UNITRENDS

Administrator Guide for Recovery Series, Recovery MAX, ION/ION+, and Unitrends Backup

10.8.6 | Document Version 1.11062024

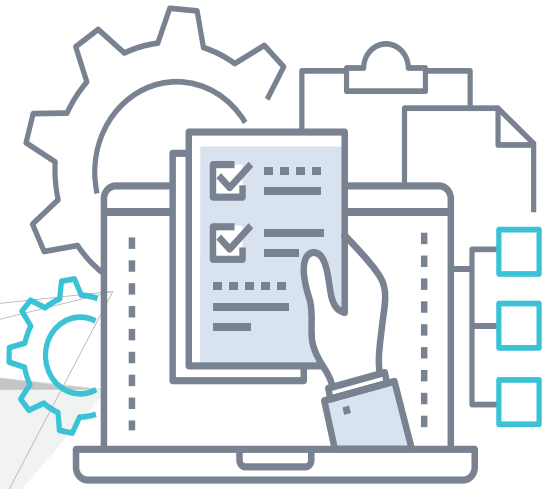

Contents

Chapter 1: Introduction

Unitrends provides comprehensive data protection and recovery solutions for any IT environment. Our innovative solutions unite protection, recoverability, and data agility delivered through an intuitive user experience.

Unitrends appliances provide enterprise-class virtual, deep virtual, physical, and unified compute protection. Use them to protect over 100 versions of servers, storage, operating systems, hypervisors, and applications, such as VMware, Hyper-V, Nutanix AHV, Citrix XenServer, NAS, SAN, Windows, Linux, SQL, SharePoint, Exchange, and Oracle.

Unitrends Recovery MAX, Recovery Series, and ION/ION+ are families of physical backup appliances for virtual and physical protection. These all-in-one solutions integrate enterprise backup and recovery software with purpose-built hardware. They can be seamlessly scaled by simply adding more appliances as your environment grows.

Unitrends Backup is an all-in-one virtual backup and recovery software-only solution. Unitrends Backup provides enterprise functionality, a common engine, and scalable data protection. This heterogeneous appliance protects systems residing on virtual, physical, and cloud-based infrastructure.

Getting started with your Unitrends appliance

Procedures in this guide assume that you have deployed your appliance and configured network settings. If you have not done these tasks, see the applicable Quick Start Guide or Deployment Guide in the following table for step-by-step instructions. Then see "Next steps" below to start protecting your environment.

Appliance type	Quick Start or Deployment Guide	Additional resources
Physical appliance lines	**Quick Start Guide**	**Site Preparation Guide (environmental specifications by appliance model)**
Gen 10 Recovery Series (models 10002–10120)	Quick Start Guide for Gen 10 Recovery Series Appliances	Site Preparation Guide for Gen 10 Recovery Series Appliances
ION and ION+ (models ION-108–IONP-832)	Quick Start Guide for ION and ION+ Appliances	Site Preparation Guide for ION and ION+ Appliances
Gen 9 Recovery Series (models 9002–9120S)	Quick Start Guide for Gen 9 Recovery Series and MAXS Appliances	Site Preparation Guide for Gen 9 Recovery Series Appliances
Gen 9 MAXS (models MAX9S-232–MAX9S-864)	Quick Start Guide for Gen 9 Recovery Series and MAXS Appliances	Site Preparation Guide for Gen 9 MAXS Appliances
Unitrends Backup virtual appliances	**Deployment Guide**	**Best practices for virtual appliances**
Unitrends Backup on VMware	Deployment Guide for Unitrends Backup on VMware	Deployment Best Practices for Unitrends Backup
Unitrends Backup on Hyper-V	Deployment Guide for Unitrends Backup on Hyper-V	Deployment Best Practices for Unitrends Backup

Appliance type	Quick Start or Deployment Guide	Additional resources
Unitrends Backup on Citrix XenServer	Deployment Guide for Unitrends Backup on Citrix XenServer	Deployment Best Practices for Unitrends Backup
Unitrends Backup on Nutanix AHV	Deployment Guide for Unitrends Backup on Nutanix AHV	Deployment Best Practices for Unitrends Backup
Unitrends Backup in Microsoft Azure	Deployment Guide for Unitrends Backup in Microsoft Azure	Deployment Best Practices for Unitrends Backup
Unitrends Backup in Amazon Web Services	Deployment Guide for Unitrends Backup in Amazon Web Services	Deployment Best Practices for Unitrends Backup

Next steps

Deployment is complete and you can get started protecting your environment:

1. Install any appliance updates (see "To install appliance updates" on page 421).
2. (Recommended) Configure encryption on the appliance (see "To configure encryption" on page 156).
3. Install the Unitrends agent on the physical machines you will protect with Unitrends backups (see "Unitrends agents" on page 361).
4. Add your assets to the appliance (see "Managing protected assets" on page 286).
5. Run backups (see "Creating backup jobs" on page 433).
6. Explore the other topics in this guide to get the most out of your Unitrends appliance!

About this Guide

This guide describes how to administer Unitrends backup and recovery solutions. Before running the procedures in this guide, you must deploy your appliance and configure network settings. For details, see "Getting started with your Unitrends appliance" on page 17.

This guide is intended for administrators and technical personnel responsible for configuring and administering Unitrends appliances, and assumes intermediate to advanced computer skills. Procedures and considerations in this guide follow best practices and requirements for the successful administration and configuration of your Unitrends backup and recovery solution.

Procedures in this guide cover supported features for Recovery Series, Recovery MAX/MAXS, and ION/ION+ appliances, and Unitrends Backup editions.

Supported features vary by appliance type and edition. To view supported features for your appliance, consult the following resources at Unitrends.com:

- Unitrends Recovery Series and Recovery Series MAX
- ION and ION+ Appliance Models

- Unitrends Backup Editions

All procedures are run from the appliance UI, unless otherwise specified. Elements in the UI are dynamic and display according to the user's role, appliance edition, environment, and type of data. For example:

- A user with the *manage* role and *backup operator* access level cannot run a recovery job. The Recover tab in the UI is disabled for this user.
- Adding backup storage does not apply to Recovery Series, Recovery MAX, and ION/ION+ appliances. The Add Storage option does not display in the Recovery Series, Recovery MAX, and ION/ION+ UIs.
- Citrix XenServer host-level backups are not supported on Unitrends Backup on Hyper-V appliances. XenServer does not display in the list of hypervisors in the Add Virtual Host dialog.

Access the UI with a Firefox or Chrome Internet browser. Internet Explorer is not supported.

Using this guide

This guide provides conceptual, procedural, and referential information for the administration of your Unitrends appliance. Unitrends recommends familiarizing yourself with this information before configuring and operating your appliance.

Feature overviews include considerations, requirements, and prerequisite information to assist in planning an effective protection strategy.

Procedures provides step-by-step instructions for performing backup and recovery operations. Instructions adhere to best practices for successful configuration and administration of your Unitrends data protection and recovery solution.

- See "Backup Administration and Procedures" on page 425 for instructions on performing, monitoring, and managing backup and backup copy jobs.
- Recovery procedures vary by backup type. For a description of each type, see "Types of data protected" on page 93. For recovery procedures, see these chapters:
  - "Recovering Host-level Backups" on page 793
  - "Recovering File-level Backups" on page 925
  - "Recovering Windows Image-level Backups" on page 1031
  - "Recovering Application Backups" on page 1147
  - "Recovering NAS Backups" on page 1121
  - "Recovering iSeries Backups" on page 1205
- Cross-references and links throughout this guide provide access to additional sources of information and assistance.

Typographical conventions

This guide uses some special typographical effects to convey certain information. Review the following for additional information:

Typographical convention	Description
Bold	Indicates one of the following: • Items you select in the UI, such as menu commands. • Text you enter in fields in the UI.
`Courier`	Indicates one of the following: • Text you enter via the command line, outside of the UI. • Output displayed by a system console, outside of the UI.
#	Sample prompt displayed before text you enter via the command-line, outside of the UI.
Greater-than symbol (**>**)	Separates sequential commands that you select or click in the UI.
Blue text	Indicates one of the following: • Link to the Unitrends website • Link to an external website • Cross reference to another section in this guide • Link to a Unitrends Knowledge Base article

Glossary of terms and acronyms

The following table describes the terms and acronyms commonly used in this document.

Term	Definition
Added disk	Applies to Unitrends Backup virtual appliances only. Virtual disk storage created on the Unitrends Backup VM's hypervisor that is added to the appliance to store backups. Also called *attached disk storage*.
Agent	Unitrends software installed on machines you wish to protect with file-level or Windows image-level backups.

Term	Definition
Agent-based asset	Physical or virtual machine that is protected by installing a Unitrends agent and running file-level or Windows image-level backups.
Appliance	The Unitrends system that backs up and recovers data. Appliance can refer to: • A physical Recovery Series, Recovery MAX, or ION/ION+ model. Consists of Unitrends hardware, Unitrends software, and additional configuration settings. • A Unitrends Backup system deployed as a virtual machine to one of the following environments: VMware, Hyper-V, Nutanix AHV, Citrix XenServer, Microsoft Azure, or Amazon Web Services. Consists of the Unitrends Backup VM, Unitrends software, attached storage, and additional configuration settings.
Application backup	Backup that captures an application's structure and data to ensure database consistency. You must install a Unitrends agent on the host asset to run application backups.
Asset	Physical and virtual machines, databases, and applications protected by the Unitrends appliance. Equivalent to the legacy term *client*. **Note:** The appliance automatically detects the virtual machines and applications on the virtual hosts and physical assets you add to the appliance.
Backup copy	Copy of a backup that is stored off-site. You can copy your backups to the following types of targets: Unitrends Cloud, a secondary Unitrends appliance, Cloud storage (managed by Amazon, AWS, Google, or Rackspace), disks, NAS devices, and other media. **IMPORTANT!** Unitrends recommends having a second copy of your backups on one of these targets in order to recover from a disaster. Backup copies that reside in the Unitrends Cloud or on a secondary Unitrends appliance are known as *hot backup copies*. Backup copies that reside on the other target types are known as cold backup copies. Equivalent to these legacy terms:

Term	Definition
	• *Replication* for hot backup copy to the Unitrends Cloud or to another Unitrends appliance. • *Archiving* for cold backup copies to other external media.
Backup group	The appliance organizes backups into groups to manage dependencies. A backup group contains a full backup and any subsequent incrementals and differentials. A backup group always starts with a full backup.
Backup mode	A backup's mode determines what data to include in the backup. Example modes: full, incremental, and differential.
Backup strategy	Combination of backups, backup copies, and other Unitrends features used to protect assets.
Backup type	Determined by the backup method used to create the backup. Examples: file-level, host-level, application.
Dashboard	A summary of the appliance's status with topic-specific tiles that capture at-a-glance data for various aspects of the appliance.
Deduplication	Specialized data compression technique that eliminates duplicate data blocks.
External storage	Applies to backup storage for Unitrends Backup appliances deployed on VMware, Hyper-V, Nutanix AHV, and Citrix XenServer only. SAN or NAS storage that is connected directly to the Unitrends Backup VM over the iSCSI, CIFS, or NFS protocol. **Notes:** • External backup storage is supported on Unitrends virtual appliances only. You cannot use external backup storage on a physical Recovery Series, Recovery MAX, or ION/ION+ appliance. • SAN or NAS storage can also be used to store backup copies. This is supported on Recovery Series, Recovery MAX, and Unitrends Backup appliances.

Term	Definition
File-level backup (formerly known as *asset-level backup*)	Backup that protects an asset's file system and operating system. You must install a Unitrends agent on the asset to run a file-level backup. (For Windows assets, you can opt to use file-level backups, image-level backups, or both backup types.)
Global menu	Toolbar across the top of the user interface that includes several menus to quickly edit global options, perform administrative tasks, and access additional resources.
Host-level backup	Backup that protects a virtual machine by leveraging hypervisor snapshots. You do not need to install a Unitrends agent on a VM to run host-level backups.
Image-level backup (Windows only)	Backup that protects a Windows asset at the disk and volume level. You must install a Unitrends agent on the asset to run an image-level backup. (You can opt to protect a Windows asset with file-level backups, image-level backups, or both backup types.)
Initial backup storage	Applies to Unitrends Backup virtual appliances only. Storage you attach to the Unitrends Backup VM that is used to store appliance configuration settings and backups. The initial backup storage must be at least 200GB in size.
Initial disk	Applies to Unitrends Backup virtual appliances only. 100GB disk used to create the Unitrends Backup VM. While deploying the virtual appliance, you select storage on the virtual host server that the installer uses to create this disk.
Instant Recovery (IR)	Process that recovers a failed or corrupted virtual machine in minutes.
iSeries backup	A backup that protects an asset's filesystem by leveraging native iSeries backup operations. You do not install an agent on the iSeries asset.
Job	Procedures performed to protect assets. Multiple job types exist, all of which can be monitored from the Active Jobs tile while in progress.
NAS backup	Backup that protects data stored on a NAS device. You do not install an agent on the NAS asset.

Term	Definition
Protected asset	Any physical machine, virtual machine, or application protected with Unitrends backups. Equivalent to the legacy term *client*.
Recovery object	Disk image created on the backup appliance during instant recovery or during file-level recovery from a host-level backup.
Recovery Point Objectives (RPOs)	Desired number of recovery points.
Recovery Time Objectives (RTOs)	Desired speed of recovery.
Replica	Virtualized copy of an asset that can immediately assume the role of that asset in case of failure. Supported for VMware virtual machines and Windows machines.
Resources	Amount of space, bandwidth, disk space, memory, etc, consumed by the job or object.
SLA Policy Automation	Unitrends feature that enables you to quickly implement a protection strategy that aligns with your business continuity plan. Simply create an SLA policy and the appliance automatically creates the backup and backup copy jobs needed for the RPO and retention settings you specified.
System load	Amount of resources being used by the system at any given time.
Tile	Topic-based sections of the dashboard.
Type	Description of both a function (backup, recover) and the storage media, such as attached virtual disk. For example, an attached disk configured as backup storage.
Unitrends Backup VM	Applies to Unitrends Backup virtual appliances only. Virtual machine created by deploying Unitrends Backup to your virtual host server.

Term	Definition
Virtual host	Host on which virtual machine assets reside. Also called a *hypervisor*.
VM replica	A virtual machine replica of a VM that is kept up-to-date by applying backups of the original VM as they run. The replica is a cold, stand-by VM that you can quickly bring online to assume the role of the original VM.
Windows replica	A virtual machine replica of a Windows asset that is kept up-to-date by applying backups of the original asset as they run. The replica is a cold, stand-by VM that you can quickly bring online to assume the role of the original Windows asset. (The Windows replica feature was formerly known as *Windows instant recovery*.)

Support for Unitrends appliances

Support is provided through the following resources:

- "Kaseya Helpdesk"
- "Contact by telephone" on page 27
- ""Show Me" interactive product tours" on page 27

Kaseya Helpdesk

The Kaseya Helpdesk is your one-stop site for help with all Kaseya product modules. Access the Kaseya Helpdesk Home page at https://helpdesk.kaseya.com/hc/en-gb.

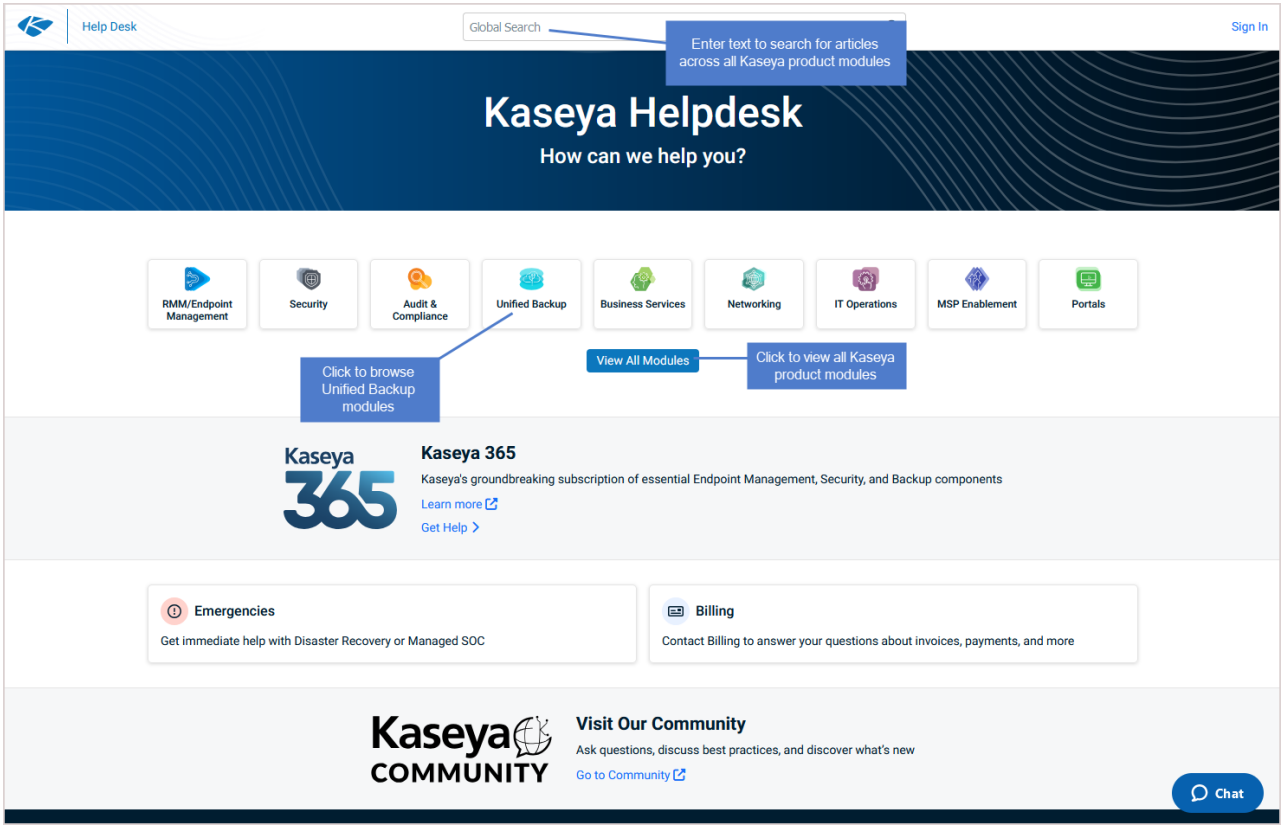

For help with your Unitrends appliance, you can go directly to the Unitrends Unified Backup section at https://helpdesk.kaseya.com/hc/en-gb#/unified_backup/unitrends_unified_backup, where you can:

- Download the Unitrends Backup virtual appliance (https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads)
- Download the latest agent releases (https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads)
- Search Knowledge Base articles – Articles in our Knowledge Base provide assistance with troubleshooting. If you encounter a problem not covered in an article, we encourage you to search the forums or post a question.
- Open a ticket with Unitrends Support.
- Chat with Unitrends Support.
- Access documentation (https://helpdesk.kaseya.com/hc/en-gb/articles/4407522416657-Documentation).

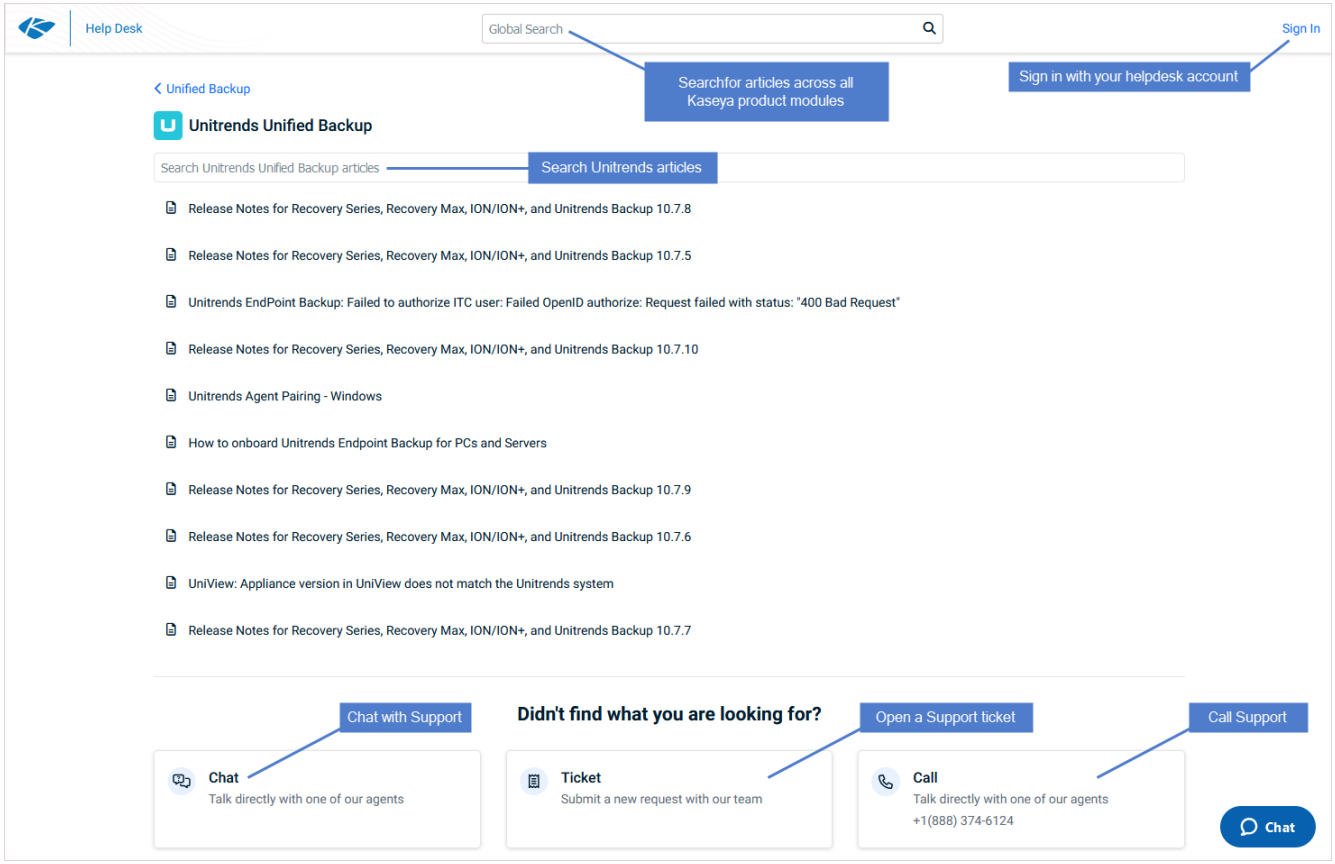

Contact by telephone

Use the following to contact Support by telephone:

- Unitrends Support North America: 1.888.374.6124
- Unitrends Support UK: +44 (0)80 8101 7687
- Unitrends Support Germany: +49 (0)89 2154822 0

You can call at any time during the hours specified in your Unitrends support service level contract. This is the recommended method for logging high priority support issues.

"Show Me" interactive product tours

We've added a new product tour experience with interactive walkthroughs of the features used to begin protecting your environment. Stay tuned for walkthroughs of additional Unitrends features in upcoming releases!

Note: To view a product tour, you must log in as a user that has the superuser, administrator, or manage role. The tours that display vary by user role.

- The product tour is launched automatically after completing the Quick Setup dialog:

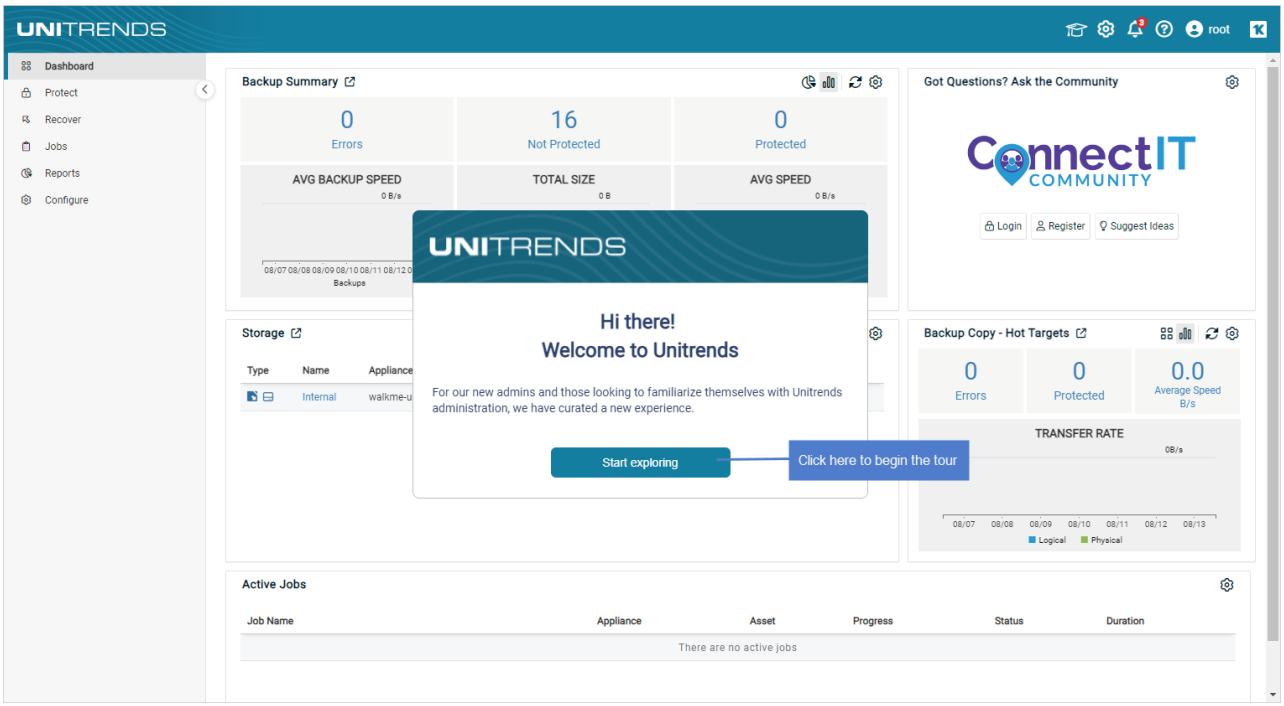

- To launch a "Show Me" tour, click and select an item under **Walkthroughs** on the **Show Me** tab:

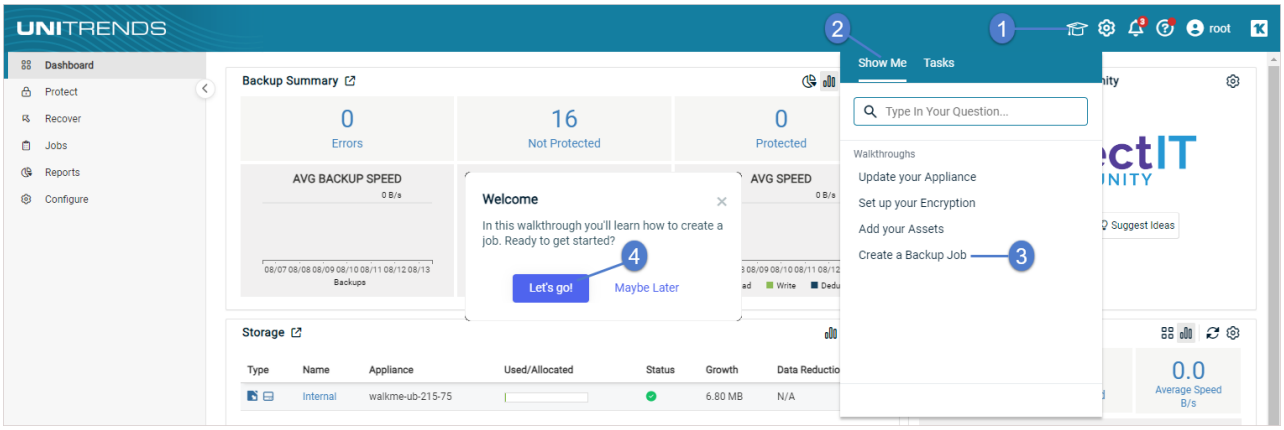

- To view product tour tasks and progress, click and select the **Tasks** tab:

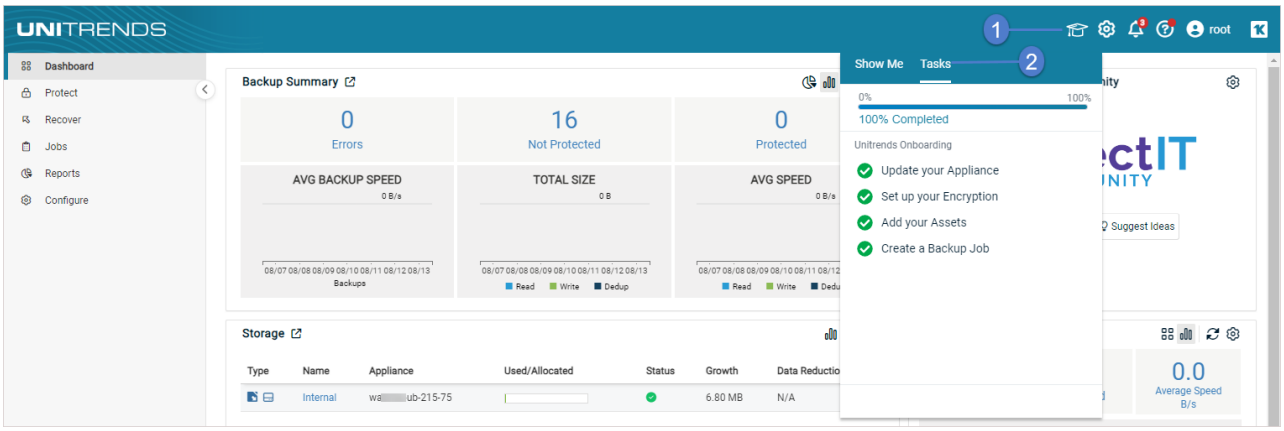

Navigating the User Interface

The user interface consists of a main dashboard and feature-specific pages and tabs for easy navigation. Only UI elements that are applicable to your environment display. For example:

- If your appliance is not receiving hot copies from another Unitrends appliance, the Copied Assets tab does not display on the Protect Page.
- If your appliance is not Enterprise Plus licensed, the Data Copy Access button and job type does not display.
- If you have not added a Hyper-V server to the appliance as a virtual host asset, the Hyper-V job type does not display.

Refer to the following topics for more on navigating the user interface:

- "Accessing the UI" on page 29
- "Global menu" on page 36
- "Dashboard" on page 39
- "Protect" on page 48
- "Recover" on page 60
- "Jobs" on page 77
- "Reports" on page 82
- "Configure" on page 84
- "Dialog help" on page 89

Accessing the UI

Use these procedures to log in to and out of the appliance UI:

- "To log in to the appliance UI"

- "To log in to the appliance UI from UniView"
- "To enable local network access if UniView is unreachable"
- "To log out of the appliance UI"

To log in to the appliance UI

1 Open a Firefox or Chrome browser and connect to your appliance by entering: ***https://<applianceIPaddress>/ui/***.

Note: If you enter ***https://<applianceIPaddress>/ui/*** and receive the message *Managed by UniView*, you must access the appliance UI from UniView. See "To log in to the appliance UI from UniView" on page 31 for details.

2 Enter the username and password of your appliance UI account. Click **Log In**.

Note: If you do not have an appliance UI account, contact your Unitrends system administrator.

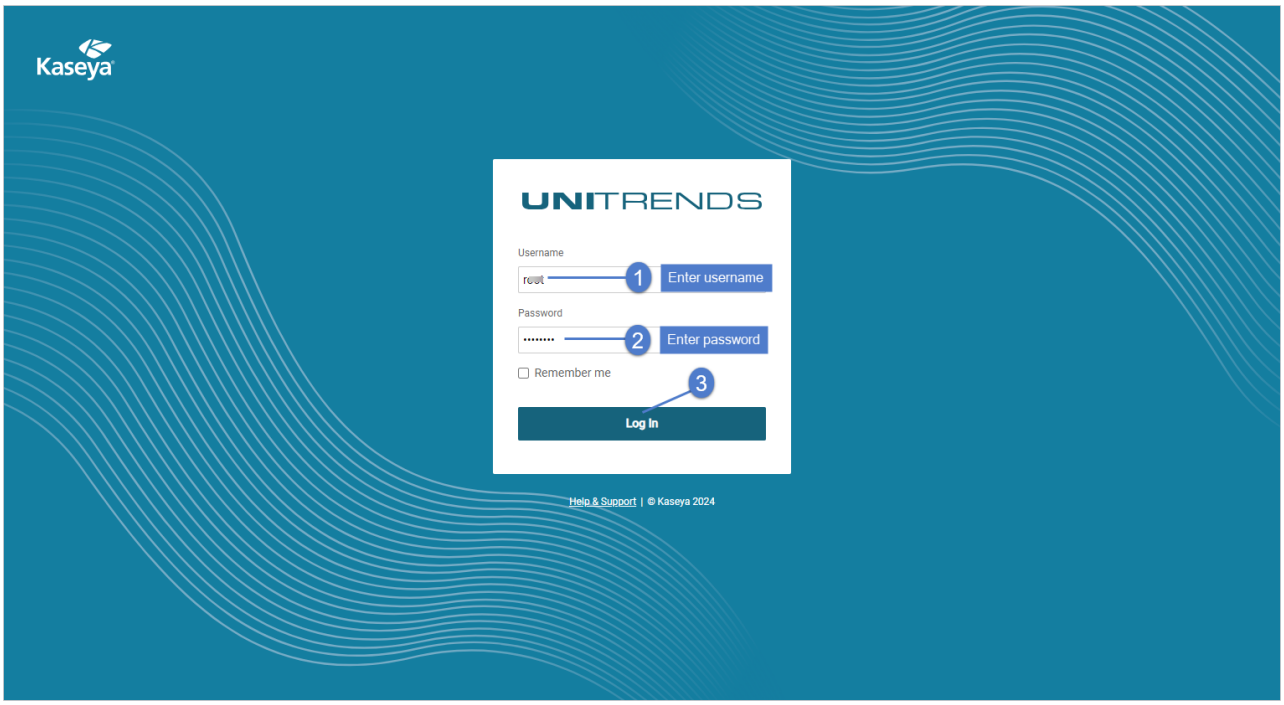

3 Upon logging in, the appliance Dashboard displays.

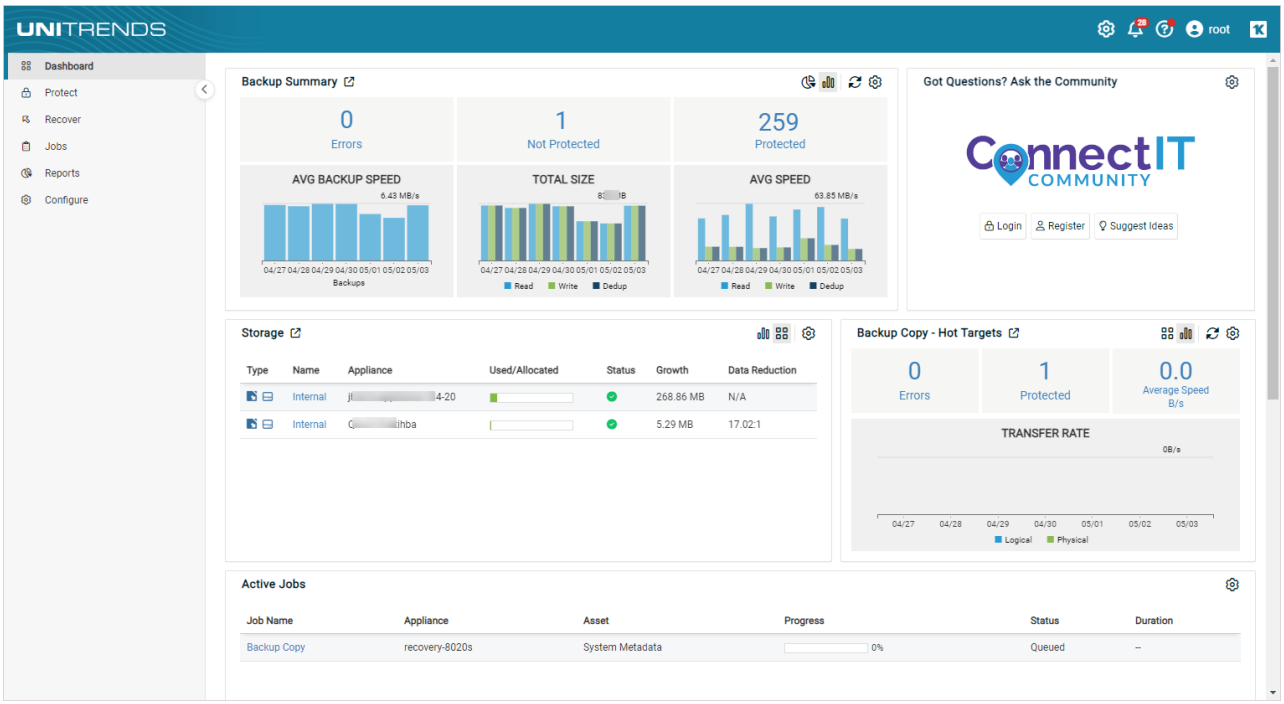

To log in to the appliance UI from UniView

For increased appliance security, the UniView Portal has a feature that blocks users from logging in directly to the appliance UI. Instead, users must connect from UniView through a proxy by using the Manage button.

If this feature has been enabled, this message displays when you attempt to log in locally:

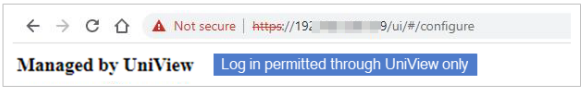

To access the appliance UI:

1 Log in to UniView.

> **Note:** If UniView is not accessible, you must access the appliance console to temporarily enable local network access. For details, see "To enable local network access if UniView is unreachable" on page 34.

2 Click **Protect** and select the **Appliances** view.

3 Locate the appliance and click its Manage button.

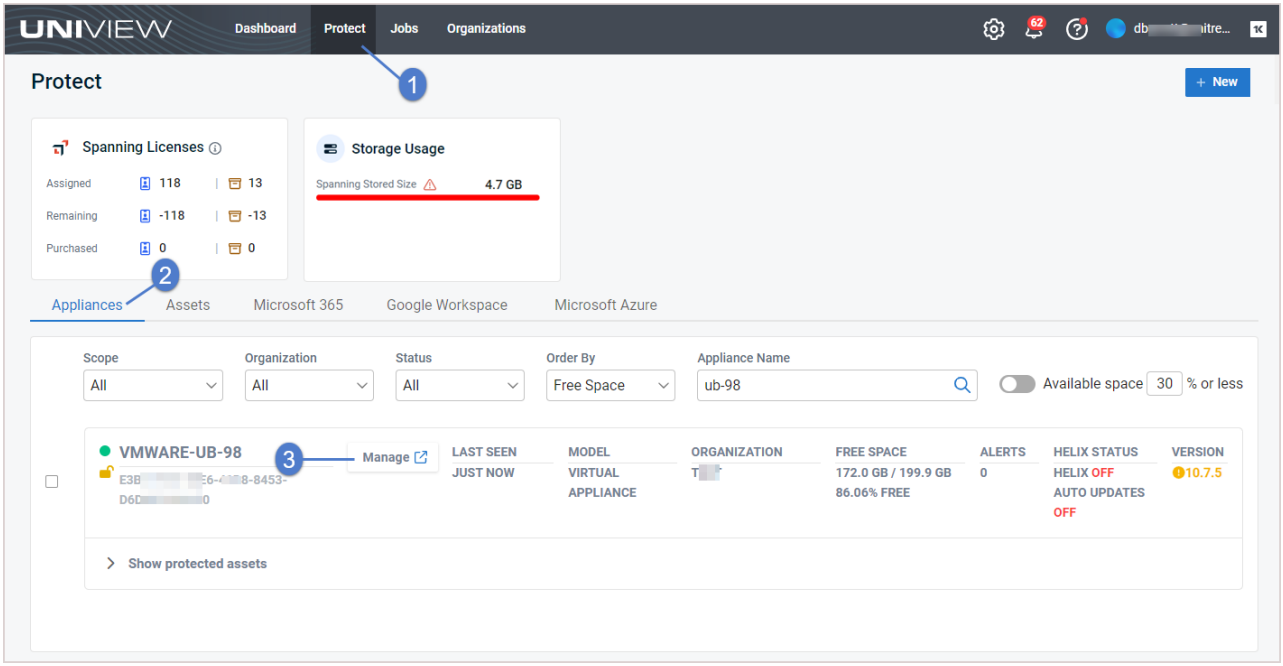

4 UniView connects to the appliance through a proxy and the appliance login page displays in a separate tab. Enter credentials and click **Log In**.

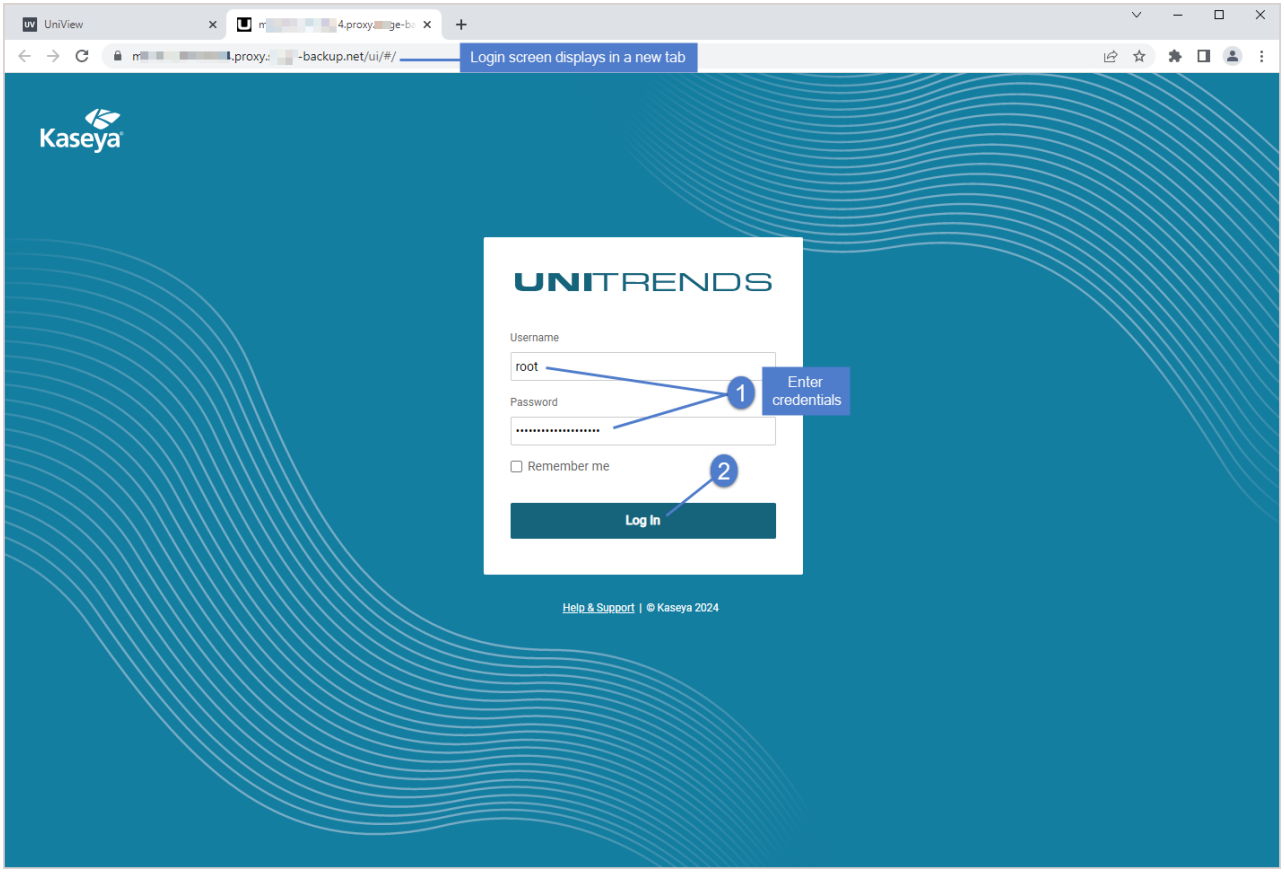

5 Upon logging in, the appliance Dashboard displays.

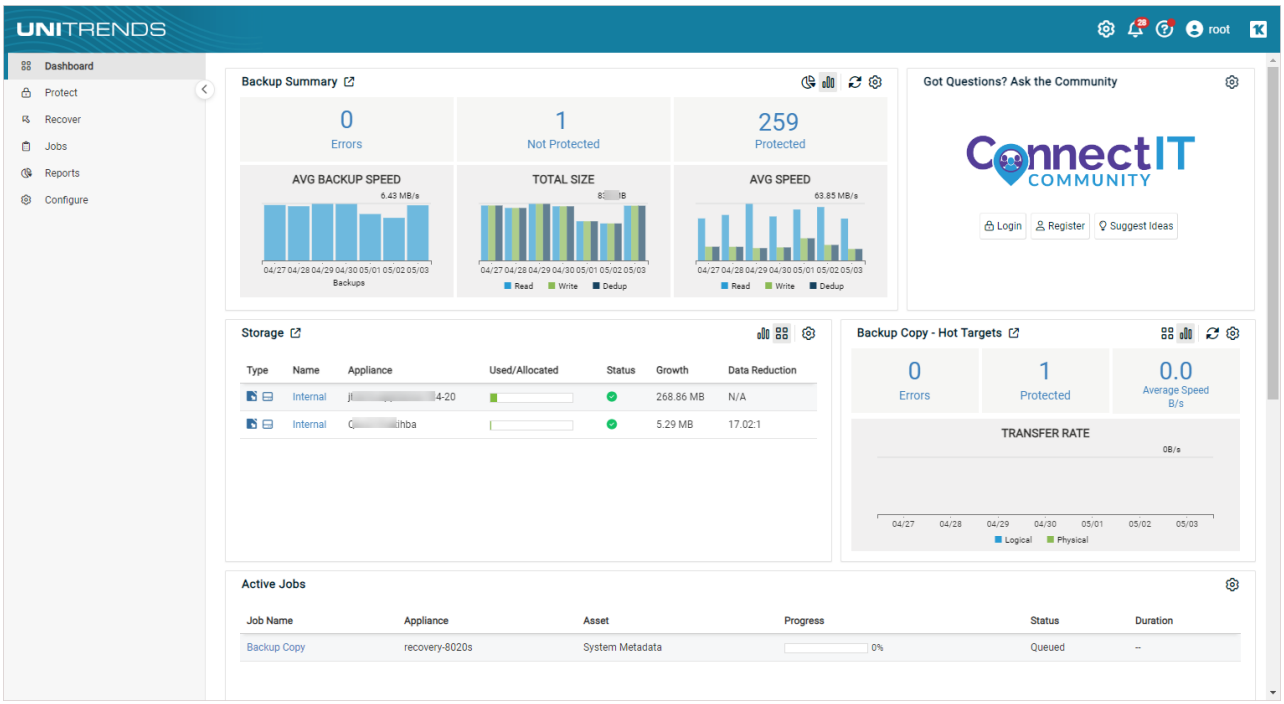

To enable local network access if UniView is unreachable

If your appliance is managed by UniView and you cannot access the UniView UI, you can temporarily enable local network access by using this procedure. When UniView is accessible, disable local network access as described in "Disable or enable local network access to an appliance" on page 170.

1 Access the appliance console by doing either of the following:

- Go to the appliance and attach a monitor, keyboard, and mouse to access the console.

OR

- If your appliance supports virtual console access with IPMI or iDRAC, access the console by using IPMI or iDRAC.

2 Enter this command in the appliance console: **dpuconfig**

You see a list of menu choices in the Unitrends Backup Console Interface. You will use the Unitrends Backup Console Interface to enable local network access.

3 On the Console Interface screen, enter **4** in the **Please enter choice** field.

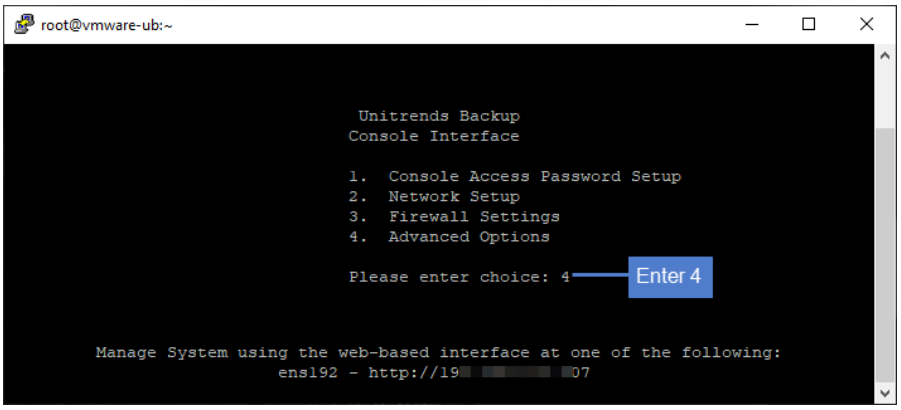

4 On the Advanced Options Menu screen, enter **2** in the **Please enter choice** field.

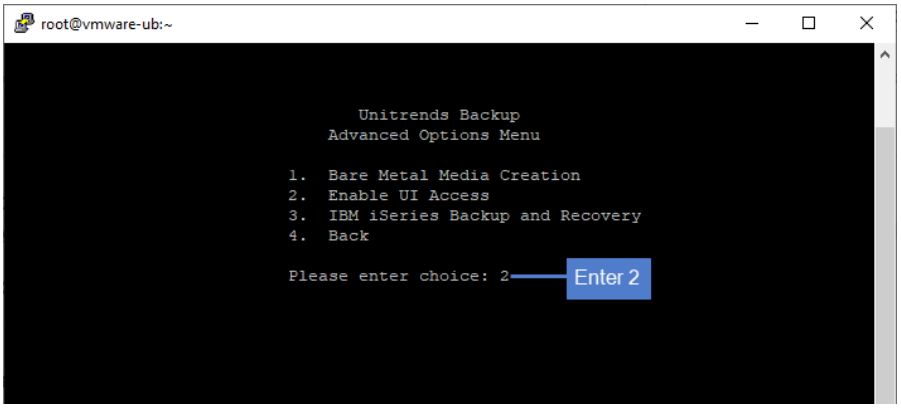

5 Local network access is enabled. Press **Enter** to exit.

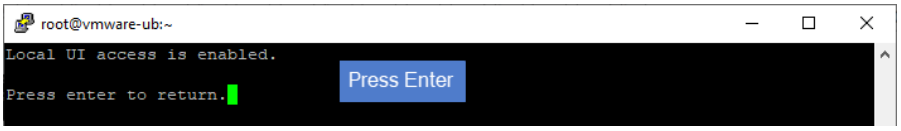

6 To exit the Advanced Options Menu, enter **4**.

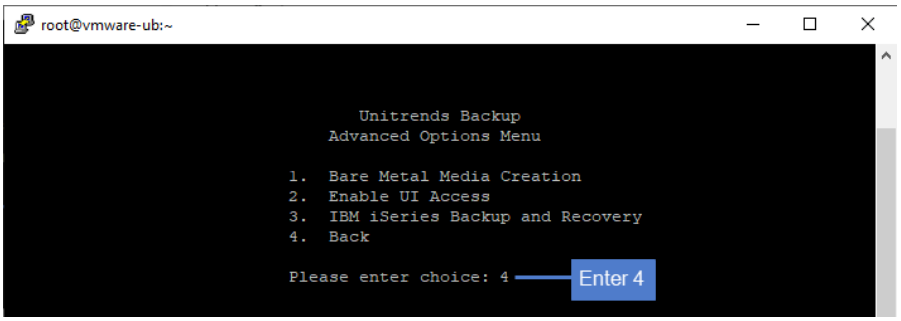

To log out of the appliance UI

1 From the Global menu, click your username and select **Logout**.

2 You are logged out of the UI and the login page displays.

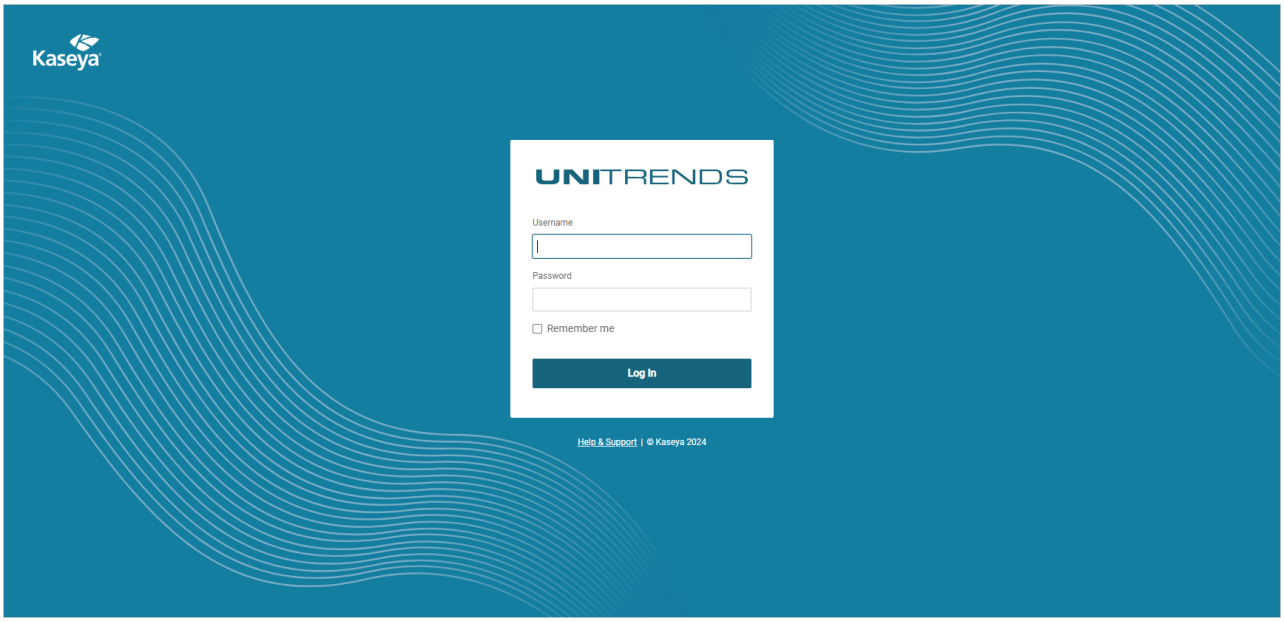

Global menu

The Global menu across the top of the user interface contains these icons to edit global options, perform administrative tasks, and access additional resources:

- "Options (Gear icon)" on page 37
- "Alerts (Bell icon)" on page 37
- "Help (question mark icon)" on page 38
- "User (Avatar icon)"
- "Kaseya (K icon)"

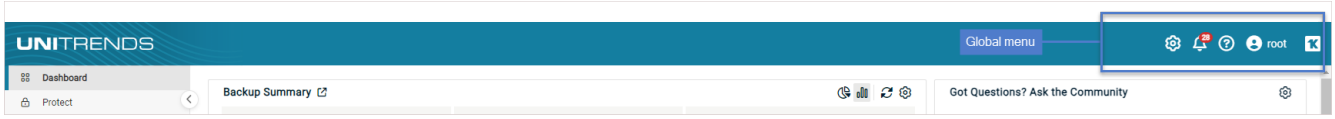

Options (Gear icon)

From the Options menu, you can select from the following options:

- Inventory Sync – Use to update the following:
  - The inventory of protected virtual machines and databases.
  - The agent pairing status of applicable agent-based assets.
- Check for updates – Use to check for appliance updates.
- Deduplication Settings – Applies to Unitrends Backup appliances only. Use to modify the appliance deduplication level.

  Note: Ransomware detection functionality requires Level 3.

- Set Language – Use to select the language for UI text.

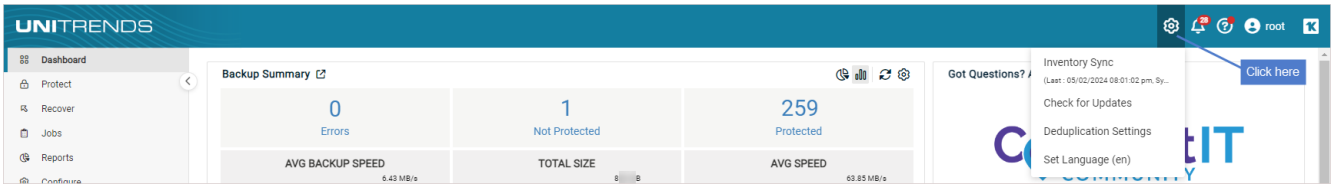

Alerts (Bell icon)

Click the Alerts icon to view the current list of alerts for the appliance. Alerts include appliance errors, warnings, and notifications. Colored icons indicate the severity level of each alert.

Clicking an alert opens the alert details. Clicking on **View More Alerts** opens the Alerts report. An alert is automatically removed once the condition has been resolved. You can manually remove all alerts by clicking the garbage can icon or delete one alert by selecting it and clicking **Dismiss Alert** in the details box.

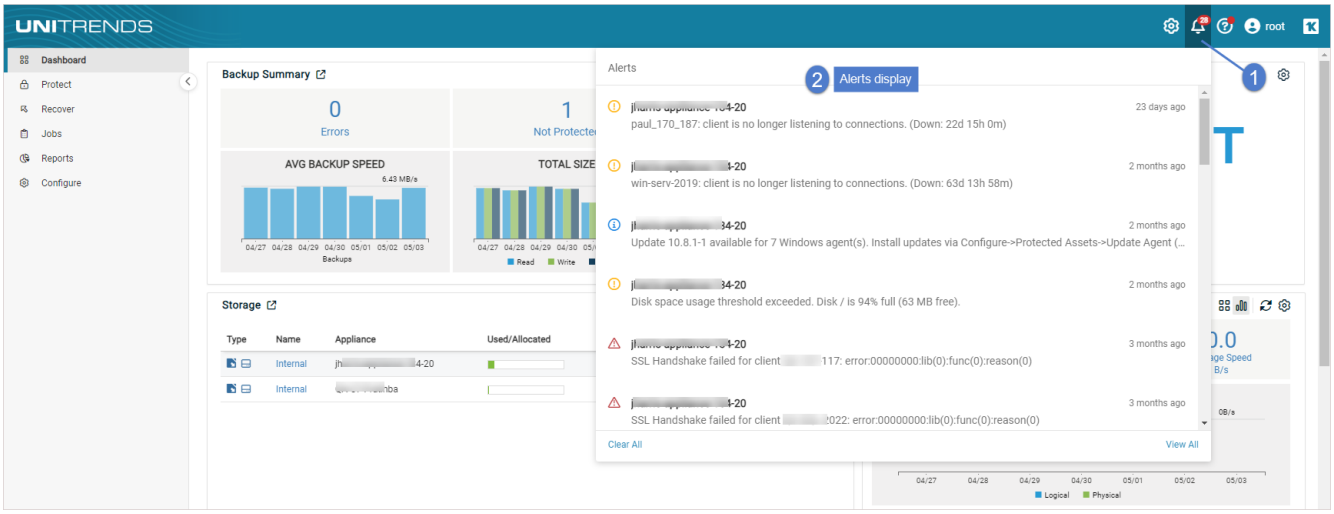

Potential ransomware infection is indicated by this icon:

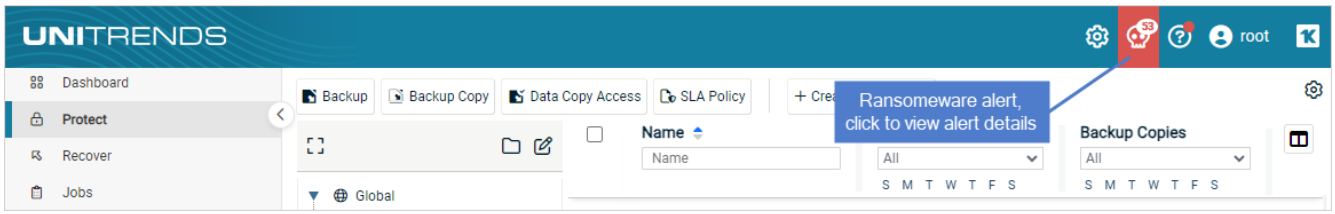

For further information on ransomware alerts, see "Ransomware detection" on page 102.

Help (question mark icon)

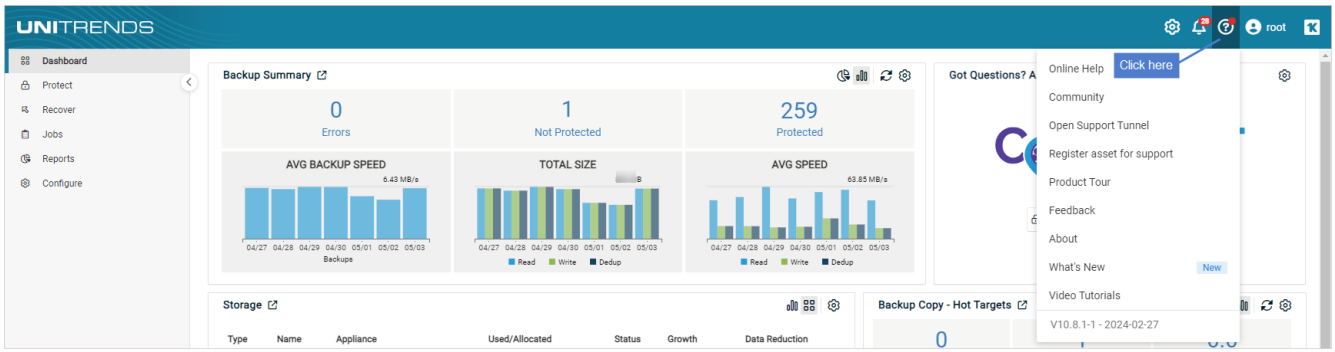

From the Help drop-down menu, you can select from the following options:

- Online Help – Displays the online help for your Unitrends appliance.
- Community - Select to access Unitrends self-help communities.
- Open Support Tunnel – Select to open a support tunnel while working with Unitrends Support. Select Close Support Tunnel when you are through working with the Support Engineer.
- Register Asset for Support – Select to register the appliance for Unitrends Support services.
- Feedback – Select to send product feedback and enhancement requests to Unitrends.
- About – Select to view appliance software, browser, and hardware information, such as appliance name, IP address, version, processor type, memory, and asset tag.
- What's New – Select to view a description of the features and fixes added in the recent releases.
- Video Tutorials – Select to access product video tutorials.

User (Avatar icon)

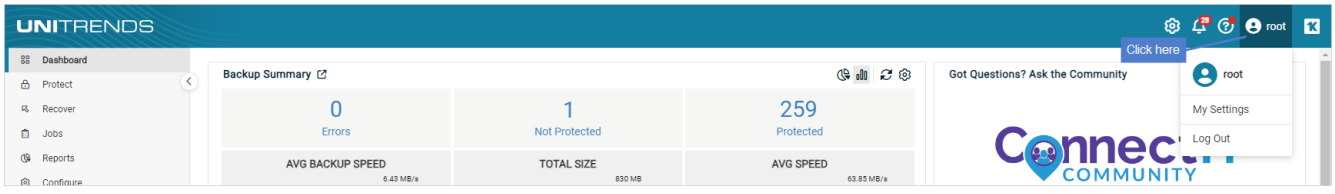

From the User (Avatar icon) menu, you can select the following options:

- User Account – Displays the name of the user that is currently logged in to the UI. The default UI user account is *root*.
- My Settings – Use to view and edit your user account details.
- Logout – Use to log out of the appliance.

Kaseya (K icon)

Click the Kaseya icon to access KaseyaOne.

Dashboard

The Dashboard provides a high-level overview of your Unitrends environment from a single pane of glass. It displays protection status for assets, and summaries for all backup storage and active jobs. In addition, you can access community forums to ask questions or get information.

The Dashboard tiles summarize information for the appliance you are logged in to, as well any managed appliances. For example, if you are logged in to an appliance that is managing two others, Backup Summary counts include jobs and assets on all three appliances.

While working with the Dashboard, you can customize the display:

- To change the layout of the dashboard, click and hold the upper region of a tile, drag it to the desired location, and release.
- – Click to run a report.
- – Click to switch to pie chart view.
- – Click to switch to bar graph view.
- – Click to switch to table view.
- – The Dashboard tiles update hourly. To update the tiles at any time, click in the top right corner of a tile. (Clicking refresh on the Backup Summary tile updates all other tiles.)
- – Click to display the tile's preferences dialog. Preferences vary by tile. For example, in the Backup Summary tile's preferences dialog you can: check or clear boxes to display or hide tiles, or click **Reset** to apply the default layout for the current release. Click **Close** to exit the dialog. In the Active Jobs tile's preferences dialog you can also set the rate at which data is refreshed.

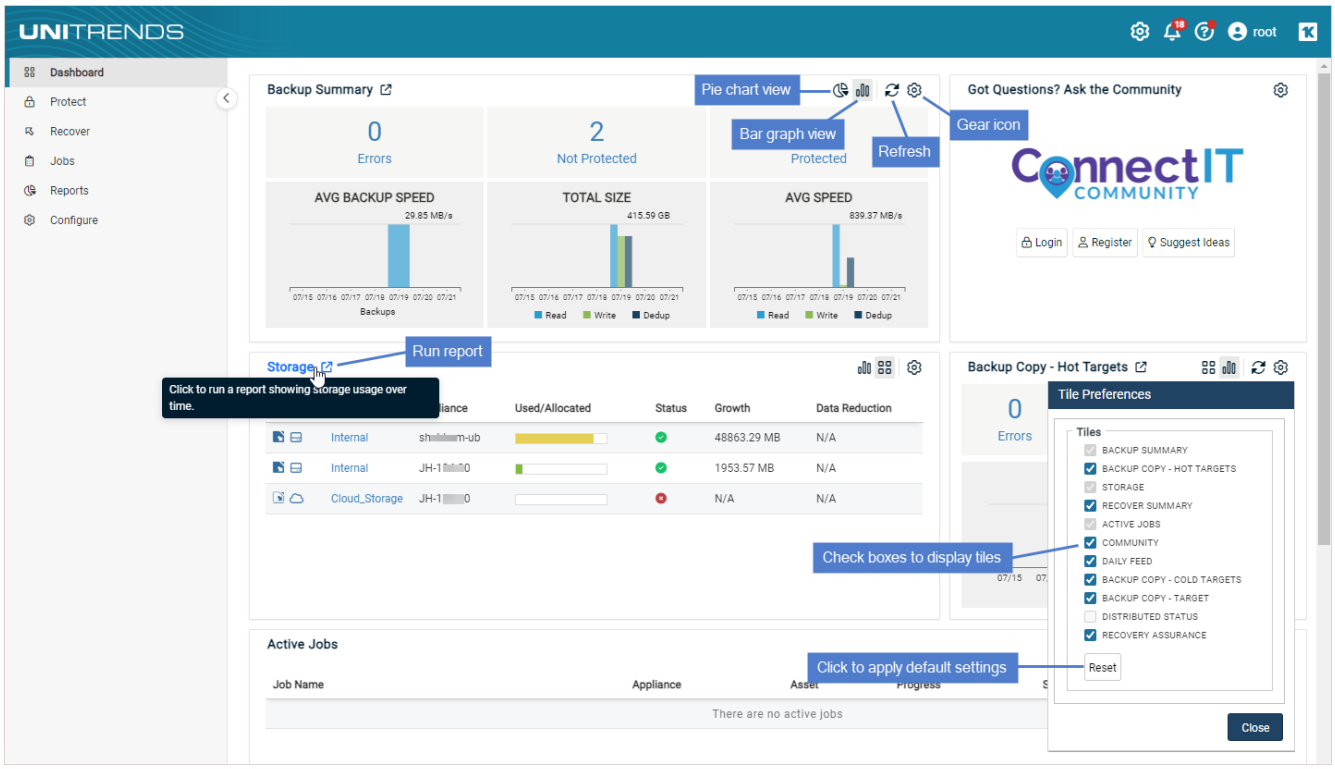

See these topics for descriptions of each tile:

- "Backup Summary tile" on page 40
- "Got Questions? Ask the Community tile" on page 42
- "Storage tile" on page 42
- "Active Jobs tile" on page 44
- "Backup Copy – Hot Targets tile" on page 44
- "Backup Copy – Cold Targets tile" on page 45
- "Backup Copy – Target tile" on page 46
- "Recover Summary tile" on page 47
- "Daily Feed tile" on page 47
- "Recovery Assurance tile" on page 48
- "Dashboard" on page 39

Backup Summary tile

This tile provides the following for backups that ran over the last seven days across all managed appliances:

Note: Unitrends appliances, virtual hosts, and VM templates are not included in the number of protected and unprotected assets.

- Bar graph view –
  - Errors – Displays all jobs, including canceled jobs, that ended in error within the last seven days. Click this number to open the "Backup Failures report" on page 1322. Bar graph shows average backup speed for each day.
  - Not protected – Displays the number of assets without a successful backup over the last seven days. Click this number to open the "Protection Summary report" on page 1331. Bar graph shows the amount of data read, written, and deduplicated each day.
  - Protected – Displays the number of assets with a valid backup over the last seven days. Click this number to open the "Protection Summary report" on page 1331. Bar graph shows average read, write, and deduplication speed for each day.

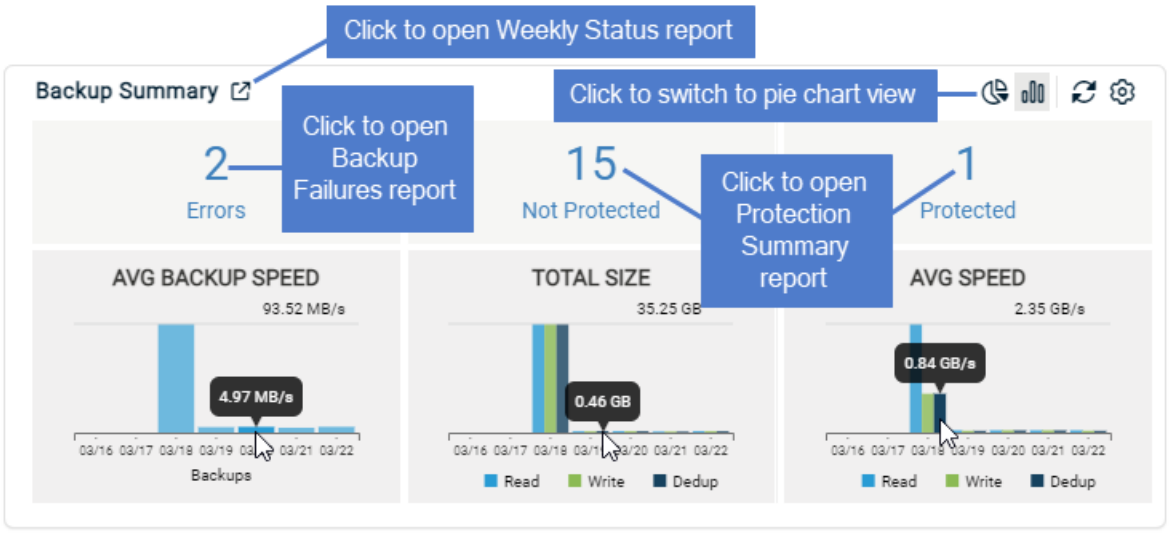

- Pie chart view –
  - Last Backup Status – Displays a chart of last backup status for all protected assets.
  - Avg Backup Speed – Average backup speed of all jobs that ran on a given day.
  - Protected VS Unprotected – Displays the number of assets with a valid backup over the last seven days (protected) versus the number of assets without a valid backup over the last seven days (unprotected).

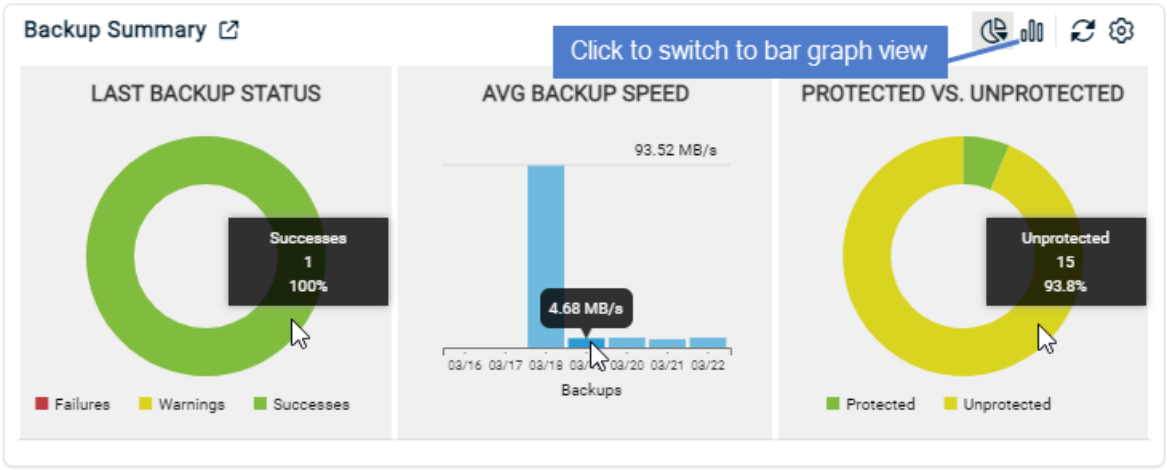

Got Questions? Ask the Community tile

From this tile, you can:

- Click **Login** to log in to the ConnectIT Community.
- Click **Register** to create a new ConnectIT Community account.
- Click **Suggest ideas** to provide feedback.

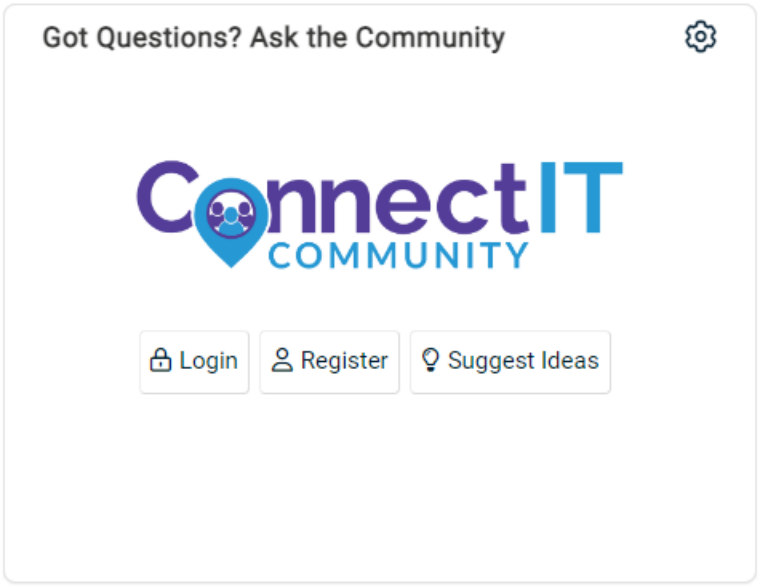

Storage tile

This tile shows these details about available and used storage for all managed appliances:

Field Name	Description
Type	Hover over an icon to display the storage type.
Name	Name of the storage.
Appliance	Name of the appliance associated with this storage.
Used/Allocated	The amount of available and used storage. Hover to display the amount of space used versus the amount still available.
Status	Hover over the icon to display the current status of the storage.
Growth	Daily average percent change (increase/decrease) in the backup data store.
Data Reduction	Data reduction ratio (backup storage/bytes written).

Table view example:

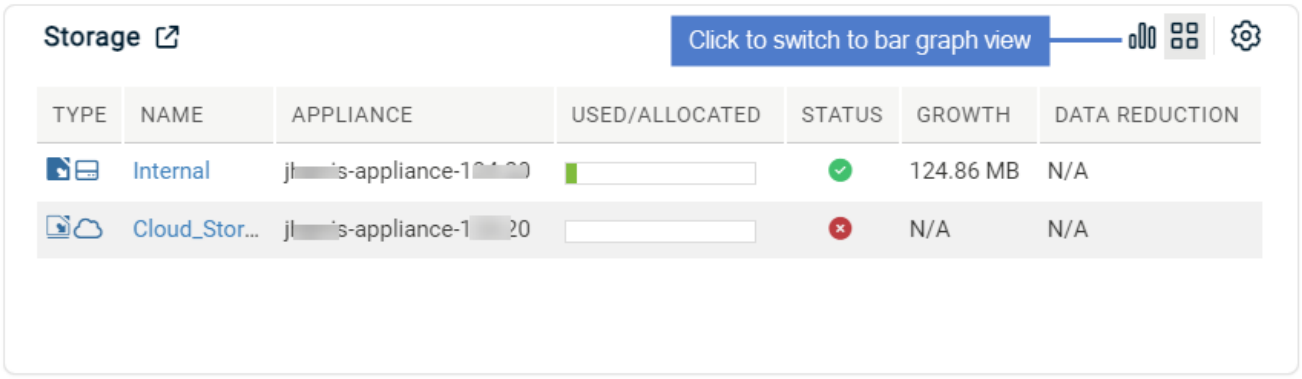

Bar graph view example:

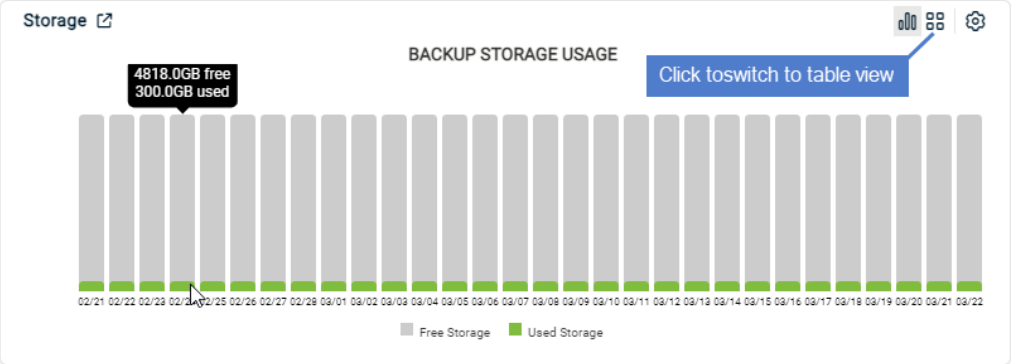

Active Jobs tile

This tile displays an at-a-glance view of all jobs currently running on all managed appliances. This view includes the job name, the appliance name, the asset being protected, a progress bar, the percent of the job completed, the current status, and the length of time the job has been running.

To filter the list, click one of the following in the lower portion of the tile: All, Backup, Recover, Backup Copy, or Data Copy Access.

When a job displays in the tile, you can click the job name to view its progress at **Jobs > Active Jobs**. The job is highlighted and details display. After a job completes, it no longer displays in the Active Jobs tile.

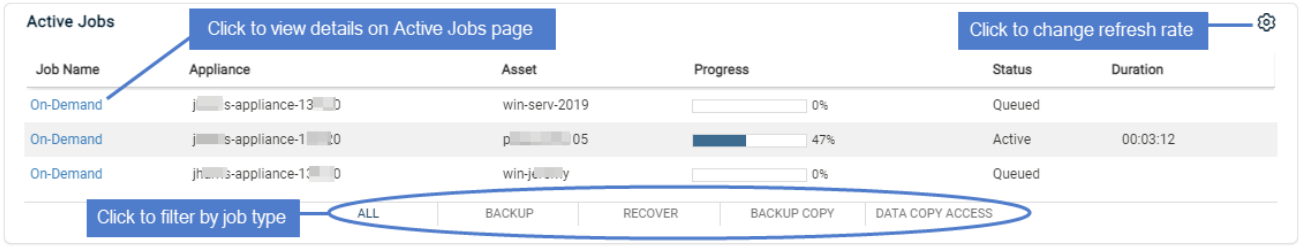

Backup Copy - Hot Targets tile

Displays data about backups that were copied from this appliance to hot backup copy targets within the last seven days. Hot copy targets include the Unitrends Cloud Unitrends appliance targets that have been added to this appliance.

This tile displays the number of errors, protected assets, and unprotected assets for all hot copy targets. A bar graph indicates the day's average backup copy transfer rate. To view details by target, click in the lower left corner. Click to return to the summary view.

- Errors – Displays the number of hot backup copy jobs that had errors within the last seven days. This includes canceled jobs that ended in error. Click this number to open the "Backup Copy - Hot Targets report" on page 1337.
- Protected – Displays the number of assets that have successful backup copies that ran within the last seven days. Click this number to open the "Protection Summary report" on page 1331.

- Average speed GB/s – Displays the average backup copy speed of all jobs that ran in last seven days. Click this number to open the "Backup Copy - Hot Targets report" on page 1337.
- Transfer Rate - Displays the average data transfer rate of backup copy jobs, by day. Only new, unique blocks are transferred to the target. Physical shows the rate for data blocks that were transferred. Logical shows the rate for logical data transferred (not actual blocks sent). Hover over a bar in the graph to display the day's average rate in GB/s.
- Bar graph view includes data for all hot targets receiving copies from this appliance:

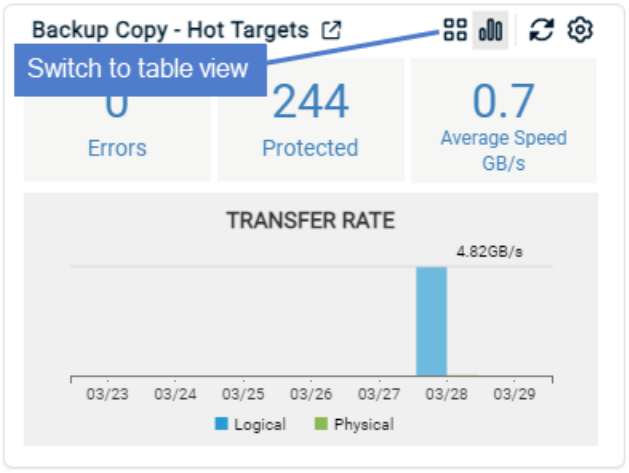

- Table view shows data by target:

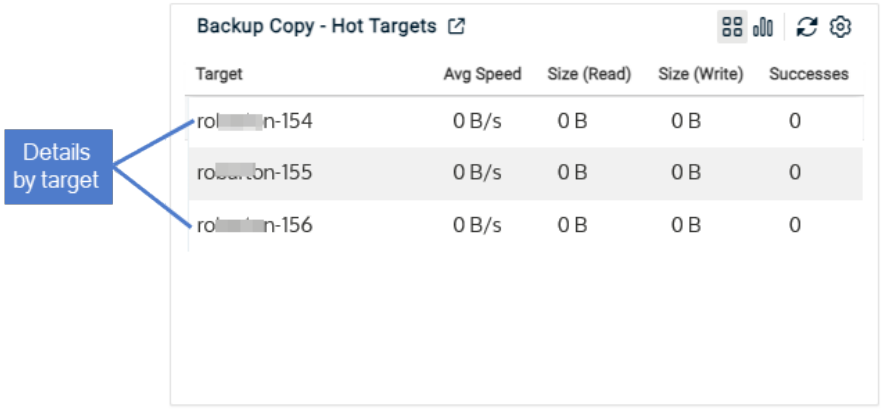

Backup Copy - Cold Targets tile

Displays data about backups that were copied to a cold backup copy target (eSATA, USB, tape, third-party cloud, attached disk, NAS, and SAN) within the last seven days. This tile displays the number of errors, protected assets, and average speed of backup copies to all cold backup copy targets. The performance graphs and average speed are calculated based on completed backup copies from appliances running version 9.0 or higher. This tile does not display data from jobs in progress or data from sources running Unitrends release 8.2 or earlier.

- Errors – Displays the number of backup copy jobs that failed over the last seven days.

- Protected – Displays the total number of protected assets and source appliances (each source appliance adds one to the count).
- Average speed B/s – Displays the average speed of completed backup copy jobs to this target from all sources.

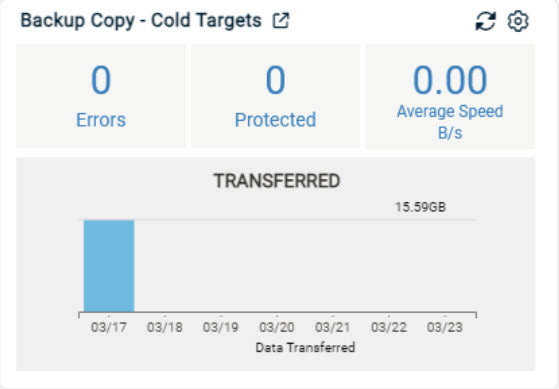

Backup Copy - Target tile

Displays data about backups that were copied to the appliance within the last seven days. Applicable only if the appliance is a hot backup copy target that is receiving backup copies from another Unitrends appliance. This tile displays the number of errors, protected assets, and average backup copy speed. Hover over a bar in the graph to display the amount of data transferred to the appliance on a given day.

- Errors – Displays the number of backup copy jobs that failed over the last seven days. Click this number to open the "Backup Copy - Hot Targets report" on page 1337.
- Protected – Displays the total number of protected assets and source appliances (each source appliance adds one to the count). Click this number to open the "Protection Summary report" on page 1331.
- Average speed MB/s – Displays the average speed of completed backup copy jobs to this target from all sources. Click this number to open the "Backup Copy - Hot Targets report" on page 1337.

Bar graph view example:

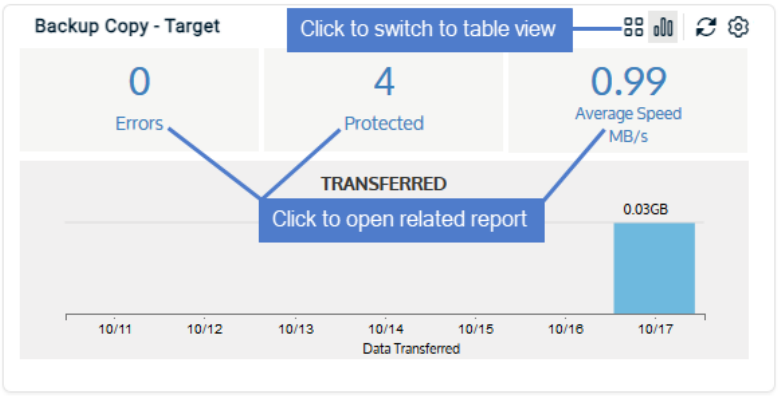

Table view example:

Recover Summary tile

This tile displays details about recent recovery jobs on all managed appliances. Hover over a bar in the graph to display the average speed of the day's recovery jobs.

- Recent restores – Displays the number of recovery jobs in the last seven days. Click this number to open the "Recovery History report" on page 1327.
- Active FLR – Displays the number of currently active file-level recovery objects. Click this number to open the File Level Recovery tab on the Recover page.
- Active IR – Displays the number of currently active instant recovery objects. Click this number to open the Instant Recovery tab on the Recover page.
- Active Replicas – Displays the number of currently active Windows and VM replicas. Click this number to open the Replicas tab on the Recover page.
- Avg speed – Displays the average speed of all recovery jobs that ran within the last seven days (current day not included). Hover over a bar in the graph to the display the average speed for a given day.

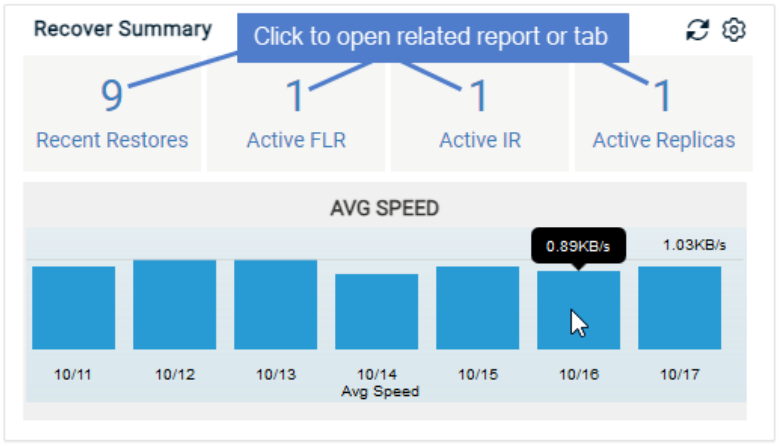

Daily Feed tile

The Daily Feed tile displays recent Tweets from Unitrends about our products and services.

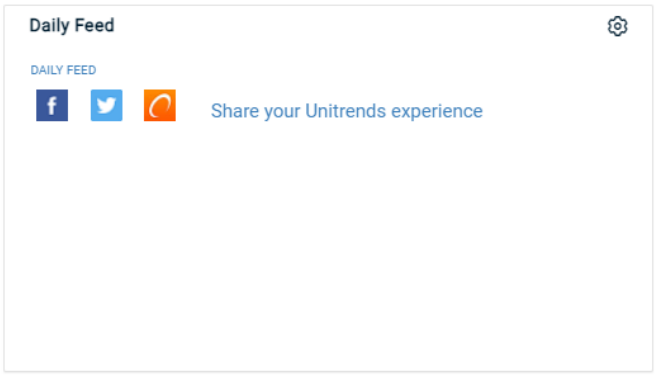

Recovery Assurance tile

This tile displays information about Data Copy Access jobs.

- RPO/RTO compliance – Displays the total number of jobs that meet RPO/RTO, the total number of jobs that fail RPO, and the total number of jobs that fail RTO. Click any of these numbers to open the "Compliance report" on page 1370.
- Recent Tests (7 Days) – Displays the total number of jobs that were successful, that completed with warnings, and that failed, over the past seven days. Click any of these numbers to open the "Recovery Assurance report" on page 1329.

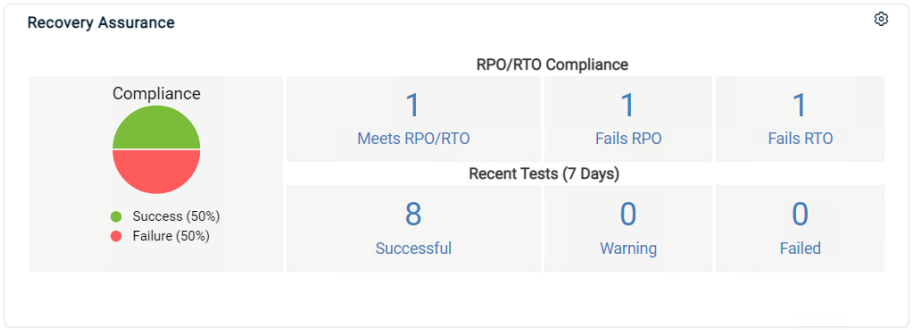

Protect

The Protect page provides status information about assets and copied assets, and enables you to create jobs to protect these assets.

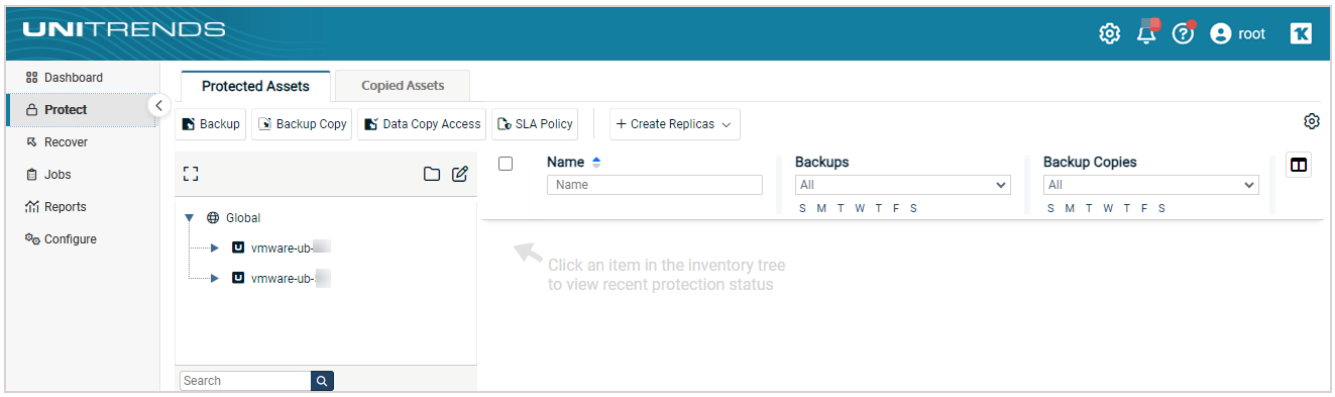

The Protect page contains the following tabs:

- "Protected Assets tab"
- "Copied Assets tab" on page 56

Notes:

Only UI elements that are applicable to your environment display. For example, if your appliance is not receiving hot copies from another Unitrends appliance, the Copied Assets tab does not display and protected assets display on the main page (you do not see a Protected Assets tab).

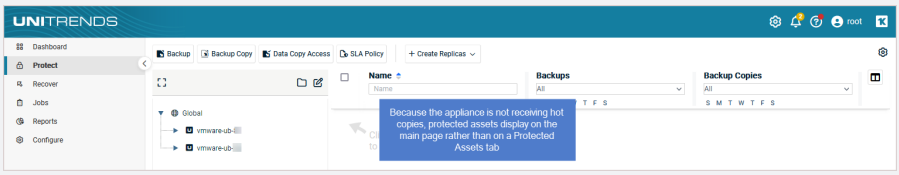

Protected Assets tab

This tab enables you to view your entire inventory of protected assets and see the status of backups and backup copies that ran over the last seven days. For more information about data protection, see the "Protection Overview" on page 91 chapter.

See the following for details on working with the Protected Assets tab:

- "Buttons" on page 49
- "Inventory tree" on page 51
- "Status table" on page 52

Buttons

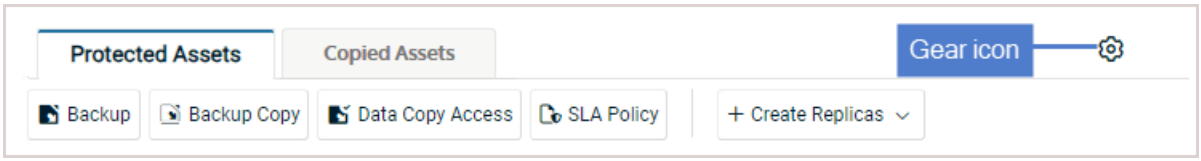

The buttons that display on the Protect Assets tab vary by environment. If you do not see a button, the feature is not applicable to your appliance. The following buttons may display:

- Backup – Opens the Create Backup Job dialog. For details on using this dialog to create a backup job, see "Creating backup jobs" on page 433.
- Backup Copy – Opens the Create Backup Copy Job dialog. For details on using this dialog to create a backup copy job, see "Creating backup copy jobs" on page 491.
- Data Copy Access – Opens the Create Data Copy Access Job dialog. For details on using this dialog to create a data copy access job, see "Recovery assurance procedures" on page 1270.
- SLA Policy – Opens the Create SLA Policy dialog. With SLA policies, the appliance automatically creates the backup and backup copy jobs needed for the assets, RPO, and retention settings specified in the policy. For details, see "Creating SLA policies" on page 536.
- Create Replicas – Click and select **Windows** to open the Create Windows Replica dialog, or click and select **VMware** to open the Create Replica VMs dialog. For details on creating replicas, see "Windows file-level replicas" on page 993 and "VM replicas" on page 876.
- Gear icon – Click to adjust the refresh rate of the Protect page.

Note: Before running jobs, it is recommended to review "Preparing for backups" on page 425 and "About creating backup and backup copy jobs" on page 426 to determine how you will implement your protection strategy.

Inventory tree

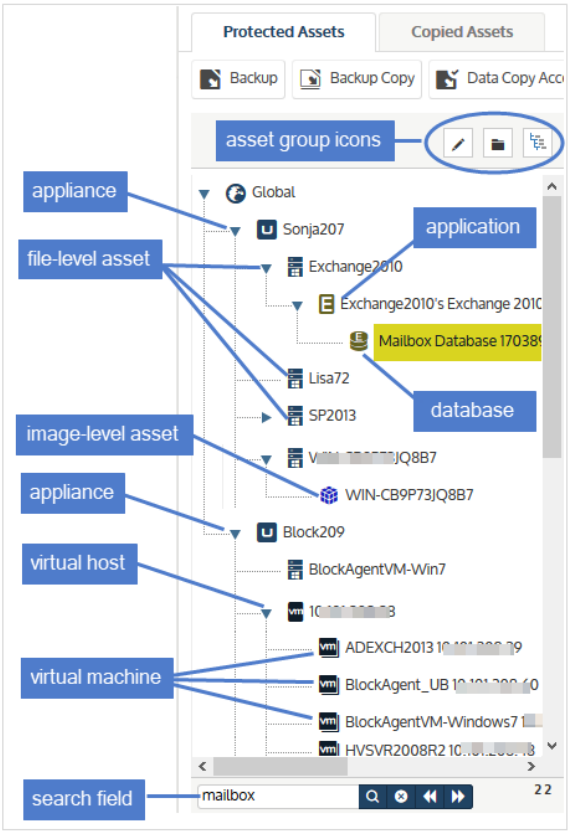

All managed appliances and their protected assets display in a tree view:

- Appliances display as top-level nodes.
- Agent-based physical and virtual assets, and virtual hosts display as second-level sub-nodes.
- Hosted VMs and applications display as sub-nodes under their host asset.

To customize the inventory tree display, you can group assets in custom folders and assign users to the groups you create. See "Grouping assets in custom folders" on page 348 for details.

Use these options while working with the inventory tree:

- To look for an asset by name, use the **Search** field below.
- To view asset groups, click the **Show Groups** icon located above the tree.
- To add, remove, or edit asset groups, click the **Manage Groups** pencil icon located above the tree. The Manage Groups icon displays only in Show Groups mode.
- To hide asset groups, click the **Hide Groups** icon located above the tree.
- To view all assets, click the **Expand All** icon located above the tree.

Status table

Selecting an appliance or asset in the inventory tree populates the status table with details about backups and backup copies that ran within the last seven days for the applicable assets. Job status icons display in the Backup and Backup Copies columns. Click a status icon to view more details. If the Backup or Backup Copies icon is gray, no job ran for the asset on that day. See these examples for details:

- "Status by appliance"
- "Status by host asset" on page 52
- "Status by single asset" on page 53
- "Filtering the status table results" on page 54
- "Viewing job details" on page 54

Status by appliance

Select an appliance to view backup and copy status for all its protected assets:

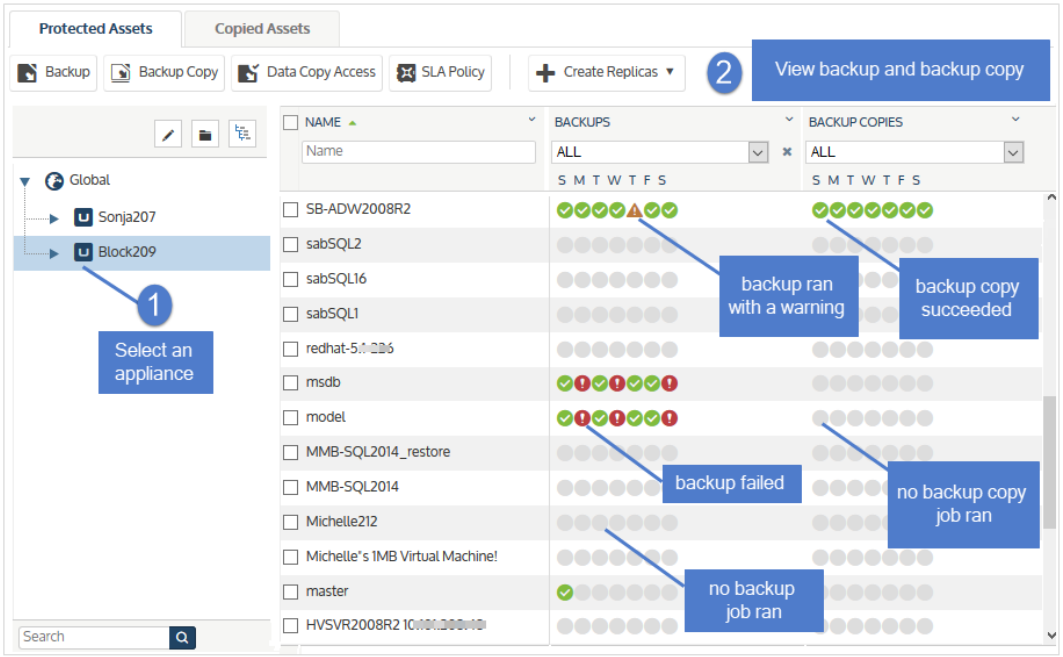

Status by host asset

Select a virtual host or application to view backup and copy status for its hosted VMs or instances:

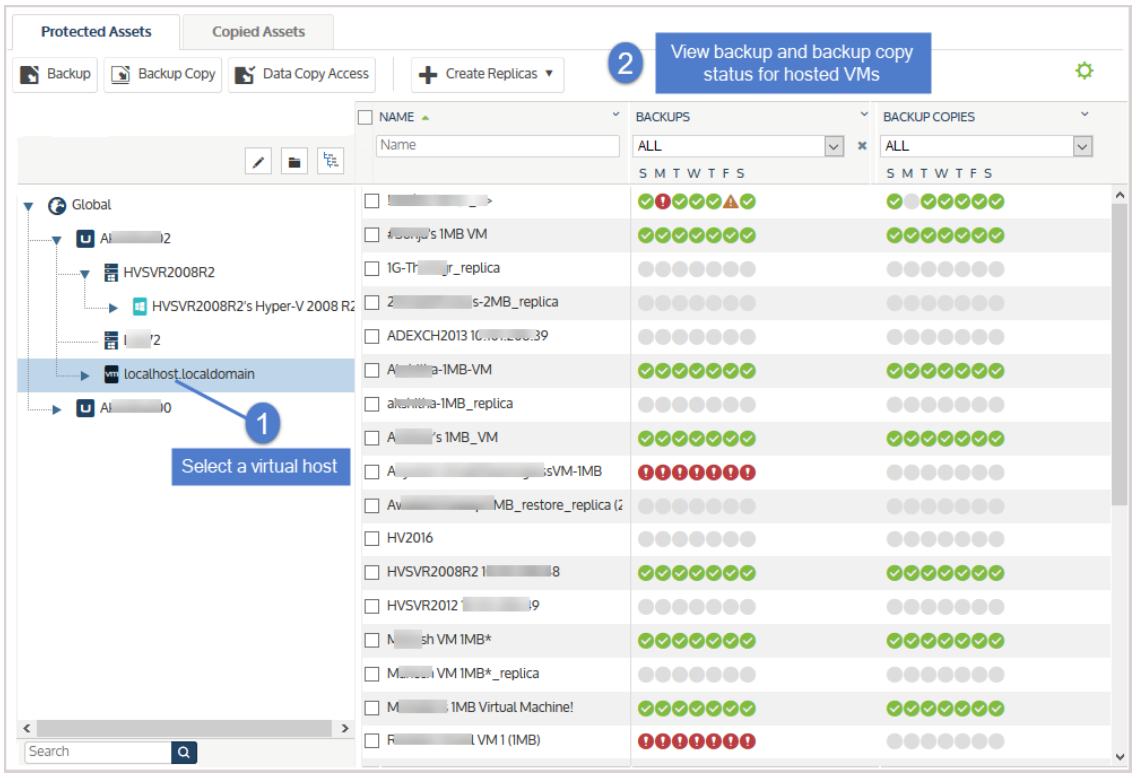

Status by single asset

Select a single agent-based asset or virtual machine to view its backup and copy status:

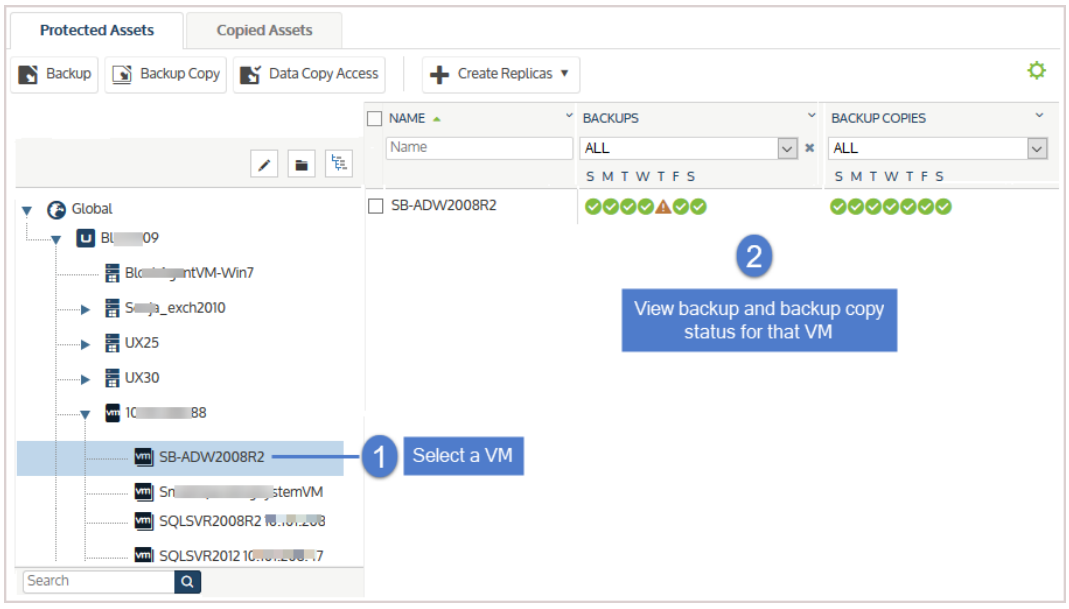

Filtering the status table results

After selecting an appliance or asset in the inventory tree, you can filter the display results by using these fields:

- Name field - Enter text to display only asset names that contain the string you entered.
- Backup drop-down - Select an item from this list to display only assets with backups that meet the condition you selected. For example, select *Failed in the last 7 days* to see only assets with backups that have failed in the last 7-day period, or *No successes in the last 7 days* to see only the assets that have not had a successful backup in the last 7-day period (this includes cases where no backup jobs have run and cases where backups have run but none were successful).
- Backup Copies drop-down - Select an item from this list to display only assets with backup copies that meet the condition you selected. For example, select *Failed in the last 7 days* to see only the assets with backup copies that have failed in the last 7-day period, or *No successes in the last 7 days* to see only the assets that have not had a successful backup copy in the last 7-day period (this includes cases where no backup copy jobs have run and cases where backup copies have run but none were successful).

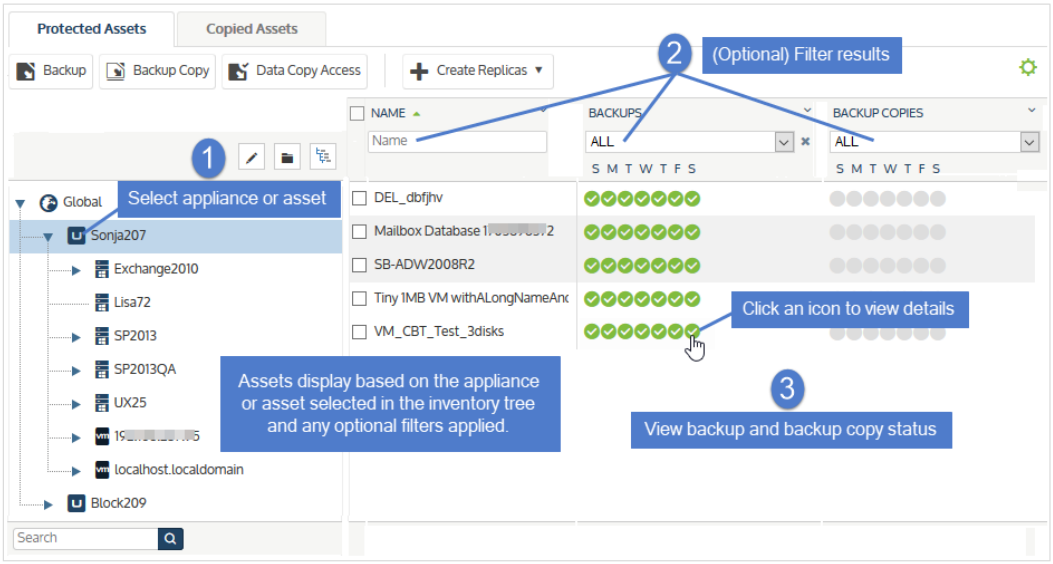

Viewing job details

To view details of an asset's backups or copies that ran on a given day:

1. Click a status icon for an overview of the backup or copy jobs that ran, by status.
2. Click **Details** to view the Backup Status dialog.

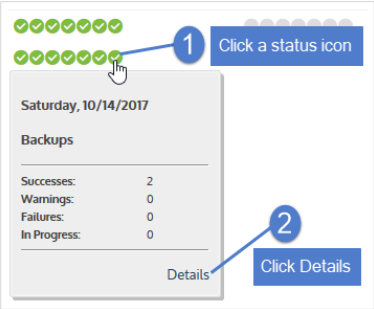

3 (Optional) Do either of the following:

- Click **Delete Failed Backups** or **Delete Failed Backup Copies** to remove failed jobs from the appliance. (This button displays only if failures occurred.)

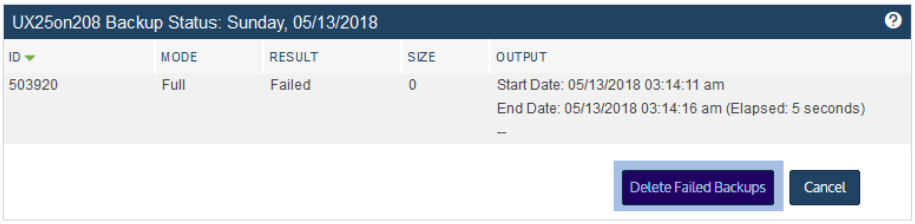

- Click **Add to Hot Backup Copy Queue** to copy successful full backups to the Unitrends Cloud or to another Unitrends appliance. (This button displays only if a hot backup copy target has been added to the appliance and a backup has not yet been copied or added to the queue.)

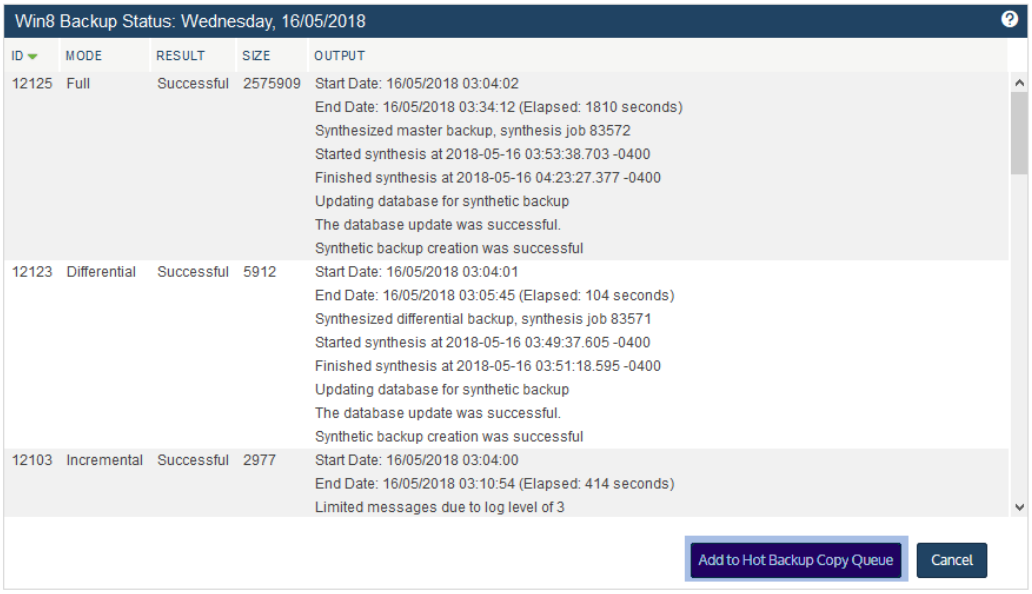

4 Click **Cancel** to exit.

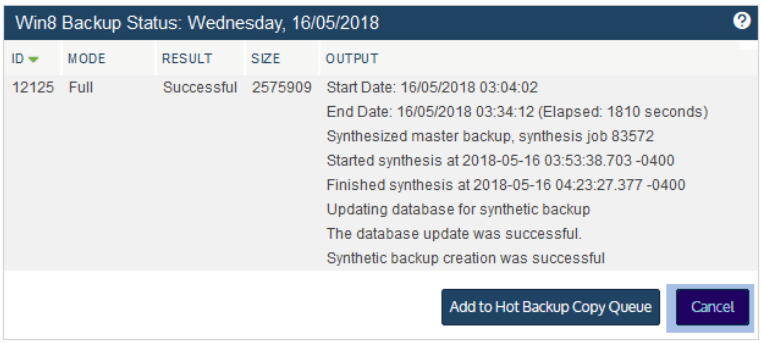

Copied Assets tab

This tab applies to appliances that are receiving hot backup copies from another Unitrends appliance. On a backup copy target appliance, the Copied Assets tab displays the status of hot backup copies received over the last seven days. See the following for details on working with the Copied Assets tab:

- "Backup Copy button"
- "Inventory tree" on page 57
- "Status table" on page 57

Backup Copy button

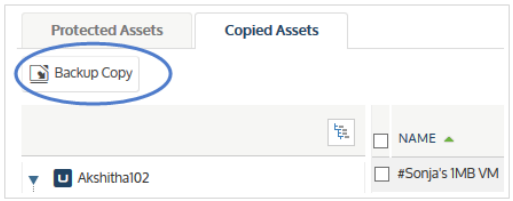

Opens the Create Backup Copy Job dialog. For details on using this dialog to create a backup copy job, see "Creating backup copy jobs" on page 491.

Inventory tree

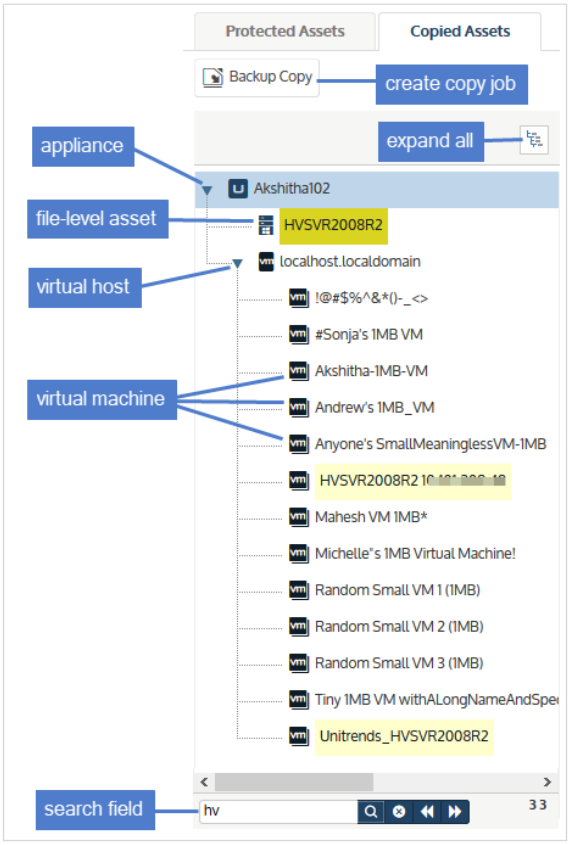

All Unitrends source appliances that are sending copies to this target display in the inventory tree, along with the assets whose backups are being copied:

- Source appliances display as top-level nodes.
- Agent-based physical and virtual assets, and virtual hosts display as second-level sub-nodes.
- Hosted VMs and applications display as sub-nodes under their host asset.

Use these options while working with the inventory tree:

- To look for an asset by name, use the **Search** field below.
- To view all assets, click the **Expand All** icon located above the tree.

Status table

Selecting an appliance or asset in the inventory tree populates the status table with details about backup copies that ran within the last seven days for the applicable assets. Job status icons display in the Backup Copies column. Click a status icon to view more details. If the Backup Copies icon is grayed-out, no copy job ran for the asset on that day.

See these topics for details:

- "Status by appliance"

- "Status by host asset" on page 58
- "Status by single asset" on page 59
- "Filtering the status table results" on page 59
- "Viewing job details" on page 54

Status by appliance

Select an appliance to view backup copy status for all its copied assets:

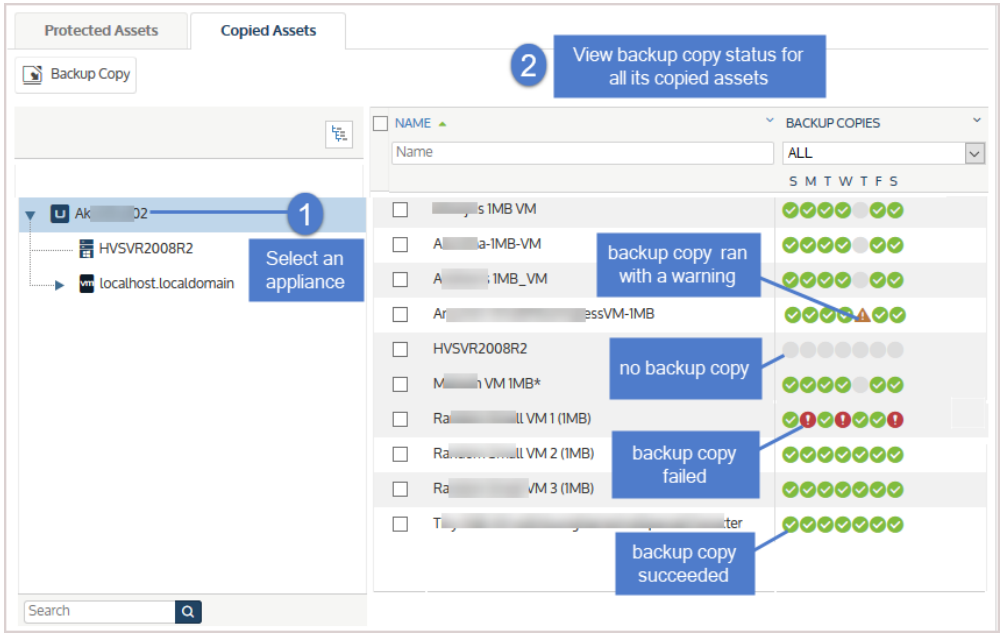

Status by host asset

Select a virtual host or application to view backup copy status for its hosted VMs or instances:

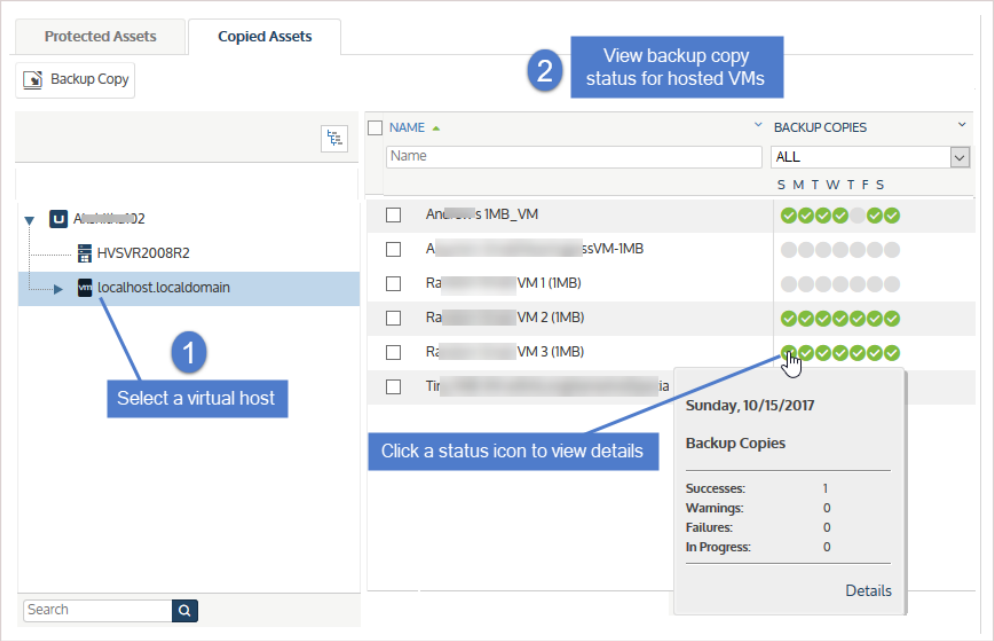

Status by single asset

Select a single agent-based asset or virtual machine to view its backup copy status:

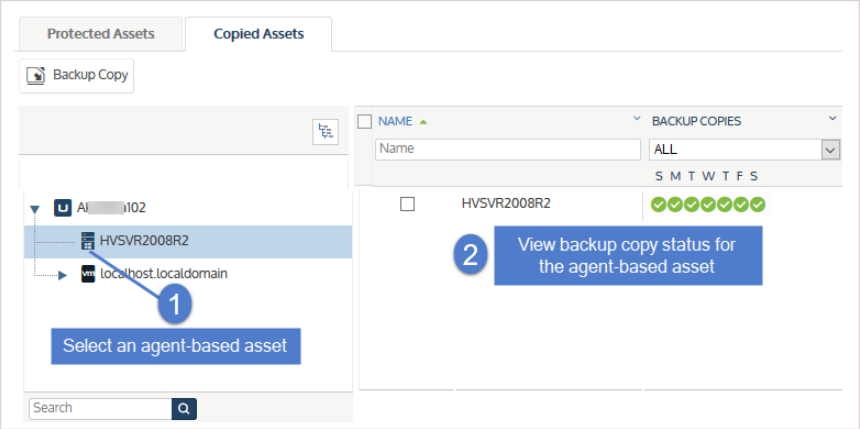

Filtering the status table results

After selecting an appliance or asset in the inventory tree, you can filter the display results by using these fields:

- Name field - Enter text to display only asset names that contain the string you entered.
- Backup Copies drop-down - Select an item from this list to display only assets with backup copies that meet the condition you selected. For example, select *Failed in the last 7 days* to see only the assets with backup copies that have failed in the last 7-day period, or *No successes in the last 7 days* to see only the assets that have not had a successful backup copy in the last 7-day period (this includes cases where no backup copy jobs have run and cases where backup copies have run but none were successful).

In this example, results are filtered to show only the assets that have not had a successful backup copy within the last seven days:

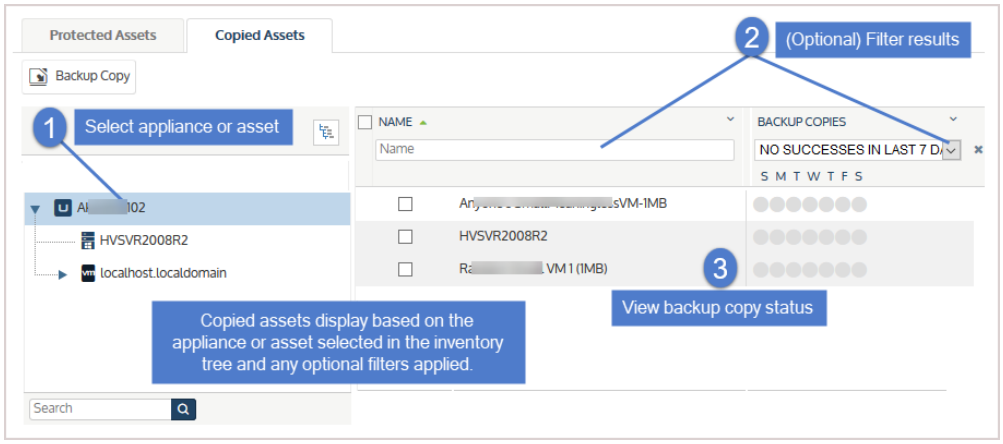

Recover

Use the Recover page to recover entire assets or individual files, to create replicas of Windows or VMware assets, and to perform instant recovery from VM or Windows image-level backups. A high-level overview of the Recover page is given below. For detailed recovery procedures, see the applicable Recovery chapter in this guide.

The Recover page contains the following tabs:

- "Backup Catalog tab"
- "File Level Recovery tab" on page 73
- "Instant Recovery tab" on page 75
- "Replicas tab" on page 76

Backup Catalog tab

The Backup Catalog lists the appliance's backups and backup copies, and contains buttons used for recovery and other management tasks.

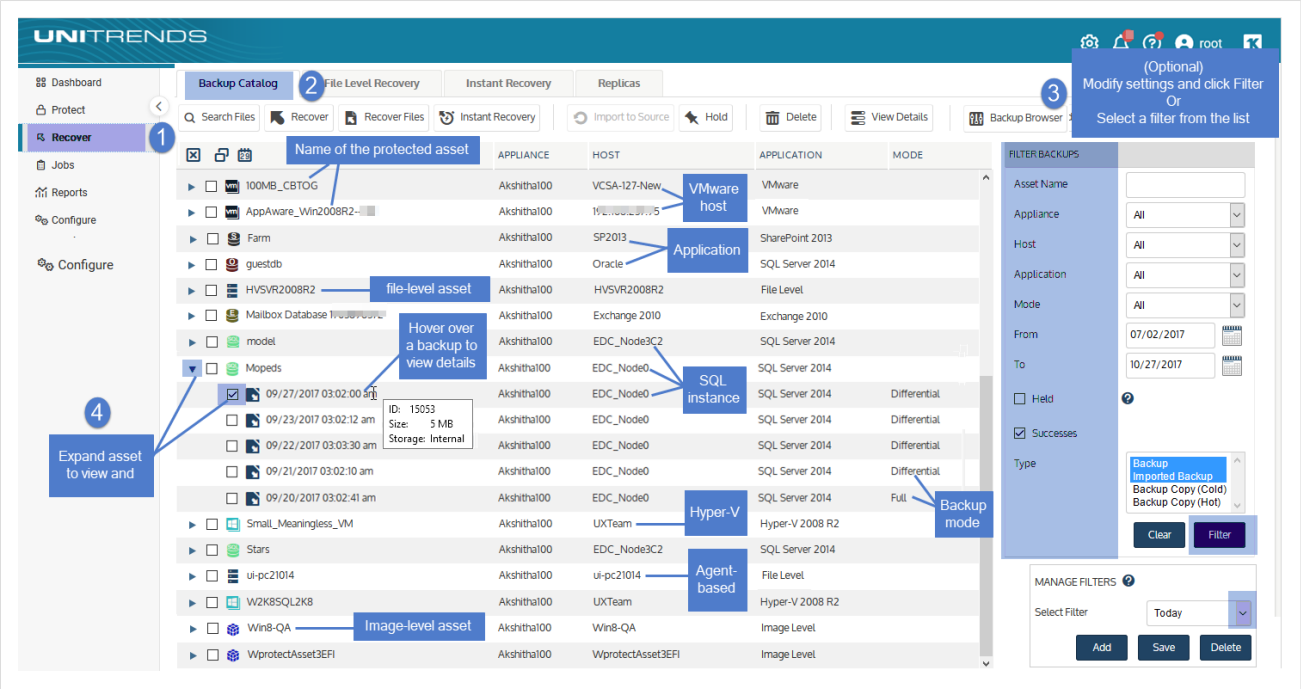

Upon accessing the Backup Catalog tab, the default filter displays associated backups and/or backup copies under each protected asset.

To change the backups and/or copies that display, you can:

- Modify Backup Catalog filter settings – Modify the display by entering filter criteria. Backups and backup copies display under the protected asset. Expand an asset to view its backups and backup copies. For details, see "Working with custom filters" on page 67.
- Search for backups by using the Backup Browser – The Backup Browser provides advanced search and filter options. Backups are not grouped under the protected asset. Search for backups by selecting an appliance and date range. Filter the display by entering search text. (Backup copies and imported backups cannot be viewed in the Backup Browser. Use the Backup Catalog instead.) For details, see "Working with the Backup Browser" on page 62.

Buttons

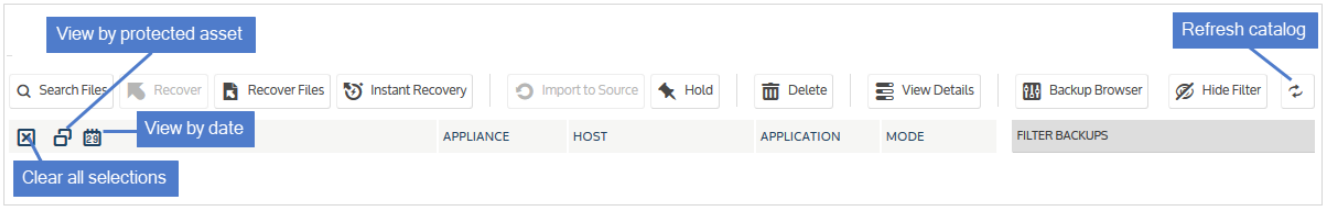

Use these buttons while working with the backup or backup copy you have selected in the Backup Catalog table:

- Search Files – Use to search for specific files in an asset's backups or backup copies and select files to recover from the search results. Supported for file-level backups and copies only.
- Recover – Use to recover an entire asset, backup, or backup copy.
- Recover Files – Use to browse the contents of a backup or backup copy and select files and/or folders to recover.
- Instant Recovery – Use to quickly recover a failed Windows, VMware, or Hyper-V asset. (See "Virtual machine instant recovery" on page 904 or "Instant recovery of Windows image-level backups" on page 1055 for details.)
- Import to Source – Use to import the selected hot or cold backup copy to the backup appliance. Once a copy is imported, you can recover from it as you do from any local backup.
- Hold – Use to place the selected backups on hold. Backups on hold cannot be removed from the appliance. To remove the hold, select the held backup and click **Unhold**.
- Delete – Use to delete the selected backups from the appliance.
- View Details – Use to view details of the selected backups and/or backup copies. (See "To view details from the Backup Catalog tab" on page 626 for details.)
- Backup Browser – Use to launch the Backup Browser, which enables advanced search and filter options. For details, see "Working with the Backup Browser".
- Hide Filter – Use to hide the Filter Backups area.
- Refresh catalog – Use to refresh the contents of the Backup Catalog tab.
- Clear all selections – Use to clear all selections on the Backup Catalog tab.
- View by protected asset – Use to view backups and/or copies by protected asset.
- View by date – Use to view backups and/or copies by date.

Working with the Backup Browser

The Backup Browser provides appliance-level search capability, advanced search options, and faster filtering performance. With the Backup Browser you can quickly search for, filter, and sort the backups stored on the appliance. You can then delete, hold, unhold, and recover backups right from the search results.

The Backup Catalog enables you to browse by asset only. Once you filter the Backup Catalog display, you must manually expand each asset to view its backups. This approach works well for asset-level searches, but can be cumbersome in some cases. For example, use the Backup Browser for these appliance-level searches:

- Search for and delete all failed backups.
- Search for all backups that have been placed on hold.
- Search for backups by application type.
- Search for all backups of assets with *test* in their names.

See these procedures for details on using the Backup Browser:

- "To search for backups in the Backup Browser"

- "To delete backups by using the Backup Browser" on page 627
- "To place backups on hold by using the Backup Browser" on page 637
- "To unhold backups by using the Backup Browser" on page 640

To search for backups in the Backup Browser

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

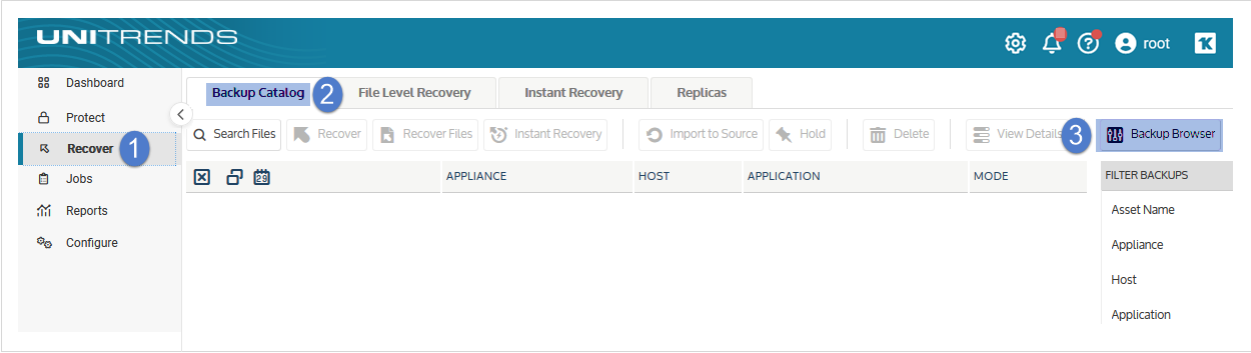

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display:

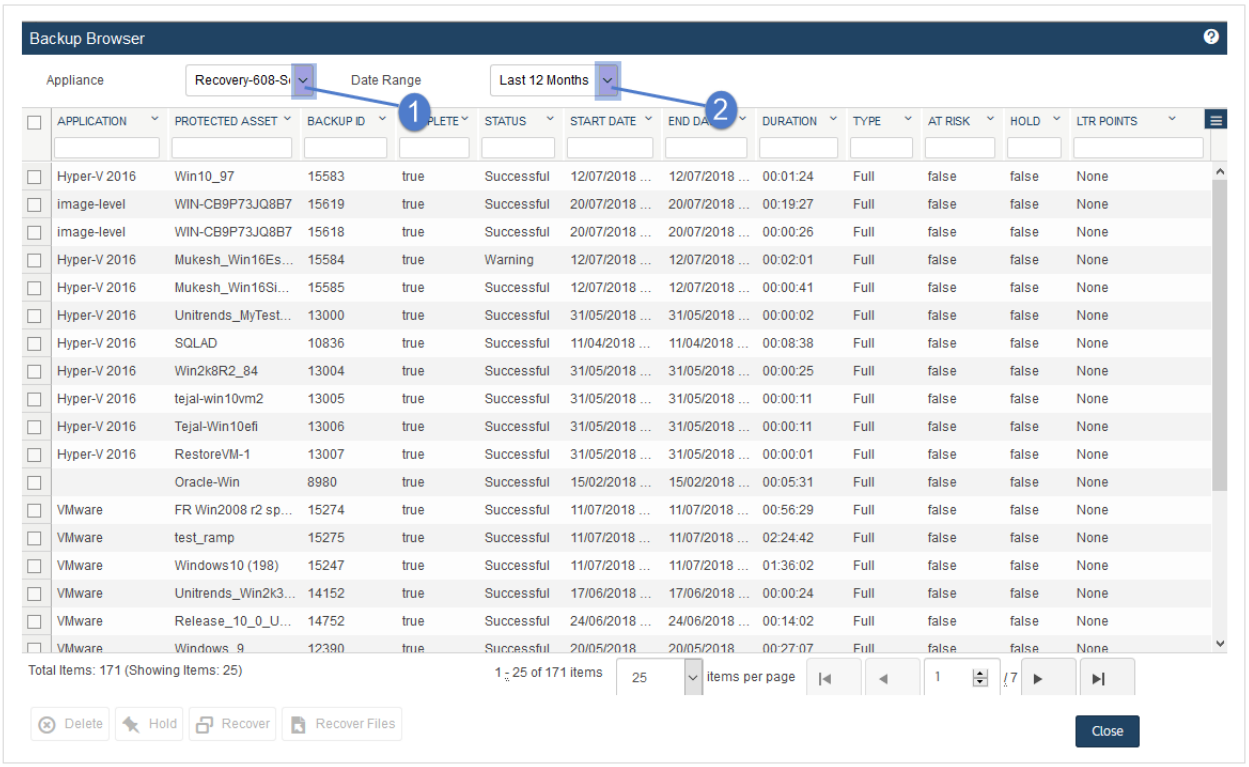

4 Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

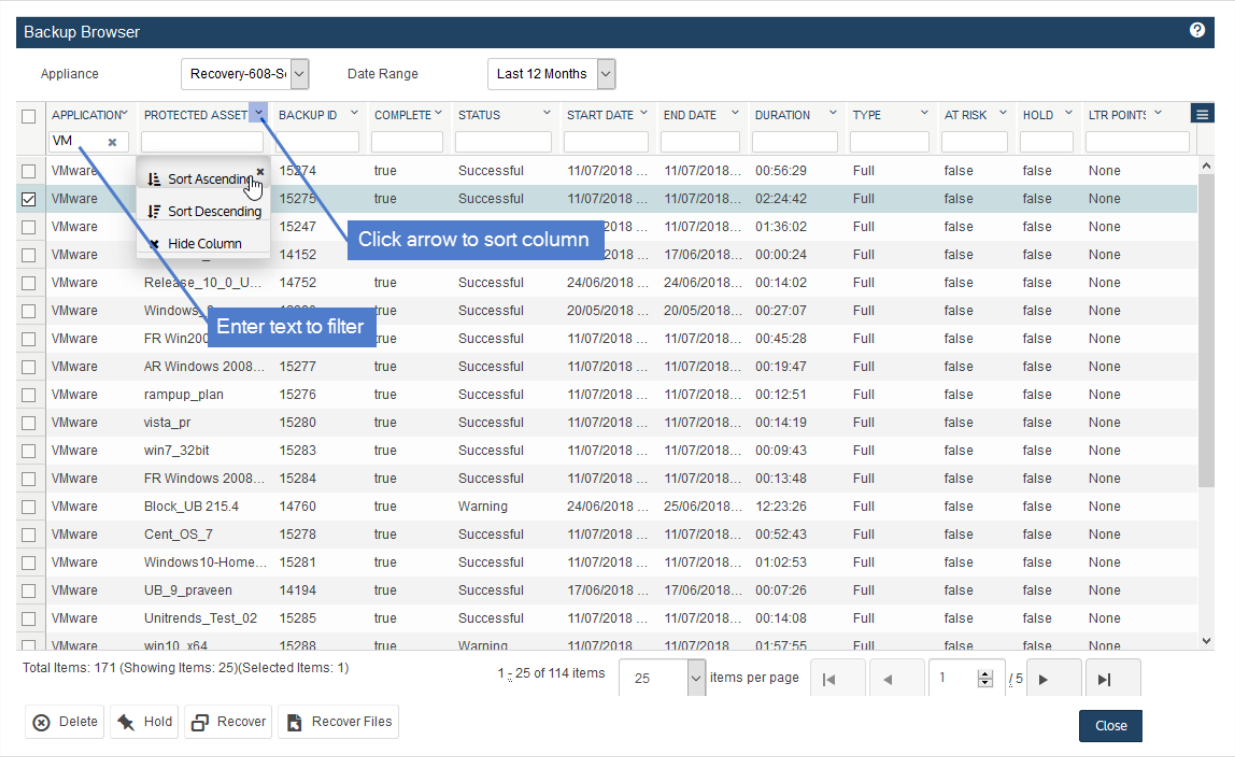

Select backups and click a button below to delete, hold, unhold, or recover.

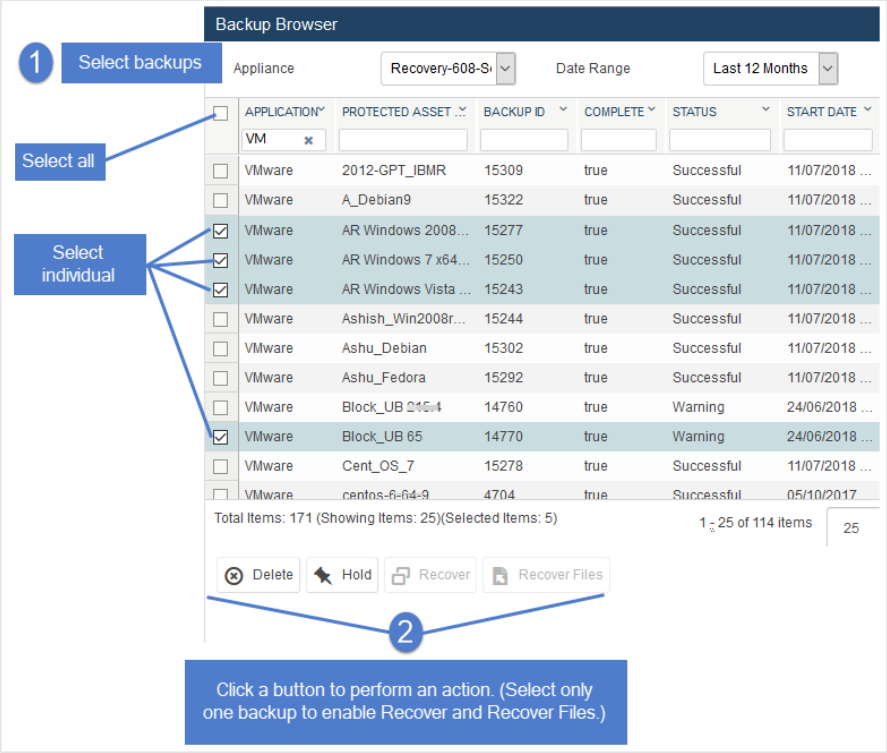

Backup Browser column descriptions

Descriptions of each Backup Browser column are given here:

Column	Description
System ID	The system-generated appliance ID.
Application	Application type: • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013. • For agent-based backups, contains *File-Level* or *Image-Level*.
Protected Asset	Name of the asset.
VM	Name of the virtual machine.

Column	Description
Database	Name of the protected database or storage group.
Instance Name	For application backups, contains the database instance name. For host-level backups, contains one of the following: • The name of the virtual machine. • The name of the Hyper-V application instance. • The name of the ESXi host if that host is being managed by a vCenter.
Backup ID	The system-generated job ID.
Complete	Indicates whether the job completed: *True* (yes) or *False* (no).
Status	Status of the backup job: *Success*, *Warning*, or *Failure*.
Start Date	The date and time at which the backup job started.
End Date	The date and time at which the backup job completed.
Duration	The amount of time it took for the job to complete, in *hh:mm:ss* format.
Type	Backup mode: Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, or Transaction (SQL only).
At Risk	Indicates if a potential ransomware infection was discovered: *True* (potentially infected) or *False* (not infected).
Hold	Indicates whether the backup has been placed on hold: *True* (yes) or *False* (no).
Encrypted	Indicates whether the backup is encrypted: *True* (yes) or *False* (no).
LTR Points	Indicates whether the backup is needed for adherence to a long-term retention (LTR) policy: *None* (no), *Daily*, *Weekly*, *Monthly*, or *Yearly*.
Storage	Storage device where the backup resides. (Default appliance storage device is *Internal*.)
Appliance	Name of the backup appliance.

Working with custom filters

Upon accessing the Backup Catalog tab, the default filter displays associated backups and/or backup copies. To change the backups and copies that display, you can modify settings manually or apply a custom filter. See these topics for details:

- "To modify the display manually"
- "To add a filter" on page 67
- "To apply a filter" on page 70
- "To assign a default filter" on page 70
- "To edit a filter" on page 71
- "To delete a filter" on page 72

To modify the display manually

1 Modify the values in the Filter Backups fields.

2 Click **Filter** to apply the new settings.

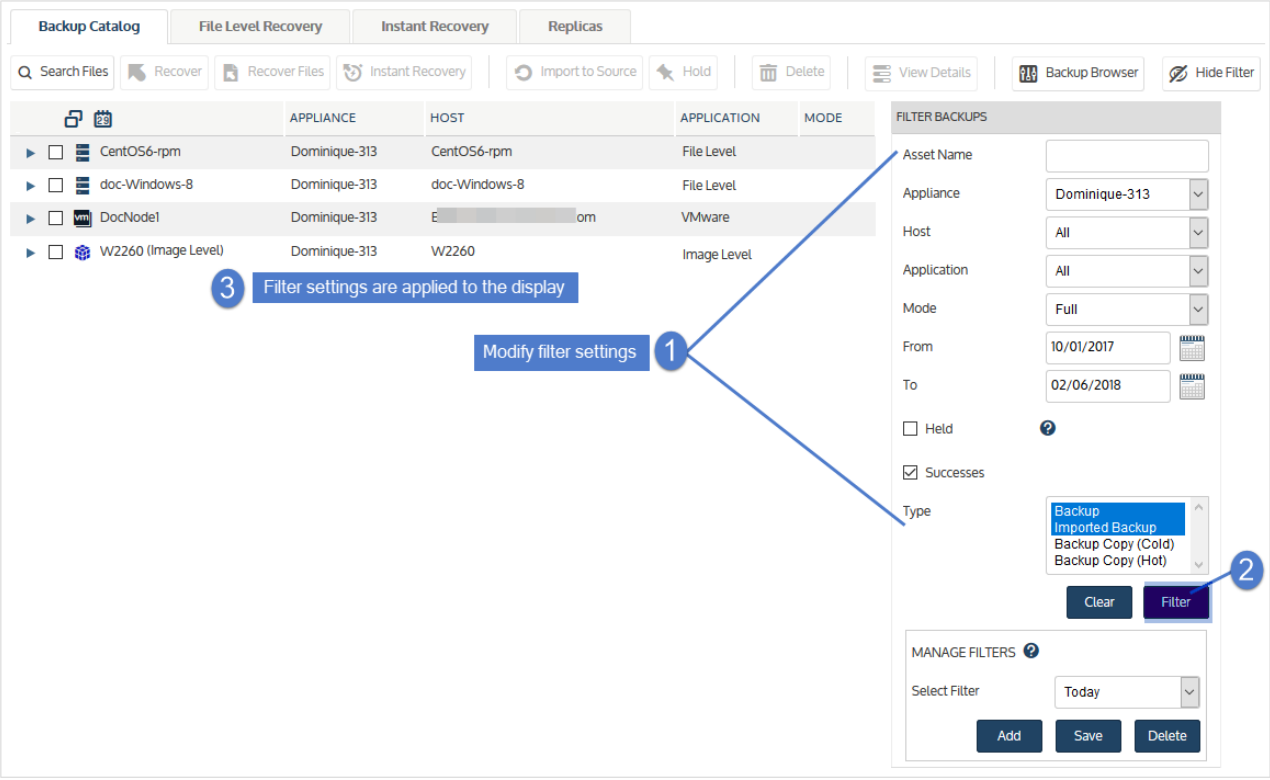

To add a filter

1 In the Manage Filters area, click **Add**.

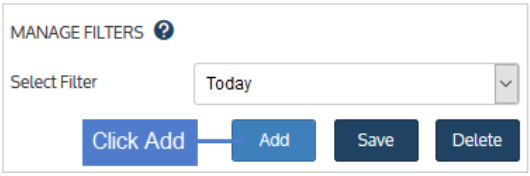

2 Enter the following:

- A unique name.
- (Optional) Check **Set as default** to automatically load this filter.

3 Click **Save**.

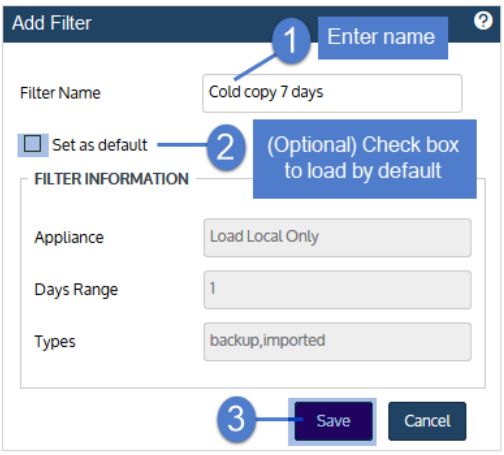

4 The new filter is selected in the Select Filter list. Modify settings in the filter fields as needed, then click **Save**.

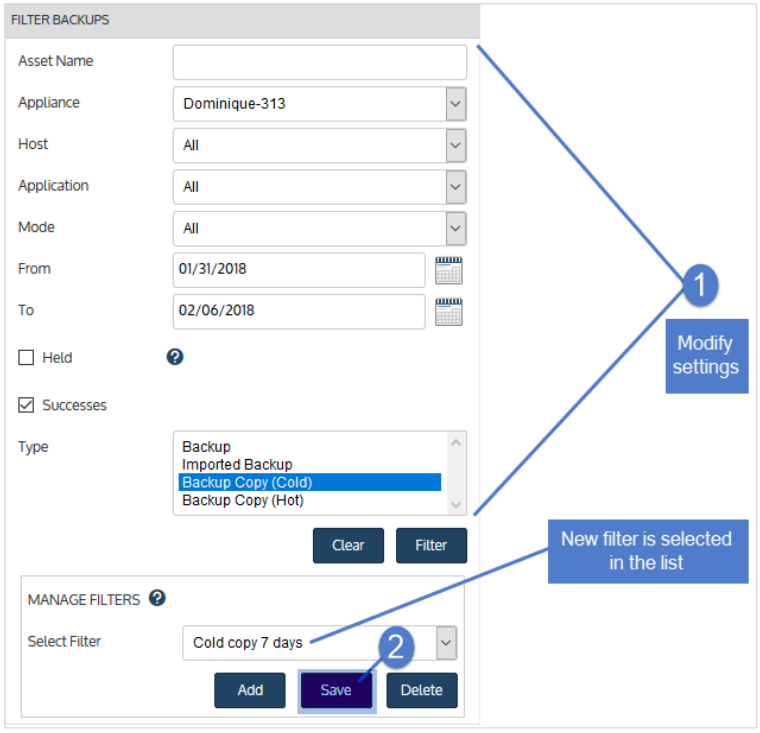

5 Review settings and click **Save** to exit.

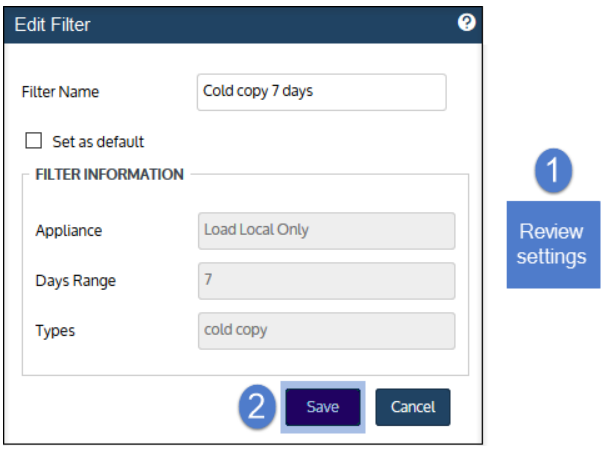

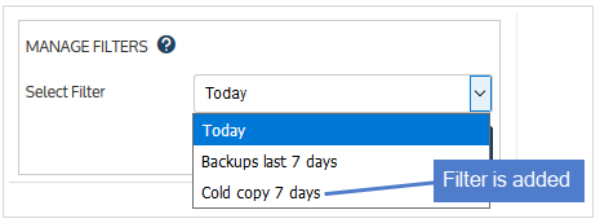

To apply a filter

Apply a custom filter by selecting it in the **Select Filter** list.

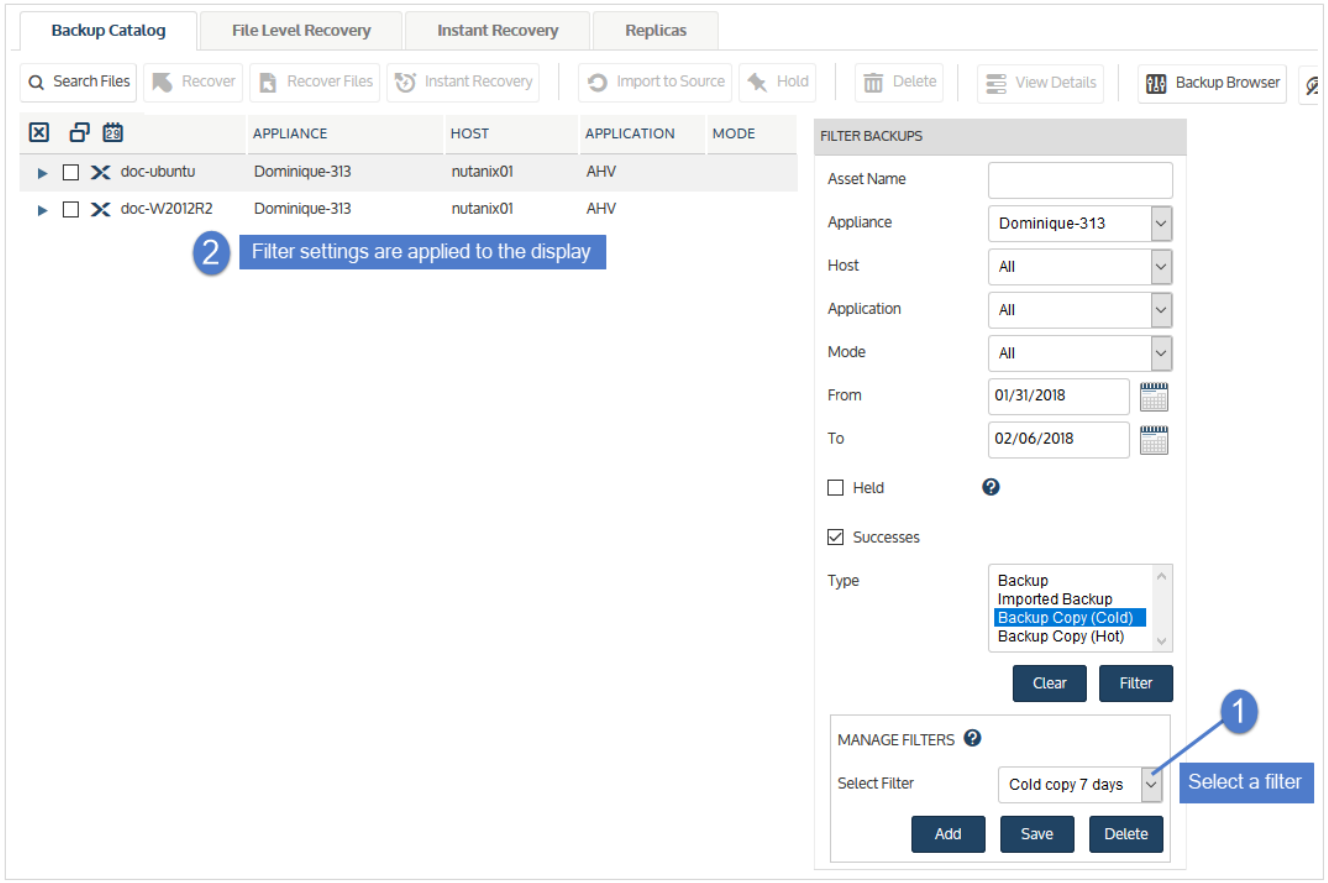

To assign a default filter

Upon accessing the Backup Catalog tab, the default filter is automatically applied, and backups and/or backup copies that meet the filter criteria display. The appliance's default filter displays today's backups and imported backup copies. To assign a different default filter, do one of the following:

- Add a new custom filter. In the Add Filter dialog, check the **Set as default** box. For details, see "To add a filter" on page 67.
- Make an existing filter the new default by using these steps:

Note: If you have already assigned a default filter and want to change your selection, simply check the **Set as default** box while creating a new filter or modifying an existing filter. This clears the **Set as default** checkbox of the previous default filter.

1. Select the filter and click **Save**.

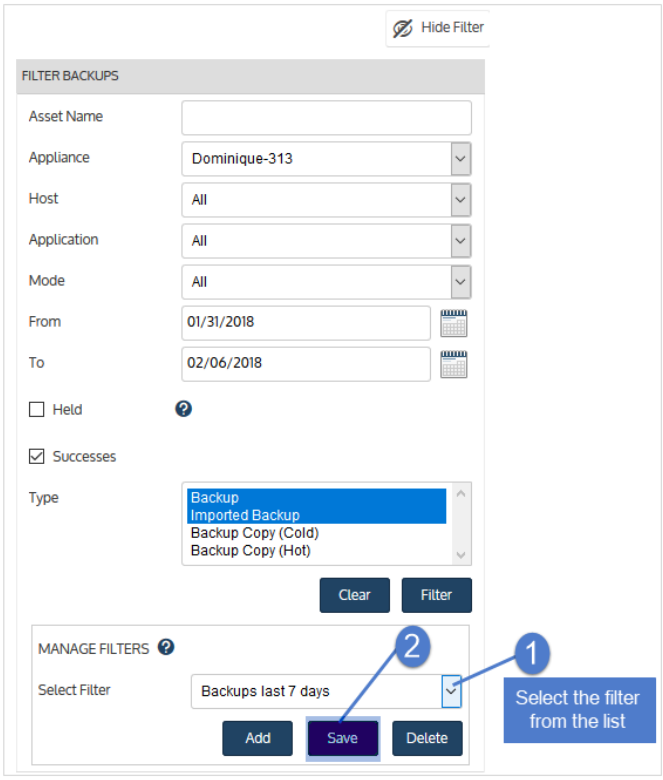

2 Check the **Set as default** box, then click **Save**.

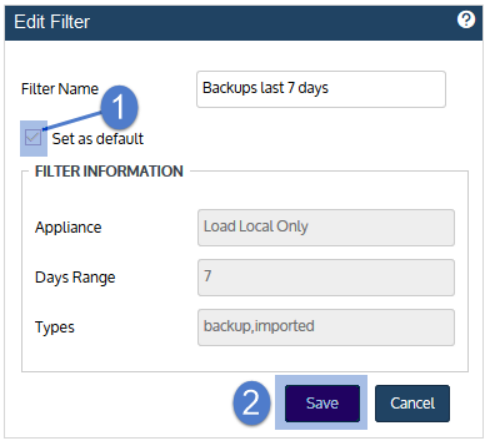

To edit a filter

1 Select the filter.

2 Modify settings in the filter fields as needed, then click **Save**.

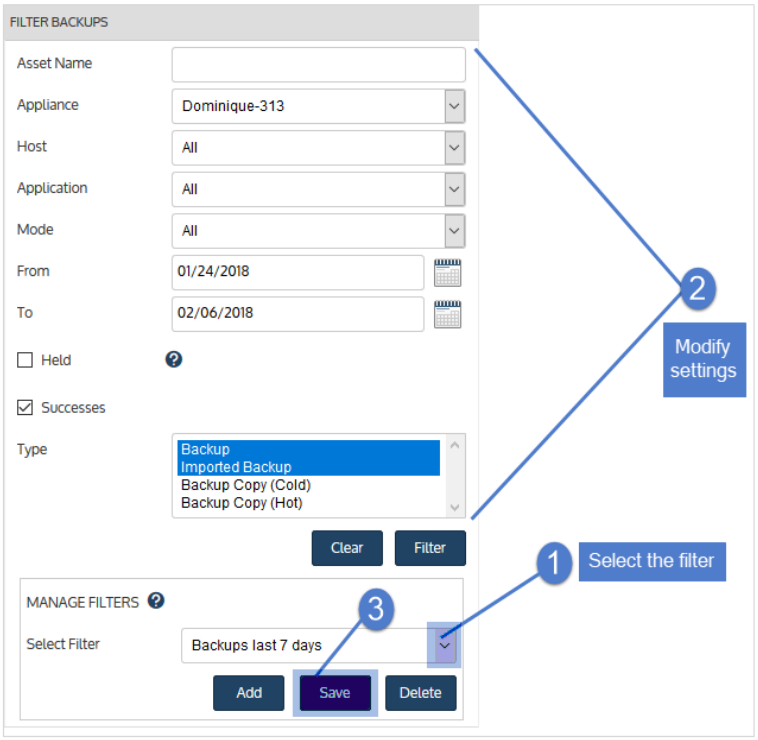

3 Review settings, modify name (optional), check **Set as default** box (optional), then click **Save**.

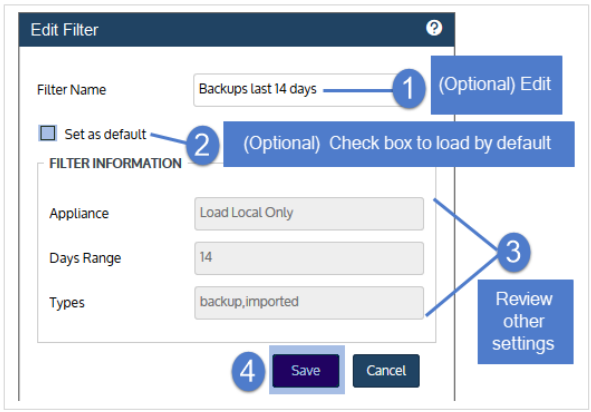

To delete a filter

1 Select the filter, then click **Delete**.

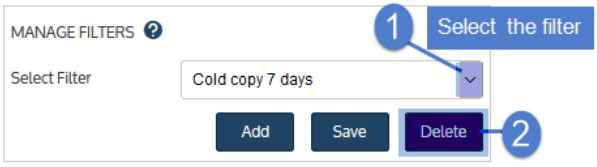

2 Click **Delete Filter** to confirm.

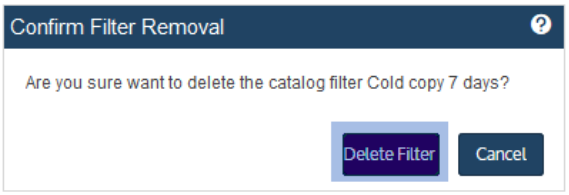

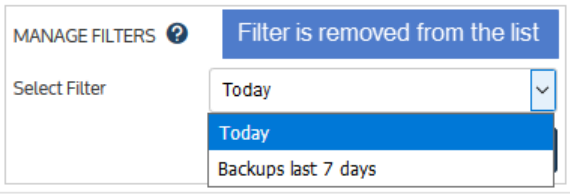

File Level Recovery tab

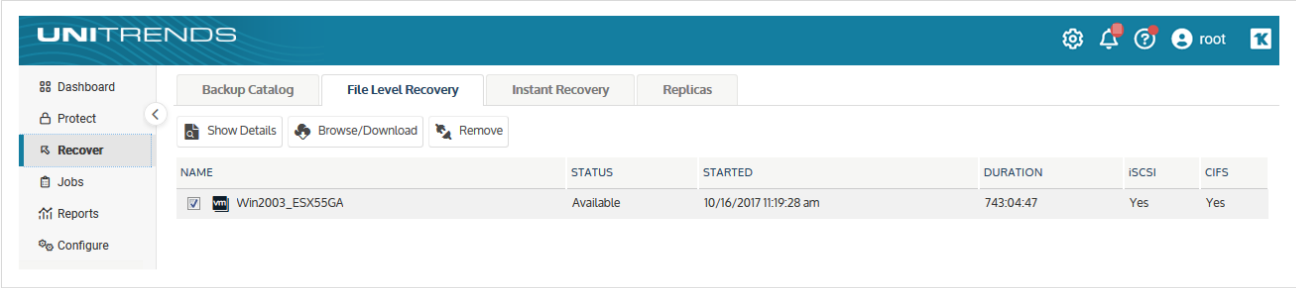

Recovery objects are created to recover files from VM host-level backups. This tab enables you to view and remove these objects. For detailed procedures, see " Recovering files from virtual machine backups" on page 808.

Use these buttons while working with objects on the File Level Recovery tab:

- Show Details – Displays the File Level Recovery Details dialog for the selected object. Details include: VM name, virtual host name, date/time the object was created, Unitrends appliance name, CIFS path to the recovery object (if applicable), iSCSI targets for recovery object (if applicable), and messages. Example:

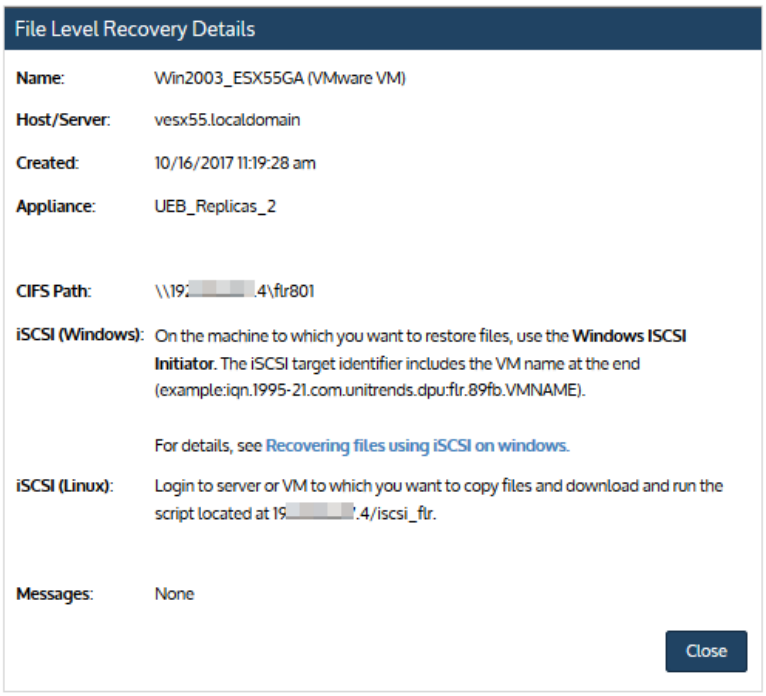

- Browse/Download – Opens a File Browser for the selected object, where you can select files and/or folders to recover. Example:

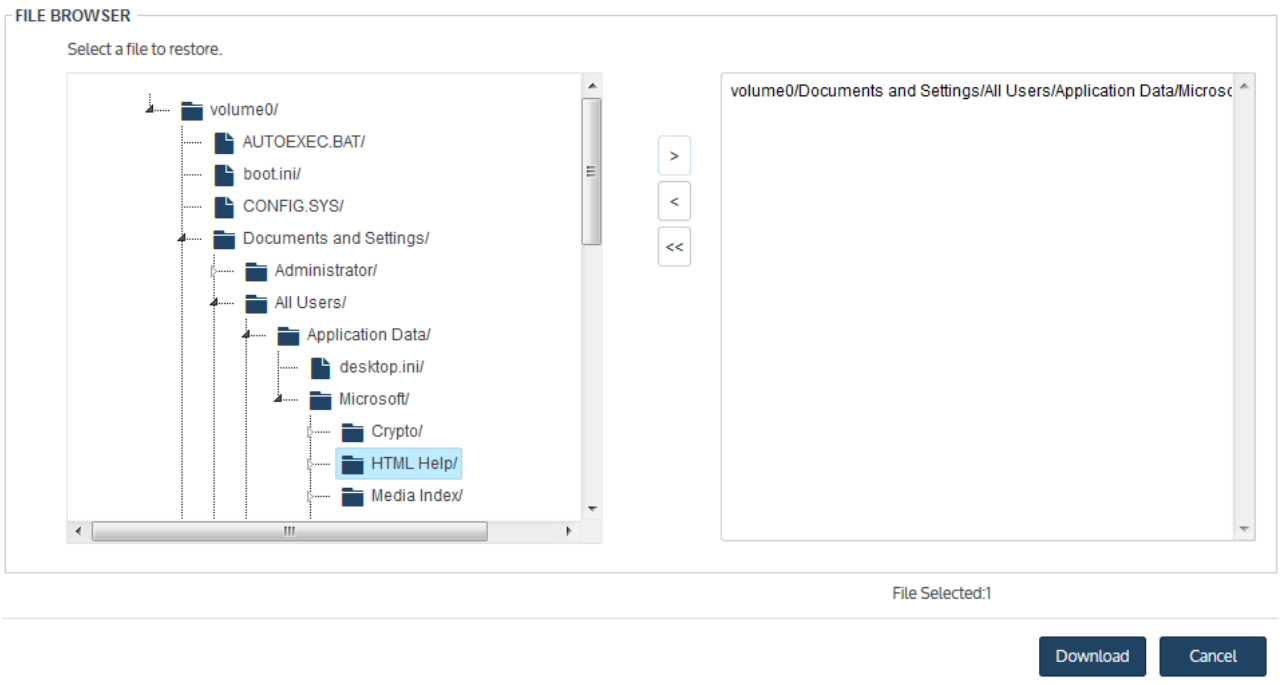

- Remove – Removes the selected file level recovery object.

Instant Recovery tab

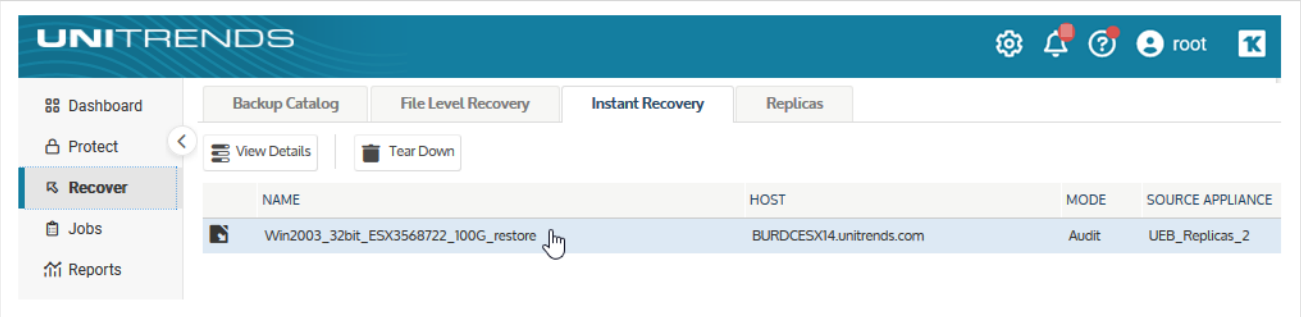

Instant recovery enables you to recover a failed or corrupted VM in minutes. To make these recovered VMs available in production very quickly, instant recovery objects run on the Unitrends appliance while data is being migrated to the recovered VM on the target virtual host. For more on instant recovery, see "Virtual machine instant recovery" on page 904.

Once you have set up a VM for instant recovery, its instant recovery object displays on the Instant Recovery tab. Use these buttons while working with instant recovery objects on this tab:

- Show Details – Displays the Details tab for the selected object. Details include:
  - Name – Name of the recovered VM.
  - Created – Date and time at which the object was created.
  - Duration – Elapsed number of days, hours, and minutes since the object was created.
  - Status – Status of the object.
  - Message – Description of the current status or most recent operation related to the object.
- Tear Down – Deletes the selected instant recovery object.

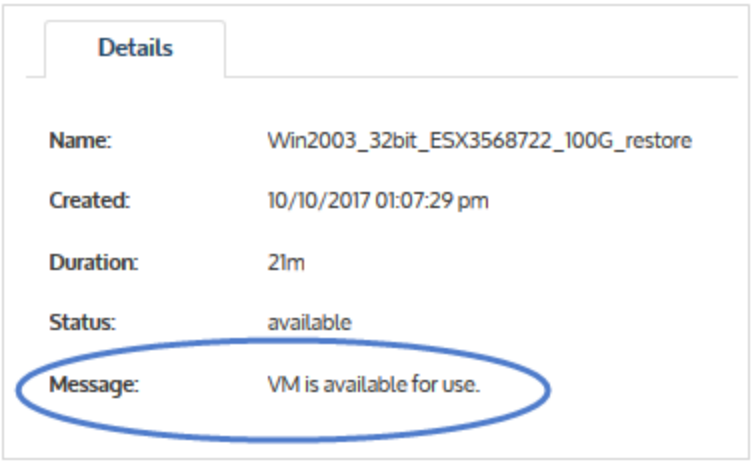

Replicas tab

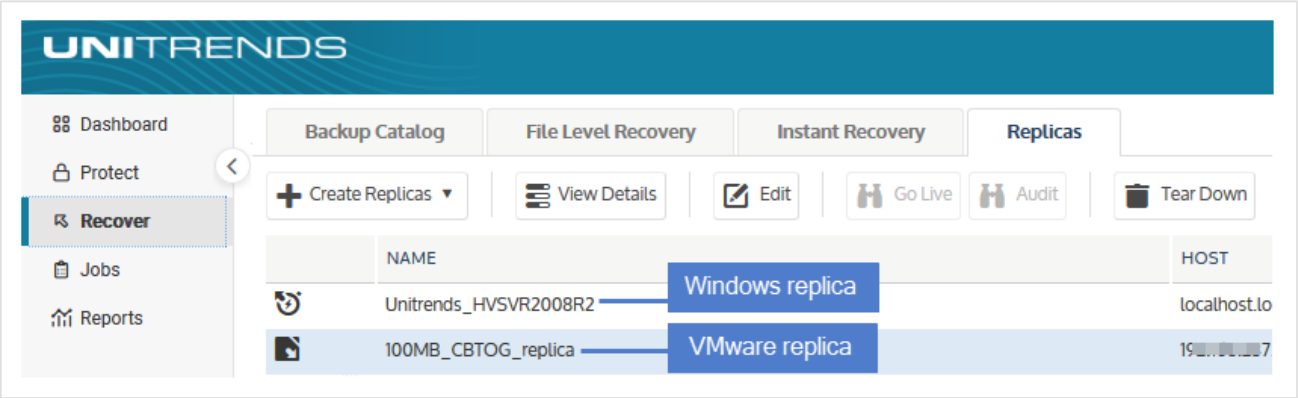

Replicas provide a quick way to recover VMware virtual machines and Windows physical assets. To use this feature, simply set up the replica by using the Create Replica VMs or Create Windows Replica dialog. The appliance then creates a virtual machine replica from the most recent backup of the original asset, and automatically applies all subsequent backups. Because the replica is continually updated, you can bring it online to immediately assume the role of original machine. For more on replicas, see "VM replicas" on page 876 and "Windows file-level replicas" on page 993.

VM replicas and Windows replicas display on the Replicas tab. Use these buttons while working with replicas on this tab:

- Create Replicas – Click to create Windows and VMware replicas.
- View Details - Displays the Details tab for the selected replica.
- Edit - Displays the Edit Replica Details dialog for the selected replica.

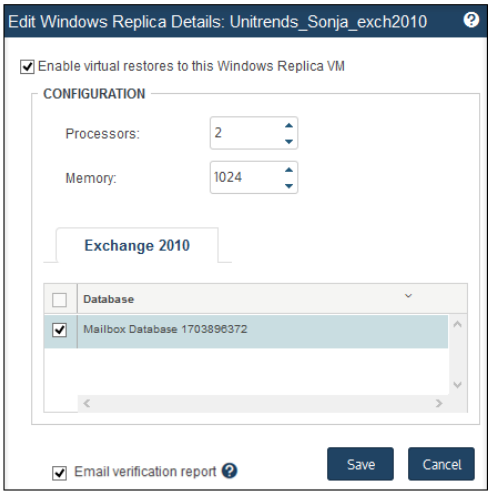

- Go Live - Boots the selected replica into live mode. (Use only if the original asset has failed and is no longer online in production.)

- Audit/End Audit- Click **Audit** to boot the selected replica into audit mode. Click **End Audit** to exit audit mode.
- Tear Down - Removes the selected replica from the appliance and, optionally, from the virtual host.

Jobs

The Jobs page enables you to create, edit, and delete jobs and view current job progress. The Jobs page contains the following tabs:

Active Jobs tab

This tab displays all currently running jobs. For detailed procedures, see "Managing active jobs" on page 607.

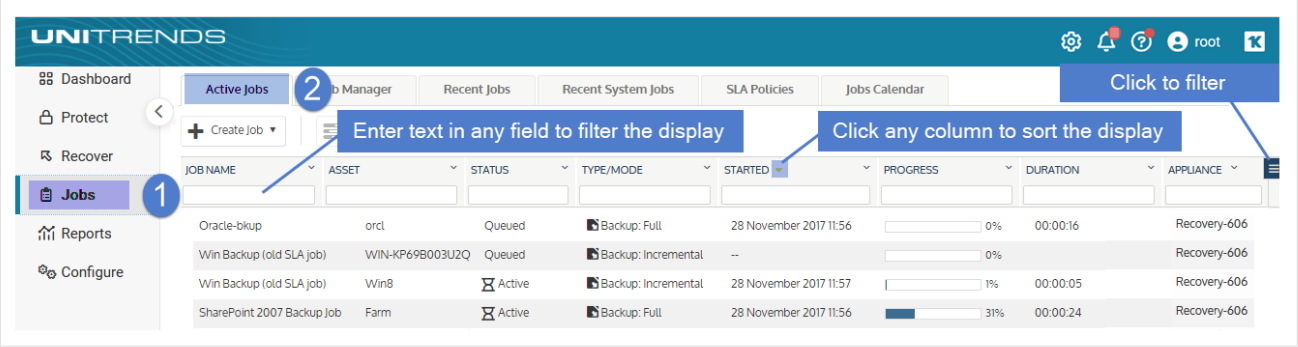

Use these buttons while working with active jobs:

- Create job - Select to open any of the following: the Create Backup Job dialog, the Create Backup Copy Job dialog, or the Create Data Copy Access Job dialog.
- View Details/Hide Details - Click to display or hide the details pane for the selected job.
- Pause - Click to pause the selected queued job. (Running jobs cannot be paused.)
- Cancel - Click to cancel the selected job.

Job Manager tab

This tab displays all scheduled jobs. Use this tab to view, create, edit, enable or disable, and delete job schedules, and to run a selected schedule on demand. If a job is disabled, it is grayed-out, does not run as scheduled, and cannot be run on demand. For detailed procedures, see "Managing scheduled jobs" on page 563.

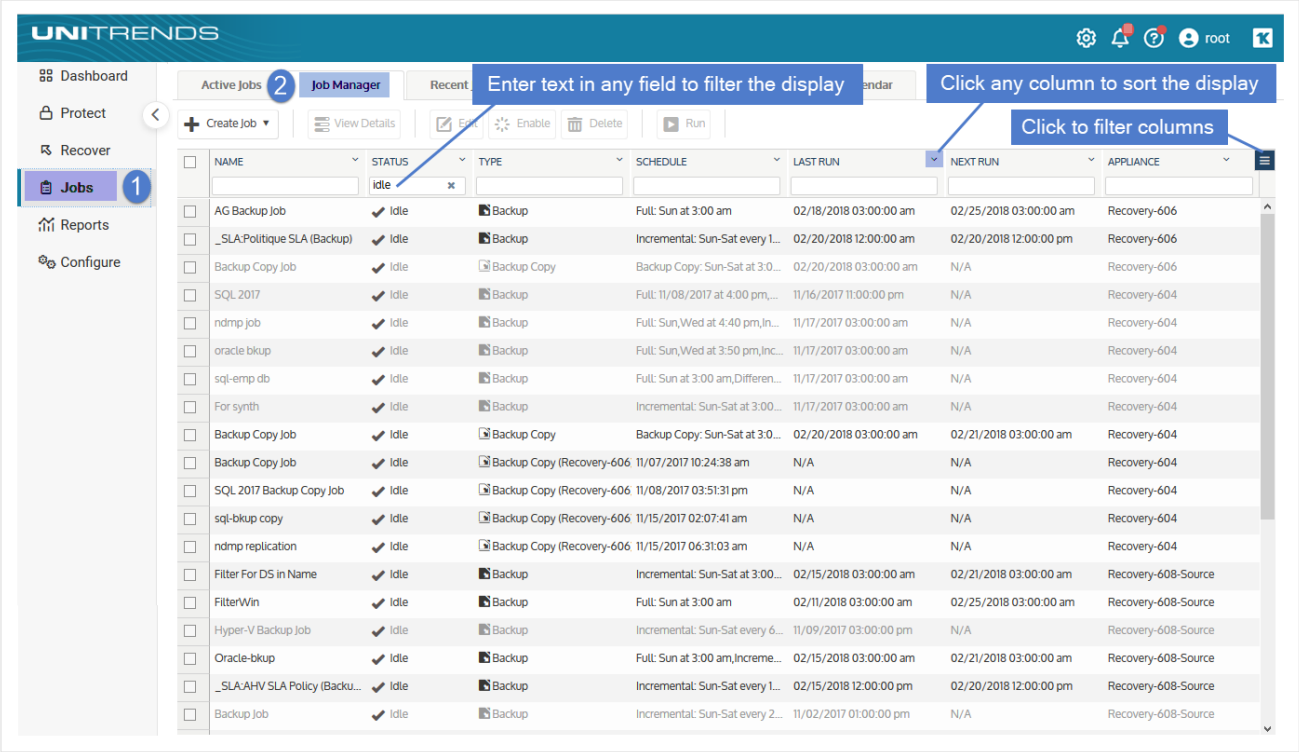

Use these buttons while working with the Job Manager:

- Create job - Select to open any of the following: the Create Backup Job dialog, the Create Backup Copy Job dialog, or the Create Data Copy Access Job dialog.
- View Details/Hide Details - Click to display or hide the details pane for the selected job.
- Edit - Click to edit the selected job.
- Disable/Enable - Click to disable or enable the selected job. An enabled job runs according to schedule and can be run on demand. A disabled job does not run.
- Delete - Click to remove the selected job.
- Run - Click to run the selected job.

Recent Jobs tab

This tab displays jobs that ran in the last seven days. For detailed procedures, see "Viewing recent jobs" on page 615.

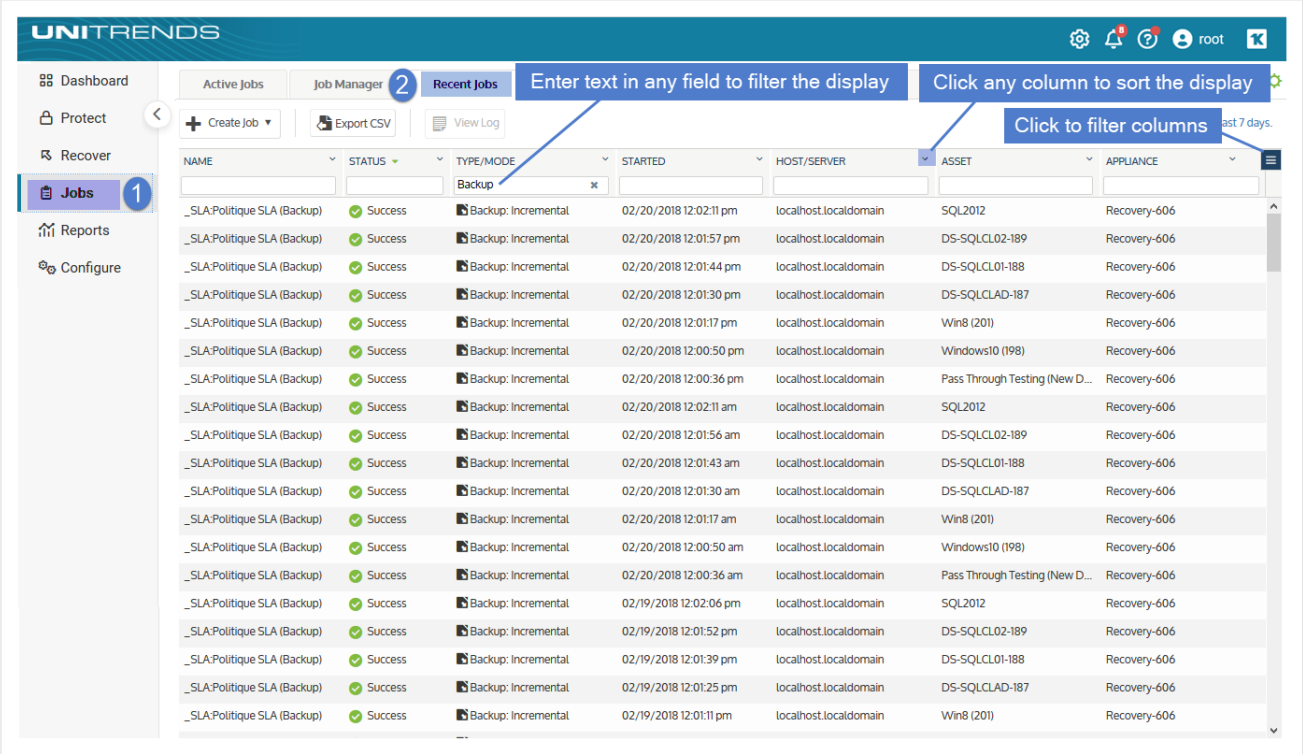

Use these buttons while working with recent jobs:

- Create job - Select to open any of the following: the Create Backup Job dialog, the Create Backup Copy Job dialog, or the Create Data Copy Access Job dialog.
- Export CSV - Click to export the job history as a CSV file.
- View log - Click to display details for the selected job.

Recent System Jobs tab

This tab displays system jobs that ran in the last seven days. For detailed procedures, see "Viewing system jobs" on page 620.

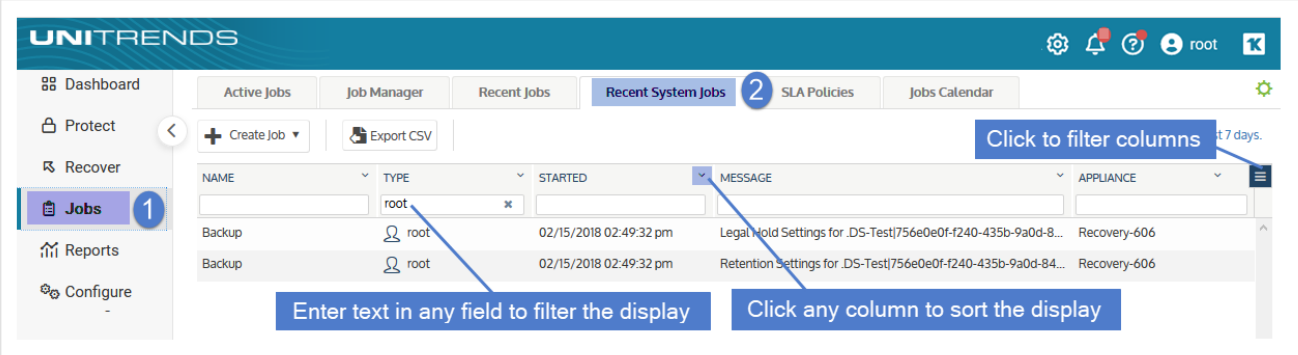

Use these buttons while working with recent system jobs:

- Create job - Select to open any of the following: the Create Backup Job dialog, the Create Backup Copy Job dialog, or the Create Data Copy Access Job dialog.
- Export CSV - Click to export the job history as a CSV file.

For more on recent system jobs, see "Viewing system jobs" on page 620.

SLA Policies tab

An SLA policy automatically creates backup and backup copy job schedules for a group of assets based on the RPO, retention, and backup copy settings you define. This tab displays all SLA policies. For detailed procedures, see "Creating SLA policies" on page 536.

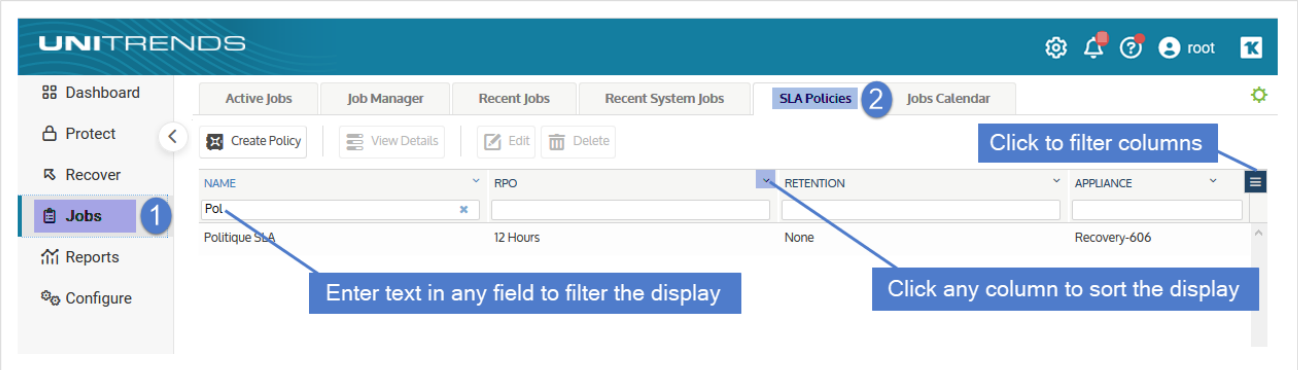

Use these buttons to view, create, edit, and delete SLA policies:

- Create Policy - Select to open the Create SLA Policy dialog.
- View Details/Hide Details- Click to display or hide the details pane for the selected policy.
- Edit - Click to edit the selected policy.
- Delete - Click to remove the selected policy.

Jobs Calendar tab

This tab provides a calendar view of scheduled jobs. Click the **month**, **day**, or **week** buttons to change the number of days that display. Click the **today**, **<**, and **>** buttons to display another month, day, or week. Click the gear icon to modify the refresh rate.

Month example:

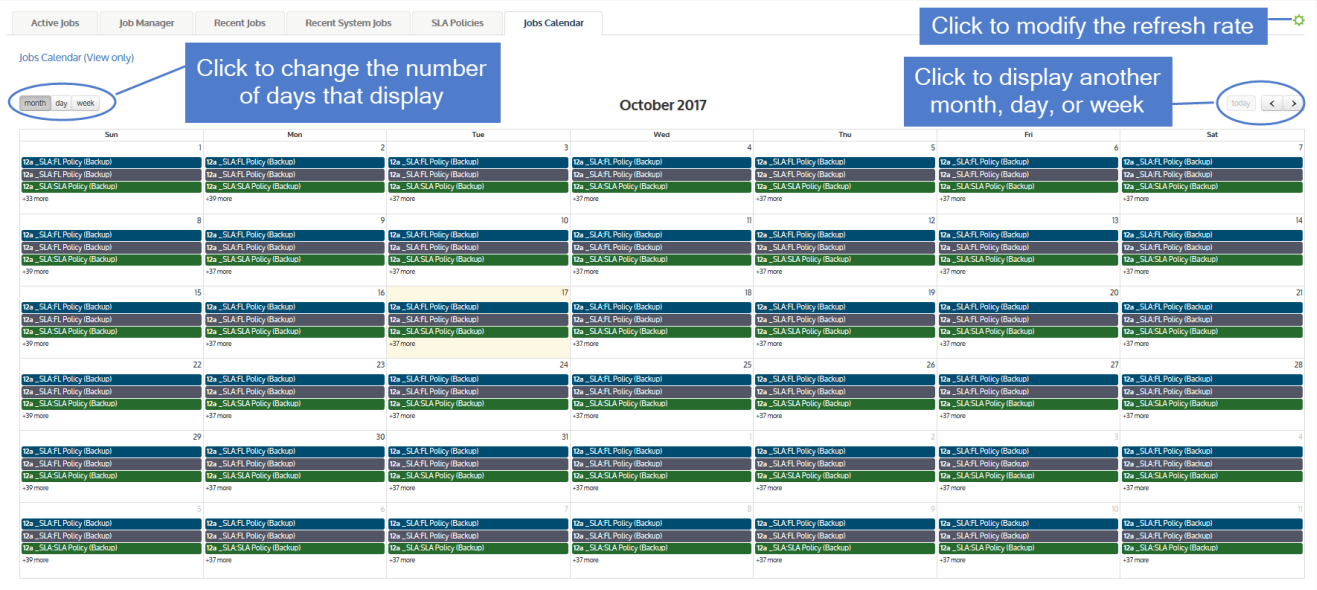

Day example:

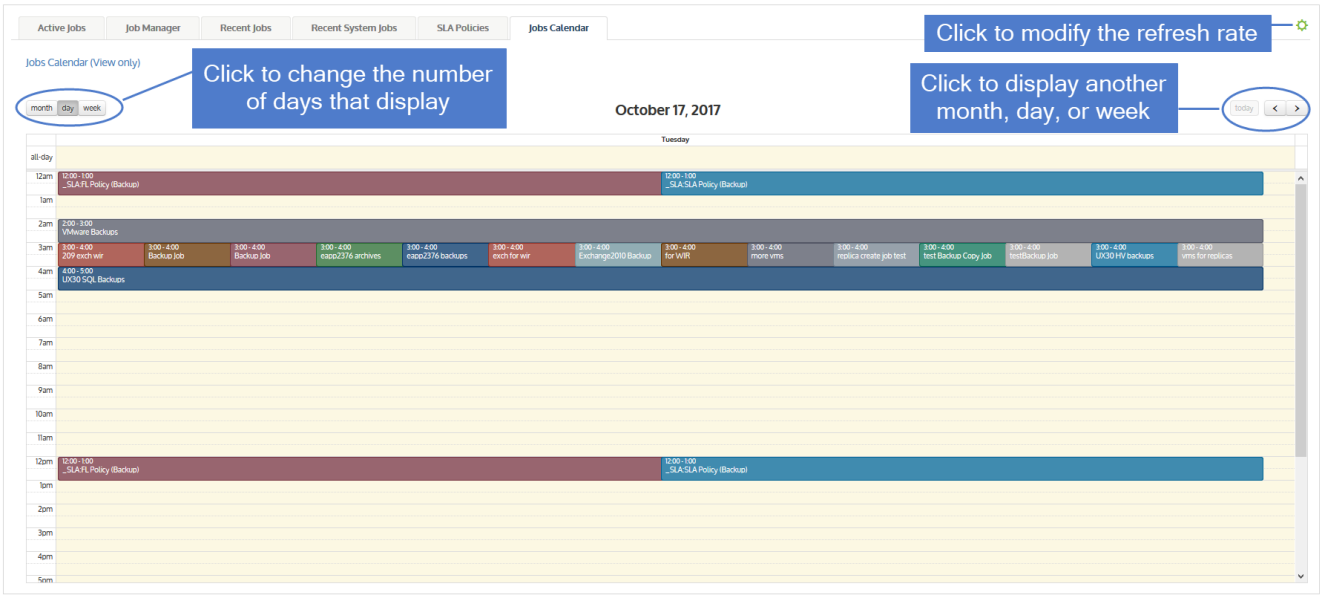

Week example:

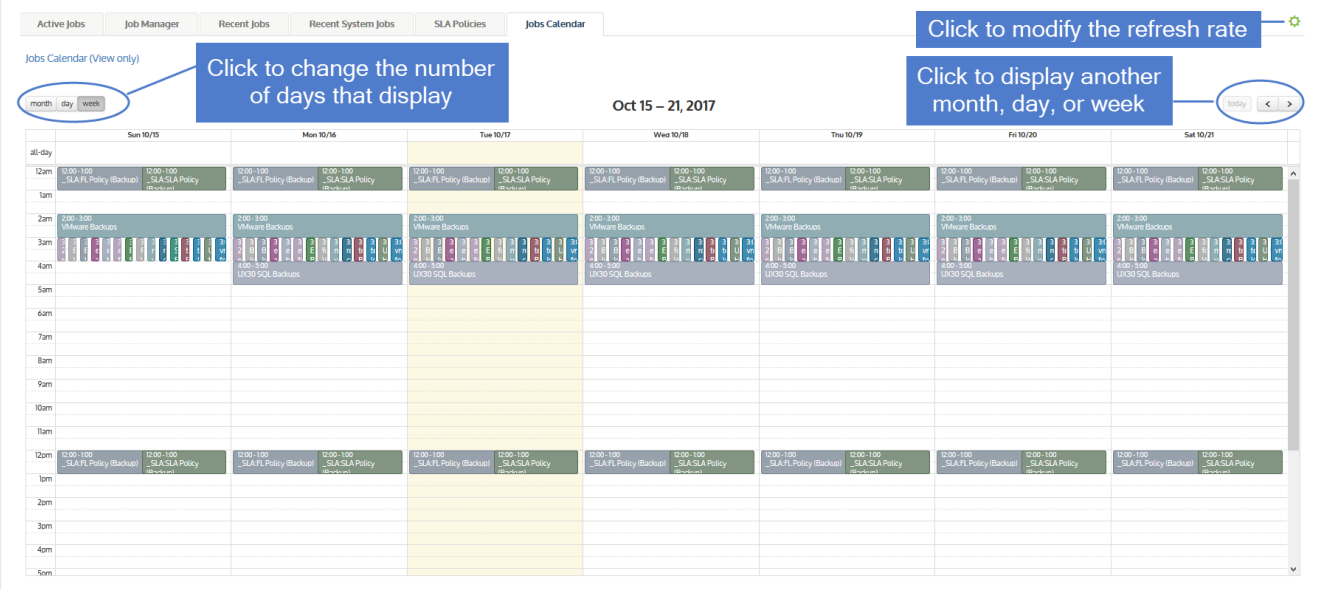

Reports

The Reports page enables you to run individual reports. The Reports page groups reports into categories. Select a category to view all the available reports. Click a report name to generate the report. Once generated, you can filter reports by any column and export them as a PDF or CSV file. For details, see "Working with reports" on page 1307.

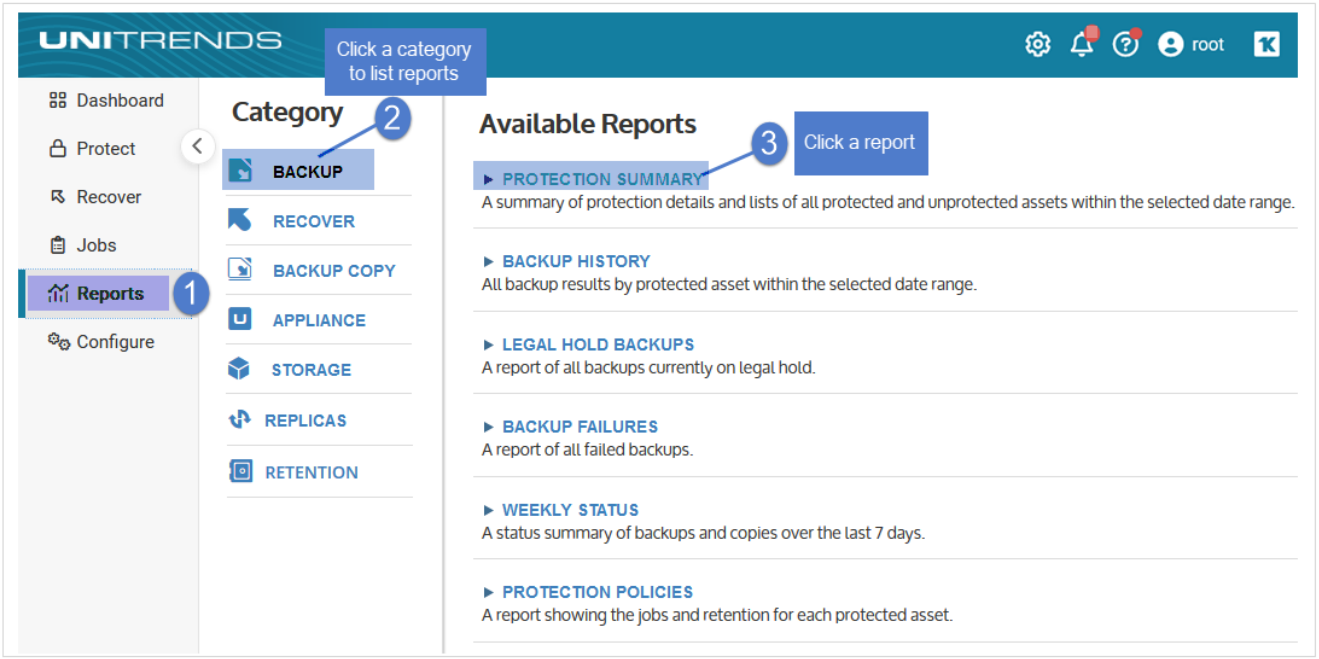

See the following for additional information on each report:

Category	Available Reports
"Backup reports" on page 1314	"Protection Summary report" on page 1331 "Backup History report" on page 1320 "Backup Failures report" on page 1322 "Weekly Status report" on page 1344 "Protection Policies report" on page 1345
"Recover reports" on page 1327	"Recovery History report" on page 1327 "Recovery Assurance report" on page 1329
"Backup Copy reports" on page 1331	"Protection Summary report" on page 1331 "Backup Copy Capacity report" on page 1336 "Backup Copy - Hot Targets report" on page 1337 "Backup Copies - Past 24 Hours report" on page 1340 "Storage Footprint report" on page 1341 "Backup Copy - Cold Targets report" on page 1342 "Weekly Status report" on page 1344 "Protection Policies report" on page 1345
"Appliance reports" on page 1347	"Update History report" on page 1347 "Capacity report" on page 1348 "Load report" on page 1351 "Alerts report" on page 1352 "Trap History report" on page 1353 "Notifications report" on page 1355
"Storage reports" on page 1356	"Storage report " on page 1356 "Data Reduction report" on page 1359
"Replicas History report" on page 1360	"Replicas History report" on page 1360
"Retention Reports" on page 1363	"Legal Hold Backups report" on page 1363 "Long-Term Retention report " on page 1365 "Min-Max Retention report" on page 1368
"Compliance report " on page 1370	"Compliance report" on page 1370
"Email reports " on page 1372	"Appliance Status report" on page 1372 "Backup Copy Hot Targets report" on page 1374 "Backup Copy Job Notifications report" on page 1376 "Compliance report" on page 1377

Category	Available Reports
	"DCA Job Notifications report" on page 1379 "Management Status report" on page 1381 "Schedule report" on page 1384 "Schedule Success report" on page 1387 "Schedule Failure report" on page 1389

Configure

Use these tabs on the Configure page to manage appliances and assets:

- "Appliances tab" on page 84
- "Protected Assets tab" on page 86
- "Copied Assets tab" on page 88

Appliances tab

From this tab you can view, add, modify, and remove appliances. Tasks are performed using the buttons across the top of the tab and the sub-tabs at the bottom. For detailed procedures, see "Appliance settings" on page 105 and "Remote Appliance Management" on page 417.

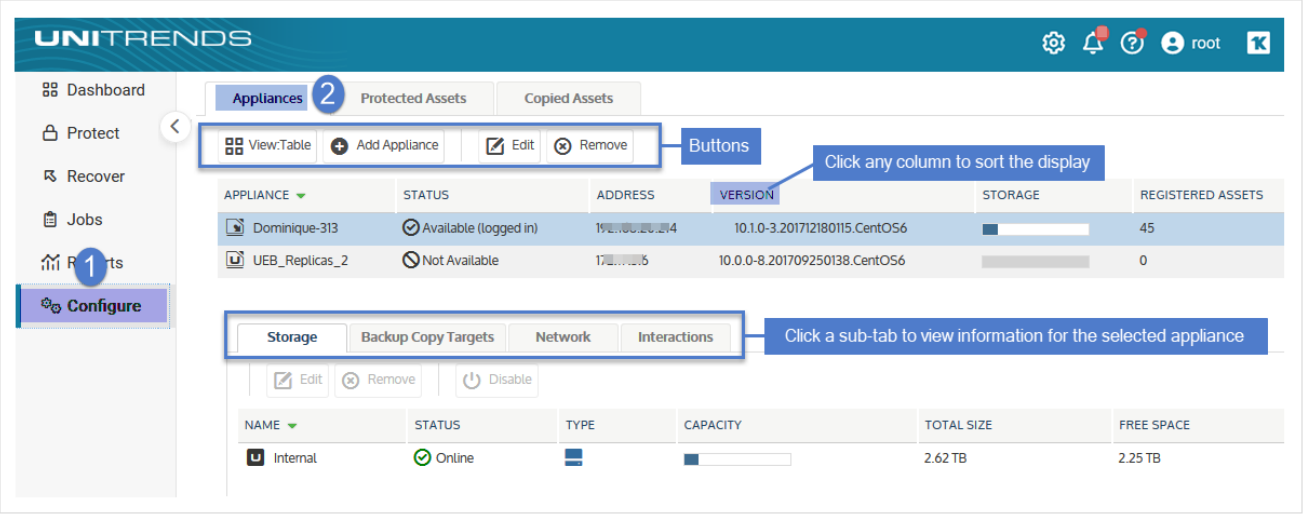

Appliance information

The Appliances tab displays the Unitrends appliance you are logged in to (called the *local appliance*), as well as any others it is managing or receiving backup copies from. The following information is provided for each appliance:

Column	Description
Appliance	Name of the Unitrends appliance.
Status	Appliance status: • *Available* indicates you can perform all management tasks for the appliance. • *Not Available* indicates the appliance is a backup copy source that cannot be managed from this UI. (To set up management for the source appliance, see "To manage remote appliances" on page 420.)
Address	Appliance IP address.
Version	Unitrends software version running on the appliance.
Storage	Total backup storage capacity. Hover to see amount used / total capacity.
Registered Assets	Number of assets that have been added to the appliance.

Appliance tab buttons

These buttons are available:

- View Table / View List – Changes the tab view. View appliances in a list or in a table.
- Add Appliance – Use to add an appliance so you can manage it from the UI of the local appliance.
- Edit – Use to edit the selected appliance. Modify various options, such as email, users, and date and time.
- Remove – Use to remove the selected appliance from the list.

Appliance sub-tabs

These sub-tabs are used to view and modify additional features of the selected appliance:

- Storage sub-tab – Use to add, edit, remove, and disable/enable backup storage. (Adding storage is supported for Unitrends Backup appliances only.) For more information, see "Backup storage" on page 196.
- Backup Copy Targets sub-tab – Use to add, edit, remove, disable/enable, erase, and scan for (discover) backup copy targets. For more information, see "Backup copy targets" on page 214.
- Network sub-tab – Use to view and edit network settings for each network adapter on the selected appliance, and to view and edit the appliance's hosts file and port security settings. For more information, see "Appliance network settings" on page 106.

- Interactions sub-tab – Use to add, edit, and remove the ConnectWise Professional Services Automation (PSA) tool, send test tickets, and view ticket history. For more information, see "ConnectWise PSA integration" on page 415.

Protected Assets tab

From this tab you can view, add, modify, and remove assets (the machines and applications you protect with your Unitrends appliance). Tasks are performed using the buttons across the top of the tab. For detailed procedures, see "Managing protected assets" on page 286.

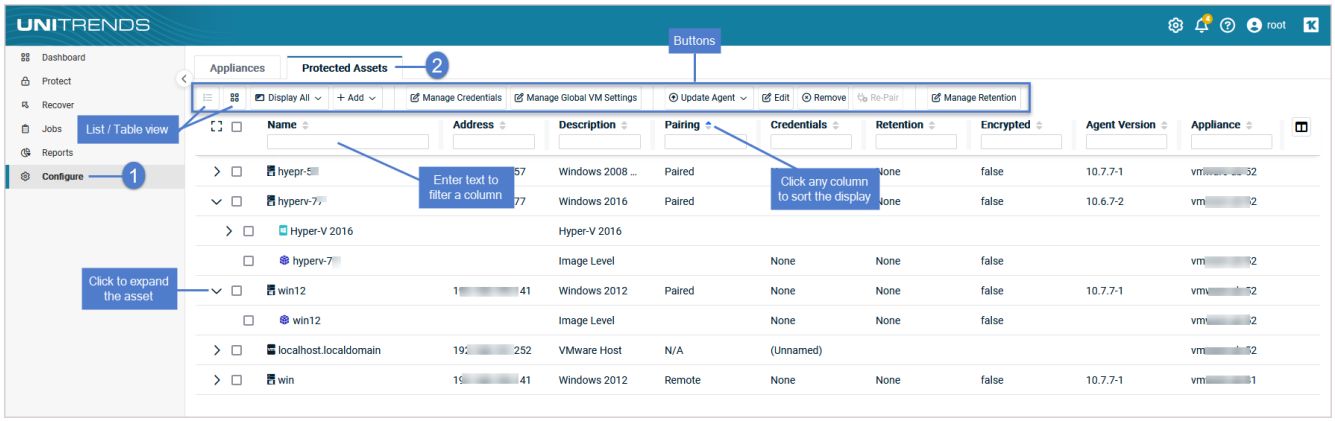

Asset information

The Protected Assets tab displays all assets that have been added to the local appliance and added to any managed appliances. Virtual hosts and physical servers display as top-level nodes in the list. To view individual virtual machines and applications, expand the virtual host or application server. The following information is provided for each asset:

Column	Description
Name	Name of the asset.
Address	IP address of the virtual host or physical asset.
Description	Description of the asset.
Pairing	Asset's secure agent pairing status. The asset's pairing status is updated when a backup runs, during an inventory sync, or any time the asset is re-saved. • To update the pairing status of all applicable assets, click and select **Inventory Sync**. • To update the pairing status of one asset, select the asset, click **Edit**, then click **Save** in the Edit Asset dialog. The asset's last pairing status displays on the Protected Assets tab. Statuses include:

Column	Description
	• Paired – The agent has been paired. • Failed – The agent is not paired. Hover to see the error message. One of these errors has occurred: – The agent pairing time window has expired. – The agent can't save pairing keys. – The agent pairing request failed. • Unsupported – Pairing is disabled or the agent pairing version is not compatible. • N/A – Not applicable for this asset. • Remote – This asset resides on a managed appliance. Log in to its appliance directly to see the asset's pairing status. **Notes:** • For agent pairing requirements, see "Windows agent requirements" on page 362 or "Requirements for secure pairing of Unitrends Linux agents" on page 391. • For agent pairing procedures, see "Secure agent pairing for Windows and Linux agents" on page 338.
Credentials	Indicates whether credentials have been assigned to the asset: • *None* indicates that no credentials have been assigned. • *Unnamed* indicates that credentials have been assigned directly to the asset, but have not been set up in the Manage Credentials dialog. • The credential name displays if the credential has been applied and set up in the Manage Credentials dialog. For details, see "Managing asset credentials" on page 322.
Retention	Asset's retention policy. *None* if no policy has been applied. For details, see "Managing retention with long-term data management" on page 328.
Encrypted	Indicates whether the asset's backups are encrypted: *True* if encrypted, *False* if not encrypted. For details, see "Encryption" on page 155.
Agent Version	Unitrends agent version running on the asset.

Column	Description
Appliance	Name of the Unitrends appliance that is protecting the asset.

Asset tab buttons

These buttons are available:

- View Table / View List – Changes the tab view. View assets in a list or in a table.
- Display All / Display Virtual / Display Physical – Use to filter the assets that display.
- Add – Use to add an asset to the local appliance or to a managed appliance. For details, see "Protected assets" on page 279.
- Manage Credentials – Use to add, edit, and delete credentials. After creating a credential, you can apply it to an asset. For details, see "Asset credentials" on page 282.
- Manage Global VM Settings - Use to choose the quiesce setting that will be applied to all newly discovered VMware, AHV, and XenServer VMs on the selected appliance. You can also opt to apply this setting to current VMs. For details, see "Quiesce settings for host-level backups" on page 283.
- Update Agent - Use to install Windows agent updates on all Windows assets or on selected Windows assets. For details, see "Push-installing the Windows agent" on page 365.
- Edit – Use to edit the selected asset. Modify various options, such as encryption, credentials, and retention. For details, see "Protected assets" on page 279.
- Remove – Use to remove the selected asset from the appliance. For details, see "Managing protected assets" on page 286.
- Manage Retention – Use to add, remove, or edit retention policies. For details, see "Configure" on page 84.
- Re-Pair – Use to pair an asset with its appliance if the asset's pairing status is *Failed* or *Unsupported*. For details, see "Secure agent pairing for Windows and Linux agents" on page 338.

Copied Assets tab

Displays only for local appliances that are receiving backup copies from another Unitrends appliance. The tab lists all assets whose backup copies are stored on the local appliance. From this tab you can view, edit, and remove copied assets by using the buttons across the top of the tab. For detailed procedures, see "Copied Assets" on page 403.

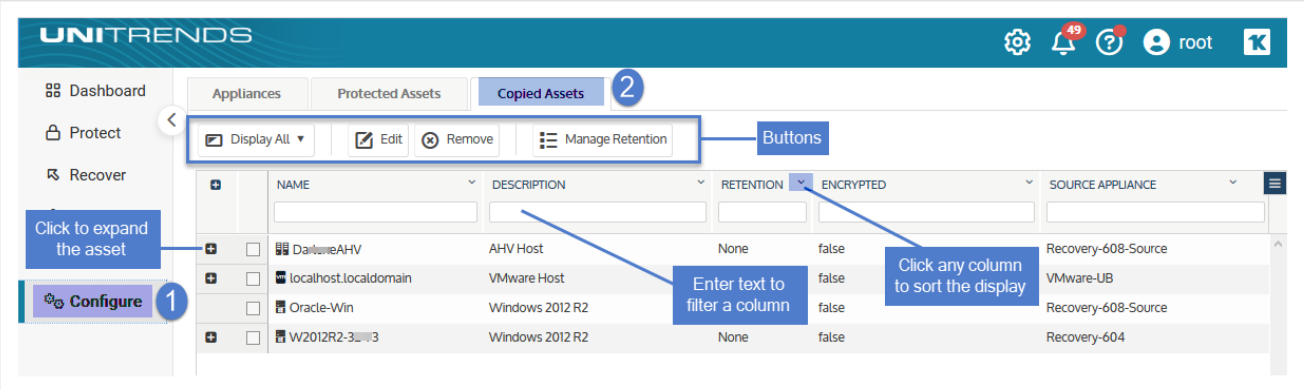

Copied asset information

The following information is provided for each copied asset:

Column	Description
Name	The name of the copied asset.
Description	Description of the asset.
Retention	The copied asset's retention policy.
Source Appliance	The name of the appliance that manages the asset and sends backup copies.

Copied asset tab buttons

These buttons are available:

- Display All / Display Virtual / Display Physical – Use to filter the assets that display.
- Edit – Use to apply a retention policy to the selected copied asset. (Supported only if the local appliance is managing the backup copy source appliance.) For details, see "Managing retention with long-term data management" on page 328.
- Remove – Use to remove the selected copied asset from the appliance. For details, see "Managing protected assets" on page 286.

Dialog help

While working with the dialogs in the UI, click the **?** in the upper-right corner for detailed instructions. Example:

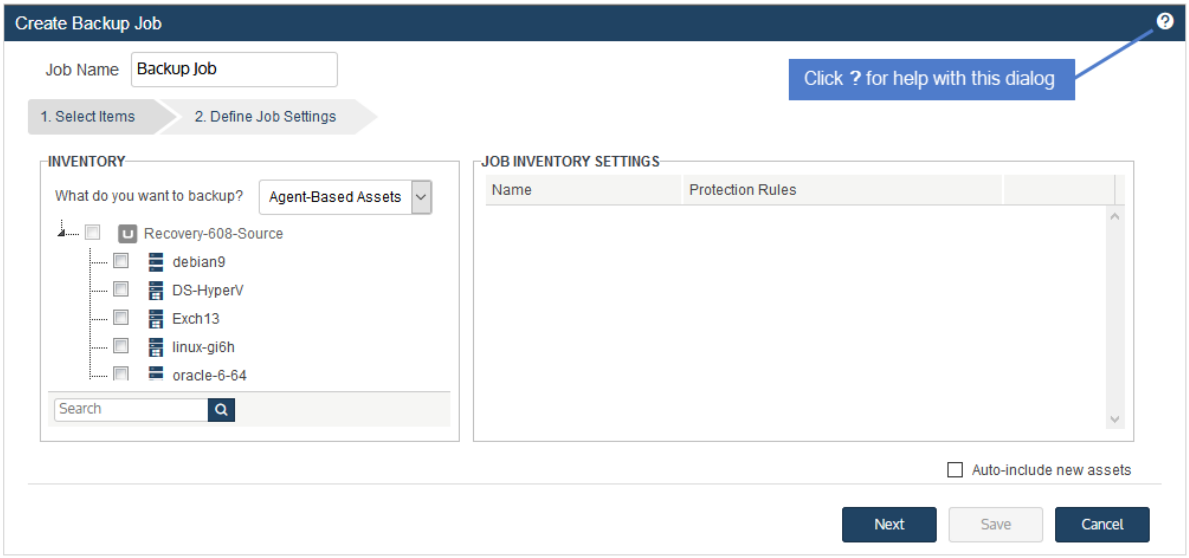

Procedures related to working with the dialog display:

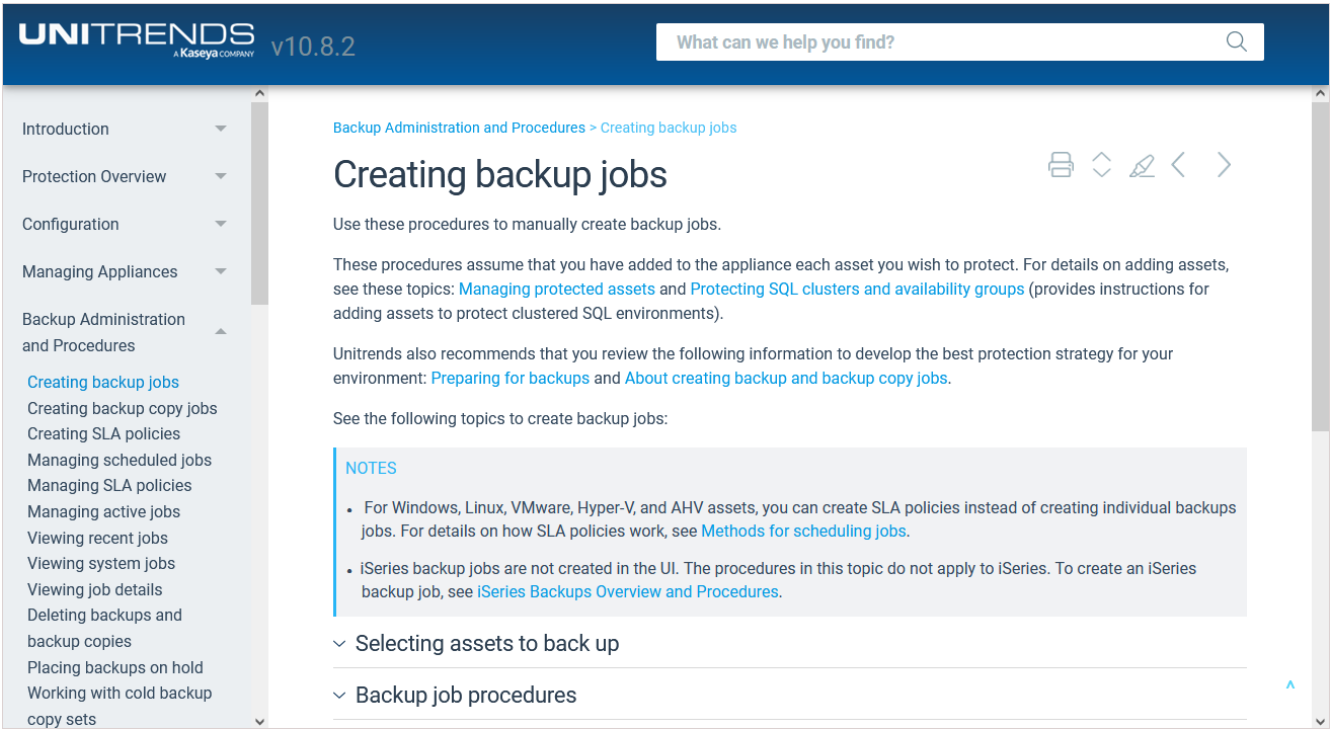

Chapter 2: Protection Overview

Unitrends Recovery Series, Recovery MAX, ION/ION+, and Unitrends Backup appliances provide comprehensive data protection for a wide range of:

- Operating systems
- Applications
- Hypervisors
- NAS devices

Any resource protected by Unitrends is called an *asset*. For a complete list of assets your appliance can protect, see the Unitrends Compatibility and Interoperability Matrix on the Unitrends website.

This chapter introduces you to Unitrends protection, providing an overview to help you determine which features will work best for your environment and your business continuity plan. For details about the protection options described here and instructions for using them, see the other applicable sections of this guide.

These key components of Unitrends data protection are described in these topics:

Data protection best practices

All data protection strategies begin with local backups on your appliance. Backups are duplicates of your data, and can run in several modes. Depending on the mode you specify, they capture all data for an asset, or a subset of data that has changed since the last backup. Each backup functions as a recovery point for the protected asset.

After you've backed up your assets, you can recover individual files, databases, file systems, entire machines, or use the instant recovery features to recover critical machines in minutes. We strongly recommend that you also make off-site copies of your local backups in order to recover from a disaster. See "Backup copies" on page 101 for details.

Customize your backup strategy to meet the recovery point objectives (RPOs) and recovery time objectives (RTOs) required for your business continuity plan. RPOs and RTOs refer to the maximum amount of data loss and downtime that you can tolerate. For example, if you can tolerate losing a day's worth of data, your RPO is one day. If you can tolerate only 30 minutes of downtime, your RTO is 30 minutes. RPOs and RTOs can vary per asset, and Unitrends offers different backup and recovery options to ensure that you meet these goals.

To meet your RPOs, use SLA policies or custom schedules to create backups at the desired frequency. To meet your RTOs, use retention policies to control the number of recovery points available on your appliance and instant recovery

to quickly spin up critical machines. Use backup copies stored on an off-site target for long-term retention and disaster recovery.

Unitrends supports a number of backup modes to ensure flexible protection policies for various types of data. A single job can use one backup mode, but your appliance can leverage multiple backup modes across various jobs. To be sure you see the full benefits of Unitrends best-in-class deduplication, be sure to run many backups. The more backups, the better the deduplication ratio.

Use the table below to choose the best mode for your environment. Modes are described in more detail in "Backup modes" on page 95.

Ranking	Backup mode	Benefits
Best	Incremental Forever	• Provides the fastest backup window after the first full backup. • Recommended for VMware, Hyper-V, and most file-level backups. • Reads the full disk once and then processes only changes going forward. • Can create schedules by using SLA policies. (All backup schedules created by SLA policies use the Incremental Forever backup mode.)
Better	Full / Incremental	• Recommended for Exchange backups. • Recommended if you want to control when full backups are taken for the purpose of backup copy management. • Recommended when you want to force a full read of all data periodically. • Inline deduplication ensures that even full backups only write changes to the backup storage.
Good	Full / Differential	• Recommended for SQL backups with additional transaction log protection for RPOs as low as one-minute. • Recommended if you want to simplify recovery of backup copies from tape at the expense of longer backup copy times compared to full / incremental. • Inline deduplication ensures that even full backups only write changes to the backup storage.
Okay	Fulls	• Recommended for Citrix XenServer backups.

Ranking	Backup mode	Benefits
		• Recommended when RPOs are very long (one week or longer). • Can be used with Incremental Forever if you only want full backups to be periodically copied to backup copy storage. • Inline deduplication ensures that even full backups only write changes to the backup storage.

Types of data protected

Unitrends protects over 100 versions of servers, storage, operating systems, hypervisors, and applications. To protect such a wide variety of assets, the appliance supports several backup methods.

The appliance runs backups based on the backup jobs or SLA policies that you create. The first step in creating a backup job is selecting the *type* of backup you want to run (for example, *File Level* or *Hyper-V*). The first step in creating an SLA policy is selecting the *type* of asset you want to protect (for example, *agent-based assets* or *VMware assets*). In both cases, the type you select determines which backup method the appliance uses and the type of backup that is created. For more on backup jobs and SLA policies, see "Preparing for backups" and "About creating backup and backup copy jobs" on page 426.

Following is a description of each Unitrends backup type.

- *File-level backups* protect an asset's file system and operating system. You must install a Unitrends agent on the asset to run file-level backups.

  Note: For Windows, you can also run bare metal backups by using the Windows bare metal agent. A bare metal backup is used for disaster recovery only. In most cases, a bare metal backup is not needed because file-level backups can be used to recover the machine (this is the recommended approach). But in some cases a bare metal backup must be used instead. To determine whether bare metal backups are needed for your asset, see "Windows Bare Metal Protection and Recovery" on page 1207.

- *Image-level backups* protect a Windows asset at the disk and volume level. You must install the Unitrends Windows agent on the asset to run an image-level backup.

  Note: You can opt to protect a Windows asset with file-level backups, image-level backups, or both backup types. The Windows agent supports both backup methods.

- *Host-level backups* protect VMware, Hyper-V, Nutanix AHV, and XenServer virtual machines by leveraging hypervisor snapshots. You do not need to install a Unitrends agent on hosted VMs.
- *Application backups* capture an application's structure and data to ensure database consistency. You must install a Unitrends agent on the host asset to run application backups.
- *NAS backups* protect data stored on a NAS device. You do not install an agent on the NAS asset.

- *iSeries backups* protect an asset's file system by leveraging native iSeries backup operations. You do not install an agent on the iSeries asset.

Backing up physical assets and hosted applications

Physical assets are protected with file-level or image-level backups. Hosted applications are protected with application backups. Physical assets are also called *agent-based assets* because a Unitrends agent must be installed on the asset to run backups.

Backing up virtual assets

For virtual assets, you can choose host-level or file-level protection. Host-level backups capture files, application data, and virtual hardware. With file-level protection, the appliance treats your VM as a physical asset to run file-level and application backups.

The table below compares the backup options for virtual assets. Host-level backups are recommended in most cases, but there are VMs for which you will want or need to use file-level protection. For considerations specific to your environment, see "Protecting VMware virtual machines with file-level backups" on page 674, "Protecting Hyper-V virtual machines with file-level backups" on page 661 , "Protecting AHV virtual machines with file-level backups" on page 698 and "Best practices and requirements for XenServer protection" on page 689 to determine which approach to take.

Host-level protection	File-level protection
Add the virtual host to your appliance and it detects all the VMs on the host. It is not necessary to install agents on VMs or add VMs to the appliance individually. This greatly simplifies protecting large virtual environments.	You must install agents on the VMs and add each one to the appliance individually.
Backups capture all data on the VMs. You can exclude entire disks (VMware, AHV, and XenServer only), but you cannot exclude files, directories, or volumes.	You can choose to protect all of the asset's data or select only particular files, directories, or volumes.
You can recover virtual machines in minutes using the VM instant recovery feature (VMware and Hyper-V only).	You can recover Windows machines in minutes using the Windows replica feature.
You can recover individual files from backups for VMs running Windows or Linux.	You can recover individual files from backups for any supported operating system. You can recover individual items from application databases.

Backups

Unitrends uses backups to create recovery points for your data. Backups are run in different modes and are organized into backup groups. Your backup strategies determine which modes you will use.

Unitrends backups fall into two general categories: local backups and backup copies. Local backups are stored on the appliance. These backups are immediately accessible and enable you to meet low RTOs. Backup copies are stored on

an off-site target. These backups are duplicates of your local backups, and are used for long-term retention and disaster recovery.

Backup modes

Backup modes determine what data to include in the backup. These modes protect all types of data and apply to all backup types (file-level, image-level, host-level, application, NAS, and iSeries).

While Unitrends supports a variety of backup modes that give you flexibility in protecting your assets, not all backup modes are supported for all assets. When manually creating a backup job for a given asset, only supported modes are available for selection. (For SLA policies, the appliance uses the incremental forever strategy to automatically create jobs. You do not select the modes to use.)

While creating backup jobs, you can select these backup modes: full, incremental, differential, selective, and bare metal (Windows only). In addition to these, the appliance automatically creates synthetic backups as needed. See the following for a description of each:

- "Full backup" on page 95
- "Incremental backup" on page 96
- "Differential backup" on page 96
- "Selective backup" on page 97
- "Windows bare metal backup" on page 97
- "Synthetic backup" on page 97

Full backup

A full backup captures all data on the asset:

- For file-level backups run with a Unitrends agent, this includes all file system and operating system data required to recover the asset. You can specify data to include or exclude from the full backup.
- For Windows image-level backups run with a Unitrends agent, this includes all disks and volumes. You can specify volumes to include or exclude from the full backup.
- For host-level backups, this includes VM metadata (configuration files) and blocks of all disks attached to the VM. For VMware, AHV, and XenServer, you can specify disks to exclude from the full backup.
- For application backups, all data is included in a full backup.
- For NAS backups, this includes all eligible data stored on the NAS device (see "NAS protection using CIFS/NFS" on page 726 or "NAS protection using NDMP" on page 727 for details on which items are automatically excluded from backup).
- For iSeries backups, this includes all eligible files, libraries, and objects (see "Requirements and considerations for iSeries protection" on page 767 for details on which items are automatically excluded from backup). The backup is of the filesystem and cannot be used to recover the asset. You can specify data to include or exclude from the full backup.
- A successful full backup must exist before a differential or incremental can run.

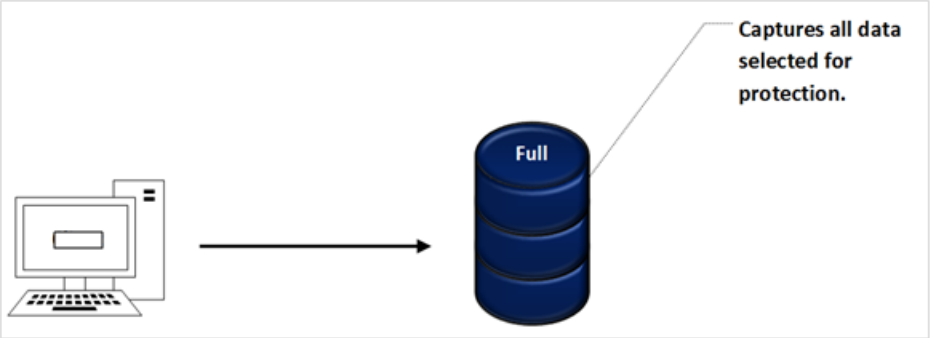

Incremental backup

An incremental captures changes in the protected data since the last successful backup (of any mode). Therefore, incremental backups are smaller and can run more quickly than full backups, but they depend on the previous backups.

The diagram below illustrates incremental backups for an asset. In this example, the incremental runs once a day, but you can schedule them more frequently if desired.

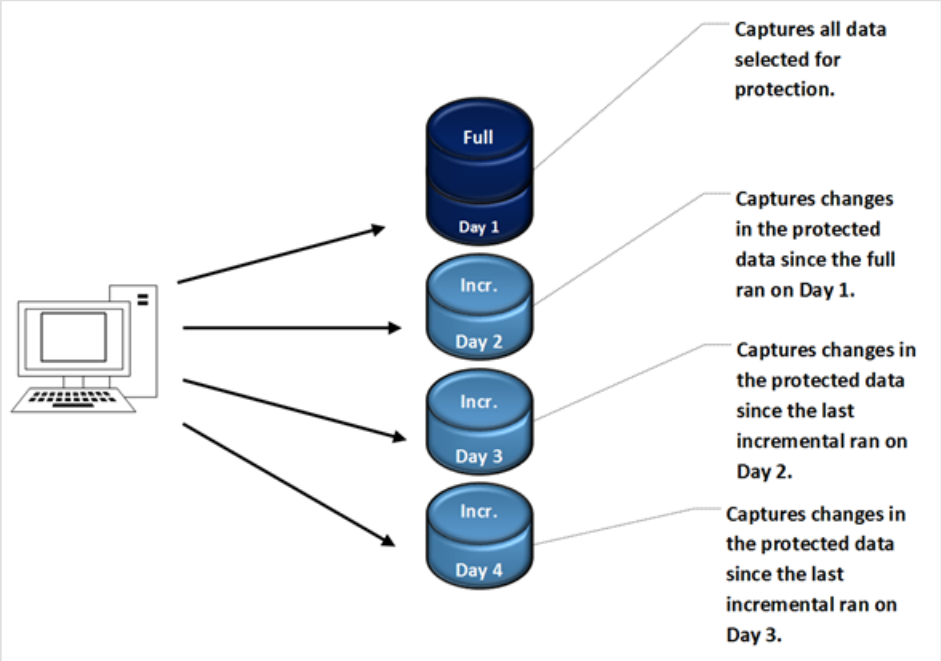

Differential backup

A differential captures changes in the protected data since the last successful full backup. The diagram below illustrates differential backups for an asset. Each differential captures all changes in the protected data since the full backup on Day 1. For example, the differential on Day 4 captures all changes since the full backup on Day 1, including the changes that were already captured by the differentials on Day 2 and Day 3.

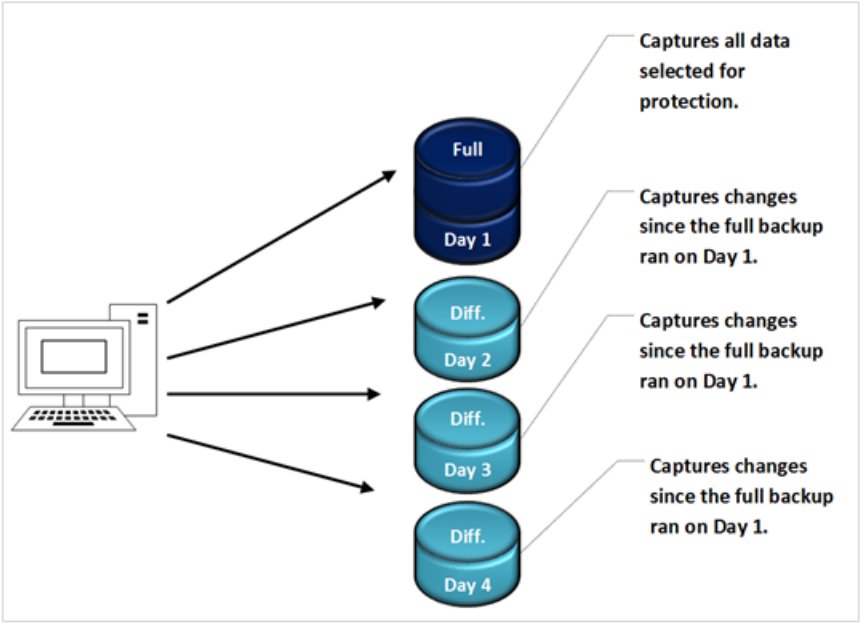

Selective backup

A selective backup is run independently of any full, differential, or incremental backup and captures only the data that you have selected. Selective backups can be used only for file-level backups.

Windows bare metal backup

A bare metal backup captures the asset's boot and critical system volumes and is used for disaster recovery only. In most cases, a bare metal backup is not needed because file-level or image-level backups can be used to recover the machine (this is the recommended approach). But in some cases a bare metal backup must be used instead. To determine whether bare metal backups are needed for your asset, see "Windows Bare Metal Protection and Recovery" on page 1207.

Synthetic backup

A synthetic backup is a full or differential backup that the Unitrends appliance synthesizes by superimposing the incremental backups on the last successful full backup. It then uses the synthesized backups to create recovery points for quick recovery. Synthetics are also used to create cold backup copies of file-level backups (file-level incrementals cannot be copied directly to cold targets).

The Unitrends appliance uses the following factors to determine when to create a synthetic backup:

- Amount of data being protected on the appliance
- Number of days from the last full backup
- Number of incremental backups since the last full backup
- Load on the appliance

Synthetic backups are created only for these backup types: file-level, image-level, and host-level backups of VMware, Hyper-V, AHV, and XenServer VMs. Synthetic backups are appliance-side only and do not impact the assets or networks. The diagram below illustrates a synthetic backup. For more information, see When are synthetic backups created.

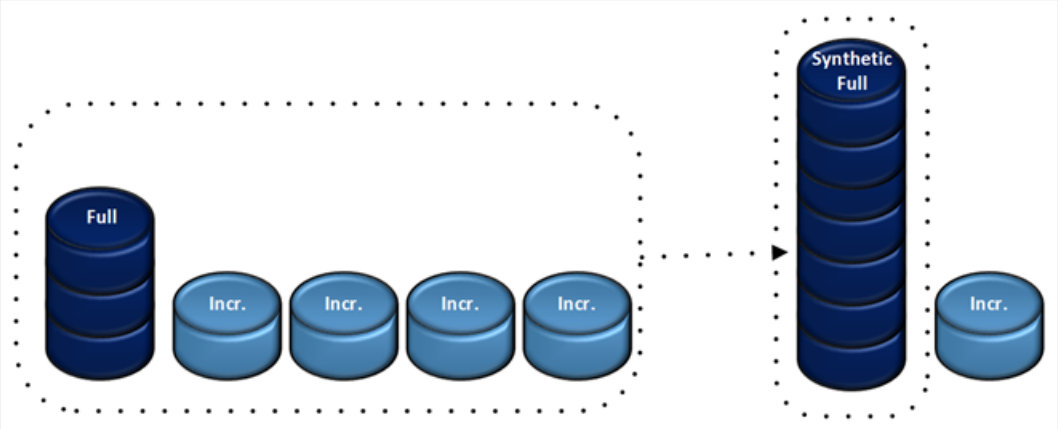

Backup groups

To protect your data, you will likely use a combination of backup modes. Your Unitrends appliance organizes an asset's backups into groups to manage any inter-dependencies between backups. The appliance creates a new group when it runs or synthesizes a full backup. Each subsequent differential or incremental forms a link in the chain of backups that constitute the group. Each link in the chain is necessary for data recovery.

Each backup is a recovery point of the asset at the point in time that the backup ran. To recover data, you select a single backup. You do not need to select any other dependent backups in the backup group. For details, see the Recovery chapters in this guide (listed in the "Recovery Overview" on page 775).

The following diagrams illustrate backup groups:

- "Incremental forever backup groups" on page 98
- "Groups with full, differential, and incremental backups" on page 99
- "Selective backup in relation to a group" on page 99

Incremental forever backup groups

The diagram below illustrates the incremental forever backup strategy for an asset. The strategy begins by automatically promoting the first scheduled incremental to a full backup. Thereafter, incremental backups run at the times specified in the job schedule. When the appliance determines a new full backup is necessary, it synthesizes a full backup and starts a new backup group.

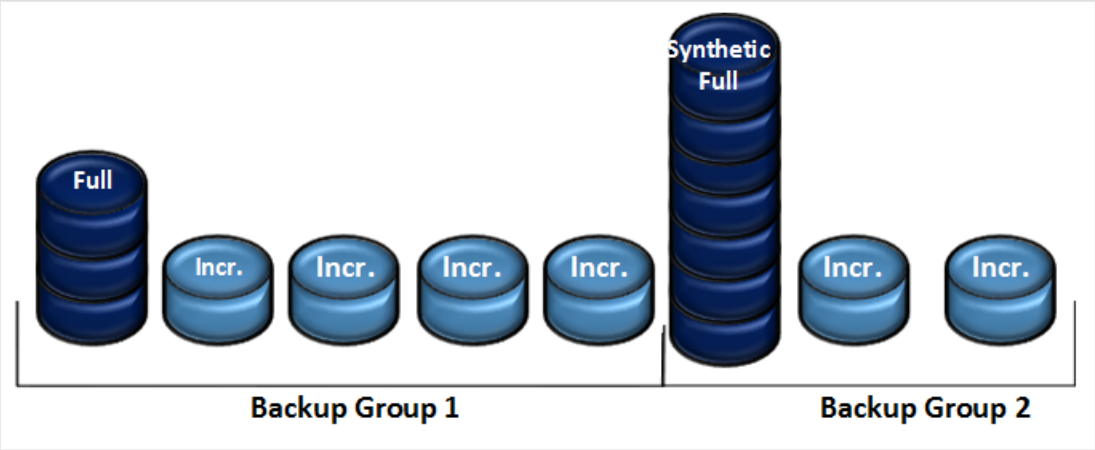

Groups with full, differential, and incremental backups

The diagram below shows two backup groups containing full, differential, and incremental backups.

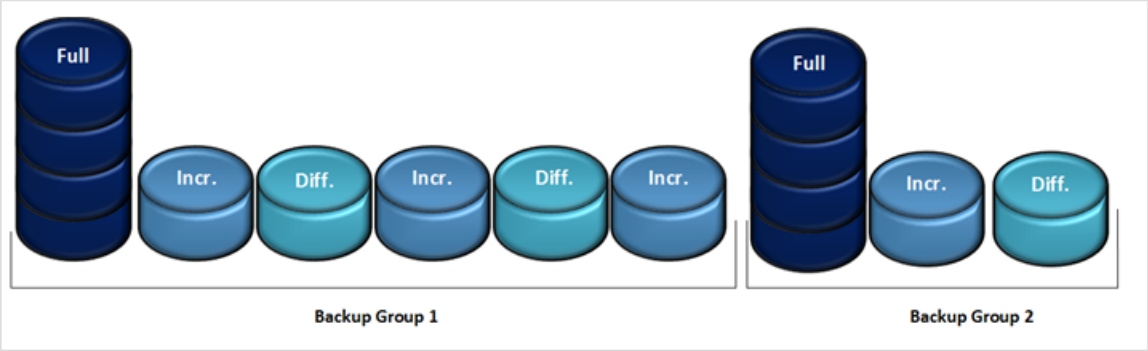

Selective backup in relation to a group

Selective backups exist independently of backup groups. The diagram below illustrates a backup group and selective backups for one asset. Both a selective backup and an incremental backup ran on Day 3. On Day 5, a differential backup and a selective backup ran. However, only the incremental and the differential belong to the group associated with the full backup run on Day 1.

Note: Selective backups are supported only for file-level protection.

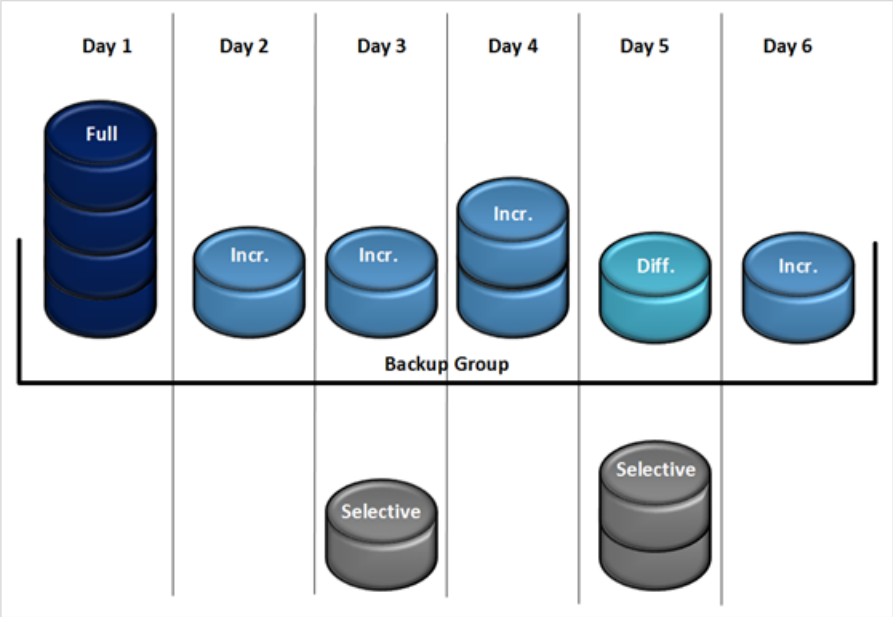

Backup strategies

A data protection strategy consists of utilizing one or more of the backup modes described above. For example, the incremental forever strategy consists of the full and incremental backup modes. Your strategy is implemented when you schedule jobs to occur at intervals and times you specify or define SLA policies to automatically create aligning schedules.

Unitrends recommends using the incremental forever backup strategy when possible, where an initial full backup is followed by incrementals at the frequency and times required to meet your RPOs.

In some cases, you will want to use a different backup strategy (such as weekly fulls with incremental or differentials). By manually creating backup schedules, you can customize your backups to fit any strategy, using the backup modes as desired. Examples of cases when you would not use the incremental forever strategy include:

- Protecting assets for which incrementals are not supported (such as Exchange and SharePoint applications, VMware hardware version 4 VMs, and VMware templates).
- Needing to control when full backups are run. (In most cases, this is not an issue since synthetics are run locally on the appliance and do not impact network or asset performance. But you may choose to schedule weekly fulls if appliance resources are taxed at certain times of the day or week.)

Storage space and backup retention

Your protection strategy should include plans for retaining the necessary local backups to meet your RPOs and RTOs. The most comprehensive strategy involves retaining recent local backups for quick recovery, and copying these backups to an offsite target for long-term retention and disaster recovery.

To create space for new backups, Unitrends appliances periodically purge older backups. Retention policies control how long backups remain on the appliance. Backups held by a policy are never purged. New backups fail if an appliance cannot purge older backups to create sufficient space.

Retention settings assure that the necessary recovery points are available on your appliance. Appliances are configured with a default backup retention policy of 30 days. This 30-day policy is applied to each protected asset.

Notes:

- The 30-day default retention policy applies to appliances imaged with release 10.7.8 or higher. This default policy does not apply to appliances that were originally imaged with an earlier release. Upgrading an appliance that was imaged with a pre-10.7.8 release does not modify its retention policies in any way.
- The 30-day default retention policy ensures that 7 daily backups and 4 weekly backups are retained for each protected asset.

You can also create your own retention policies to hold backups for a specified number of days. You can create multiple policies and customize them to achieve different RPOs and RTOs for your assets. See "Managing retention with long-term data management" on page 328 for details.

The amount of total backup storage capacity on the appliance varies by appliance type:

- Recovery Series, Recovery MAX, and ION/ION+ physical appliances come with a set amount of backup storage. You cannot add backup storage to the appliance.
- Unitrends Backup virtual appliances are deployed as virtual machines. During deployment, the initial backup storage was created using either a virtual attached disk, a SAN LUN, or a NAS share. After initial deployment, you can add more backup storage as desired. See "About adding backup storage to a Unitrends Backup appliance" on page 199 for details.

This storage capacity is used to store local backups, for VM instant recovery, and for Windows replicas. To use the instant recovery or Windows replicas features, you must reserve a portion of this storage to be used for instant recovery write space. For more on storage, see "Backup storage" on page 196.

Backup copies

Backup copies are duplicates of your backups that are stored off-site. Unitrends recommends having a second copy of your backups in order to recover from a disaster. You can copy your backups to the following types of targets:

- Unitrends Cloud
- A secondary Unitrends appliance
- Cloud storage (managed by Amazon, AWS, Google, or Rackspace)
- Disks, NAS devices, and other media that can be stored off-site

To copy backups, you add the backup copy target to your Unitrends appliance, then create a job or SLA policy that defines which backups to copy and other options. Backups are then copied according to the settings you defined. Once the backup copy target is full, the appliance does one of the following:

- If the *Delete older backup data to free space* option is selected in the backup copy job, removes older backup copies to make room for new ones.
- If the *Fail backup copy job and send alert* option is selected in the backup copy job, fails the backup copy job without removing older copies or copying any new backups.

Depending on the type of target selected, the appliance creates either hot or cold backup copies. Hot backup copies reside on the Unitrends Cloud or on a secondary appliance. Cold backup copies reside on cloud storage managed by other various storage providers and on other backup copy media that can be stored offsite.

Backup copies support the same recovery operations as local backups. However, because recovering data from copies requires additional steps, local backups should be used whenever possible to meet low RTOs.

For instructions on creating and managing backup copy targets and jobs, see the following topics:

- "Backup copy targets" on page 214 for procedures used to configure and add backup copy targets to your Unitrends appliance.
- "Backup Administration and Procedures" on page 425 for procedures used to create and manage backup copy jobs.

Recovery

After successfully backing up your assets, you have different options for recovering individual files, databases, or entire assets. All backup recovery options can be performed using local backups or backup copies. However, if you are using a cold backup copy for the recovery, you must first import the backup copy to a Unitrends appliance.

For critical virtual machines and Windows physical servers, you can set up a virtual replica that can be spun up in minutes in the event of asset failure. This replica performs just like the original asset, so production downtime is reduced to just minutes. With Windows replicas, the replica continues performing the role of the failed asset until you can get a new physical Windows server deployed. With VMware and Hyper-V instant recovery, you can use the replica until you deploy a new VM or keep using the replica VM itself.

For more on these recovery options, see:

- "Recovering Host-level Backups" on page 793
- "Recovering File-level Backups" on page 925
- "Recovering Windows Image-level Backups" on page 1031
- "Recovering Application Backups" on page 1147
- "Recovering NAS Backups" on page 1121
- "Recovering iSeries Backups" on page 1205

Ransomware detection

Unitrends uses a series of predictive analytics to identify possible ransomware infection on your protected assets. Ransomware attacks are commonly associated with aggressive rewrite activity and large influxes of randomized data.

If both of these metrics suddenly exceed baseline levels established by the appliance, an alert is triggered that can be viewed in the "Global menu" on page 36 at the top of the UI:

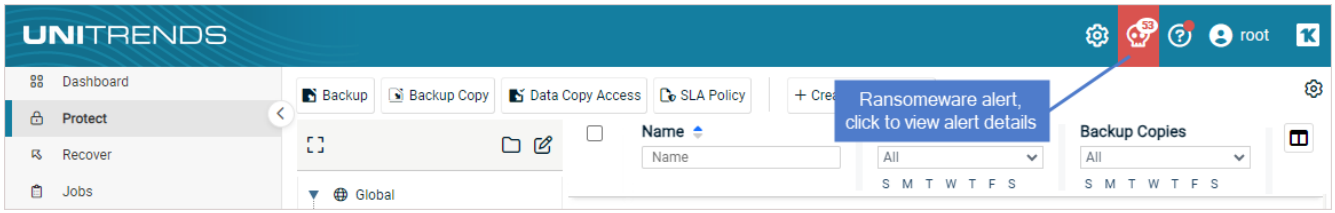

Click the ransomeware alert to view details. Unitrends recommends that you confirm the presence of malware on the identified asset and recover to the most recent safe backup. The "Backup History report" on page 1320 identifies any backups that are at risk of infection.

Notes:

- Ransomware detection functionality requires deduplication level 3. For further instructions on configuring deduplication levels, see "Configure deduplication settings on your Unitrends Backup appliance" on page 193.
- Sudden changes in the volume of compressed, encrypted, or media data in your protected environment may trigger ransomware alerts.

This page is intentionally left blank.

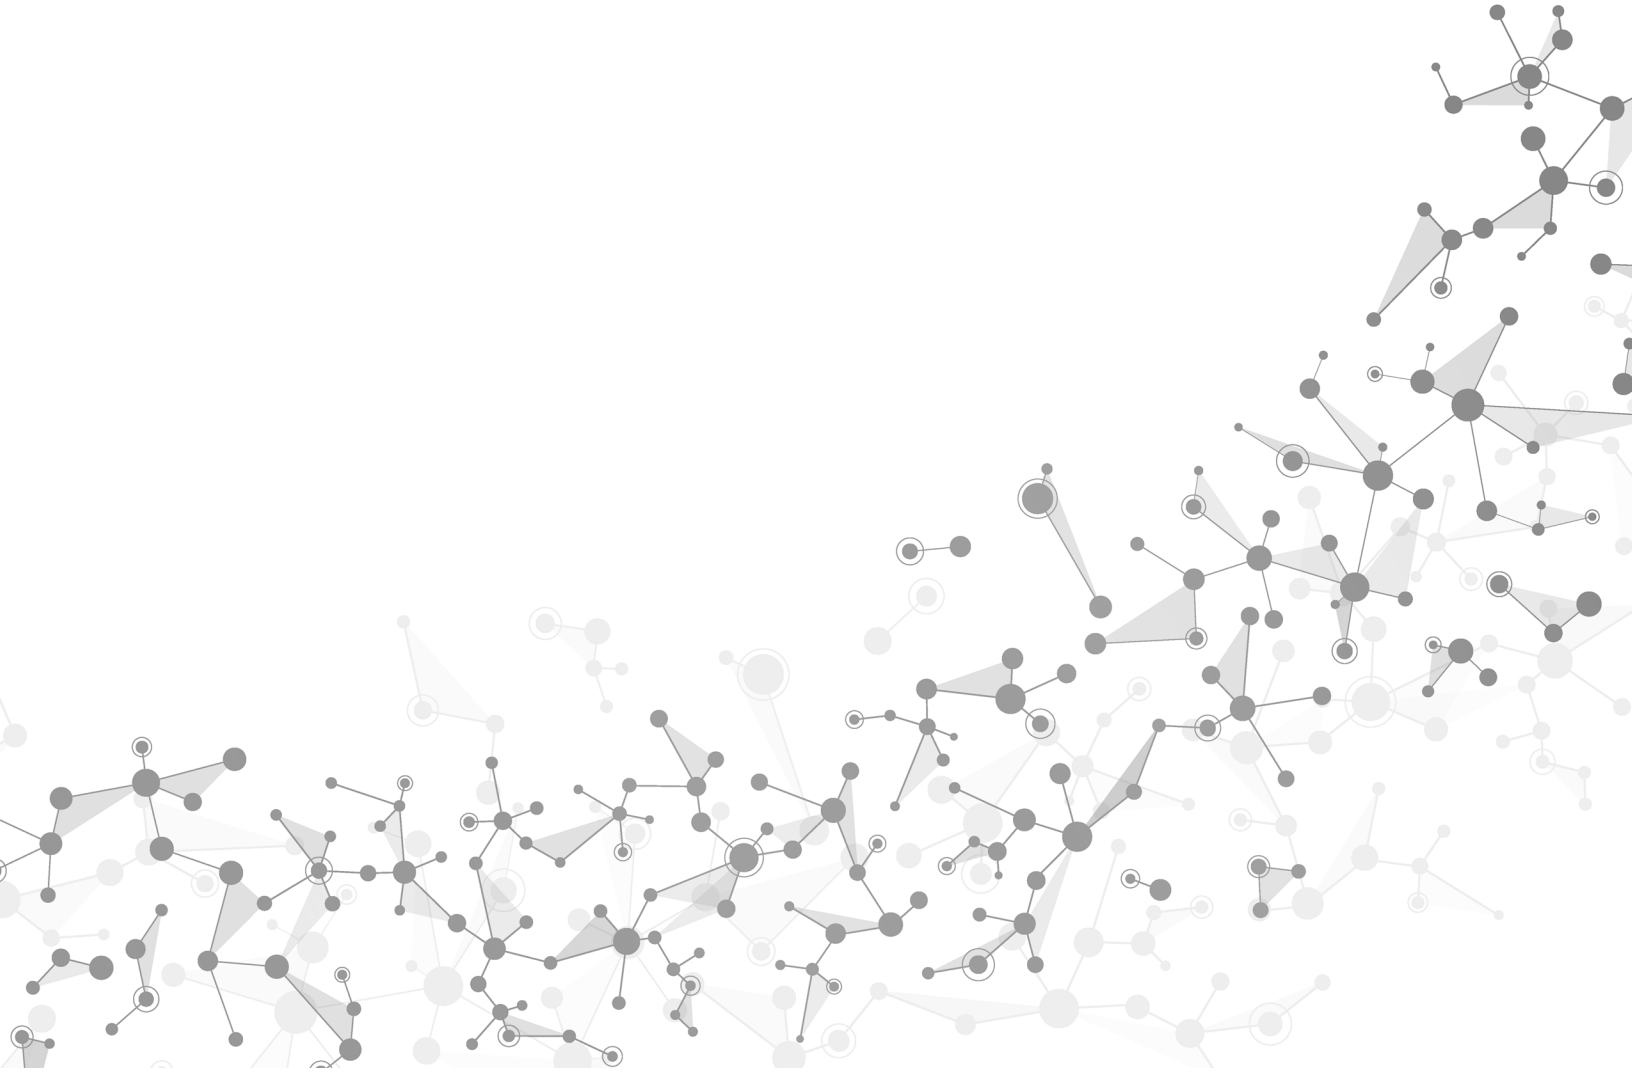

Chapter 3: Configuration

See these topics for instructions on performing administrative and configuration tasks for your appliances and protected assets.

Appliance settings

Use these procedures to configure appliance settings.

- "Configure deduplication settings on your Unitrends Backup appliance" on page 193
- "Set appliance language" on page 194
- "Appliance Samba share" on page 195

Note: To configure a Recovery Series, Recovery MAX, or ION+ physical appliance for SAN-direct VMware backups, see "VMware SAN-direct backups" on page 683. (SAN-direct backups are not supported on Unitrends Backup or ION appliances.)

Appliance network settings

There are several addresses you should permit for all deployments. All of these ports are outgoing connections from the Unitrends appliance. We do not require incoming NAT of ports or exposing the unit to a public IP, only outgoing communication from a local source Unitrends appliance is needed.

IMPORTANT! Never expose the appliance Web UI or SSH connections to open external ports. Doing so may void your support agreement until the appliance can be secured properly. Never deploy the Unitrends appliance on a public IP. All incoming ports to a Unitrends appliance must be firewall protected. Privately operated hot backup copy targets should be deployed in such a way as to secure the VPN connection to only trusted source external IPs.

During deployment, these network settings were configured for the appliance: IP address, subnet, gateway, primary DNS, hosts file, and open ports. Settings were either configured manually (if using a static IP address) or by DHCP.

- The IP address and subnet enable communication between the appliance and other machines on your network.
- The gateway enables communication between the appliance and machines on different subnets.
- Appliance DNS settings are required for the following:
  - To connect the appliance to the Internet.
  - To add assets using only their hostnames (rather than by fully qualified domain names).
  - To add backup copy cloud targets to the appliance.
  - To update your appliance from the UI.
  - To access the Unitrends Community forums from the UI.
- The hosts file enables communication between the appliance and its protected assets without using DNS. (But DNS is required for other features and must be set up on the appliance.) During deployment, the hosts file is created and contains an entry for the appliance itself. Additional entries are automatically added to this file any time you add an asset to the appliance or configure a secure tunnel connection (for backup copy to the Unitrends Cloud or to another Unitrends appliance). In most cases it is not necessary to modify this file.
- Port security controls which ports are open on the appliance. By default, the appliance is configured with all ports open (port security is set to *None Open All*). Other port security levels are available and you can close ports by applying one of these other levels.

You can modify the network settings described above as needed. See the following for details:

- "To view or edit network settings" to modify IP address, subnet, gateway, or DNS settings.
- "To view or edit the hosts file"
- "To view or edit port security settings" on page 113

In addition to the standard appliance network settings, additional ports must be open to connect to the Internet and to copy backups to a hot backup copy target. See "Additional port requirements" on page 115 for details.

To view or edit network settings

Notes:

- Before changing network settings, you should verify that no jobs are running. Changing network settings while a job is running causes the job to fail.
- Do not edit the network settings of an adapter that is configured to use DHCP unless you intend to assign it a static IP address.
- If the appliance is being used as a backup copy target, you must assign a static IP address. DHCP is not supported for appliance backup copy targets.
- If you change the IP address, you will no longer be able to access the appliance from a web browser using the previous IP address. To avoid losing web access to your appliance, make sure to assign it valid network settings and to make a note of these new settings.

1. On the **Configure > Appliances** page, select the appliance.
2. Click the **Network** tab below.
3. Select the adapter and click **Edit**.

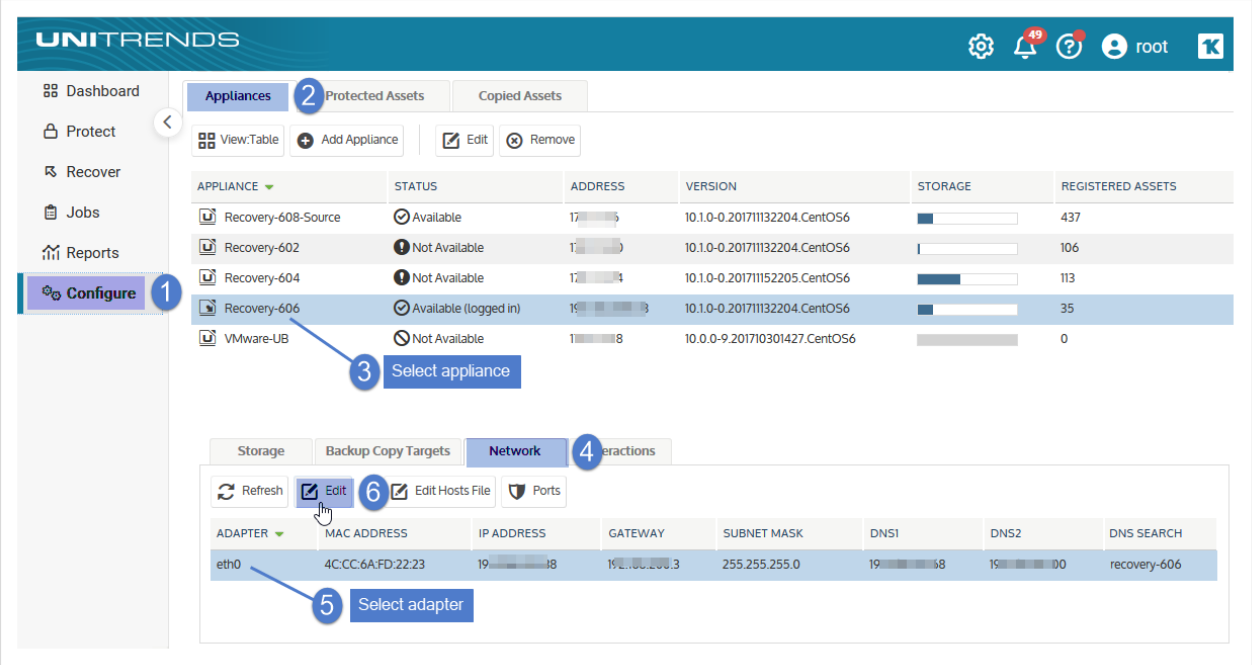

4 Enter the changes and click **Save**. See below for details.

- The options displayed in the Edit Network Adapter dialog differ by whether the appliance is running on CentOS 7 or CentOS 6. To check the Recovery OS Version, click on **? > About**, as shown here:

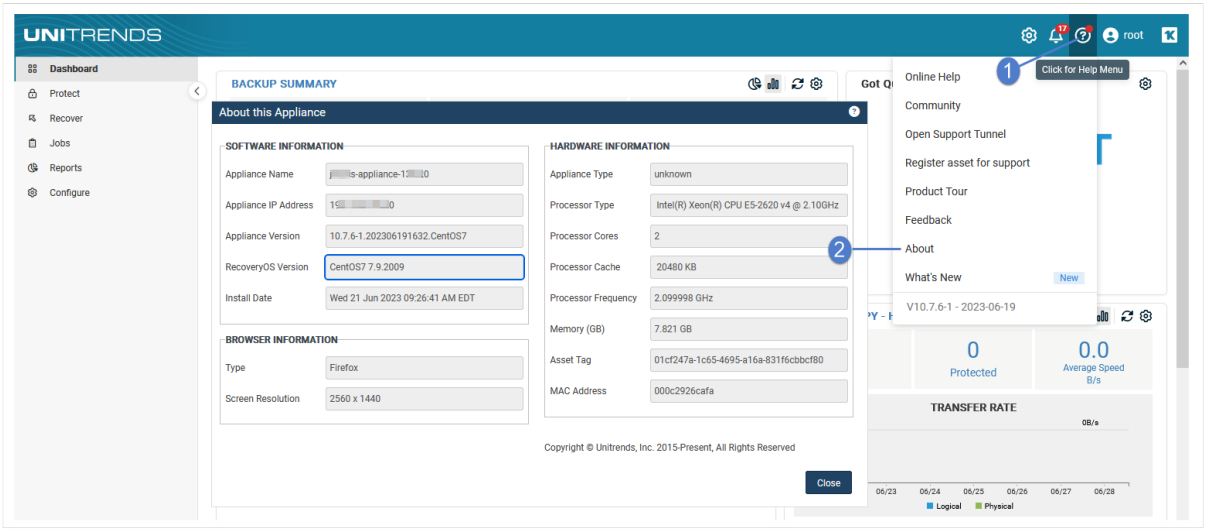

CentOS 7 appliances

For appliances running on CentOS 7, you can modify the following:

- To disable the adapter, select **Off** from the Configure Adapter Type list.

- If DHCP is available in your environment, you can switch from a static IP address to DHCP by selecting **DHCP** from the Configure Adapter Type list.
- To switch from DHCP to a static IP address, select **Static** from the Configure Adapter Type list and enter an IP Address, Gateway, Subnet, and DNS Settings.

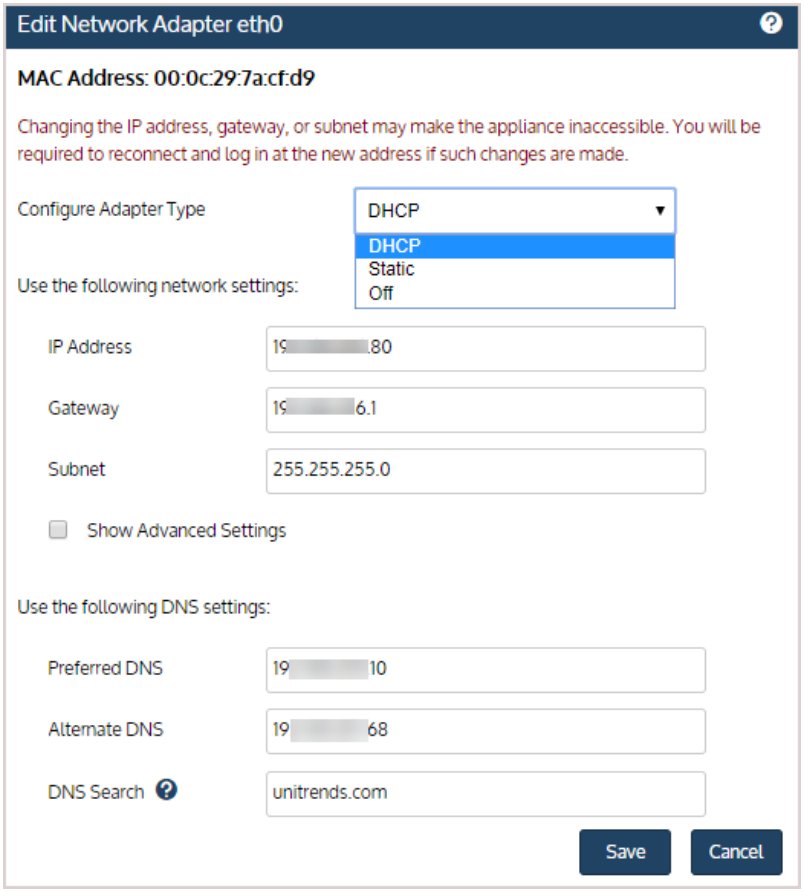

CentOS 6 appliances

For appliances running on CentOS 6, you can modify the fields shown below:

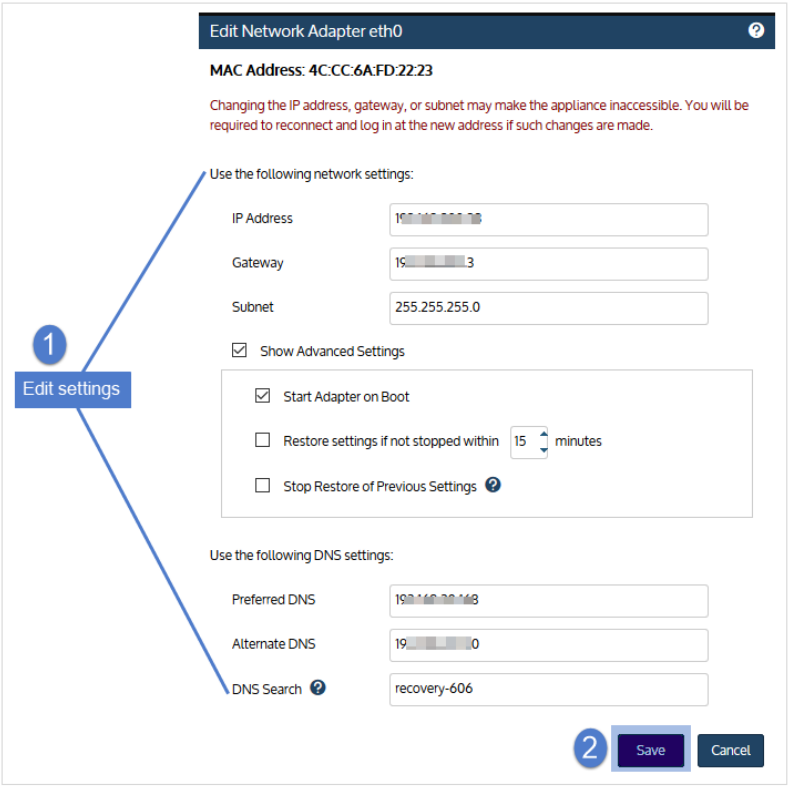

To view or edit the hosts file

1. On the **Configure > Appliances** page, select the appliance.
2. Click the **Network** tab below.
3. Select the adapter and click **Edit Hosts File**.

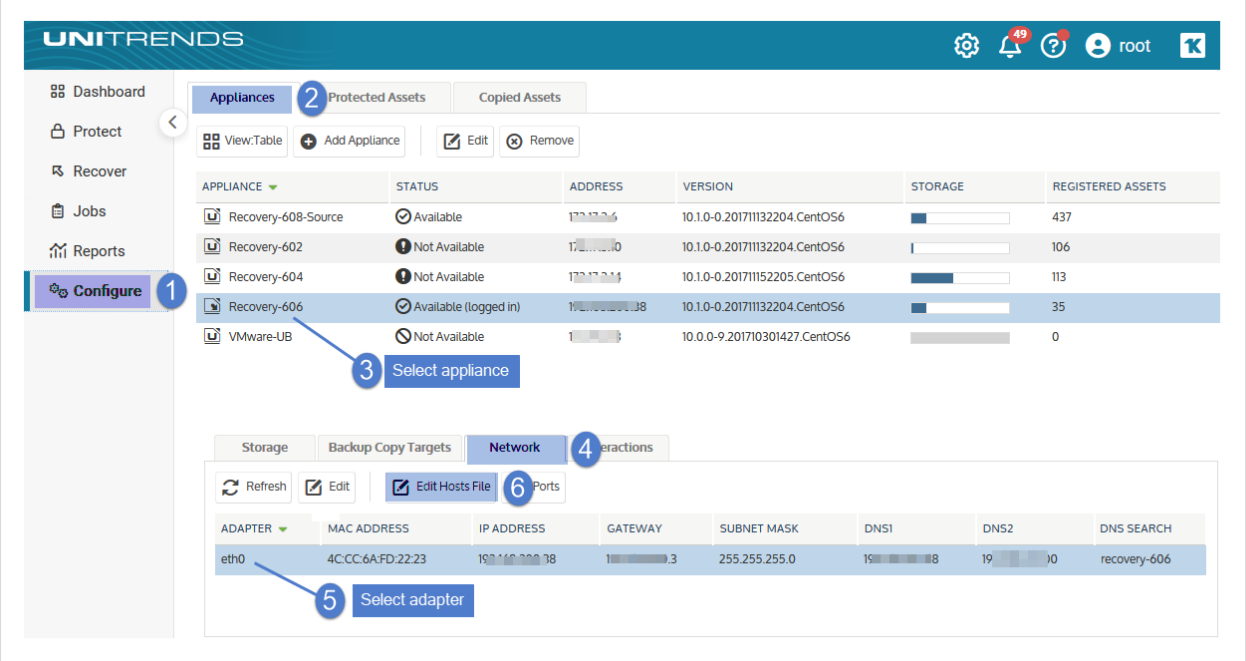

4 Do any of following:

- To add an asset, click **Add**, enter all applicable information, and click **Save**:

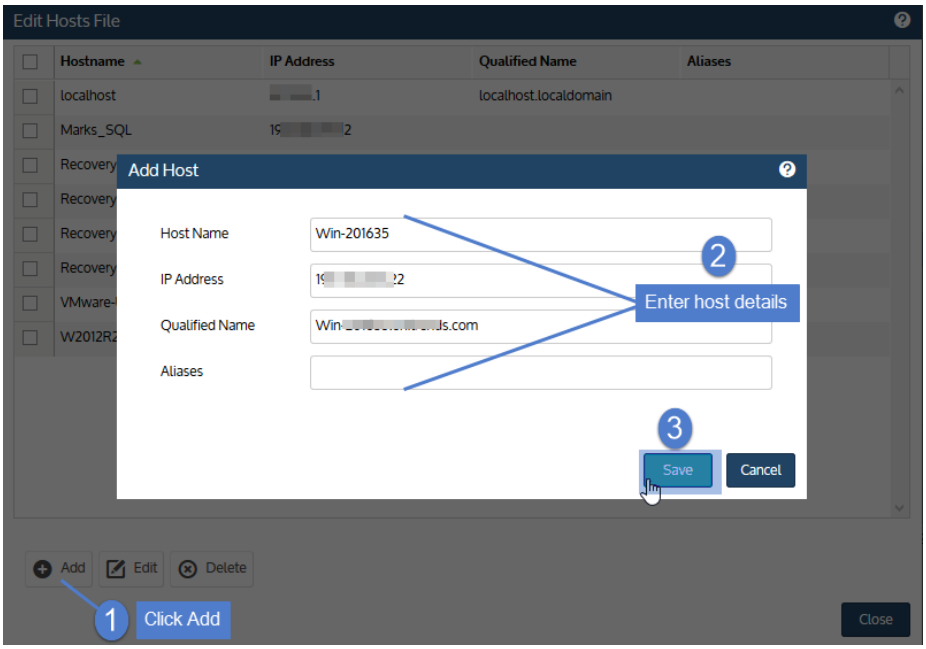

- To edit an asset, select it in the list, click **Edit**, modify information, and click **Save**:

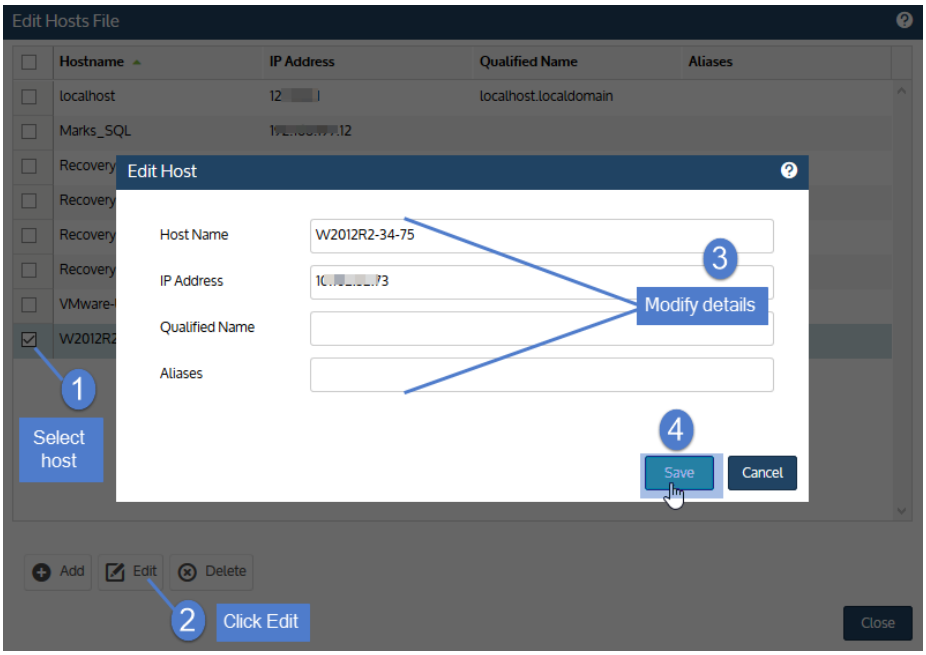

- To delete an asset, select it in the list, click **Delete**, check the **I understand the risks...** box, and click **Delete Host**:

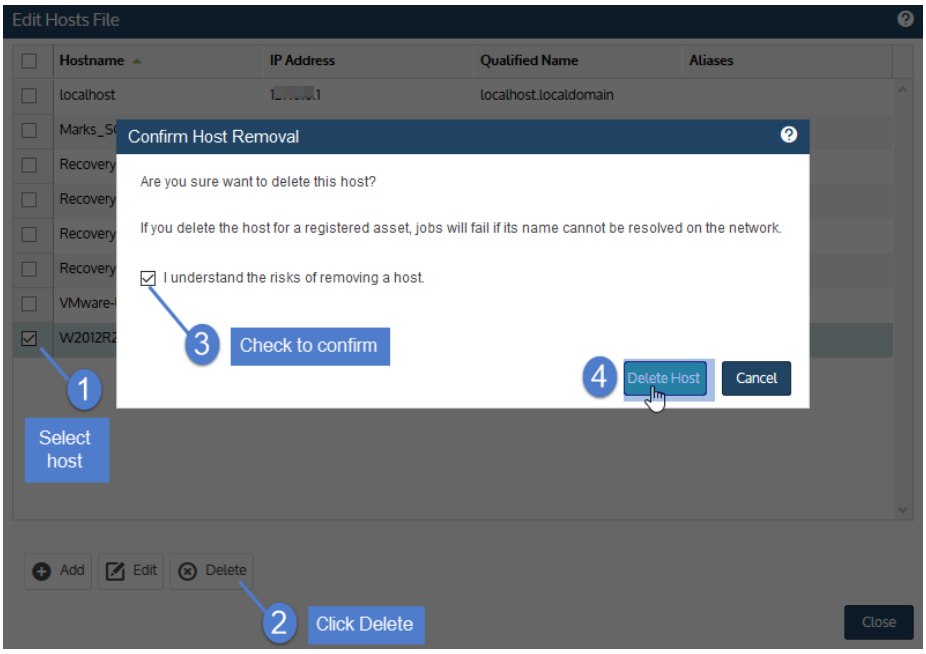

5 Review your changes and click **Close** to exit.

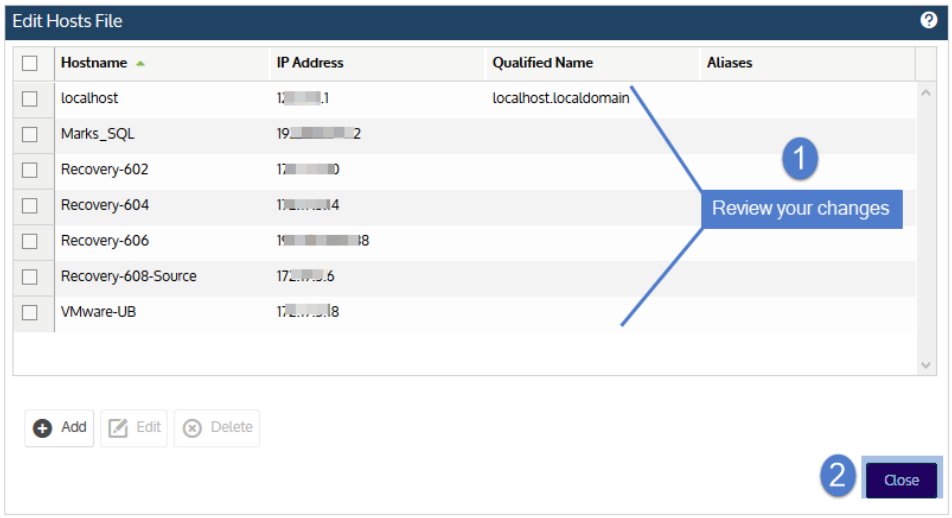

To view or edit port security settings

Note: If this appliance is managed by UniView (e.g., users can only log in to the appliance UI from UniView), you cannot modify port security settings. The Ports tab is disable and you receive a *Managed by UniView* message if you attempt to click the tab.

1. Log in to the appliance UI.

   You must log in directly to the appliance. You cannot change the port security settings of a managed appliance.
2. On the **Configure > Appliances** page, select the appliance.
3. Click the **Network** tab below.
4. Select the adapter and click **Ports**.

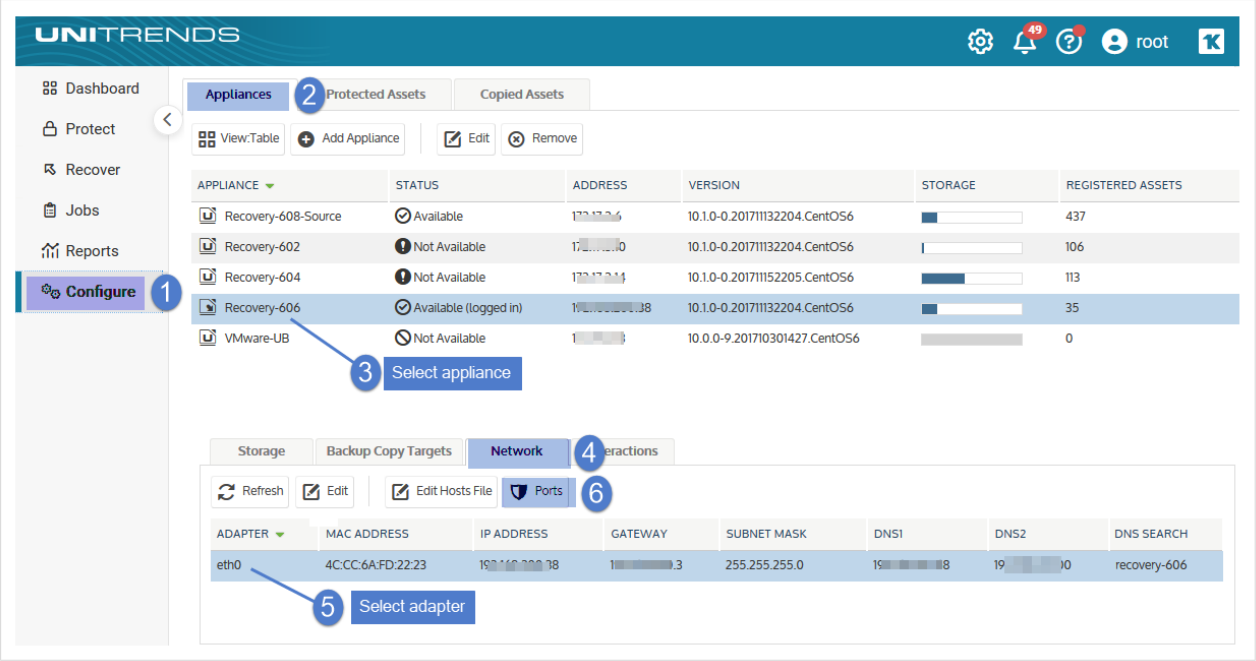

5 View the Port Security area to see the current port security setting and to determine which setting you want to apply. Click each option to see the associated closed ports:

- **None (Open All)** opens all ports.
- **Low**, **Medium**, and **High** closes the ports listed in the table above.

6 Select a **Port Security** option and click **Save**.

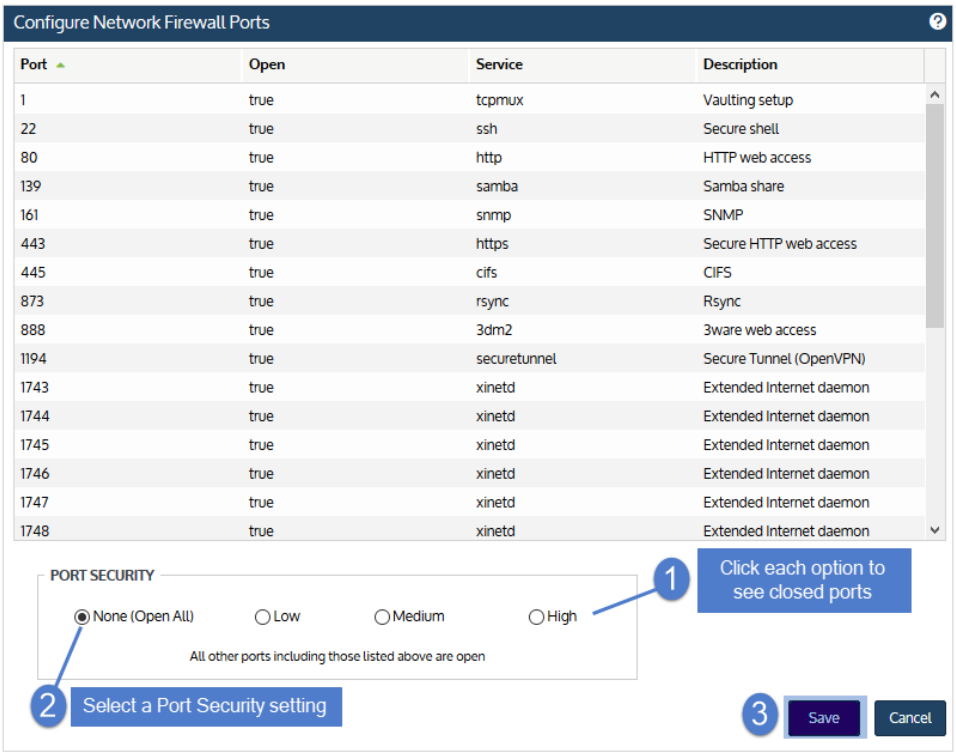

Additional port requirements

Additional ports must be open for connectivity to the Internet and for connectivity to any hot backup copy target. See the following for details:

Note: Unitrends does not officially support backup through firewalls. For details, see this KB article: Backup fails through Router, DMZ, or Firewall.

- "Connectivity between the appliance and the Internet"
- "Connectivity between the appliance and a hot backup copy target"

Connectivity between the appliance and the Internet

Task	Port, Protocol, and Rule	Destination	Notes
Backup and backup copy operations	443/HTTPS Outbound from the Unitrends appliance	kaseyagroup-appliance-registry.jfrog.io	A secure docker container registry required to update backup and backup copy components.
Product Updates	443/HTTPS Outbound from	repo.unitrends.com	repo.unitrends.com is used by the Unitrends appliance

Task	Port, Protocol, and Rule	Destination	Notes
	the Unitrends appliance 22/SFTP Outbound from the Unitrends appliance	sftp.unitrends.com	to perform software updates. sftp.unitrends.com is used to collect files related to active support tickets.
Remote Support	443/HTTPS Outbound from the Unitrends appliance	support-itivity.unitrends.com	Used for opening a remote tunnel to the Unitrends support team.
Proactive Monitoring	161/UDP Outbound from the Unitrends appliance 161/TCP Outbound from the Unitrends appliance 162/UDP Outbound from the Unitrends appliance 162/TCP Outbound from the Unitrends appliance	notifications.unitrends.com	Used for SNMP trap collection for all proactive monitoring.

Connectivity between the appliance and a hot backup copy target

Task	Port, Protocol, and Rule	Destination	Notes
Backup copy to the Unitrends Cloud or your Unitrends target appliance.	The OpenVPN port provided by Unitrends **Or** The port number you have configured for the secure tunnel connection to the backup copy target appliance must be open Outbound for the TCP and UDP protocols. Port 443 must also be open Outbond for the UCP protocol.	For Unitrends Cloud, the public-facing IP address provided by Unitrends. Target appliance hostname and IP	Used for copying data to the Unitrends Cloud or your Unitrends target appliance.

Email reporting

To configure email reporting

1 On the **Configure** > **Appliances** page, select the appliance and click **Edit**.

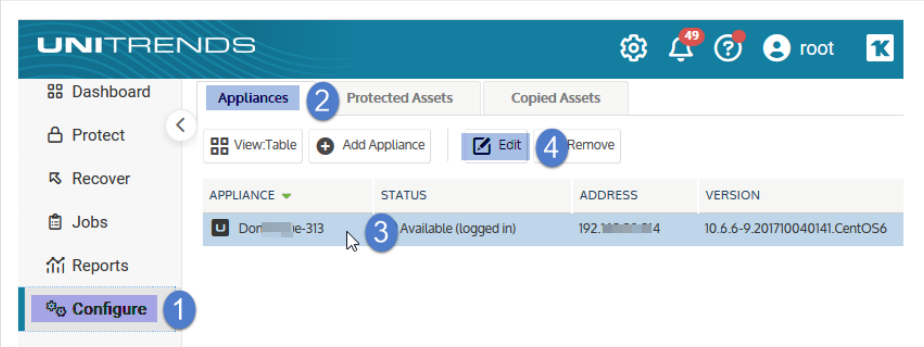

2 Select the **Email** tab.

3 Check the **Enable email reporting** option.

4 Enter the fully qualified SMTP server name or its IP address.

> **Note:** If a DNS record has not been configured, you must use the IP address of the SMTP server.

5 If you have an externally-hosted SMTP server that requires authentication, check the **Server requires credentials** option and enter username and password credentials.

IMPORTANT!

- When using a non-local mail server or an internal SMTP relay configuration, we recommend using an authenticated mail user to prevent filtering issues (for example, cases where alerts are not sent to specific recipients due to filtering rules applied to unauthenticated connections or defined in the mail domain policy). Use a mail user service account that is exempt from routine password change to prevent email from being blocked or delayed.
- Google security requirements – As of May 30, 2022, Google no longer supports the use of third-party apps or devices which ask you to sign in to your Google Account using only your username and password (for details, see Less secure apps & your Google Account). To use Gmail as an SMTP relay, you must:
  - Create an application password for the Unitrends appliance in the Google Account:
    - Access the Google account and go to **Manage Google Account > Security > App passwords**.
    - From the Select app list, select **Other (*Custom name*)**. Enter a custom name (e.g., *Unitrends*).
    - Click **Generate**. The Google app password is generated and displays on the Generated app password page.
  - On the Email tab of the Unitrends Edit Appliance dialog, check **Server requires credentials** and enter the Google username and the 16-digit, Google-generated app password.

6 To test the SMTP configuration, click **Send a Test Email**. Enter your email address and click **Send**. The appliance sends a test email. (If you do not receive the email, check the SMTP settings.)

7 Enter an email recipient and select the email report types that this recipient will receive.

(Optional) Click **+** to add more recipients. For further information on email report types, see "Email reports " on page 1372.

- Appliance — Appliance Status, Replication Activity (if hot backup copy is configured), and Compliance (if data copy access jobs are scheduled) reports are sent to the recipient if this option is selected.
- Jobs — The Schedule report is sent to the recipient if this option is selected.
- Failures — Failure reports are sent to the recipient if this option is selected.

8 Click **Save**.

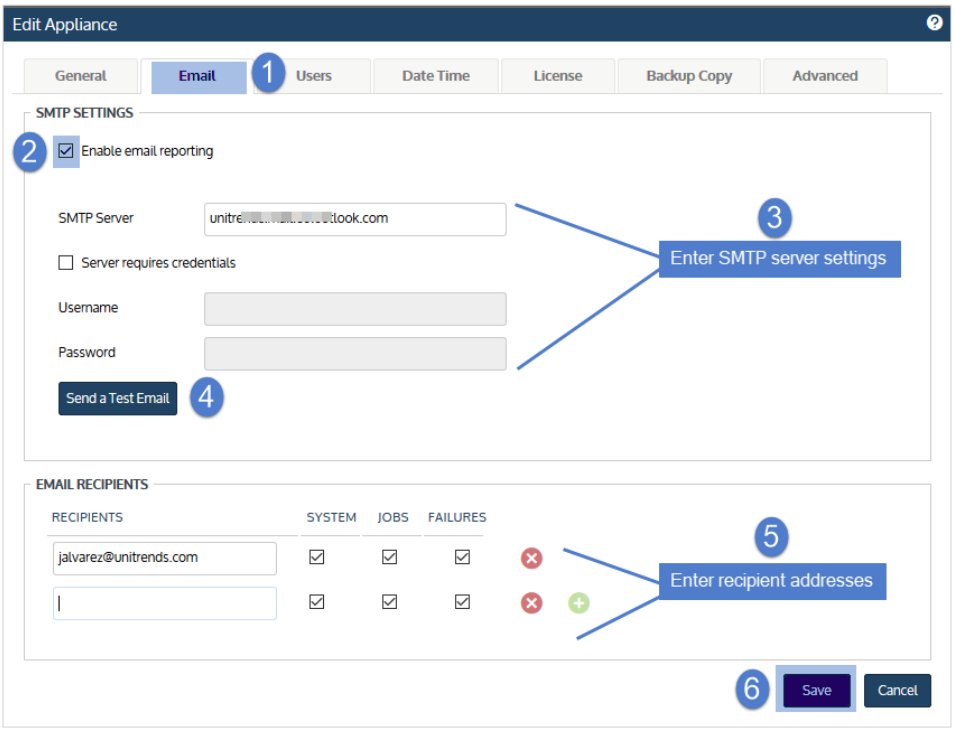

Users and roles

Unitrends appliances are managed and monitored from the User Interface (UI). A user account is required to access the UI. By default, a superuser named *root* is created on the appliance, but you can create additional user accounts. You can set up users on the appliance itself or use Active Directory (AD) authentication.

A user's role determines what the user can access on the appliance. Roles are assigned to users in these ways:

- When you create a user on the appliance, you assign a role to the user.
- If you are using AD authentication, you add Unitrends AD user groups to your AD domain (as AD security groups) and assign users to these groups. Each Unitrends AD user group comes configured with a Unitrends role that is applied to each user in the group.

See these topics for details on how role-based access works:

- "Roles and access levels"
- "Additional considerations" on page 123

Use these procedures to manage users created on the appliance:

- "To add a user" on page 123
- "To edit a user" on page 131
- "To remove a user" on page 133

Use these procedures to manage Active Directory (AD) users:

Roles and access levels

Unitrends' self service role-based access control model enables you to restrict a user's access at the appliance, asset, and task level. Each user account is assigned a role that defines the types of operations the user can perform on the appliance. In addition, the Manage role can be further customized by applying an access level and other options. User roles and access levels are described in the following table.

Role	Access level	Description
Monitor	not applicable	A user with this role is only able to view the status of operations, such as jobs, and to run reports. The user cannot create or start jobs or configure the appliance in any way.
Manage	No Restrictions	A user with this role and access level can: • View all settings and run reports. • Perform all backup, backup copy, and recovery tasks (create, modify, run, cancel, and delete jobs, delete backups, recover backups, etc.). • Perform all other management and configuration tasks other than creating or modifying users. (The user can add, modify, and delete assets, managed appliances and backup copy targets, modify network, storage, and email settings, edit their password, etc.).

Role	Access level	Description
		Note: If using Active Directory authentication, users in the Unitrends-Manage AD group have the No Restrictions access level. You cannot modify the access level of a Unitrends AD group. However, you can add a Unitrends-Manage AD user to the UI and apply a different access level. For details, see "To add a role to a user in the Unitrends-Manage AD group" on page 138.
Manage	Backup/Recovery Operator	A user with this role and access level can: • Run reports. • View jobs. • Run, cancel, suspend, or resume existing backup and backup copy jobs. (The user cannot create, modify, or delete backup or backup copy jobs.) • Recover from backups and backup copies. • Cancel active recovery jobs. You can further restrict recovery operations by using: • Edit Scope to specify the assets the user can access. – The user can recover from backups and backup copies of in-scope assets only. Backups and backup copies of other assets do not display in the UI and cannot be accessed by the user. – The user can cancel active recovery jobs of in-scope assets only. Recovery jobs of other assets display in the UI but the user cannot cancel these jobs.

Role	Access level	Description
		Note: VMware vApps and resource pools cannot be excluded from a user's scope. If the user has access to any hosted VM on an ESXi server, the user can access its hosted vApps and resource pools. • Edit Options to restrict what and how the user can recover from backups and backup copies. (For example, recover only files, recover only to the original asset, etc.).
Manage	Backup Operator	A user with this role and access level can: • Run reports. • View jobs. • Run, cancel, suspend, or resume existing backup and backup copy jobs. (The user cannot create, modify, or delete backup or backup copy jobs.)
Manage	Recovery Operator	A user with this role and access level can: • Run reports. • View jobs. • Recover from backups and backup copies. • Cancel active recovery jobs. You can further restrict recovery operations by using: • Edit Scope to specify the assets the user can access. – The user can recover from backups and backup copies of in-scope assets only. Backups and backup copies of other assets do not display in the UI and cannot be accessed by the user. – The user can cancel active recovery jobs of in-scope assets only. Recovery jobs of other assets display in the UI but the user cannot

Role	Access level	Description
		cancel these jobs. **Note:** VMware vApps and resource pools cannot be excluded from a user's scope. If the user has access to any hosted VM on an ESXi server, the user can access its hosted vApps and resource pools. • Edit Options to restrict what and how the user can recover from backups and backup copies. (For example, recover only files, recover only to the original asset, etc.).
Superuser or Administrator	not applicable	A user with this role can perform all operations, including adding, editing, and deleting users.

Additional considerations

Review these additional considerations before managing users:

- User accounts can only be used to log in to the appliance for which they were created. Users are not shared across Unitrends appliances. To log in to another appliance, the user must be created directly on that appliance.
- Once a user logs in, they can access the local appliance and any appliances that are managed by the local appliance. The user's role determines the operations they can perform on the local and managed appliances.
- To add, modify, or delete users, you must be logged in to the UI as a user that has the *administrator* or *superuser* role. Users with *monitor* or *manage* roles can only see their own user account.
- To add a user, you must supply a username, password, and role for the new user.
- Once you set up users, you can assign them to asset groups to customize how protected assets are grouped and displayed for the user in the UI. For details, see "Grouping assets in custom folders" on page 348.
- In most cases, you access the appliance through the UI by entering UI user and password credentials. If you are an advanced user and need command line access, you can use a terminal emulator, such as PuTTY, to connect to the appliance using operating system account credentials. The appliance's OS *root* user is the only account that can be used for command line access. Use caution when performing tasks from the appliance command line. Before using the command line, check the Support Toolbox. Many lower-level appliance tasks can be run from this handy interface.

To add a user

1. Log in to the appliance as a user that has the administrator or superuser role.
2. On the **Configure > Appliances** page, select the appliance and click **Edit**.

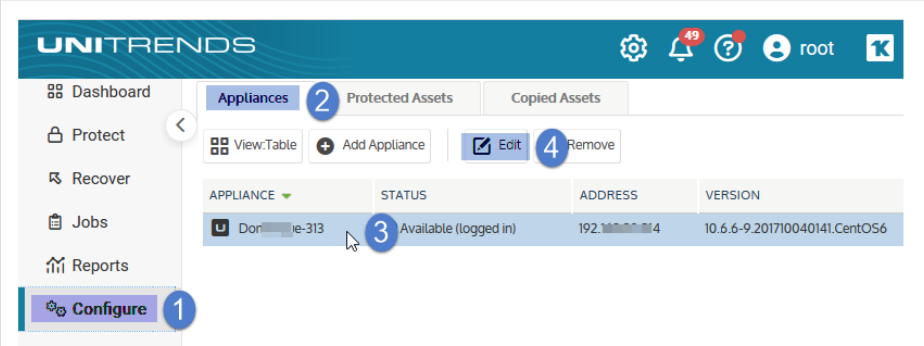

3 Select the **Users** tab.

4 Click **Add**.

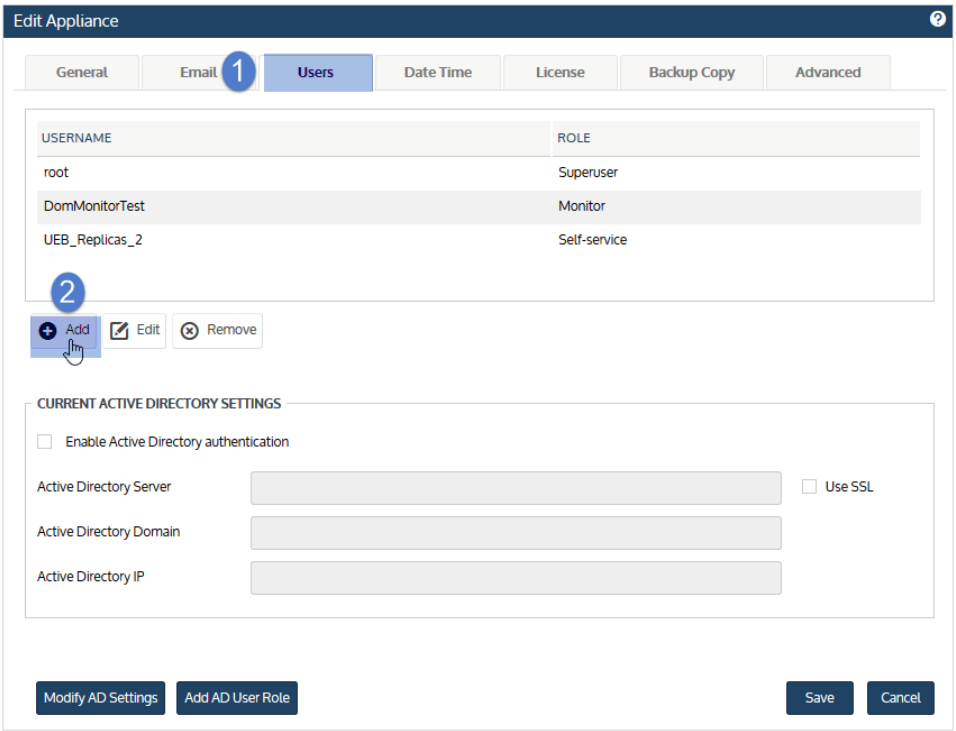

5 In the Add User area, enter the **Username**, **Password**, and **Confirm Password**, then select a **Role**. For a description of each role, see "Roles and access levels" on page 120.

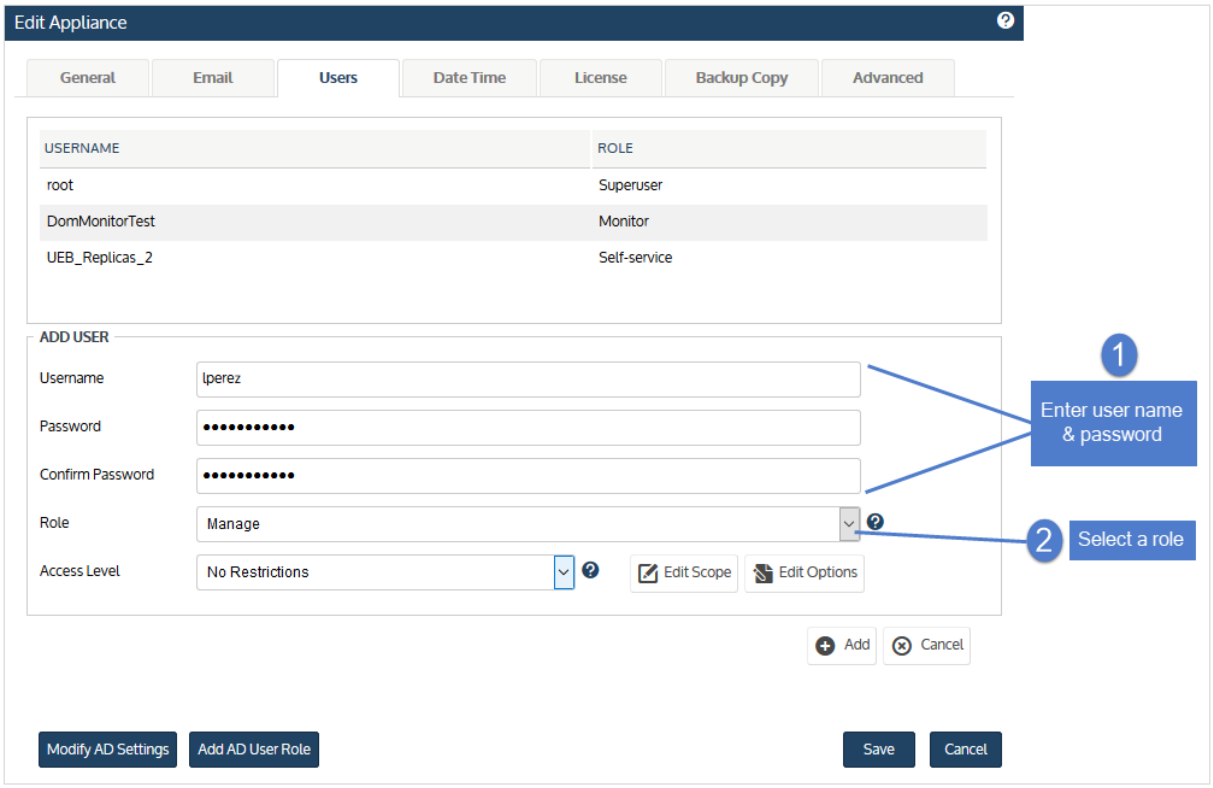

6 Do one of the following:

- If you selected the Manage role, continue with step 7.
- If you did NOT select the Manage role, skip to step 11.

7 Select an Access Level. For a description of each access level, see "Roles and access levels" on page 120.

8 Do one of the following:

- If you selected the Recovery Operator or Backup/Recovery Operator access level, continue with step 9.
- If you selected the No Restrictions or Backup Operator access level, skip to step 11.

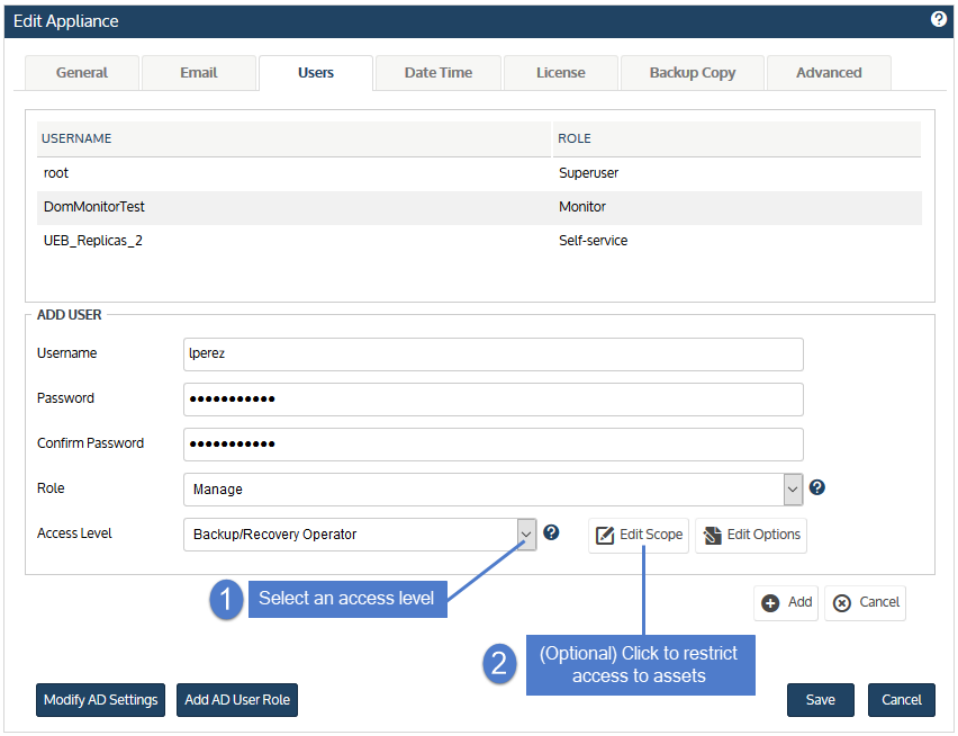

9 (Optional) Restrict the assets the user can access for recovery tasks. Skip this step if you want the user to have access to all assets. (For details on how restricting assets works, see "Roles and access levels" on page 120.)

Restrict access to assets by doing these steps:

- Click **Edit Scope**.
- In the Inventory area, check boxes to select each asset the user will have access to. Selected assets display in the Recovery Scope Assets area.
  - You can expand nodes in the inventory tree to view and select hosted applications and VMs.
  - You can select a virtual host or application instance to enable access to all of its hosted VMs or databases. Note that as new VMs and databases are added they are not automatically added to the user's scope. To add them, edit the AD user role to discover the new VMs or databases.

  Note: VMware vApps and resource pools cannot be excluded from a user's scope. If the user has access to any hosted VM on an ESXi server, the user can access its hosted vApps and resource pools (even if you do NOT select them in the inventory tree to add them to the Recovery Scope Assets list).
- Click **Save**.

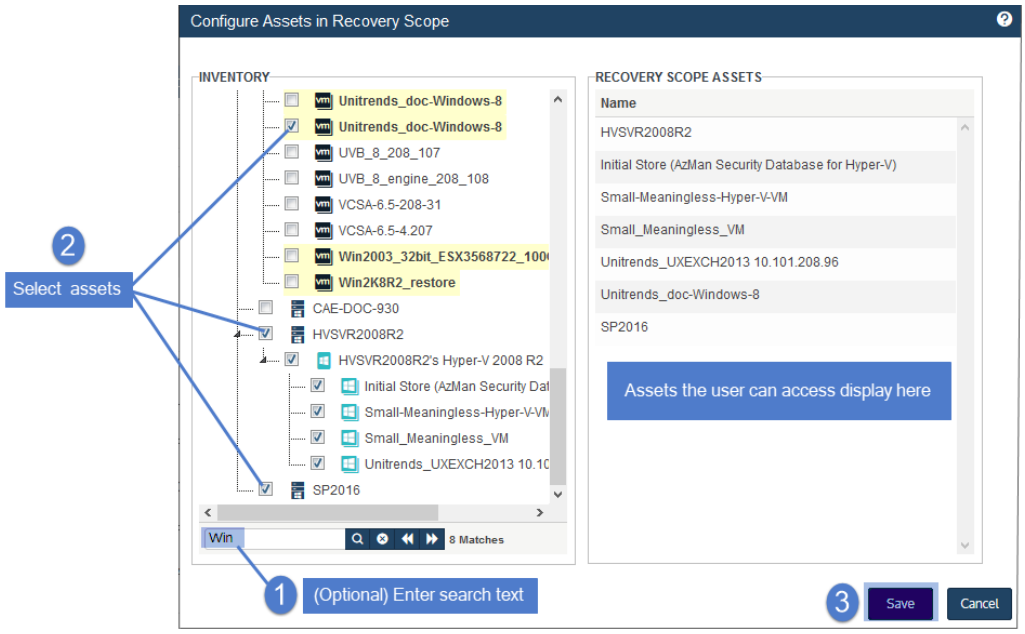

10 (Optional) Restrict the user's recovery options. Skip this step if you want the user to have the ability to perform all recovery operations.

Restrict recovery options by doing these steps:

- Click **Edit Options**.
- In the Recovery Options dialog, select one or more options, then click **Save**. Options are described in the following table.

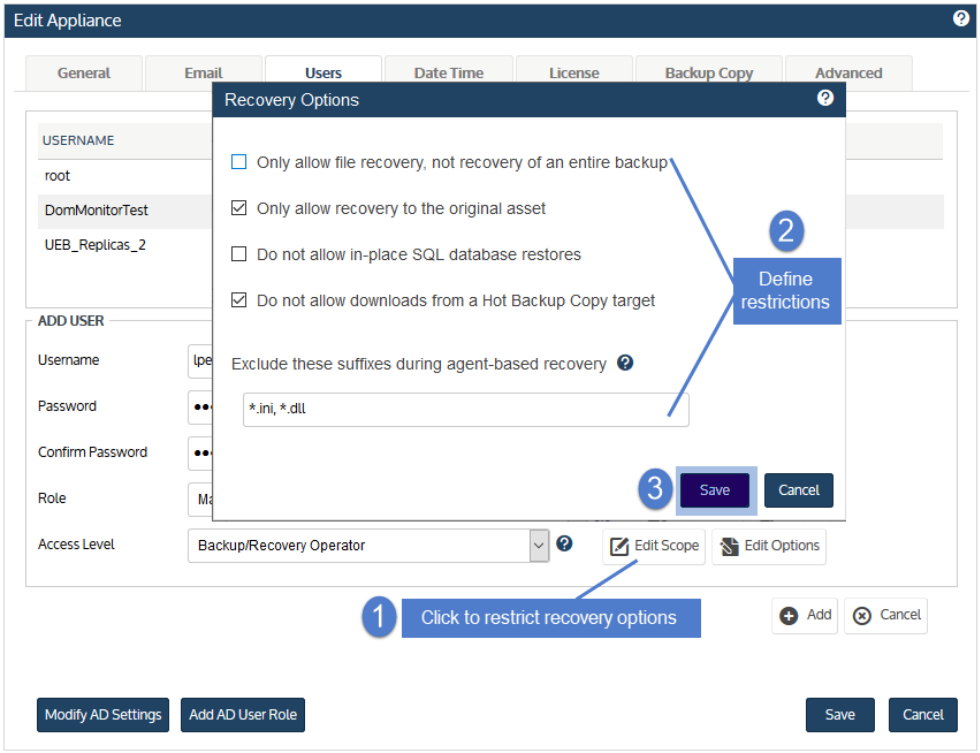

Recovery option	Description
Only allow file recovery, not recovery of an entire backup	Check this box to allow the user to recover only selected files from a backup or backup copy. The user is not able to recover any of the following: • An entire backup. • An entire virtual machine from a host-level backup. • A virtual machine by using instant recovery. • A Windows asset by using Windows replicas. • An asset by using bare metal recovery.
Only allow recovery to the original asset	Check this box to allow recovery to the original asset only. The user is not able to recover to a different asset.
Do not allow in-place SQL database restores	Check this box to prevent SQL database restores to the original location. The user can only recover SQL backups to an alternate location. The user is not able to recover any of the following; • System databases (master, model, and msdb).

Recovery option	Description
	• Stretch databases.
Do not allow downloads from a Hot Backup Copy target	Check this box to prevent the user from recovering from backup copies that reside in the Unitrends Cloud or on a remote hot backup copy target. The user can recover from local hot backup copies only (hot copies that reside on the appliance they are logged in to). The user is not able to: • Import a hot backup copy from a remote target to the local appliance. • Download files from a hot backup copy that resides on a remote target.
Exclude these suffixes during agent-based file level recovery	Applies to recovering files from file-level backups only. This restriction does not apply to recovering entire file-level backups or recovering from other backup types (such as host-level or application backups). Enter a comma-separated list of file extensions to prevent the user from recovering files of these types. You must include the * wildcard before each extension. For example, enter the following to exclude ZIP files, HTML files, and executables: `*.zip, *.html, *.exe`

11 Click **Add** to add the new user.

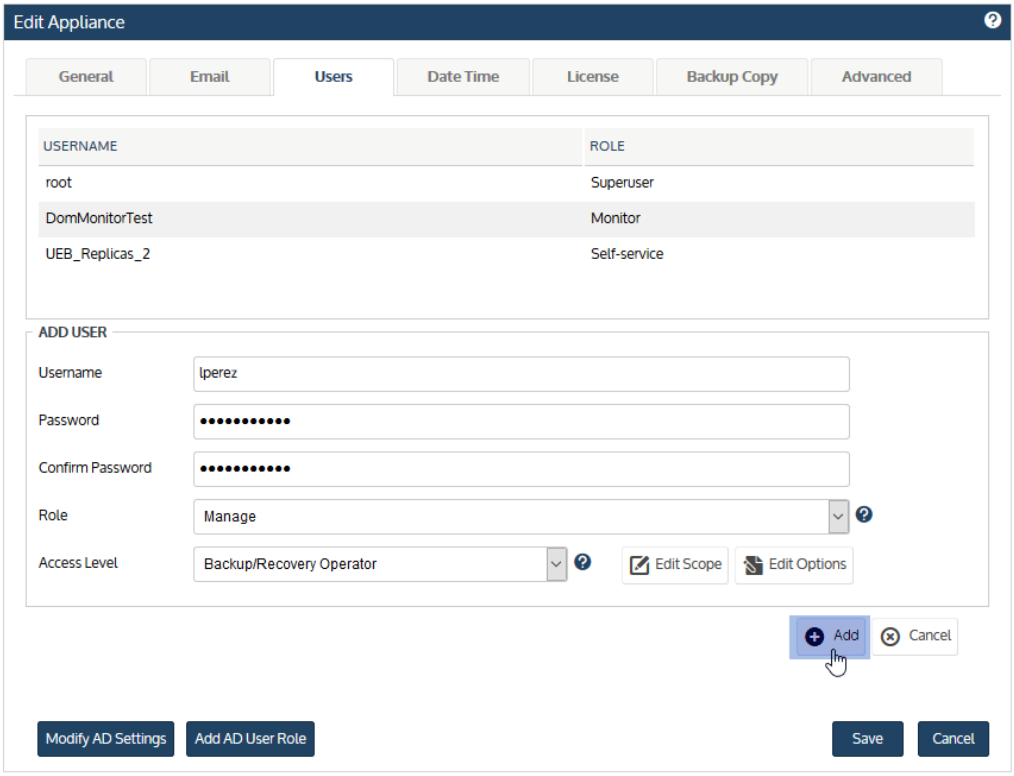

12 Click **Save**.

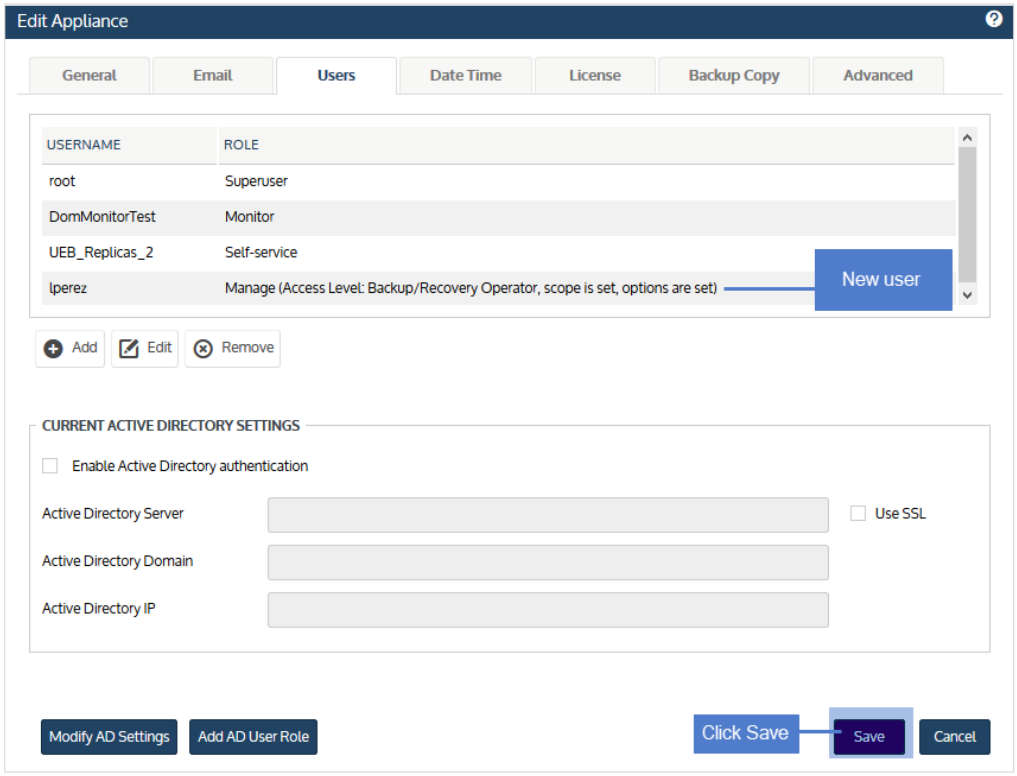

To edit a user

1. Log in to the appliance as a user that has the administrator or superuser role.
2. On the **Configure > Appliances** page, select the appliance and click **Edit**.

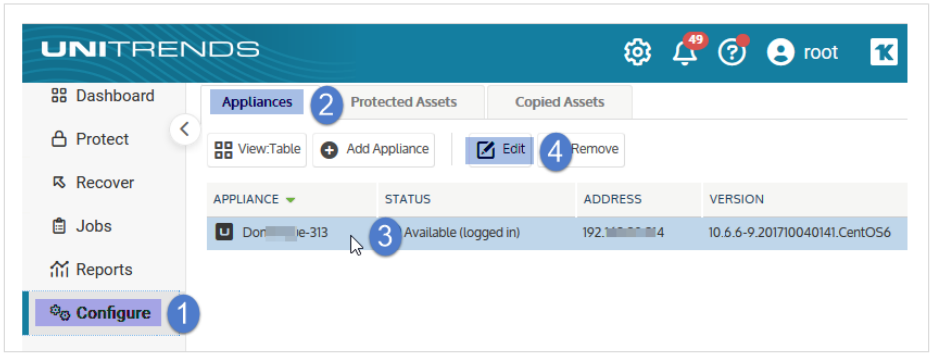

3. Select the **Users** tab.
4. Select a user in list, then click **Edit**.

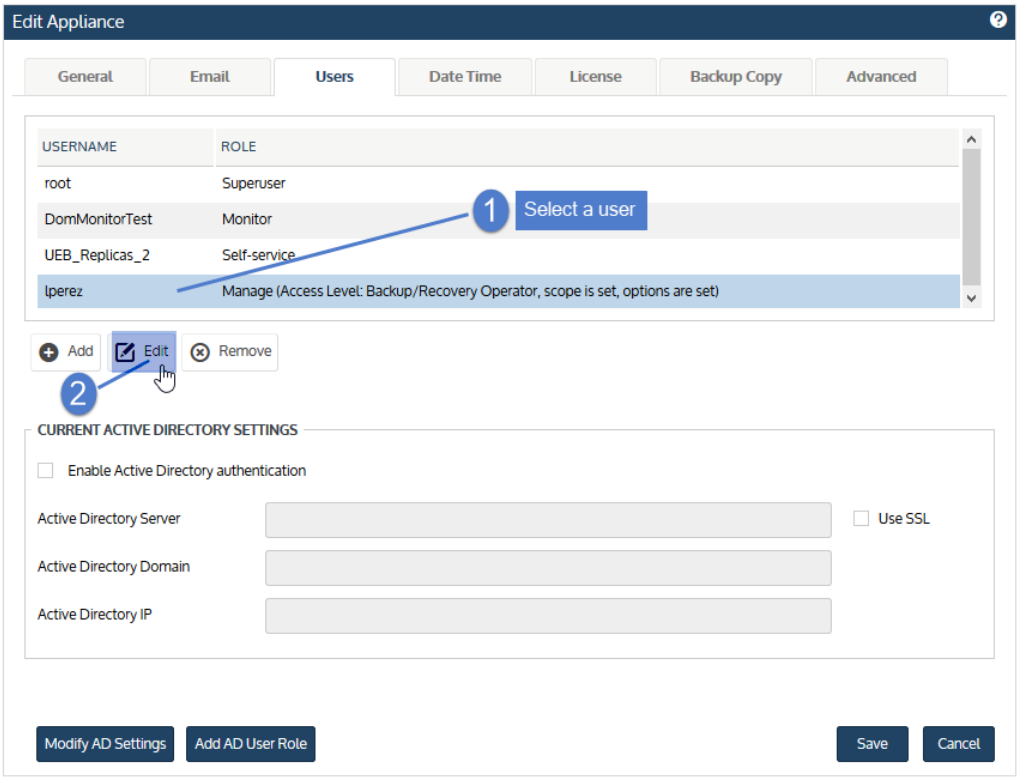

5 Modify fields, then click **Update**. (For details on each setting, see "To add a user" on page 123.)

6 Click **Save**.

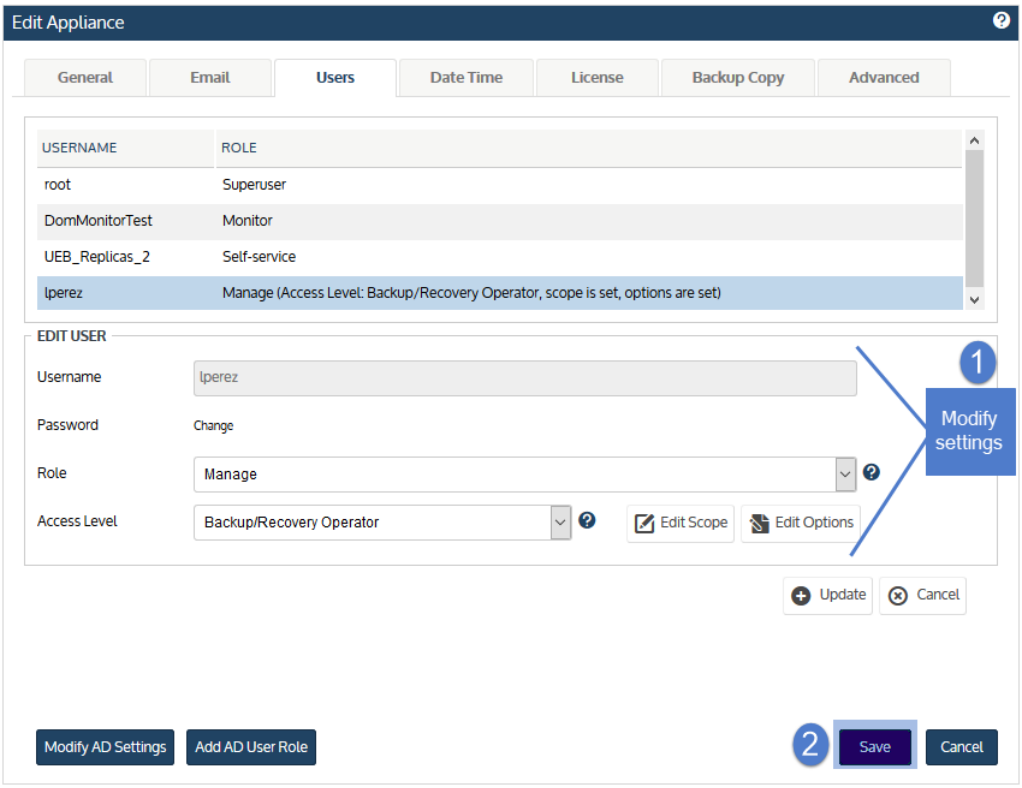

To remove a user

1. Log in to the appliance as a user that has the administrator or superuser role.
2. On the **Configure > Appliances** page, select the appliance and click **Edit**.

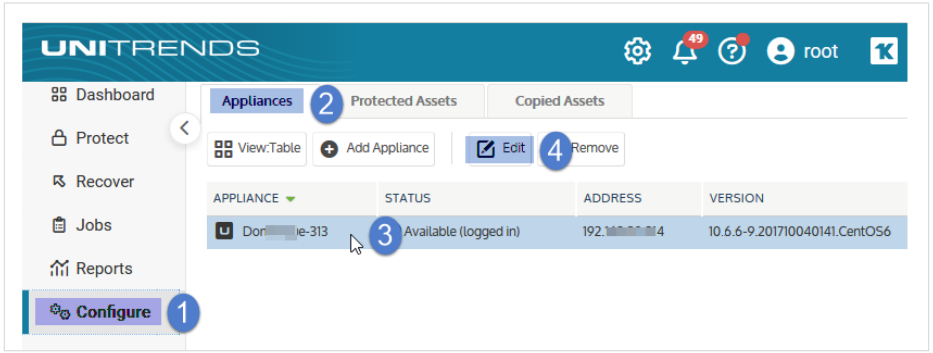

3. Select the **Users** tab.
4. Select a user in list.
5. Click **Remove**, then **Confirm**. The user is removed.

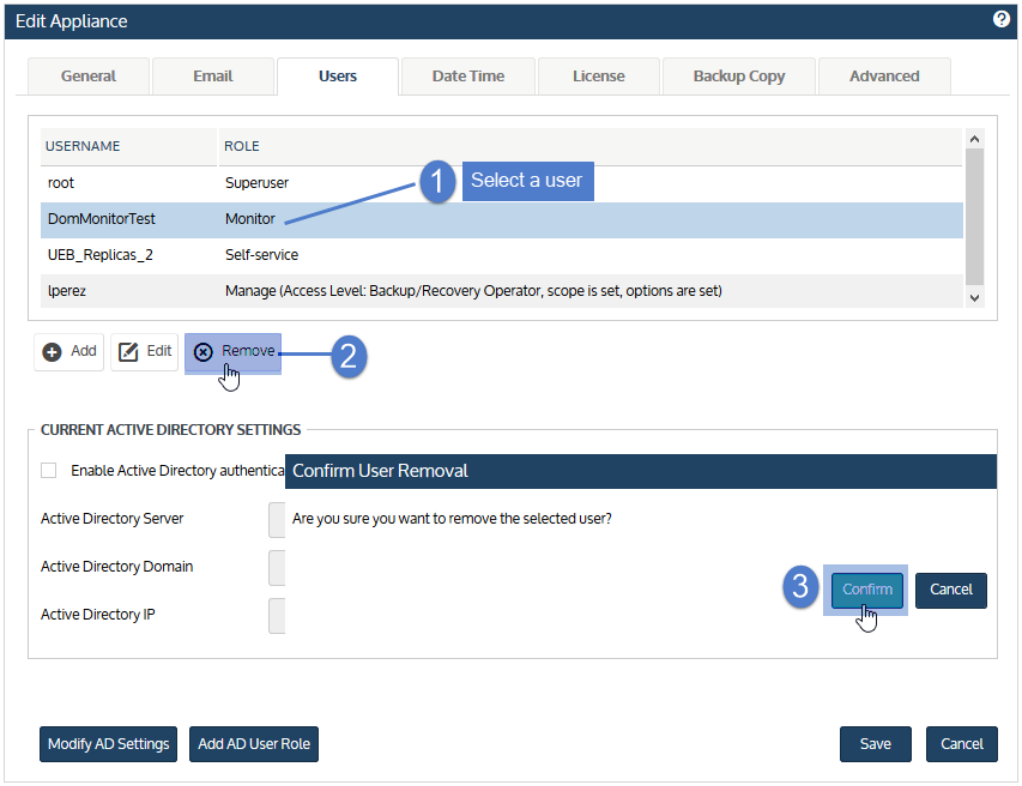
Edit Appliance
General
Email
Users
Date Time
License
Backup Copy
Advanced
USERNAME
ROLE
root
Superuser
DomMonitorTest
Monitor
UEB_Replicas_2
Self-service
lperez
Manage (Access Level: Backup/Recovery Operator, scope is set, options are set)
1
Select a user
Add
Edit
Remove
2
CURRENT ACTIVE DIRECTORY SETTINGS
Enable Active Directory authentica
Active Directory Server
Active Directory Domain
Active Directory IP
Confirm User Removal
Are you sure you want to remove the selected user?
3
Confirm
Cancel
Modify AD Settings
Add AD User Role
Save
Cancel

6 Click **Save**.

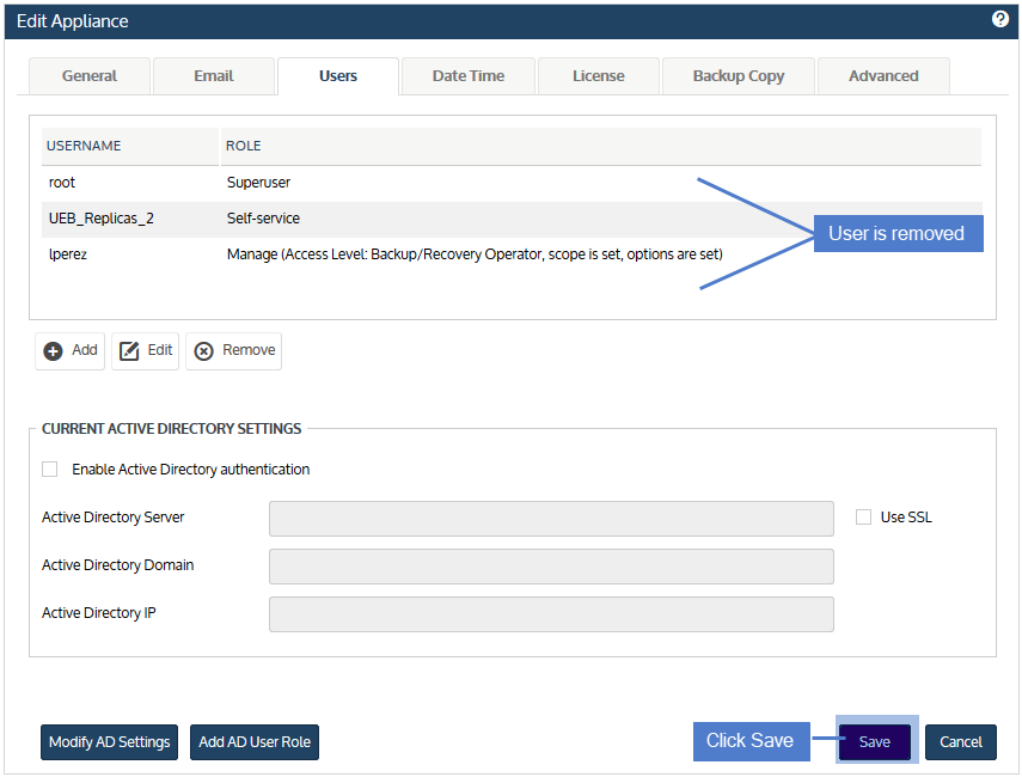

To set up Active Directory authentication

You can use Active Directory (AD) domain credentials for Unitrends user accounts. Set Unitrends users up as members of specified AD domains and they can access the appliance without being added as users on the appliance itself.

Note: AD authentication is implemented at the UI and Apache component level. The Unitrends operating system is not joined to the AD domain.

The AD security group to which a user belongs determines which features that user can view and utilize. Users are granted one of the following roles: *monitor*, *manage*, *administrator*, or *superuser*.

Use these steps to set up AD authentication:

1 Create the following security groups in your Active Directory domain. For a description of each role, see "Roles and access levels" on page 120.

AD security group	Description
Unitrends-Superuser	Members of this security group are granted the *superuser* role in the Unitrends UI.
Unitrends-	Members of this security group or BULTIN\administrators are granted the

AD security group	Description
Admin	*administrator* role in the Unitrends UI.
Unitrends-Manage	Members of this security group are granted the *manage* role and *no restrictions* access level in the Unitrends UI. These users can view statuses and reports, run jobs, and perform other management tasks, such as adding or modifying assets and retention settings. You can restrict a user's backup and recovery options by adding the AD user to the UI and applying a different access level. For details, see "To add a role to a user in the Unitrends-Manage AD group" on page 138.
Unitrends-Monitor	Members of this security group are granted the *monitor* role in the Unitrends UI. These users are only able to view the status of completed jobs and run reports. They cannot run jobs, view running jobs, or configure the appliance in any way.

Note: You may name these security groups to suit your environment. If you use your own names, be sure to enter those names when you configure AD authentication on the appliance. Unitrends AD user group names in your AD domain must match the names you enter below in step 8.

2 Add users to the Unitrends AD security groups as desired.

Users who are not BULTIN\administrators must be assigned to a Unitrends group to log in to the UI using AD authentication.

Note: Add users to the groups only. Do not add groups. Nested grouping is not a Microsoft best practice and may cause undesirable results.

3 Do one of the following:

- Create a DNS entry for the AD server with reverse lookup configured, then skip to step 8.
- Continue with step 4 to add the AD server to the Unitrends appliance's hosts file.

4 Log in to the appliance UI as a user that has the administrator or superuser role.

5 On the **Configure > Appliances** page, select the appliance and click the **Network** tab below.

6 Click **Edit Hosts File**.

7 Add the Active Directory server to the appliance hosts file:

Note: This host entry must be added before you configure the appliance for AD authentication.

- Click **Add**.
- Enter the Host Name of the AD server that manages the Active Directory domain.
- Enter the IP Address of the AD server.
- For Qualified Name, enter the Active Directory domain only. Do not include the server name.

- Click **Save**.
- Example: for an AD server called *SERVER_AD* whose IP address is *192.168.111.75* and AD domain is *company_domain.com*, enter the following:
  - **SERVER_AD** in the Host Name field
  - **192.168.111.75** in the IP Address field
  - **company_domain.com** in the Qualified Name field

8 Configure the appliance for AD authentication. Click **Modify AD Settings** and enter the following in the Current Active Directory Settings area:

Field	Action
Enable Active Directory Authentication	Check this box to start using AD authentication, or leave unchecked to start using AD authentication at a later time.
Use SSL	The Use SSL option is not used.
Active Directory Server	Enter the hostname of the AD server that manages the Active Directory Domain. If left blank, the appliance populates this field using the hosts file entry. If you are using DNS and did not add the AD server to the hosts file, be sure to enter the hostname here. This field is limited to 15 characters.
Active Directory Domain	Enter the name of the AD domain. Do not include the AD server name. For example, *ad_domain.company_domain.com*. This name must be present in the appliance hosts file or resolvable through DNS.
Active Directory IP	Displays the IP of the AD server if you added it to the appliance hosts file. (No IP displays if you are using DNS.)
Unitrends user groups	Unitrends AD user groups display below the Active Directory IP. If you opted to enter different user group names above in step 1, modify the names here to match those that you entered. User group names in these fields must match the names in your AD domain.

9 Click **Save**.

Users in the Unitrends AD groups can now log in to the appliance UI by using their AD credentials. Note that AD users do not display on the Users tab in the Unitrends UI (unless you add an AD user role for the user).

Users in the Unitrends-Manage group are granted the *no restrictions* access level. You can restrict a user's backup and recovery options by adding the AD user to the UI and applying a different access level. For details, see "To add a role to a user in the Unitrends-Manage AD group" on page 138.

To log in using AD authentication

This procedure assumes you have set up the Unitrends user account in Active Directory and have configured AD authentication as described in "To set up Active Directory authentication".

1 Connect to the appliance by directing a Chrome or Firefox browser to:

```
https://<appliance IP address>ui/
```

2 On the Login page, enter the AD domain and user name in either of the following formats:

*ad_domain**ad_username* or *ad_username*@*ad_domain.company_domain*.

For example, for user *jsmith* on AD domain *accounting* and company domain *americanaccountants.com*, enter:

accounting\jsmith

or

jsmith@accounting.americanaccountants.com

3 Enter the password for this AD user.

4 Click **Login**.

To add a role to a user in the Unitrends-Manage AD group

Once you have set up AD authentication, users that have been added to the Unitrends AD security groups in your AD domain can log in to the appliance and perform operations that are enabled for their group. No further setup is required.

For users in the Unitrends-Manage security group, you have the option to apply additional restrictions by adding AD user roles. Once you add a user role, the AD user displays in the list of users on the Users tab. (AD users that do not have roles defined do not display in the Unitrends UI.)

Note: Adding a user role is supported for AD users with the *manage* role only (users in the Unitrends-Manage group, if you are using the default Unitrends group names). Because the appliance cannot determine a user's role in your AD domain, the UI does not prevent assigning roles to AD users that have the *monitor*, *administrator*, or *superuser* roles. Be aware that the access level you assign by doing this procedure only applies if the user has the manage role in your AD domain. If not, the access level is not used and the user's role (*monitor*, *administrator*, or *superuser*) determines the operations they can perform.

Use this procedure to add an AD user role:

1 Log in to the appliance as a user that has the administrator or superuser role.

2 On the **Configure > Appliances** page, select the appliance and click **Edit**.

3 On the Edit Appliance dialog, select the **Users** tab.

4 Click **Add AD User Role** at the bottom of the tab.

5 Enter the AD **Username**. Do not include the AD domain or company domain. For example, for user *jose.perez* on AD domain *accounting* and company domain *americanaccountants.com*, enter: **jose.perez**

6 Select an **Access Level**. For a description of each access level, see "Roles and access levels" on page 120.

7 Do one of the following:

- If you selected the Recovery Operator or Backup/Recovery Operator access level, continue with step 8.
- If you selected the No Restrictions or Backup Operator access level, skip to step 10.

8 (Optional) Restrict the assets the user can access for recovery tasks. Skip this step if you want the user to have access to all assets. (For details on how restricting assets works, see "Roles and access levels" on page 120.)

Restrict access to assets by doing these steps:

- Click **Edit Scope**.
- In the Inventory area, check boxes to select each asset the user will have access to. Selected assets display in the Recovery Scope Assets area.
  - You can expand nodes in the inventory tree to view and select hosted applications and VMs.
  - You can select a virtual host or application instance to enable access to all of its hosted VMs or databases. Note that as new VMs and databases are added they are not automatically added to the user's scope. To add them, edit the AD user role to discover the new VMs or databases.

  Note: VMware vApps and resource pools cannot be excluded from a user's scope. If the user has access to any hosted VM on an ESXi server, the user can access its hosted vApps and resource pools (even if you do NOT select them in the inventory tree to add them to the Recovery Scope Assets list).
- Click **Save**.

9 (Optional) Restrict the user's recovery options. Skip this step if you want the user to have the ability to perform all recovery operations.

Restrict recovery options by doing these steps:

- Click **Edit Options**.
- In the Recovery Options dialog, select one or more options, then click **Save**. Options are described in the following table.

Recovery option	Description
Only allow file recovery, not recovery of an entire backup	Check this box to allow the user to recover only selected files from a backup or backup copy. The user is not able to recover any of the following: • An entire backup. • An entire virtual machine from a host-level backup. • A virtual machine by using instant recovery. • A Windows asset by using Windows replicas.

Recovery option	Description
	• An asset by using bare metal recovery.
Only allow recovery to the original asset	Check this box to allow recovery to the original asset only. The user is not able to recover to a different asset.
Do not allow in-place SQL database restores	Check this box to prevent SQL database restores to the original location. The user can only recover SQL backups to an alternate location. The user is not able to recover any of the following; • System databases (master, model, and msdb). • Stretch databases.
Do not allow downloads from a Hot Backup Copy target	Check this box to prevent the user from recovering from backup copies that reside in the Unitrends Cloud or on a remote hot backup copy target. The user can recover from local hot backup copies only (hot copies that reside on the appliance they are logged in to). The user is not able to: • Import a hot backup copy from a remote target to the local appliance. • Download files from a hot backup copy that resides on a remote target.
Exclude these suffixes during agent-based file level recovery	Applies to recovering files from file-level backups only. This restriction does not apply to recovering entire file-level backups or recovering from other backup types (such as host-level or application backups). Enter a comma-separated list of file extensions to prevent the user from recovering files of these types. You must include the * wildcard before each extension. For example, enter the following to exclude ZIP files, HTML files, and executables: `*.zip, *.html, *.exe`

10 Click **Add** to add the role. The AD user role displays in the list of users above.

11 Click **Save** to exit the Edit Appliance dialog.

To edit an AD user role

1 Log in to the appliance as a user that has the administrator or superuser role.

2 On the **Configure > Appliances** page, select the appliance and click **Edit**.

3 On the Edit Appliance dialog, select the **Users** tab. AD users that have roles defined display in the user list.

4 Select the AD user in the list, then click **Edit**.

5 Modify user role settings, then click **Update**. (For details on each setting, see "To add a role to a user in the Unitrends-Manage AD group" on page 138.)

6 Click **Save** to exit the Edit Appliance dialog.

To remove an AD user role

1 Log in to the appliance as a user that has the administrator or superuser role.

2 On the **Configure > Appliances** page, select the appliance and click **Edit**.

3 On the Edit Appliance dialog, select the **Users** tab. AD users that have roles defined display in the user list.

4 Select the AD user.

5 Click **Remove**, then **Confirm**. The user role is removed.

6 Click **Save** to exit the Edit Appliance dialog.

To add an AD user

Add the user to one of the Unitrends AD security groups in your Active Directory domain.

To remove an AD user

Remove the user from its Unitrends AD security group in your Active Directory domain.

To modify AD settings

Use this procedure to enable or disable AD authentication, modify AD server settings, and modify Unitrends AD security group names.

> **Note:** Unitrends AD security group names entered in the Active Directory Settings dialog must match the names defined in your AD domain. Only modify the Unitrends AD group names if these names have changed in your AD environment.

1 Log in to the appliance as a user that has the administrator or superuser role.

2 On the **Configure > Appliances** page, select the appliance and click **Edit**.

3 On the Edit Appliance dialog, select the **Users** tab.

4 Click **Modify AD Settings** at the bottom of the tab.

5 Modify information as needed and click **Save**. For details on each setting, see "To set up Active Directory authentication" on page 135.

Passwords

Users log in to the appliance UI by supplying UI user and password credentials. These UI users are created as described in "Users and roles" on page 119. In addition to these UI users, the appliance also has an operating system account for command line access.

Use these procedures to manage passwords:

- "Change your UI password"

- "Change a UI user password"
- "Change the appliance operating system password"

Change your UI password

1 Log in to the UI. Your username displays in the Global menu in the upper-right portion of the UI.

2 Click your username, then select **My Settings**.

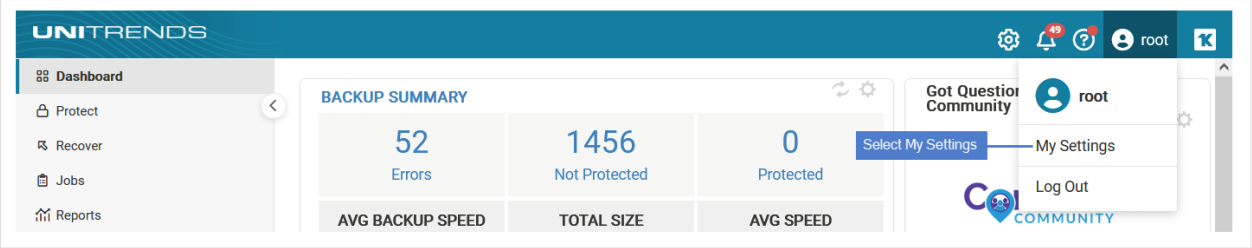

3 Click **Change Password**.

4 Enter the **Existing password**, **New password**, and **Confirm password**.

5 Click **Save**.

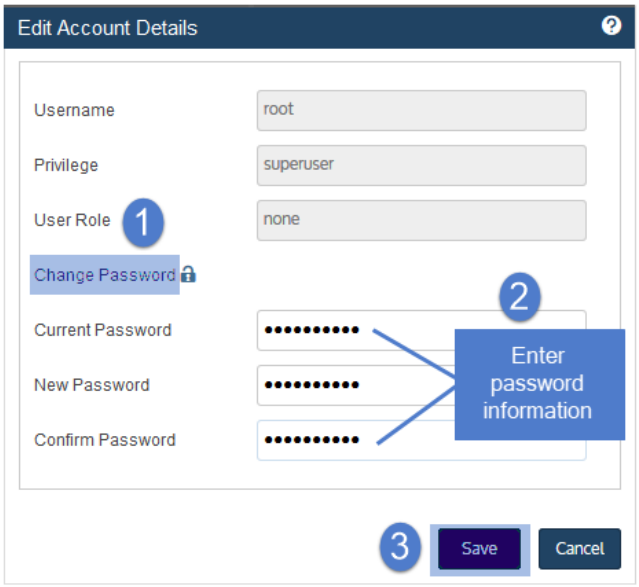

Change a UI user password

1 Log in to the UI as a user that has the superuser or administrator role.

2 On the **Configure > Appliances** page, select the appliance and click **Edit**.

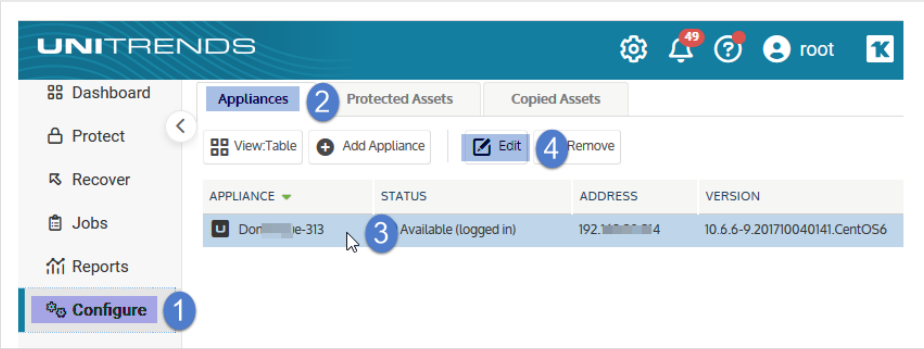

3 Select the **Users** tab.

4 Select a user in list, then click **Edit**.

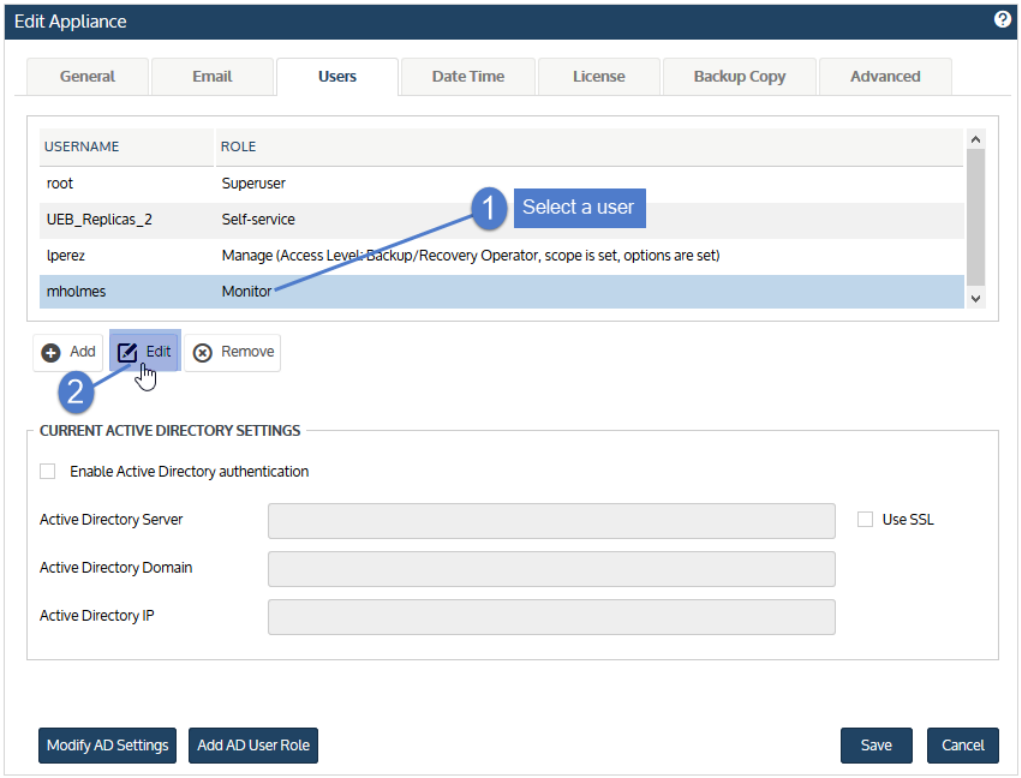

5 Click **Change** (to the right of Password).

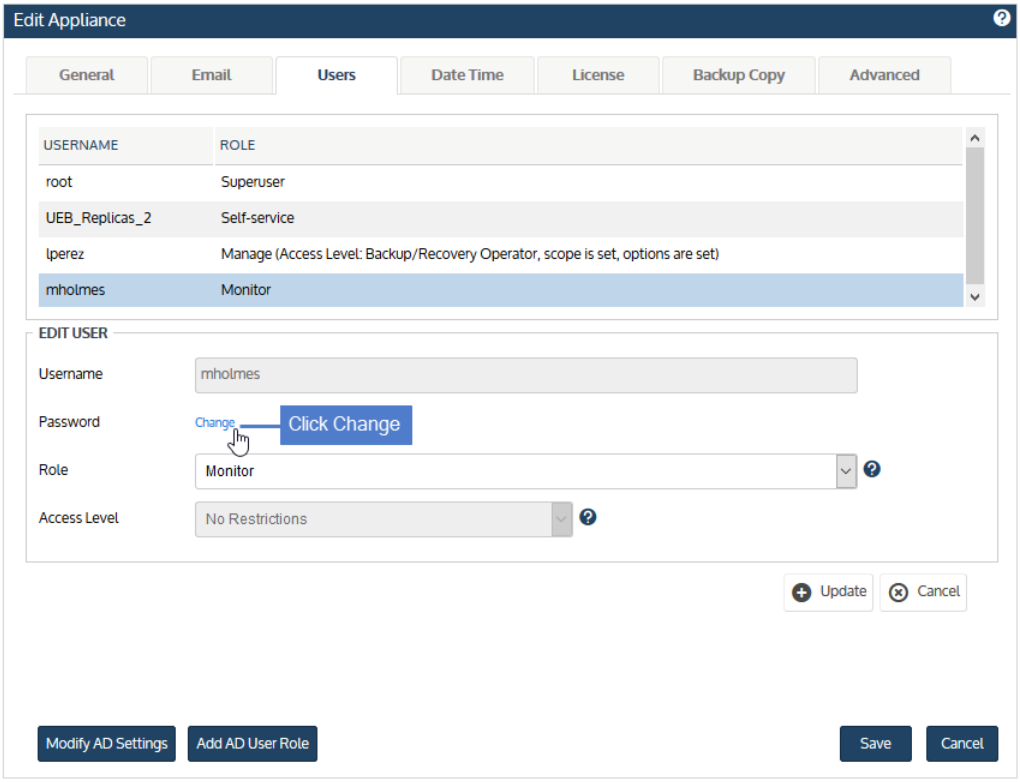

6 Enter the **Current Password**, **New Password** and **Confirm Password**.

7 Click **Update**.

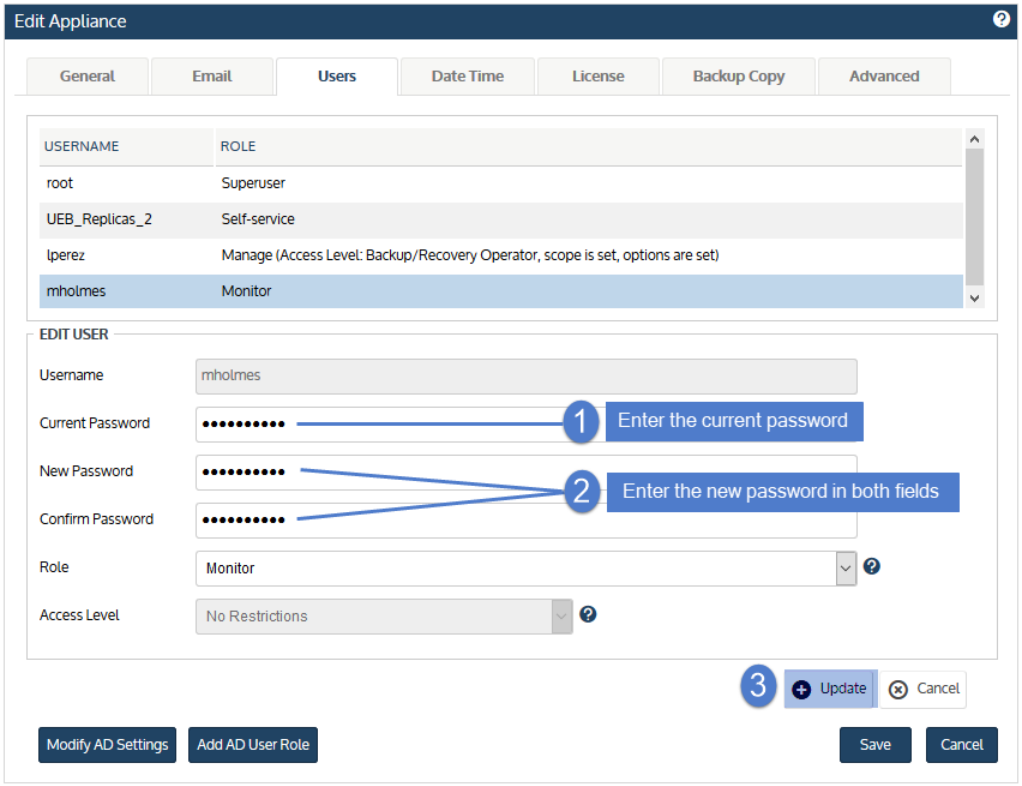
Edit Appliance
General
Email
Users
Date Time
License
Backup Copy
Advanced
USERNAME
ROLE
root
Superuser
UEB_Replicas_2
Self-service
lperez
Manage (Access Level: Backup/Recovery Operator, scope is set, options are set)
mholmes
Monitor
EDIT USER
Username
mholmes
Current Password
1
Enter the current password
New Password
2
Enter the new password in both fields
Confirm Password
Role
Monitor
Access Level
No Restrictions
3
Update
Cancel
Modify AD Settings
Add AD User Role
Save
Cancel

8 Click **Save**.

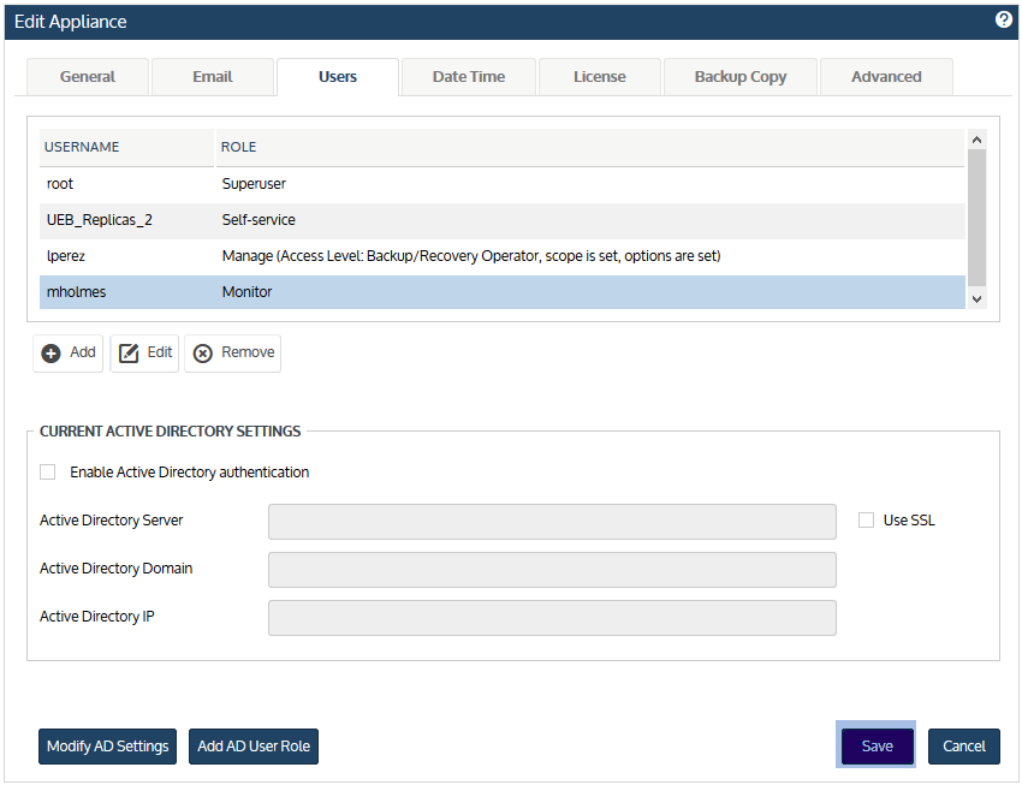

Change the appliance operating system password

In most cases, you access the appliance through the UI by entering UI user and password credentials. If you are an advanced user and need command line access, you can use a terminal emulator, such as PuTTY, to connect to the appliance using operating system account credentials (user *root*). Use caution when performing tasks from the appliance command line. Before using the command line, check the Support Toolbox. Many lower-level appliance tasks can be run from this handy interface.

To change the operating system password

1 On the **Configure > Appliances** page, select the appliance and click **Edit**.

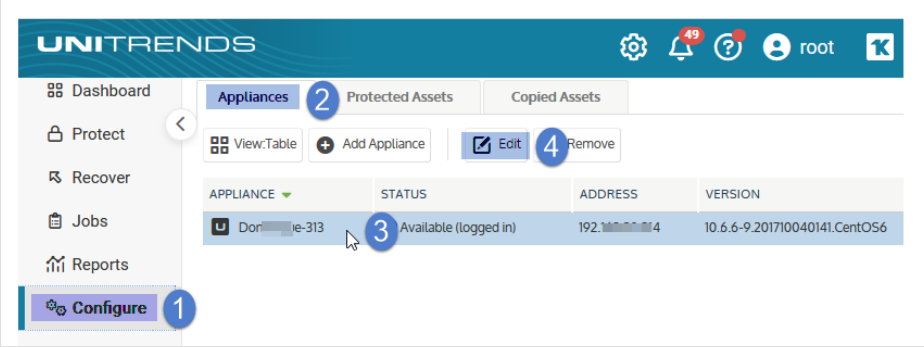

2 Click **Advanced** and select **OS Password**.

3 Enter the current and new passwords, then click **Save**.

Note: For increased security, ensure that the password you enter is different than your UI user password.

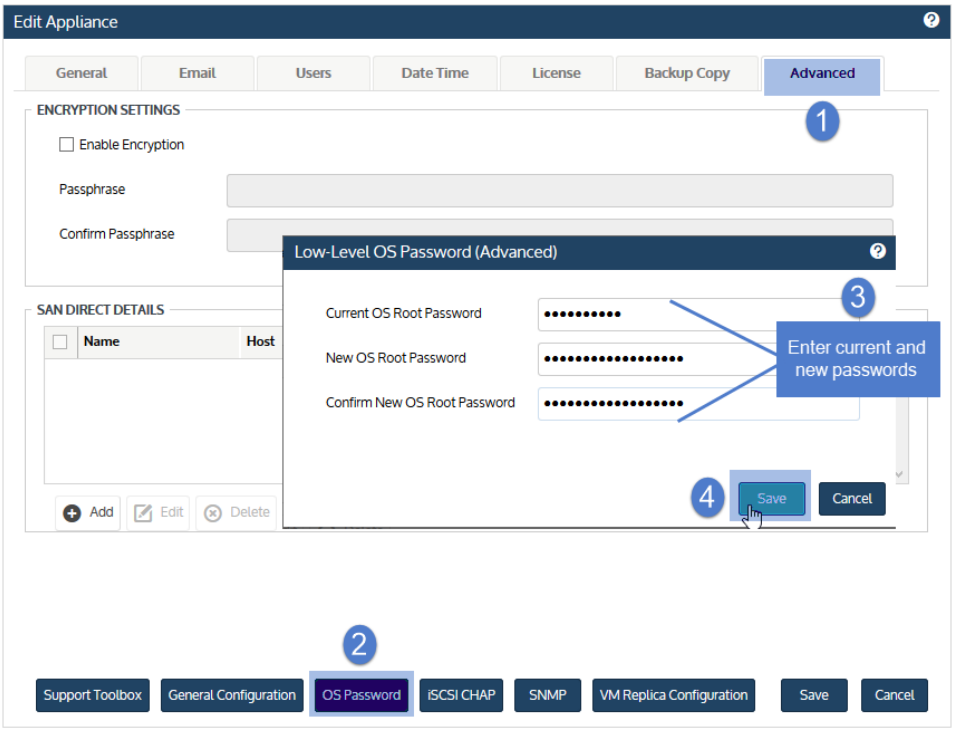

Date and time settings

During deployment, date and time settings were configured for the appliance. You can edit these settings as needed. You can manually set the date and time or sync to an NTP server. To use an NTP server, you will need to supply its address. Edit these settings from the **Date Time** tab of the **Configure > Appliances > Edit > Edit Appliance** page.

To edit date and time settings

1 On the **Configure > Appliances** page, select the appliance and click **Edit**.

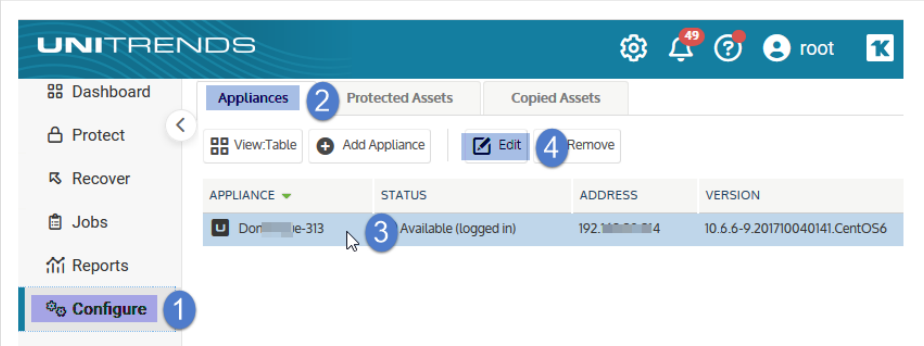

2 Select the **Date Time** tab.

3 Do one of the following:

- Select **Enter a date and time for your appliance**, specify a Date, Time, and Timezone, then click **Save**:

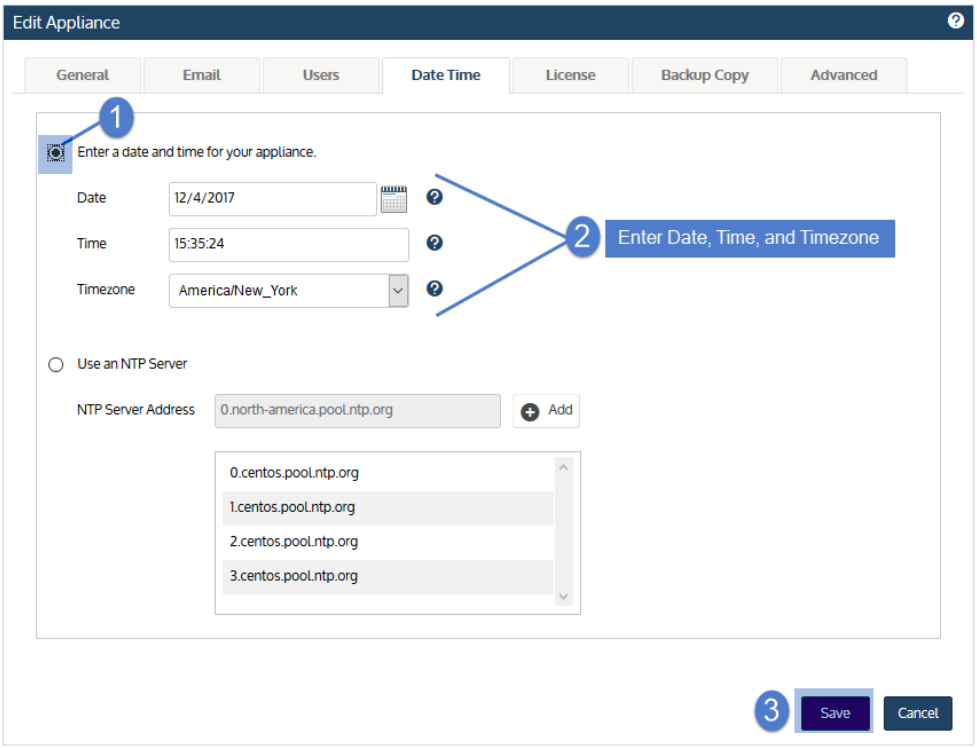

- Select **Use an NTP Server** to sync the appliance date and time to one of the NTP servers in the list. (Or you can opt to add your own NTP server by entering its NTP Server Address and clicking **Add**.) Click **Save** to sync to the NTP server:

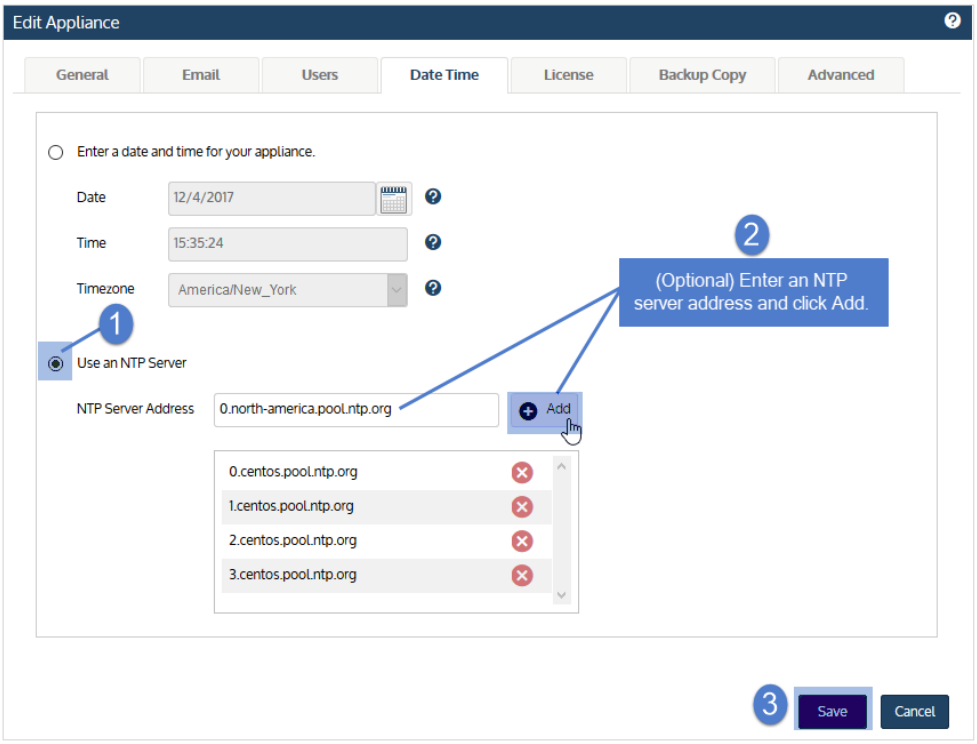

License settings

You can add or modify the appliance license as needed.

Note: Applying a license stops all running jobs.

Licensing procedures for physical Recovery Series, Recovery MAX, and ION/ION+ appliances differ from those for virtual Unitrends Backup appliances:

- Recovery Series, Recovery MAX, and ION/ION+ appliances ship fully licensed. It is likely you will never need to modify this license unless directed to do so by Unitrends Support. If you need to update a license, apply the license you receive from Unitrends.
- Unitrends Backup appliances deploy without a license. Register the license as described in "To register a Unitrends Backup appliance" on page 149, then apply the license you receive from Unitrends as described in "To license a Unitrends Backup appliance" on page 151.

To register a Unitrends Backup appliance

1 On the **Configure > Appliances** page, select the appliance and click **Edit**.

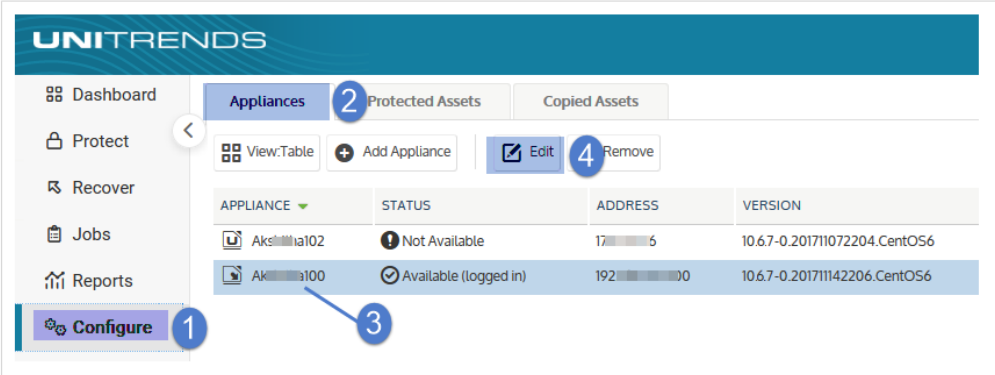

2 Select the **License** tab and click **Update**. The Registration Center displays.

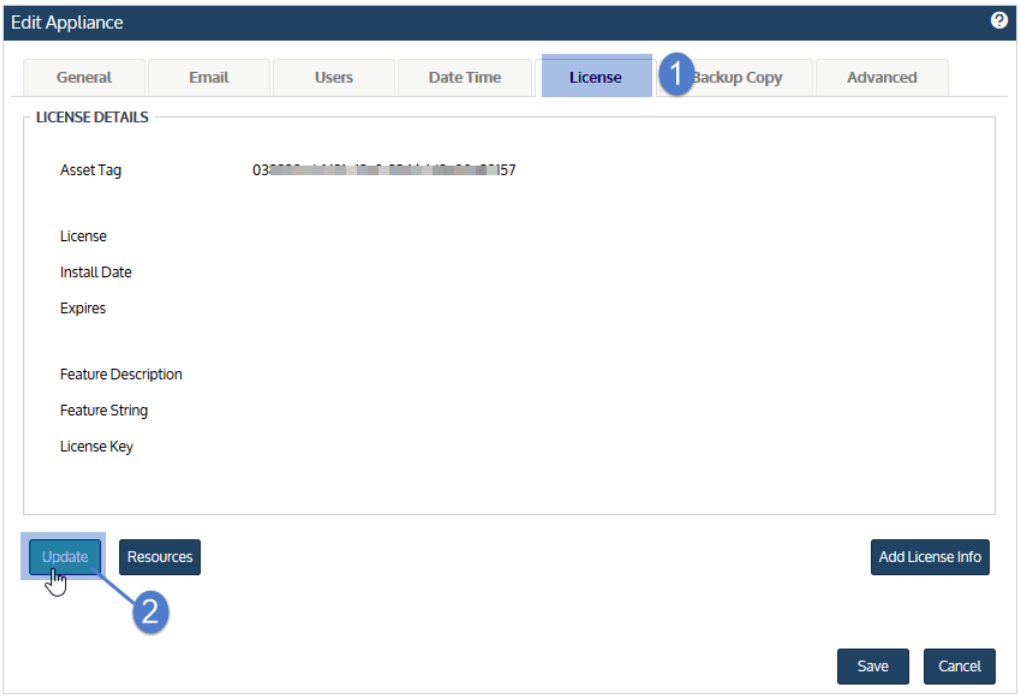

3 Select one of the following:

Selection	Description
Start my free trial	Submit this form to start your free 30-day trial.
Activate my purchase	Enter your email address and activation code. You license key will be emailed to the address you enter here.
Activate with	Enter your promotional code to register your product and receive your license

Selection	Description
promo code	key.
Request a quote	Request a license quote.

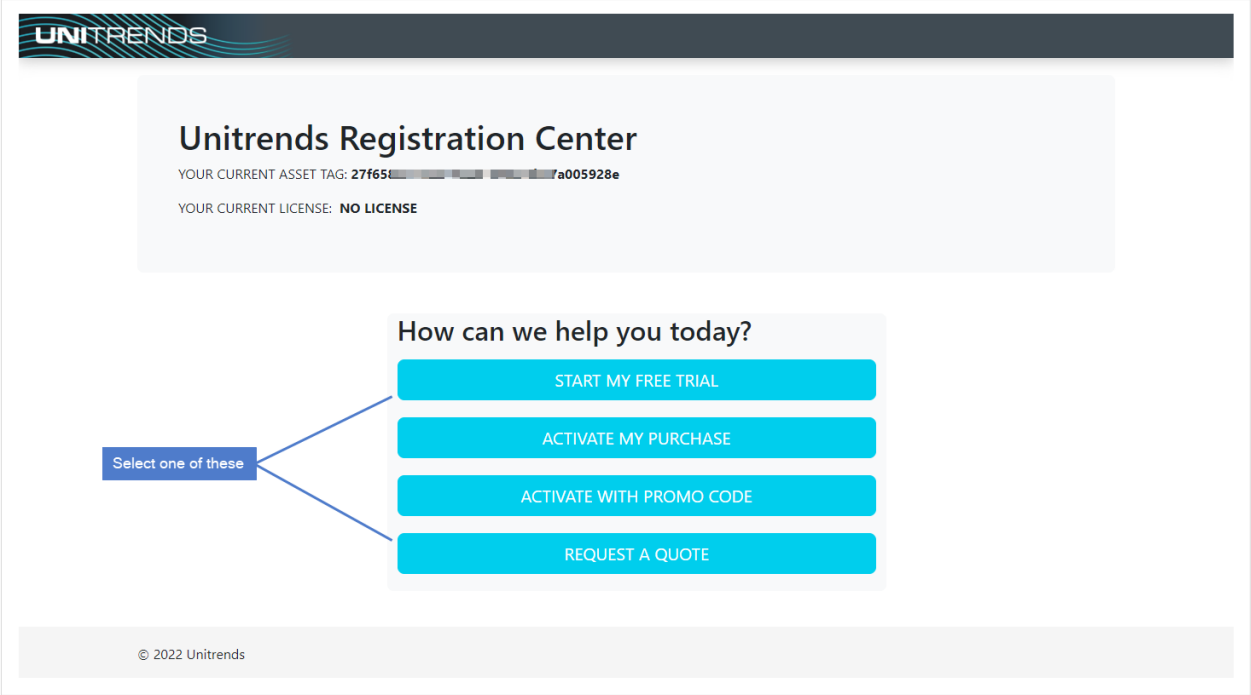

4 Complete and submit the applicable form.

Once you have purchased a license, Unitrends sends an email containing license details. Use the next procedure to apply this license information to the appliance.

To license a Unitrends Backup appliance

Use these steps to enter license information you have received from Unitrends.

1 On the **Configure > Appliances** page, select the appliance and click **Edit**.

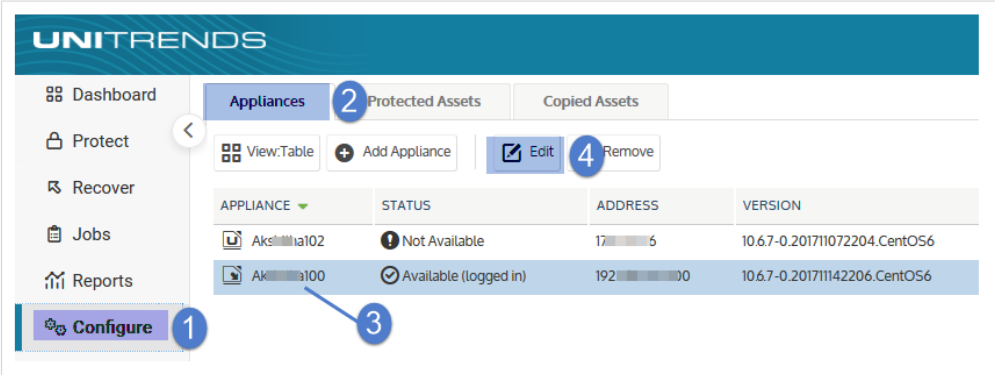

2 Select the **License** tab and click **Add License Info**.

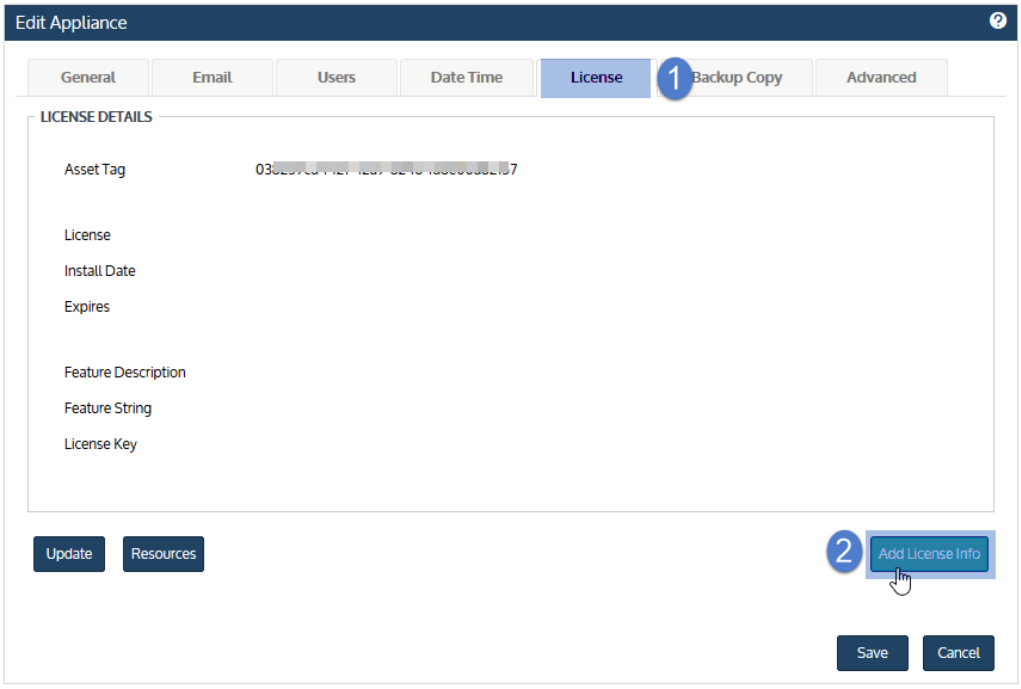

3 Enter the **License Key**, **Expiration Date**, and **Feature String**.

4 Click **Save**. The license is applied.

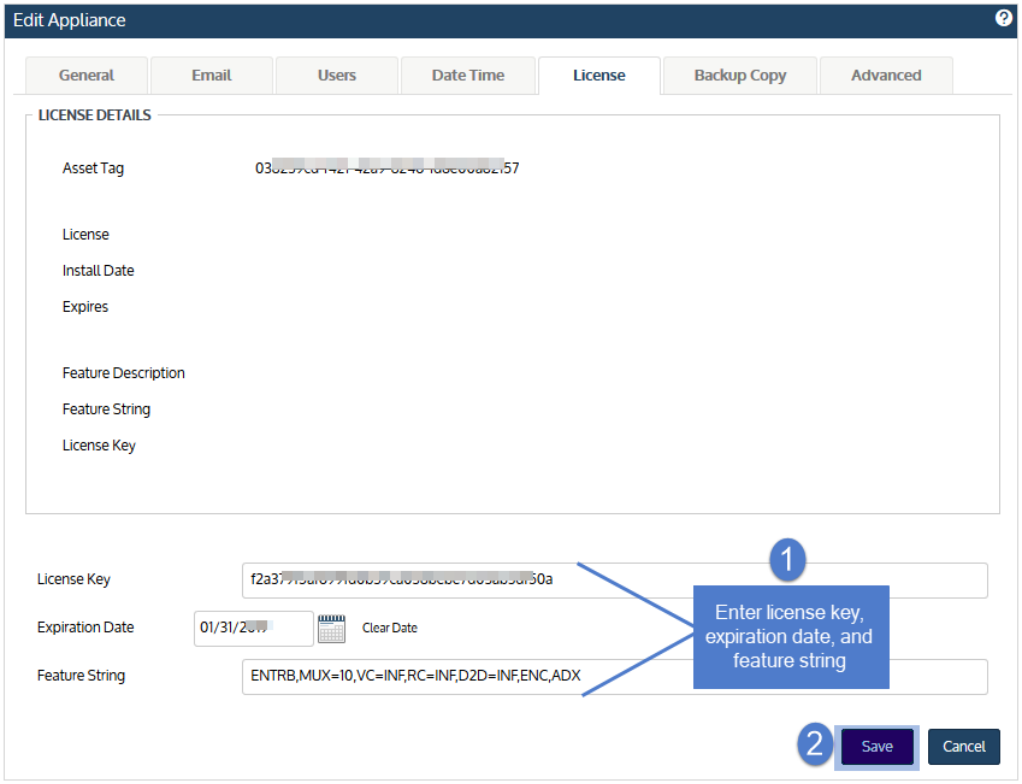
Edit Appliance
General
Email
Users
Date Time
License
Backup Copy
Advanced
LICENSE DETAILS
Asset Tag
License
Install Date
Expires
Feature Description
Feature String
License Key
License Key
Expiration Date
Clear Date
Feature String
ENTRB,MUX=10,VC=INF,RC=INF,D2D=INF,ENC,ADX
1
Enter license key, expiration date, and feature string
2
Save
Cancel

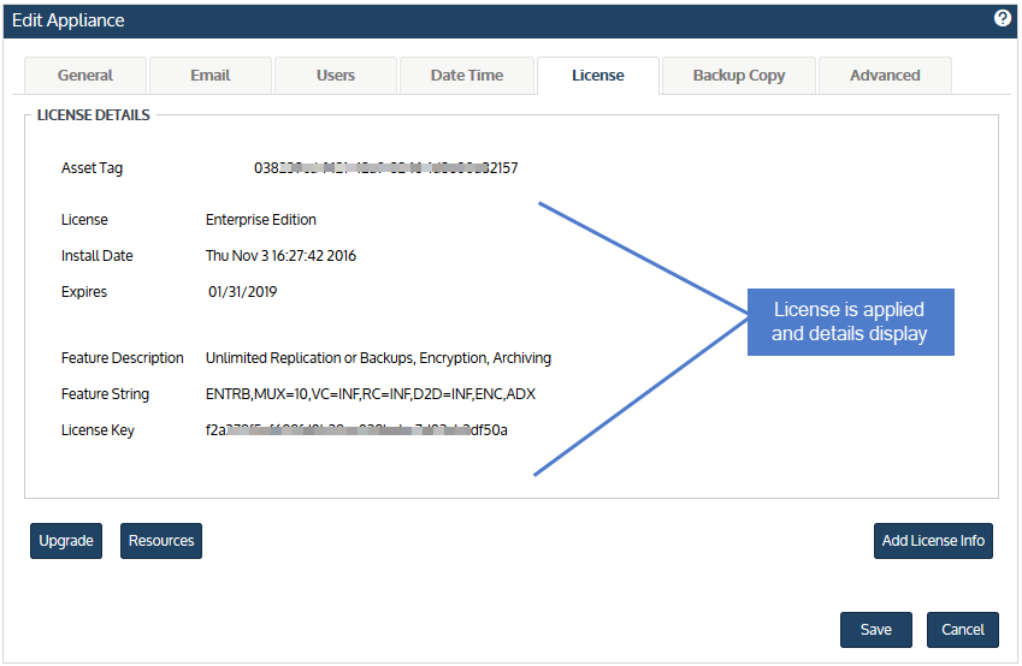
Edit Appliance
General
Email
Users
Date Time
License
Backup Copy
Advanced
LICENSE DETAILS
Asset Tag
License
Enterprise Edition
Install Date
Thu Nov 3 16:27:42 2016
Expires
01/31/2019
Feature Description
Unlimited Replication or Backups, Encryption, Archiving
Feature String
ENTRB,MUX=10,VC=INF,RC=INF,D2D=INF,ENC,ADX
License Key
License is applied and details display
Upgrade
Resources
Add License Info
Save
Cancel

To view license resources for Unitrends Backup appliances

Use these steps to see the number of application, server, socket, VM, and workstation resources you can protect with your license, and the appliance's current resource usage.

1 On the **Configure > Appliances** page, select the appliance and click **Edit**.

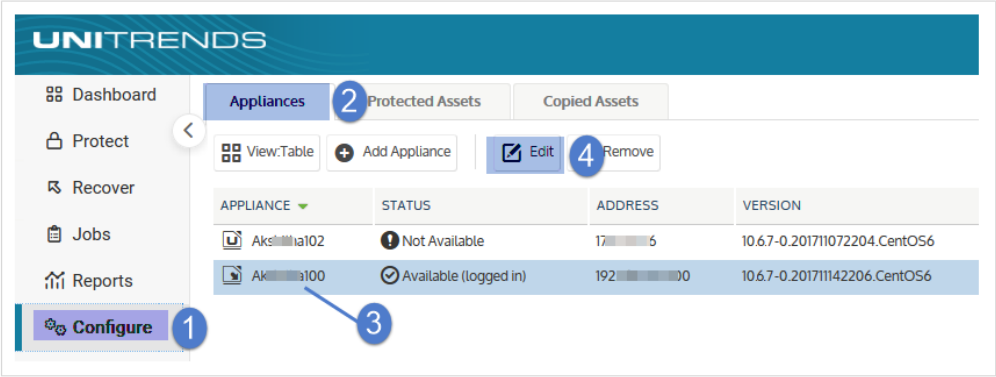

2 Select the **License** tab and click **Resources**.

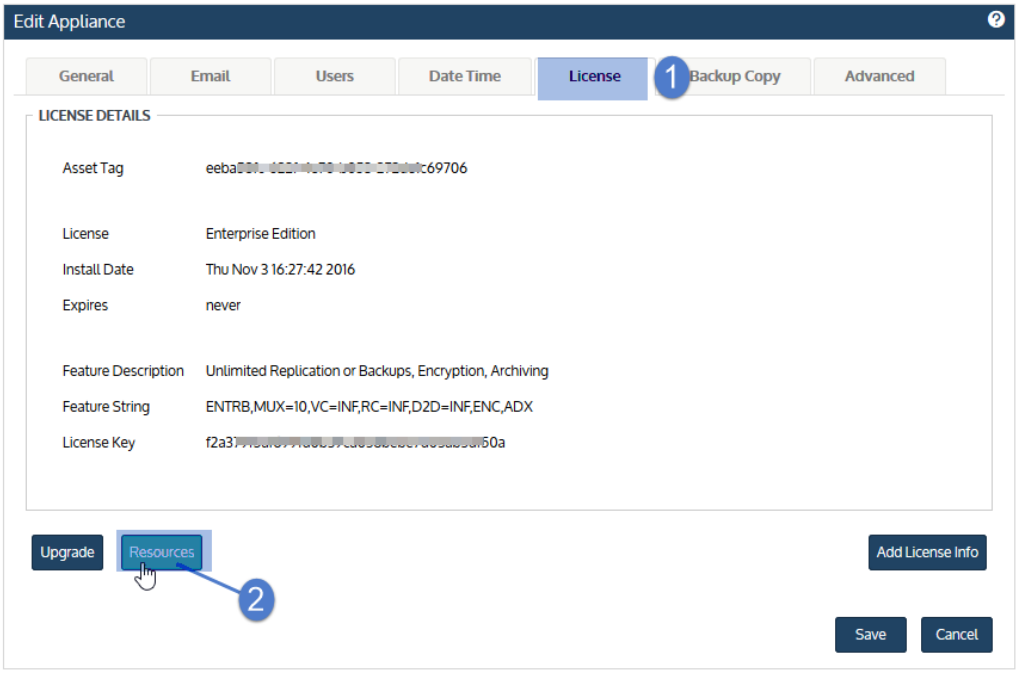

3 Resource limits and usage display.

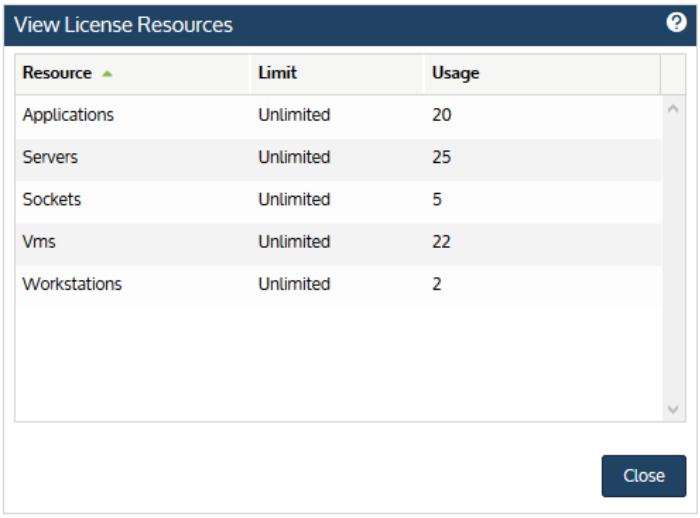

4 Click **Close** to exit.

Encryption

Use encryption to protect data from unauthorized access and theft. All data remains encrypted until a request is made to recover the data. If the correct passphrases are in place, recovery proceeds without administrator involvement.

Unitrends encryption provides:

- Encryption at the asset level.
- The ability to manage and change passphrases.
- Backup, backup copy, and recovery of encrypted data.

To set up encryption, you provide a passphrase and save the master key file. Once you configure encryption, you can encrypt backups by asset (**Configure > Protected Assets > Edit Asset**). See these topics for details:

- "Encryption considerations and limitations"
- "To configure encryption" on page 156
- "To change the encryption passphrase" on page 159

Encryption considerations and limitations

The following encryption limitations apply:

- Encryption slightly degrades performance for backups, backup copies, and recovery. Use encryption only if you really need to hide your data.
- Make sure to keep the passphrase secure. If you forget the passphrase, there is no way to recover it or recover any encrypted backups.

- Once you have enabled encryption for an asset, that asset's subsequent backups are encrypted. Encryption takes place during backup jobs. When unencrypted backups run on an appliance before you configure encryption, those backups remain unencrypted.
- Small Form Factors (SFF) do not support encryption.
- The following backup types are not encrypted:
  - Legacy MS Exchange Information Store backups
  - CEP brick-level backups
  - Any data stored on the appliance via Samba or NFS

To configure encryption

1 On the **Configure > Appliances** page, select the appliance and click **Edit**.

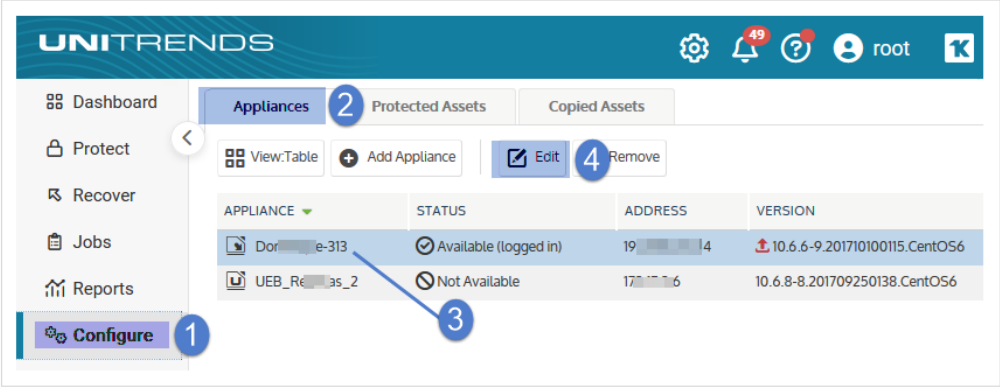

2 Select the **Advanced** tab.

3 Check **Enable Encryption**.

4 Enter a **Passphrase** and **Confirm Passphrase**.

IMPORTANT!	Be sure to keep the passphrase secure. If you forget the passphrase there is no way to recover it.

5 Click **Save**.

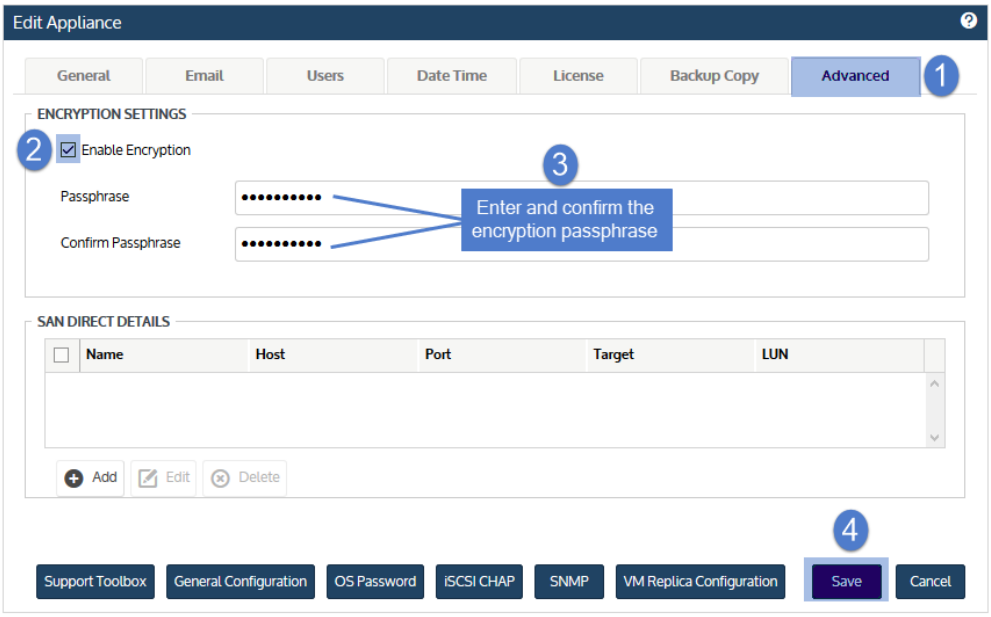

6 Return to the Edit Appliance dialog.

7 Select the **Advanced** tab and click **Save Master Key File**.

8 (Optional) If your appliance has a CD drive, you can save the key file directly to a CD. Insert a CD into your appliance's CD drive. (If no CD is inserted, the key file is saved to the appliance's samba share.)

9 Click **Continue**.

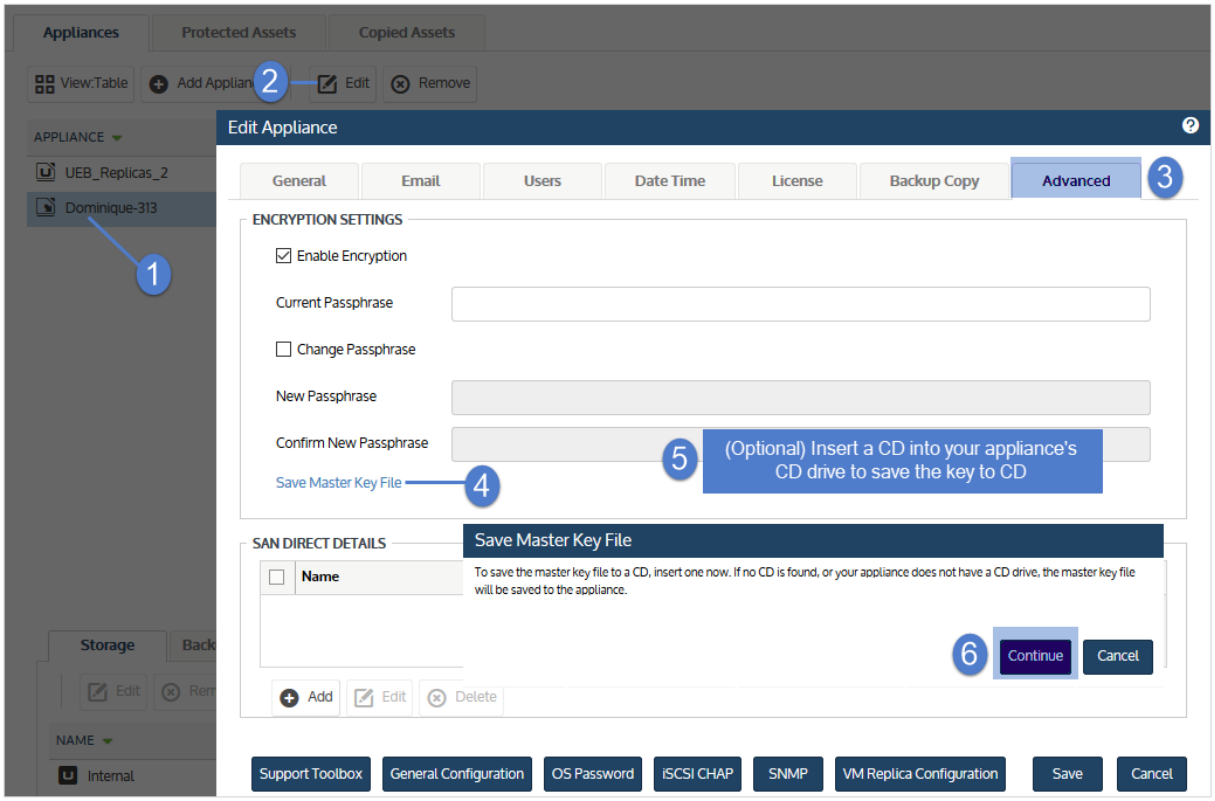

10 You receive a message indicating the master key file was saved to the appliance's samba share or to CD. Click **OK**.

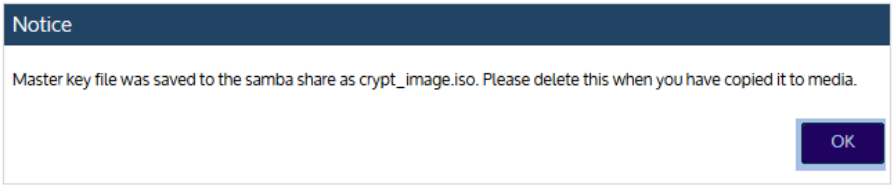

IMPORTANT! Be sure to keep the master key file secure. If you ever need to perform disaster recovery of the appliance, you will need this key to access any encrypted backups.

11 If you saved the key to the appliance's samba share, do these steps:

- Log in to a Windows workstation as an administrator with full system access.
- Launch File Explorer and enter the following path to access the master key file on the Unitrends appliance: *\\ApplianceIP\samba*
- Copy the master key file, called *crypt_image.iso*, to removable media and store it in a safe location.
- Once you have copied the key to removable media, delete *crypt_image.iso* from *\\ApplianceIP\samba* for increased security.

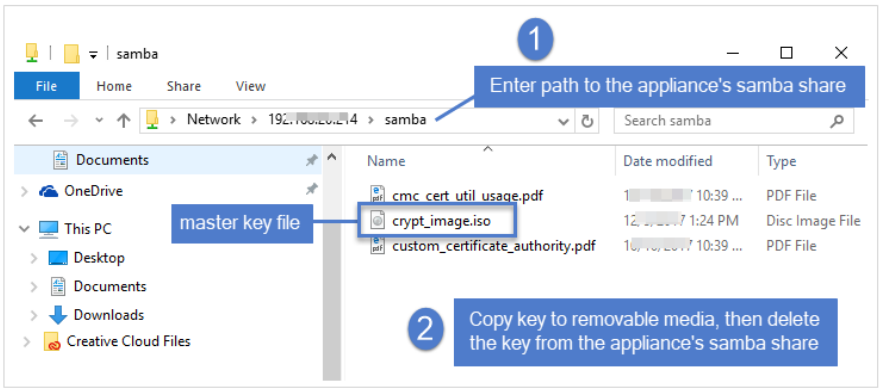

12 Enable encryption for each asset whose backups you wish to encrypt. To enable encryption for an asset, go to **Configure > Protected Assets**. Select the asset, click **Edit**, select **Encrypt Backups**, then **Save**.

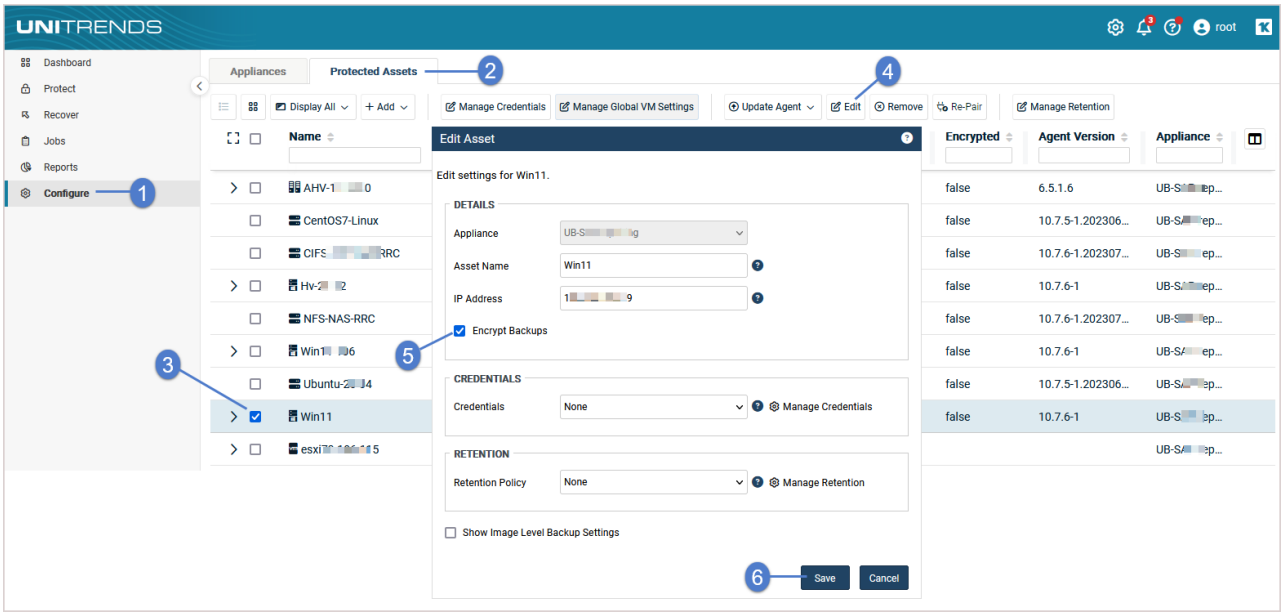

To change the encryption passphrase

1 On the **Configure > Appliances** page, select the appliance and click **Edit**.

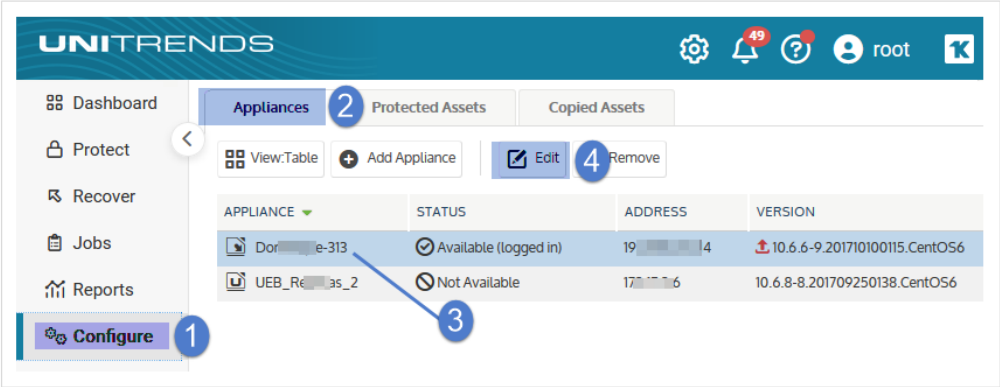

2 Select the **Advanced** tab.

3 Enter the **Current Passphrase**.

4 Check **Change Passphrase** and enter a **New Passphrase** and **Confirm New Passphrase**.

> **IMPORTANT!** Be sure to keep the passphrase secure. If you forget the passphrase there is no way to recover it.

5 Click **Save Master Key File**.

6 (Optional) If your appliance has a CD drive, you can save the key file directly to a CD. Insert a CD into your appliance's CD drive. (If no CD is inserted, the key file is saved to the appliance's samba share.)

7 Click **Continue**.

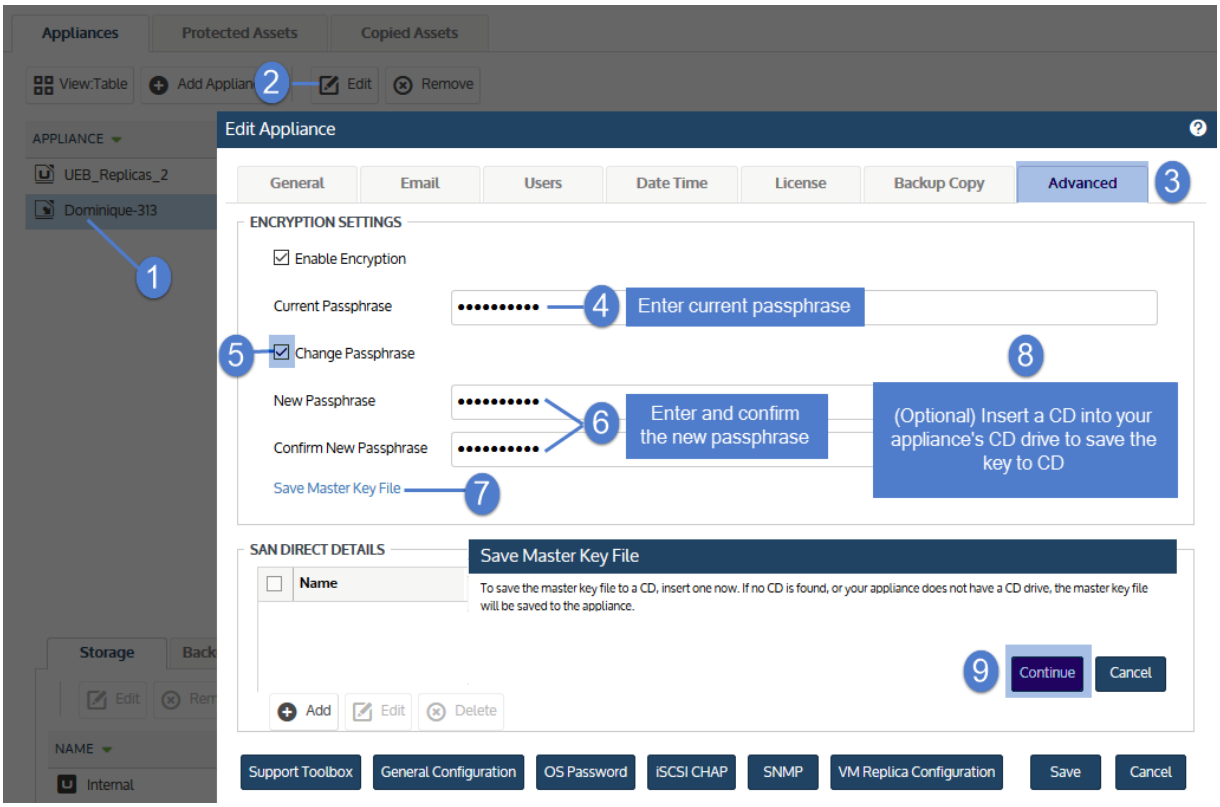

8 You receive a message indicating the master key file was saved to the appliance's samba share or to CD. Click **OK**.

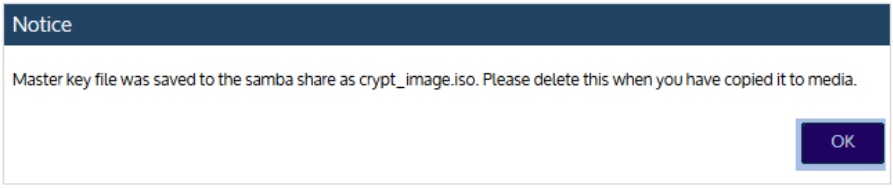

IMPORTANT! Be sure to keep the master key file secure. If you ever need to perform disaster recovery of the appliance, you will need this key to access any encrypted backups.

9 Click **Save** in the Edit Appliance dialog.

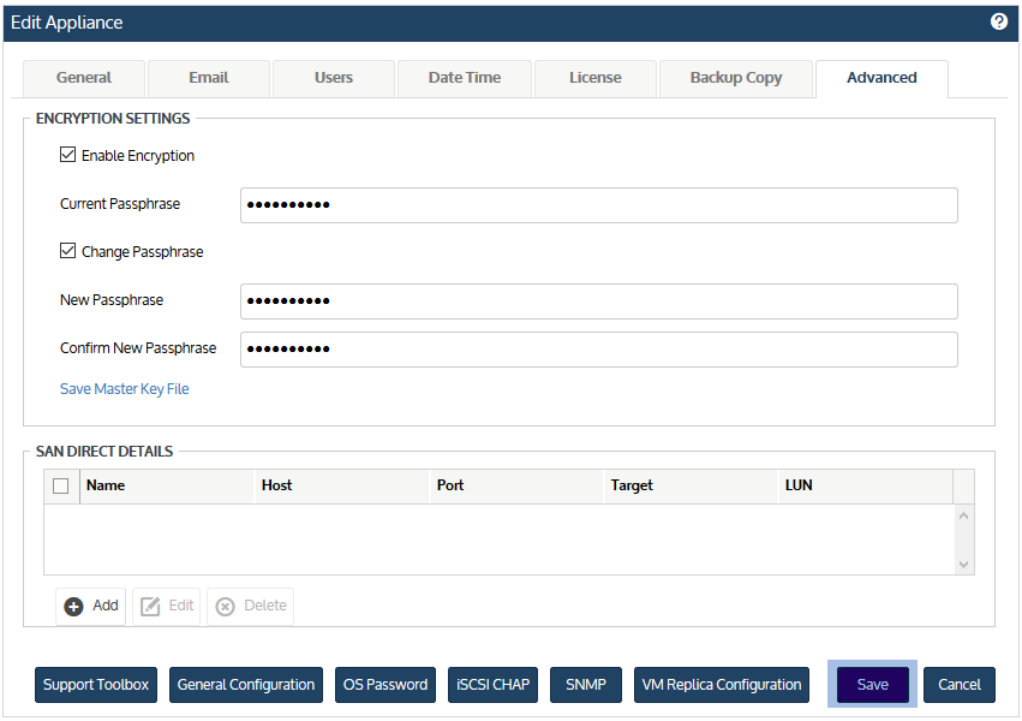

10 If you saved the key to the appliance's samba share, do these steps:

- Log in to a Windows workstation as an administrator with full system access.
- Launch File Explorer and enter the following path to access the master key file on the Unitrends appliance: *\\ApplianceIP\samba*
- Copy the master key file, called *crypt_image.iso*, to removable media and store it in a safe location.
- Once you have copied the key to removable media, delete *crypt_image.iso* from *\\ApplianceIP\samba* for increased security.

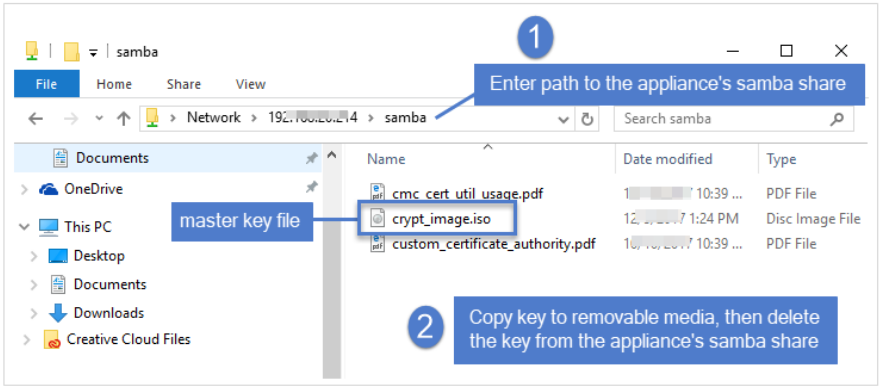

Add the appliance to your UniView Portal

Note: After adding the appliance to UniView, it is recommended that you disable local network access for increased security. Once local network access is disabled, users must access the appliance UI from UniView (as described in "To log in to the appliance UI from UniView" on page 31). To disable local network access, see "Disable or enable local network access to an appliance" on page 170.

1 Log in to your UniView Portal:

- Open a Firefox or Chrome browser and enter ***https://login.backup.net/*** to access the login page.

- Enter the backup.net homerealm that was provided to you by the UniView Portal Onboarding team. Click **Next**.

- Enter the username and password of your UniView Portal account. Click **Log In**.

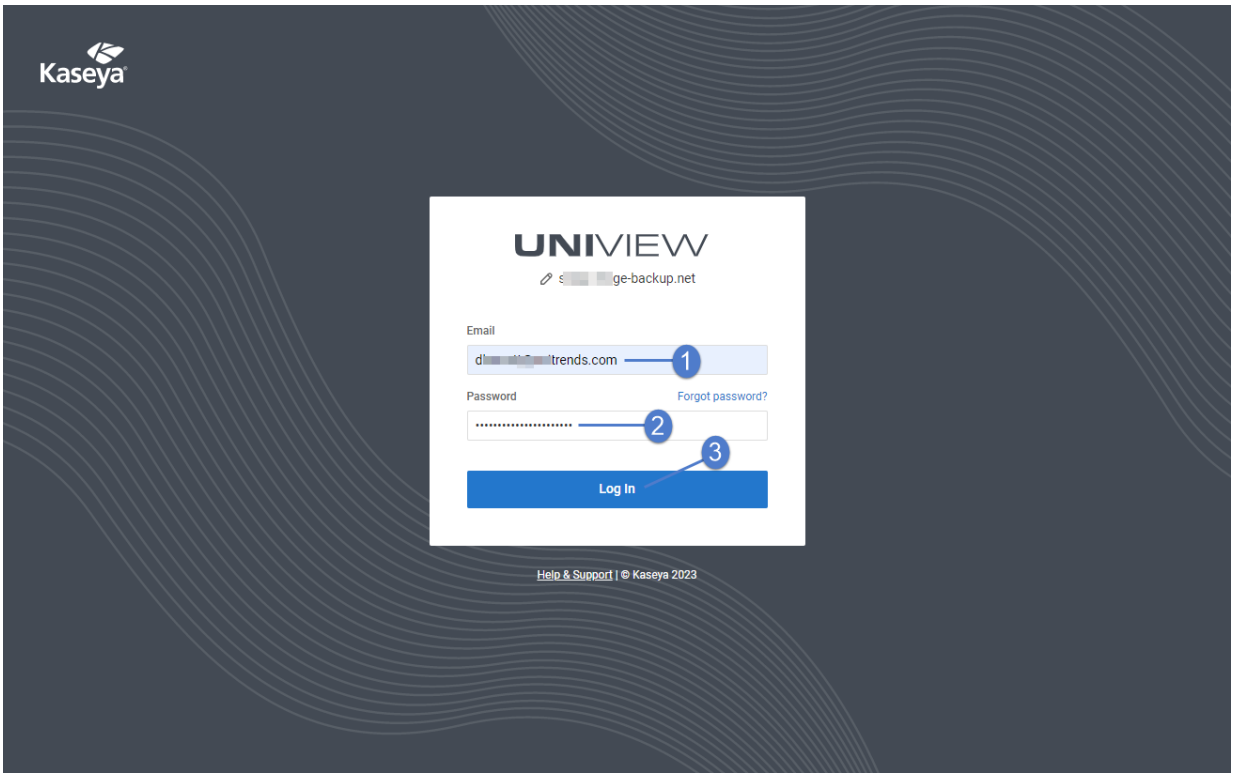

- Enter your two-factor authentication (2FA) code, then click **Verify**. You can obtain the code from your authenticator app or use a recovery code.

  Note: Only use a recovery code if you have lost your IOS or Android device, or cannot access your authenticator application for some other reason.

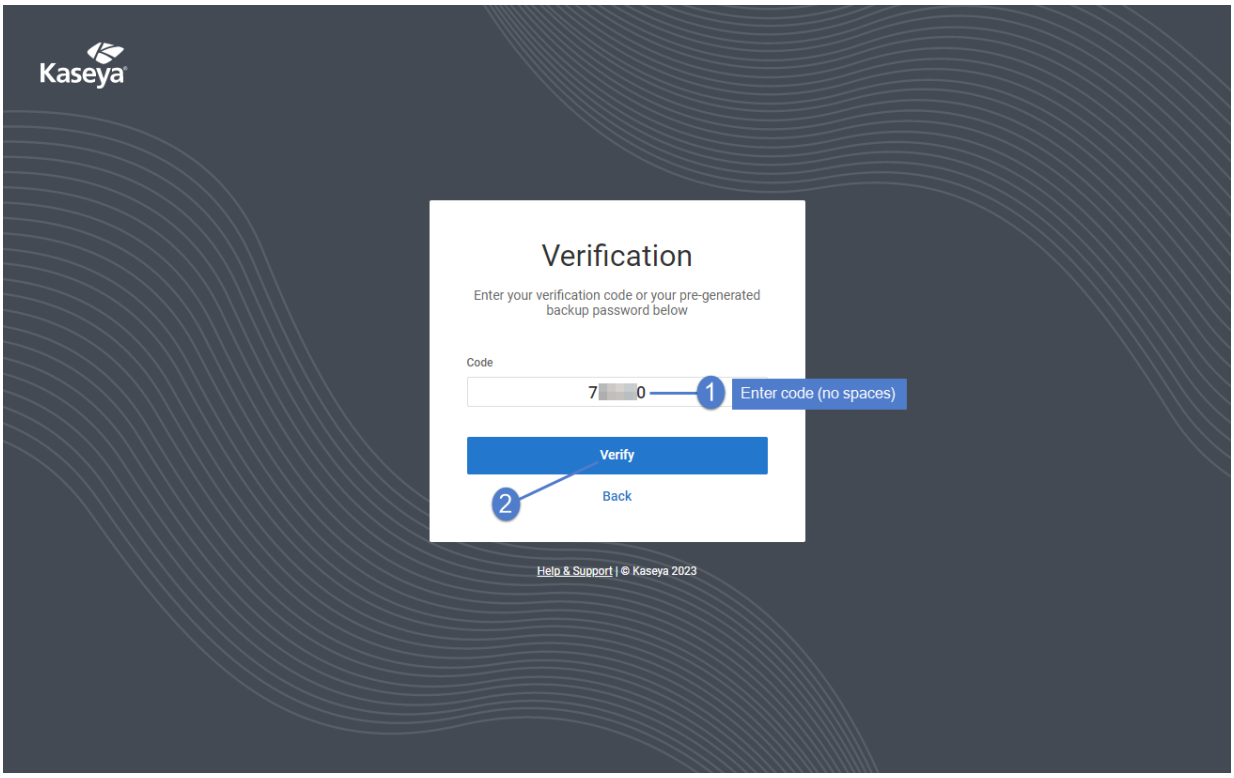

- Upon logging in, the portal Dashboard displays.

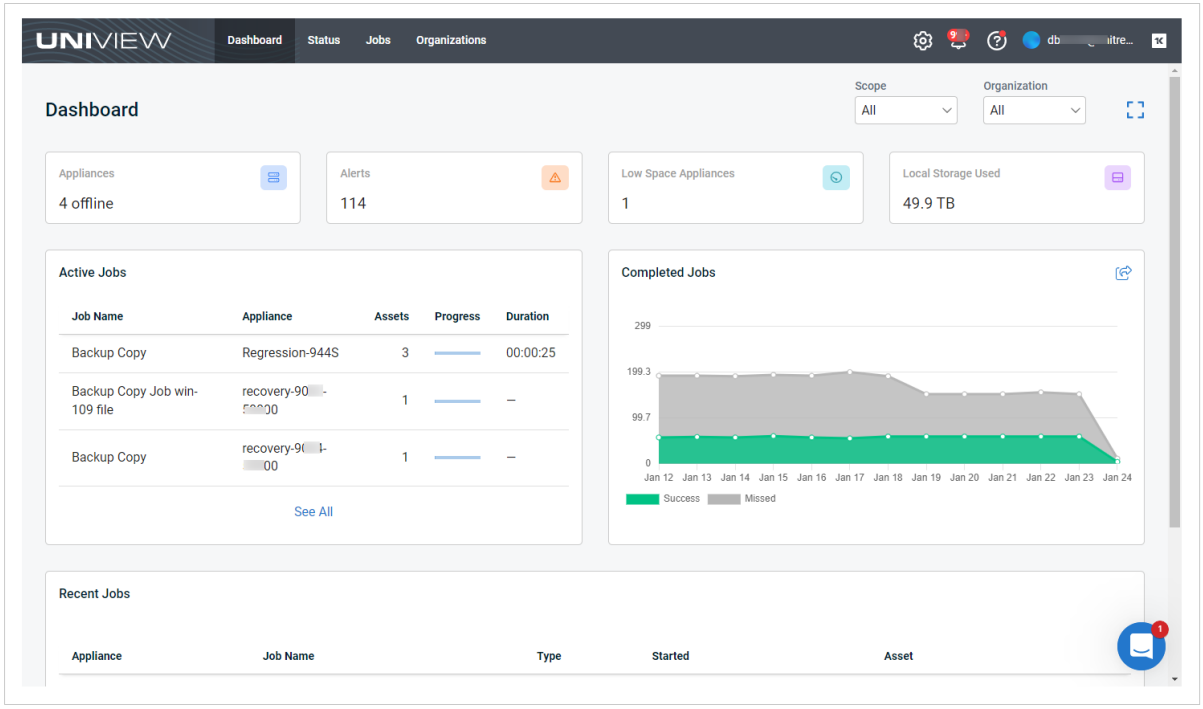

2 In the Appliances view, click **New +**.

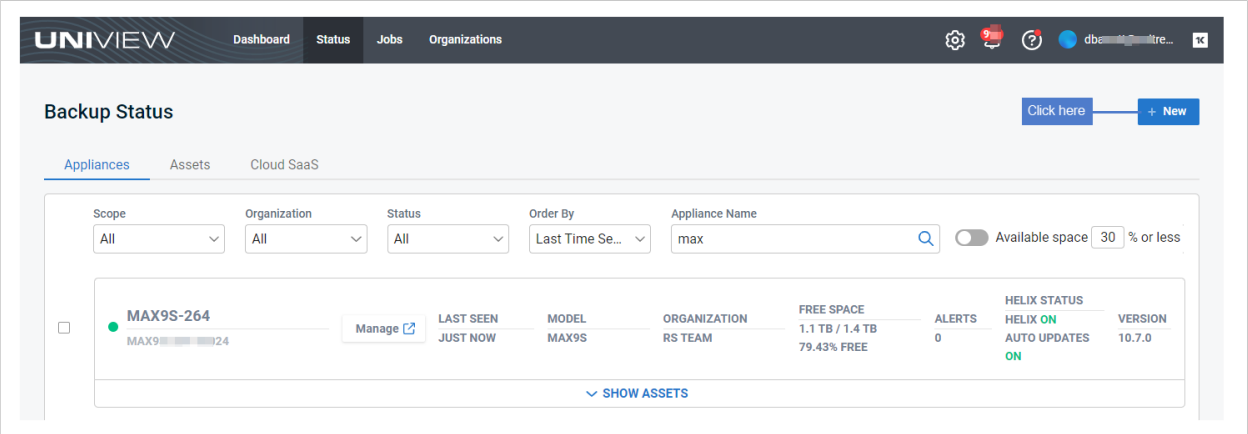

3 Do one of the following:

- Select an organization from the Name list.

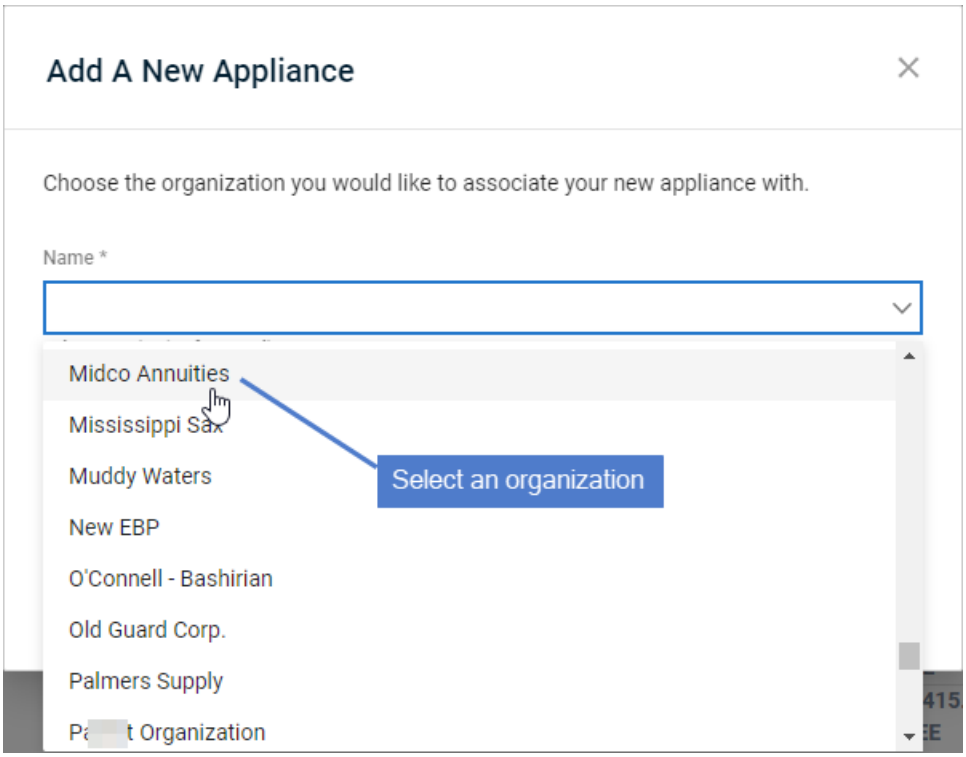

OR

- Click **Add New Organization** and do one of the following:

- If you have integrated your PSA system ((ConnectWise Manage, Autotask, or BMS), use the Import Organizations dialog to select an organization and scope.
- If you have not integrated a PSA system, use the Add Organization dialog to enter the organization name, select a scope, and click **Save**.

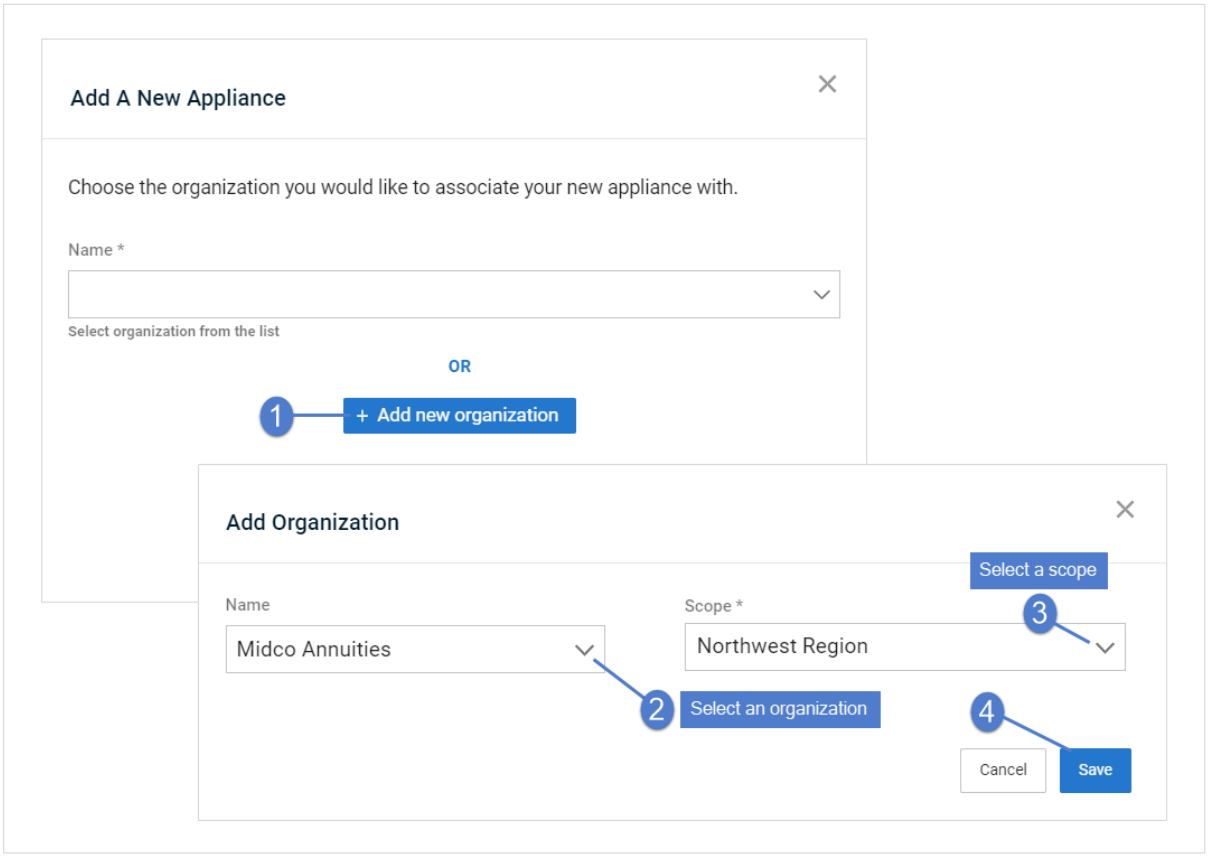

4 Click **Next**.

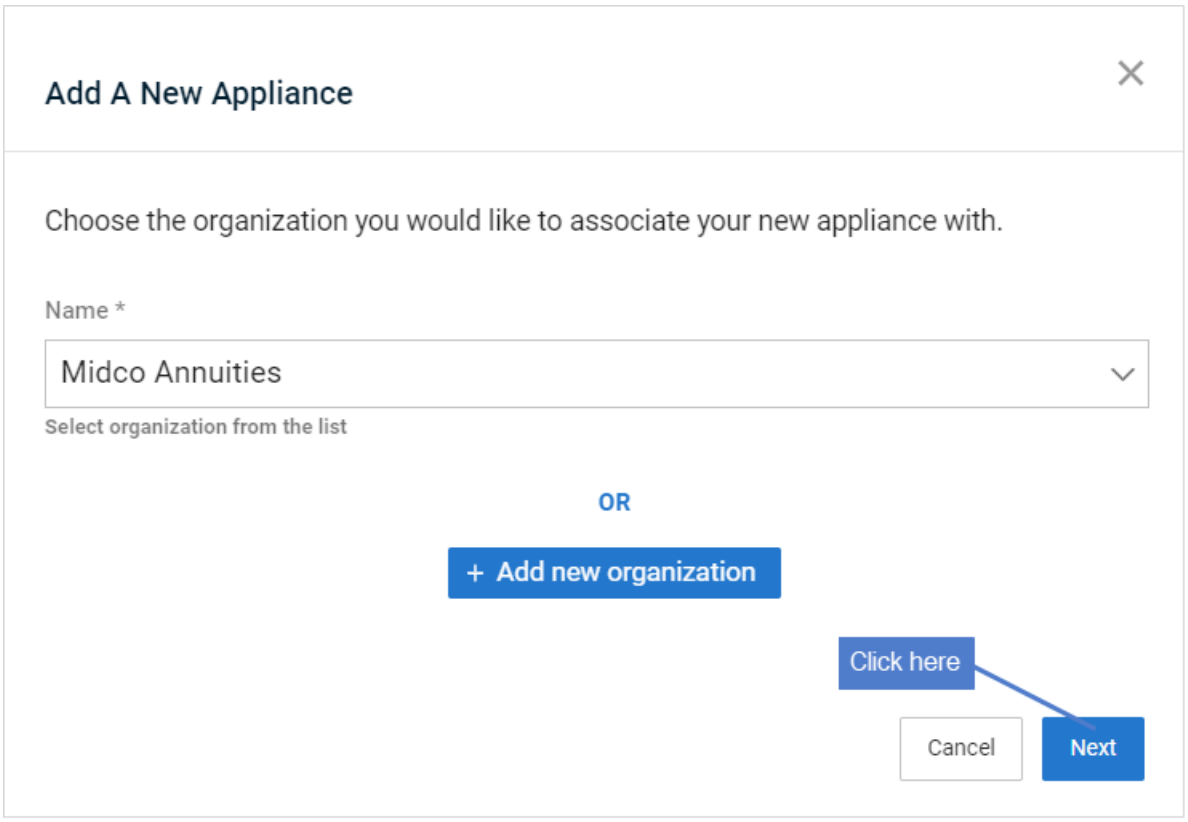

5 Copy the backup.net ID that displays. (Leave this dialog open. Do *not* click Done.)

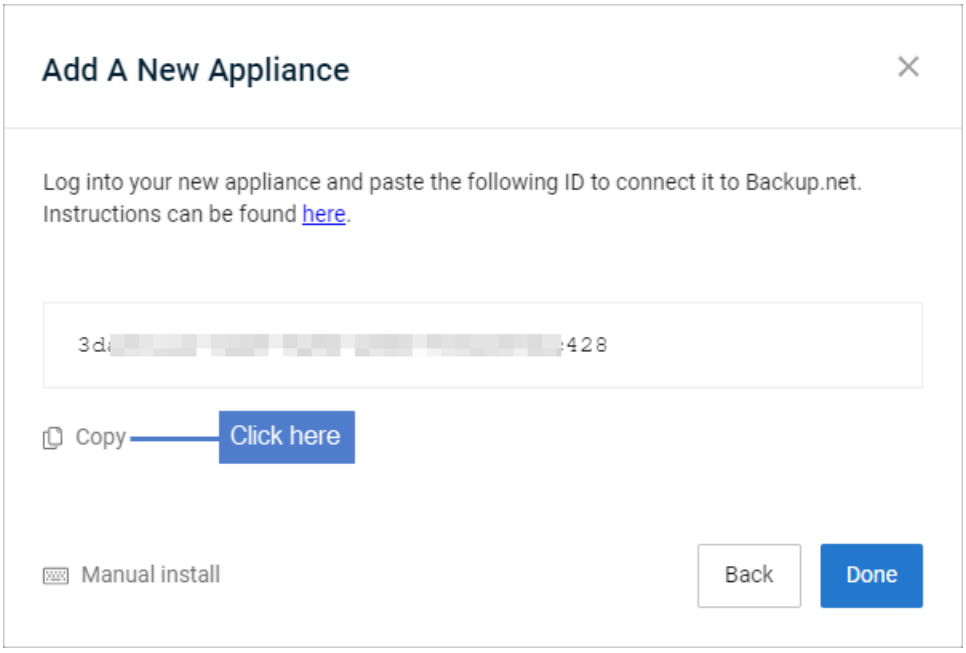

6 Log in to the Unitrends appliance. In the Edit Appliance dialog, click **Backup.net Config**, then paste the ID into the Backup.net ID field. Click **Save**. Click **OK**.

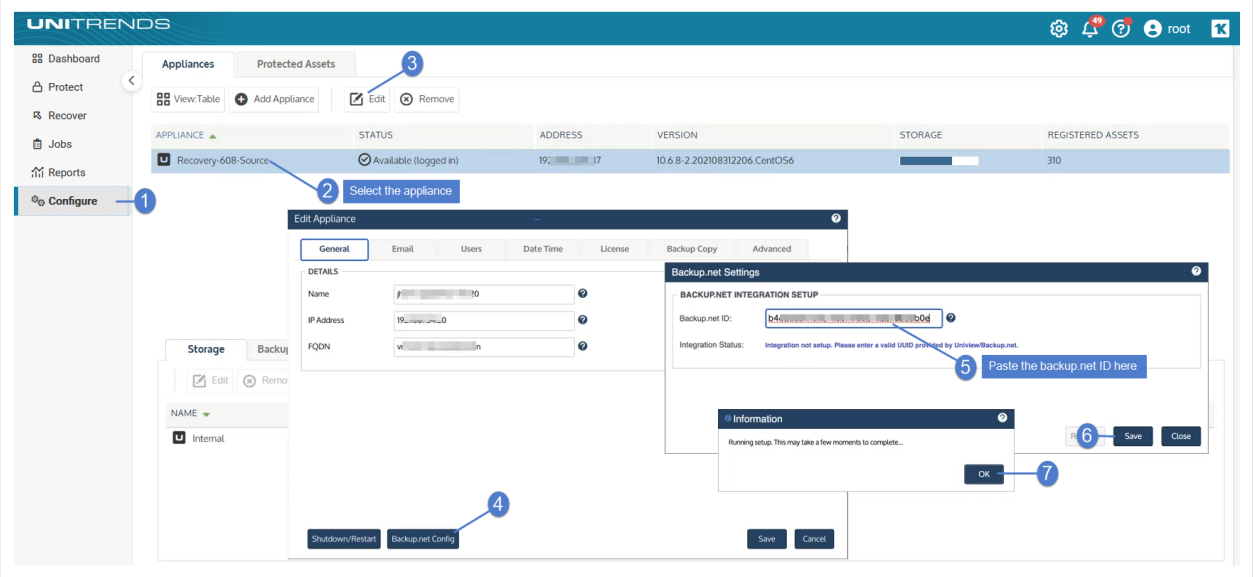

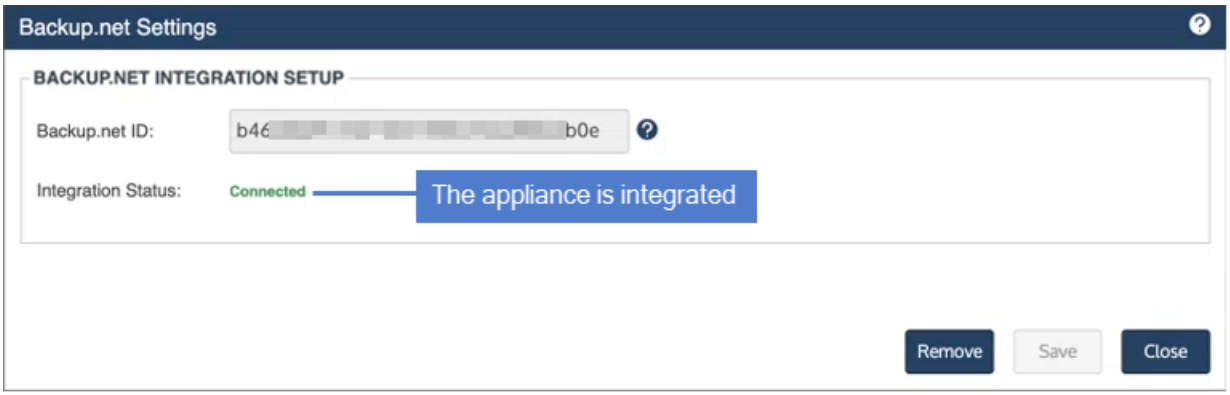

7 Return to the UniView Portal. Click **Done**.

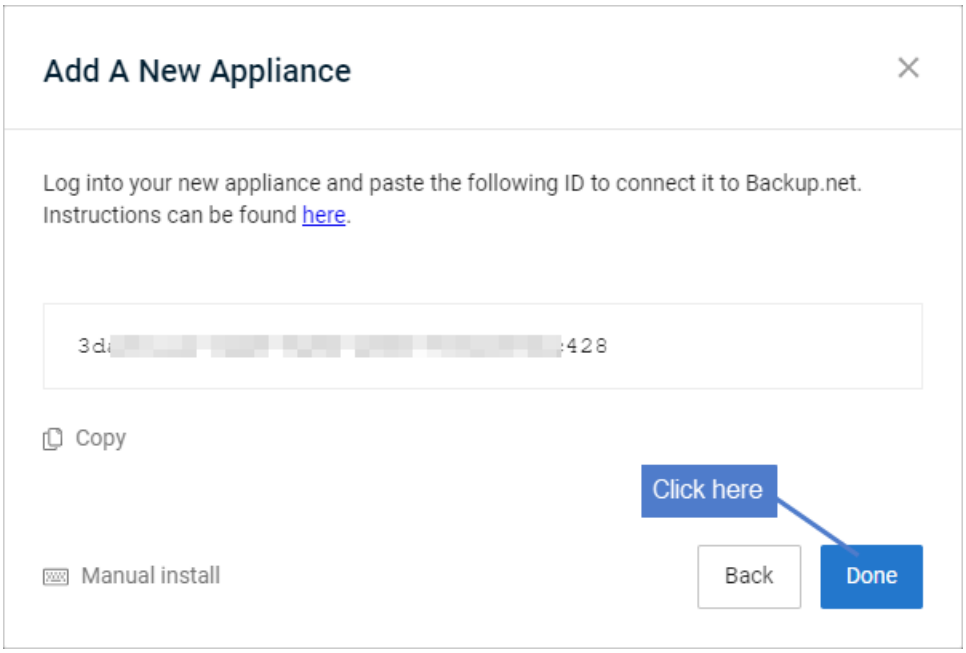

The appliance is added and displays in the Appliances list.

Notes:

- It can take a few minutes for the appliance to display in the list. If needed, refresh the page.
- For increased appliance security, the UniView Portal has a feature that blocks users from logging in directly to the appliance UI. Once local network access has been disabled, users must connect to the appliance from UniView. To use this feature, see "Disable or enable local network access to an appliance" on page 170.

Disable or enable local network access to an appliance

UniView allows users to restrict local access to the Unitrends appliance on the local network. The appliance UI and management functions can still be accessed through UniView. Disabling local access enforces MFA, significantly reduces potential security exposure, and allows admins greater access controls through roles and scopes in UniView.

Once local network access has been disabled, users must connect to the appliance from UniView (as described in "To log in to the appliance UI from UniView").

Consider the following before disabling local network access:

- To enable or disable local network access, you must log in to UniView as a Superuser, Admin, or Manage user. (UniView users with Monitor access cannot enable or disable local network access.)
- To enable or disable local network access, the Unitrends appliance must be running version 10.7.2 or higher.
- Hot backup copy to a Unitrends appliance target – To add a Unitrends appliance backup copy target to the appliance, local network access must be enabled on the backup copy target appliance. If needed, use the

procedure below to enable local network access on the target appliance before adding the hot backup copy target. Once the target has been added, use the procedure below to disable local network access.

- iSeries protection – To protect your iSeries platform, you must access the appliance directly from the local network. Do NOT disable local network access if your appliance is protecting an iSeries environment.
- Appliance disaster recovery (DR) – Local network access must be enabled on the DR target appliance. Once you have recovered the configuration and last backups from the failed appliance, use this procedure to disable local network access on the target appliance.

To disable or enable local network access

1 In the Appliances view, click the appliance.

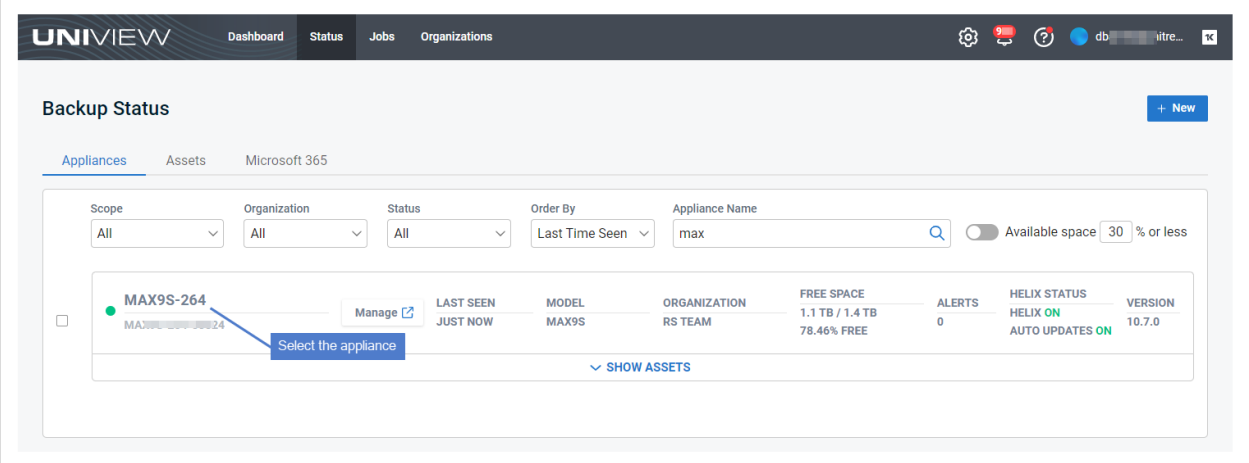

2 On the Appliance Details page, scroll down to the Settings section:

- indicates local network access is enabled (users can access the appliance UI by entering *https://<applianceIPaddress>/ui/* in a browser on the local network).
- indicates local network access is disabled (users must access the appliance UI by logging in to UniView and clicking the appliance's Manage button, as described in "To log in to the appliance UI from UniView").

3 Do one of the following:

- Click to disable local network access.
- Click to enable local network access.

Notes:

- If the toggle is disabled ( ), you are using Monitor user credentials and cannot enable or disable local network access.
- It may take a minute or two to disable or enable access. During this transition time, the or toggle is disabled ( ).

- If the appliance is not running version 10.7.2 or higher, this feature is not supported and you see this message:

  Appliance version not supported
  Please upgrade your appliance to enable this feature

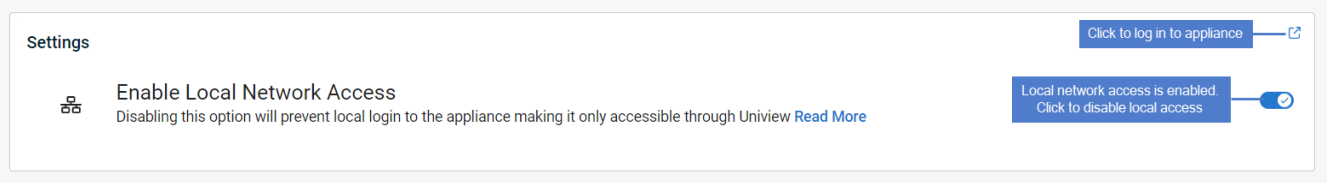

Remove the appliance from your UniView Portal

IMPORTANT! If this appliance is managed by UniView (e.g., users can only log in to the appliance UI from UniView), local network access to the appliance is disabled. Before removing the appliance from the UniView Portal, you must enable local network access if users will still need access the appliance UI. For details, see "Disable or enable local network access to an appliance" on page 170.

1 Log in to your UniView Portal:

- Open a Firefox or Chrome browser and enter ***https://login.backup.net/*** to access the login page.

- Enter the backup.net homerealm that was provided to you by the UniView Portal Onboarding team. Click **Next**.

- Enter the username and password of your UniView Portal account. Click **Log In**.

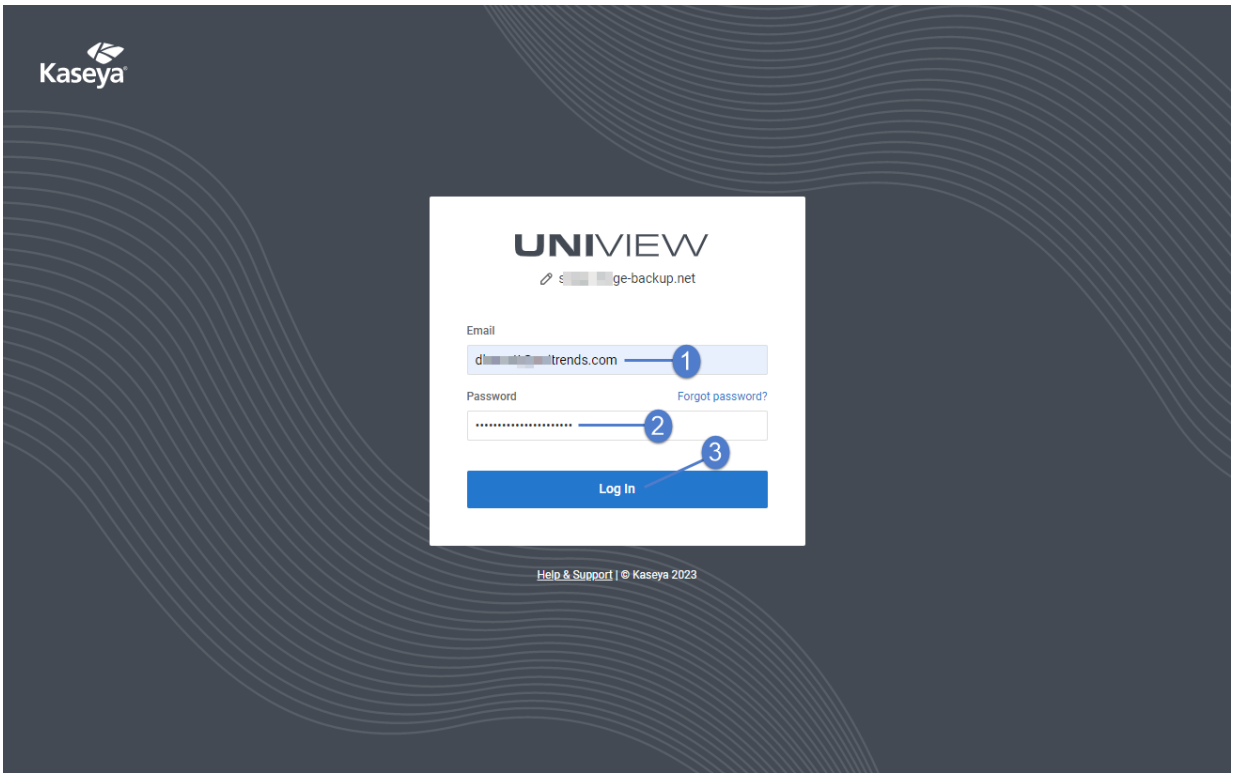

- Enter your two-factor authentication (2FA) code, then click **Verify**. You can obtain the code from your authenticator app or use a recovery code.

  Note: Only use a recovery code if you have lost your IOS or Android device, or cannot access your authenticator application for some other reason.

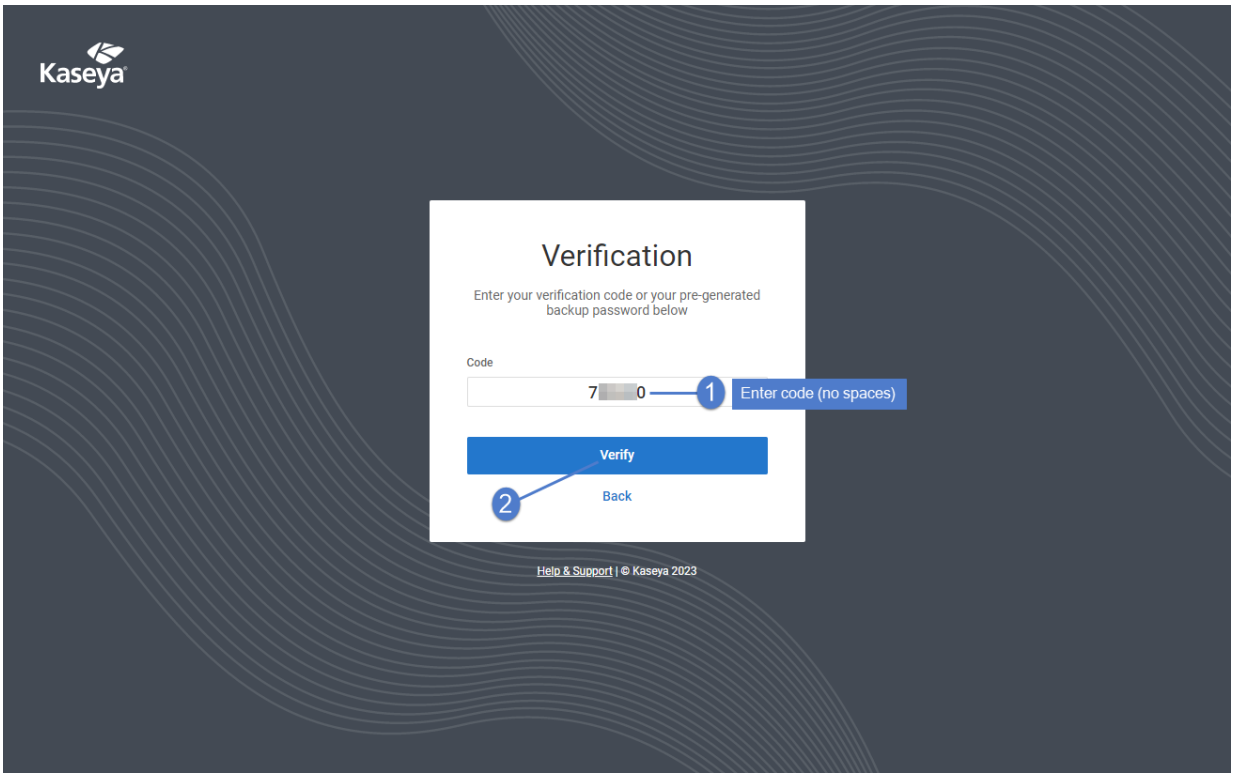

- Upon logging in, the portal Dashboard displays.

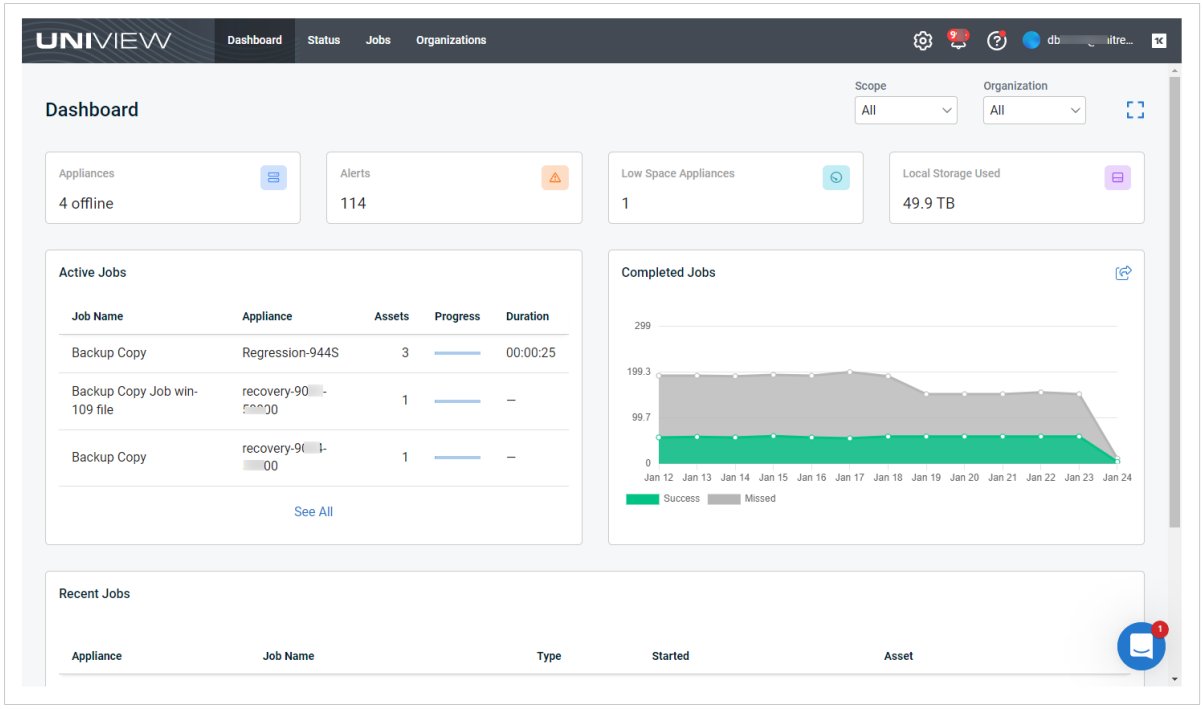

2 In the Appliances view, click the appliance.

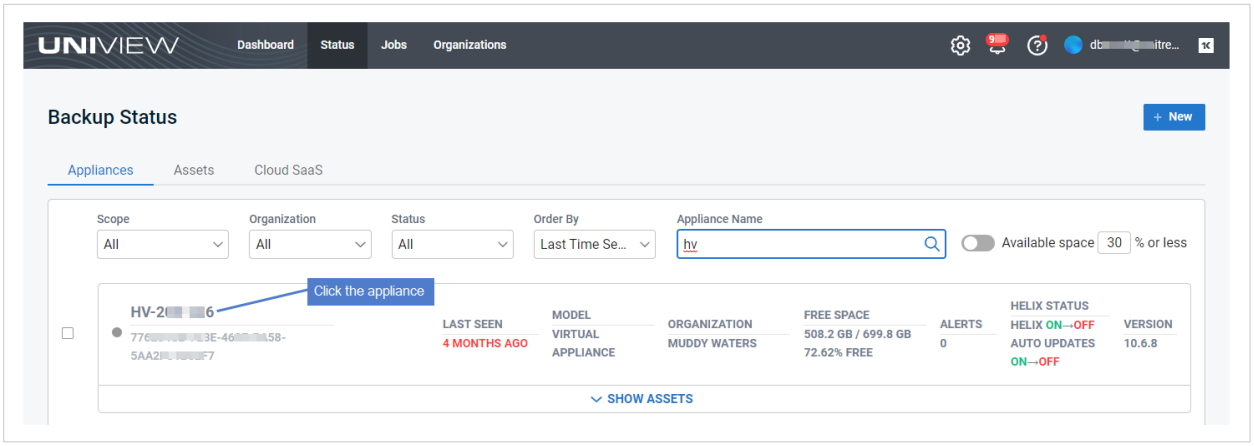

3 On the Appliance Detail page, click . Type DELETE and click the **Delete** button.

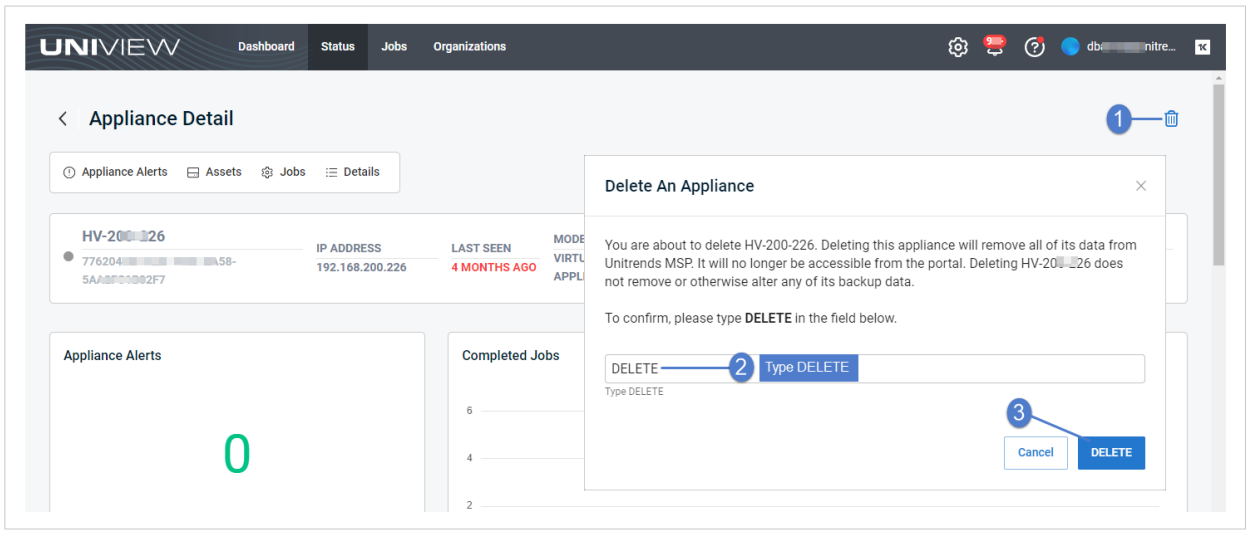

4 Log in to the Unitrends appliance. In the Edit Appliance dialog, click **Backup.net Config**, then click **Remove**. Click **OK**.

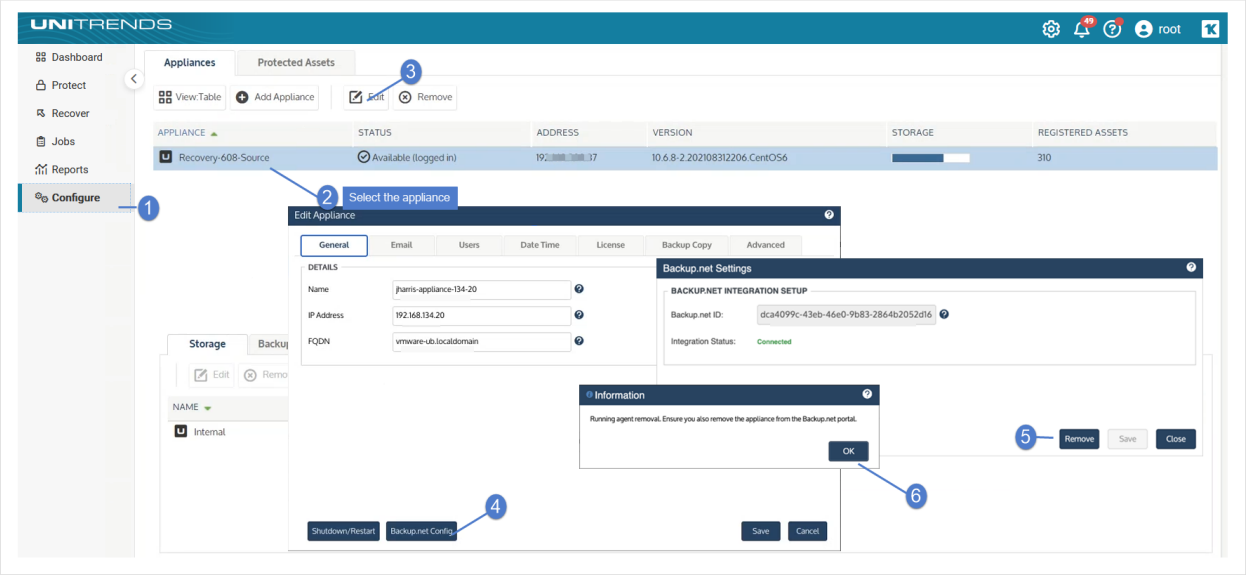

VM replica configuration

The VM replicas feature creates standby replicas of critical VMs that can be brought online in seconds. During replica setup, the user specifies the number of recovery point snapshots to retain with the replica VM on the hypervisor.

By default, the appliance supports retention of up to 14 replica snapshots. You can modify the maximum number allowed by the Unitrends appliance by doing these steps:

1 On the **Configure > Appliances** page, select the appliance and click **Edit**.

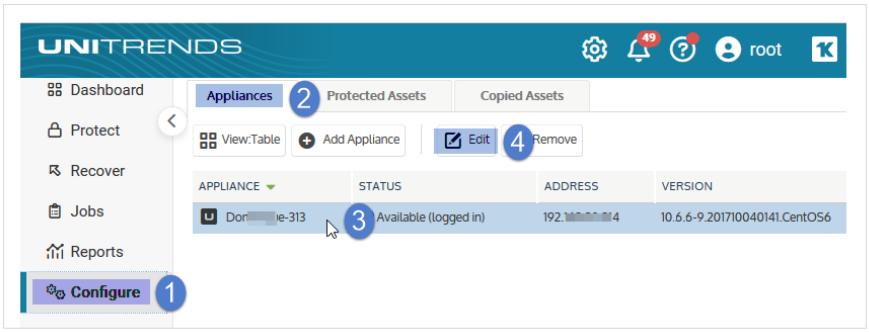

2 In the Edit Appliance dialog, click **Advanced** and select **VM Replica Configuration**.

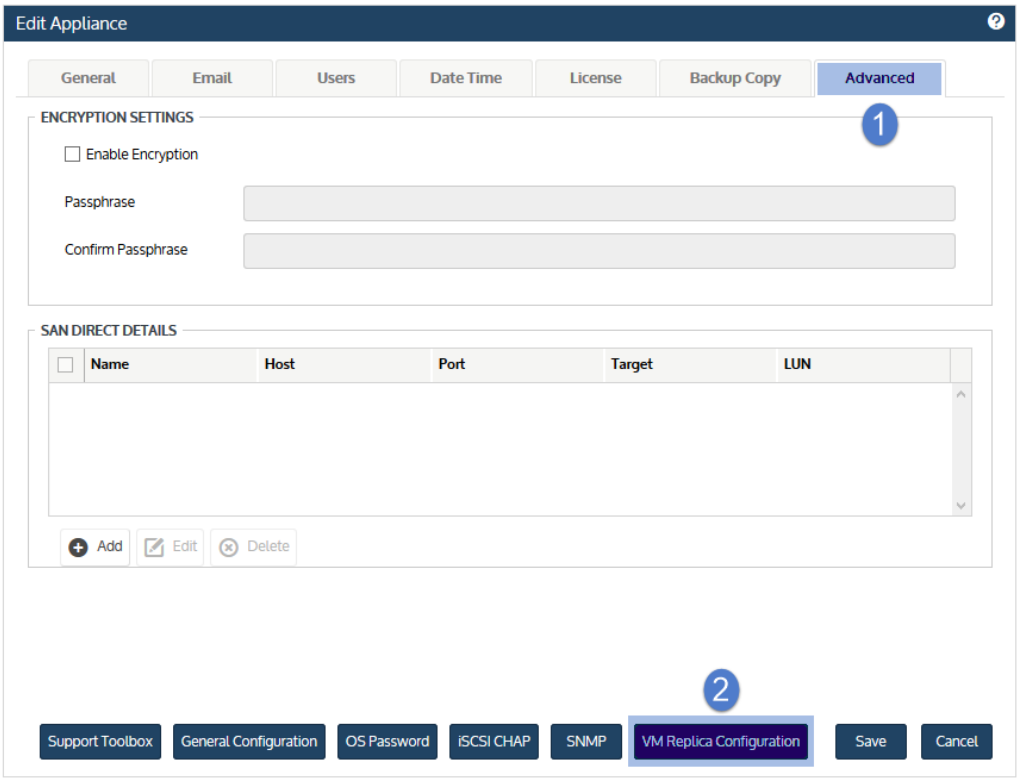

3 Modify the number of Maximum Recovery Points. Supported values are 1-28.

4 Click **Save**.

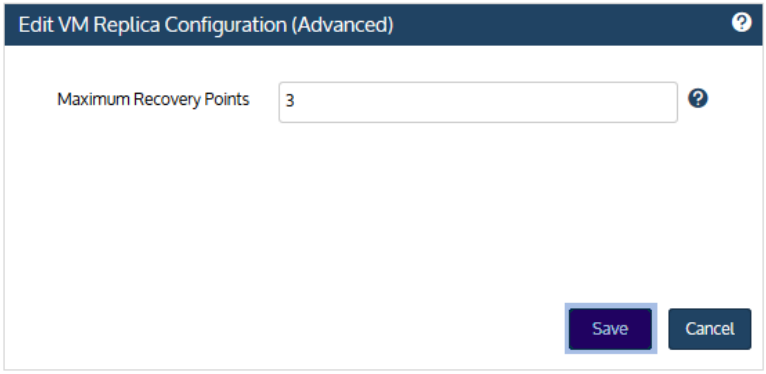

SNMP trap notifications

You can configure your appliance to send system and application-specific alerts to your network management server using the SNMP protocol. Alerts are delivered as incoming trap messages to the network management application. This enables you to quickly identify and respond to hardware or software conditions that require action.

Through the use of the Unitrends SNMP agent and MIB, you can configure alerts to be sent to your own Remote Monitoring and Management (RMM) software.

SNMP agent requirements

To use the Unitrends agent:

- The appliance must be running version 9.0.0-12 or higher.
- The Unitrends SNMP agent supports SNMP gets with SNMP version 1, 2c, and 3.

To set up SNMP trap notifications

1. If you will be using SNMP V3, configure the username and password from the command line as follows:
   - Using a terminal emulator, such as PuTTY, connect to the appliance using the following:
     - Appliance IP address
     - Port 22
     - SSH connection type
   - Log in as user *root*.
   - Enter the following command to configure the SNMP V3 username and password:

     ```
     # /usr/bp/bin/cmc_snmpd user create<snmp_user><snmp_passwd>
     ```

     The script defaults to authorization type MD5 and privacy/encryption of DES.
2. Log in to the appliance UI.
3. On the **Configure > Appliances** page, select the appliance and click **Edit**.

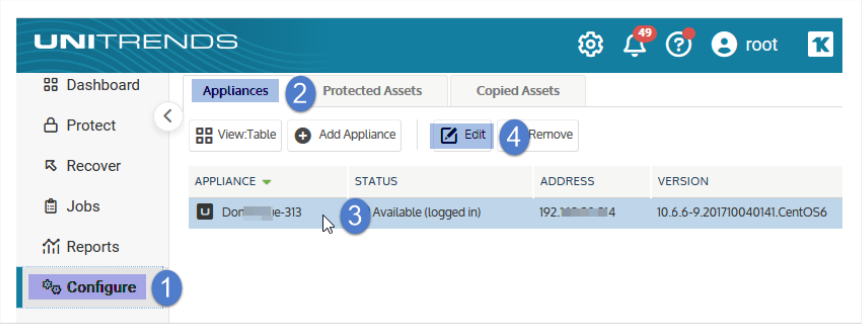

4. Click **Advanced** and select **SNMP**.

> **Note:** Your appliance comes configured with a default destination of *notifications.unitrends.com*. Various system alerts are sent to this address to enable Unitrends to proactively resolve problems, if and when they arise. For example, if a disk drive is failing, Unitrends receives a trap and dispatches a warranty request on the failed component (if the appliance support contract is up-to-date). This destination must remain in place for proactive monitoring to continue.

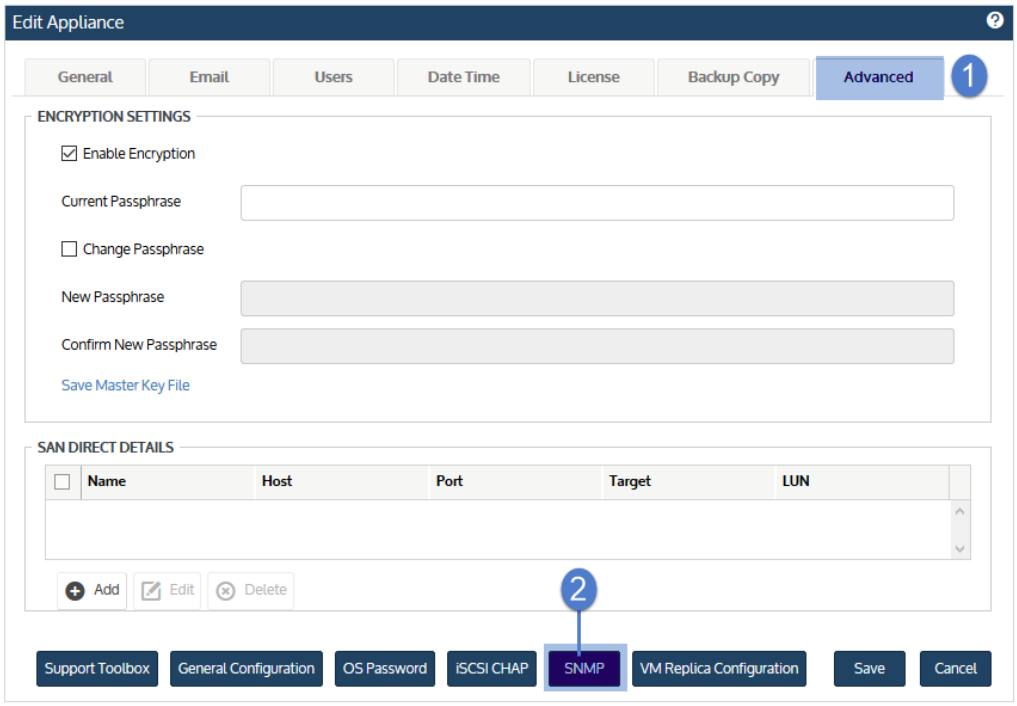

5 Click **Download MIB** and save the *unitrends.mib.txt* file.

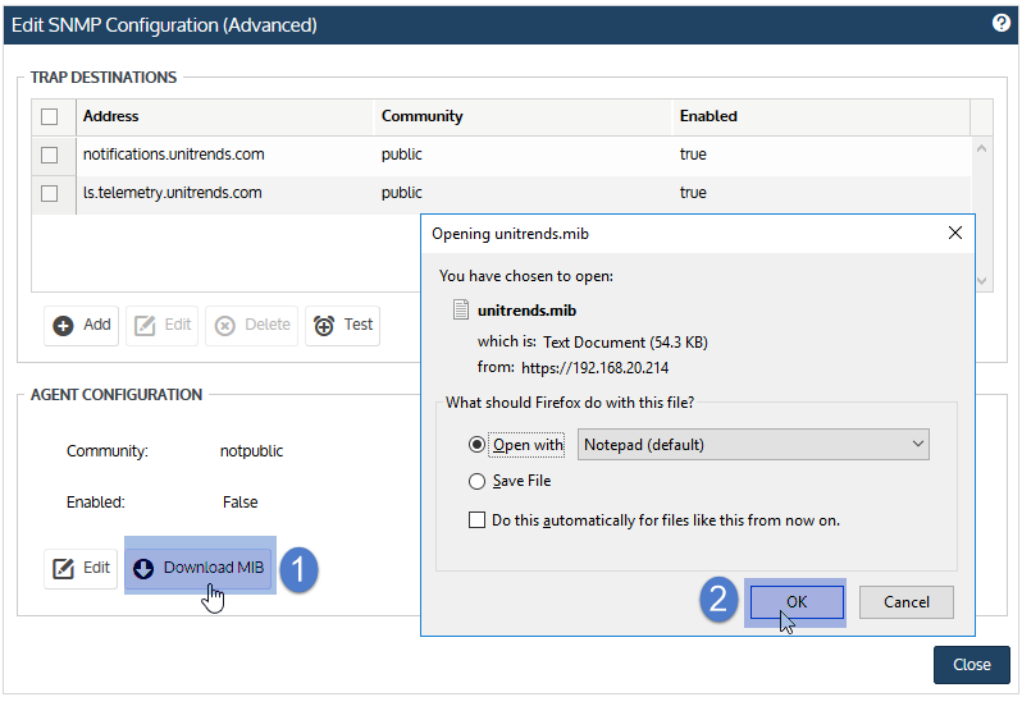

6 Install *unitrends.mib.txt* in your RMM environment. The file is also available at *http://<Unitrends appliance IP>/snmp/*.

> **Note:** You will also need the Net-SNMP MIBs. These come standard in most RMM software. If you need them, they are available at *http://<Unitrends appliance IP>/snmp/*.

7 Return to the appliance UI. In the Trap Destinations area, click **Add**.

8 Enter the Destination address and Community, check **Enabled**, and click **Save**.

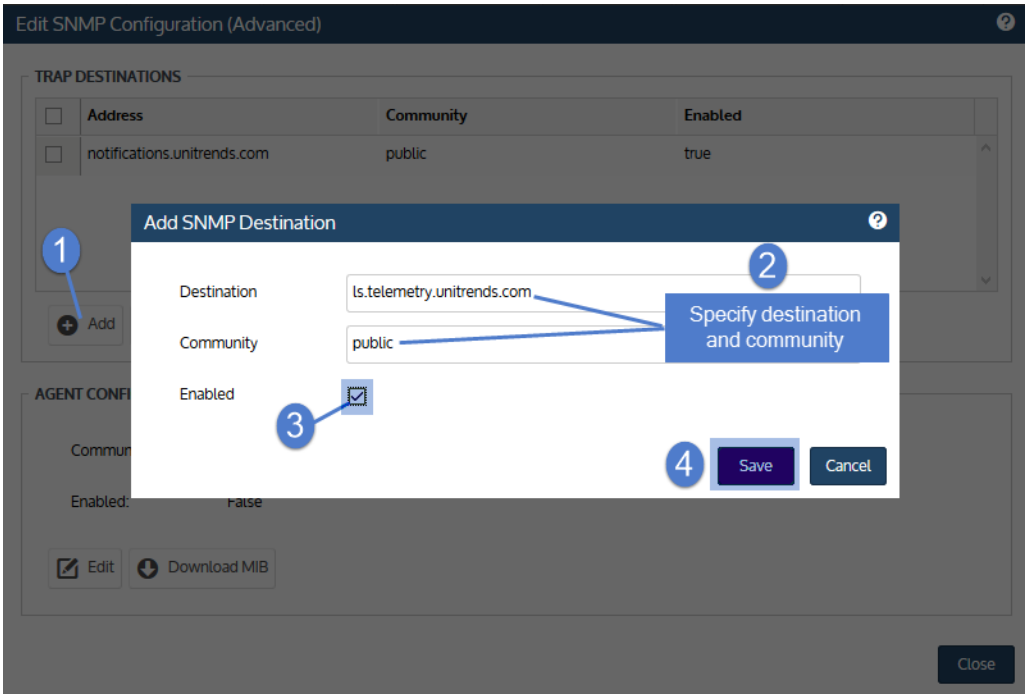

9 Click **Test**. A test trap is sent to all destinations. You see a *Success* message if the destination you configured is operational.

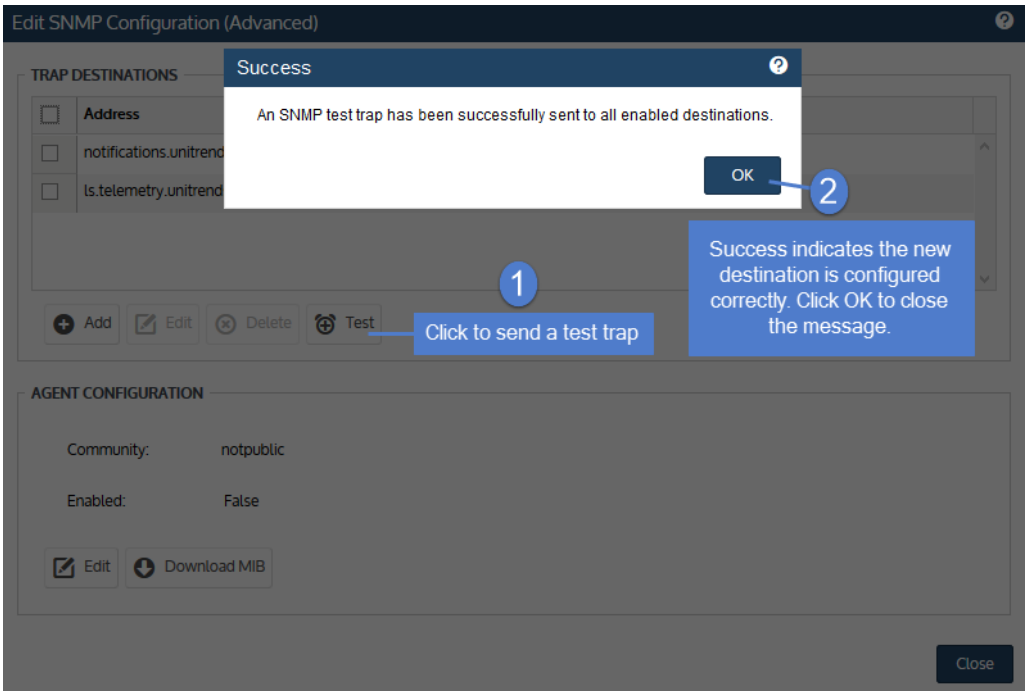

CHAP authentication for iSCSI connections

Unitrends supports Challenge Handshake Authentication Protocol (CHAP) for iSCSI connections to external storage:

Note: CHAP authentication is used for iSCSI connections to external backup storage and backup copy targets only. CHAP is NOT used to recover files from host-level backups over iSCSI.

- You can configure the iSCSI connection with CHAP before configuring CHAP on the target storage. Once the target is configured, CHAP authentication is enforced.
- If CHAP has not been configured on the target storage, the appliance detects this and gains access without CHAP authentication, even if CHAP has been enabled on the Unitrends appliance.
- If CHAP has been configured on the storage target, you must enable CHAP authentication on the Unitrends appliance. If not, any attempt to add the target to or access the target from the Unitrends appliance fails.
- A single CHAP username and password is used by the Unitrends appliance. Therefore, all of its CHAP-enabled iSCSI targets must be configured with this username and password.
- CHAP is supported from the initiator (Unitrends appliance) to the target only. Mutual (bi-directional) CHAP is not supported.
- CHAP authentication occurs upon first log in to the target. Subsequent operations on the target succeed, without further authentication, for the duration of the iSCSI session or until the target sends a random challenge request.

To configure iSCSI CHAP authentication

1 Log in to the appliance UI.

2 On the **Configure > Appliances** page, select the appliance and click **Edit**.

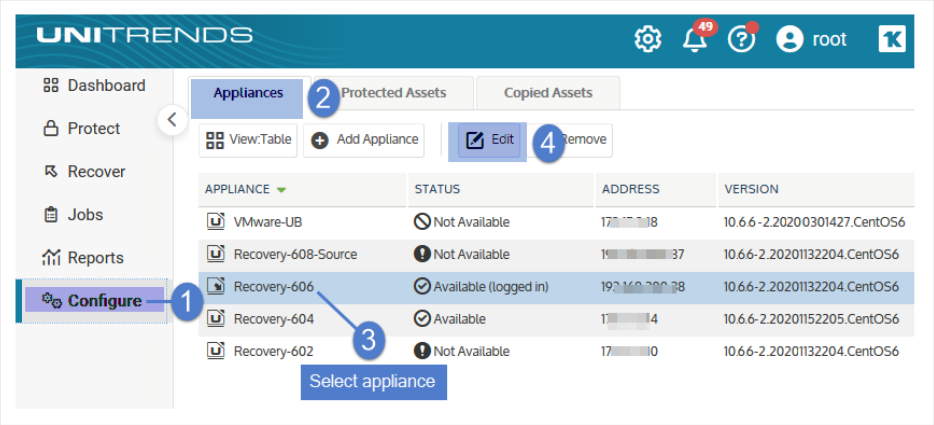

3 Click **Advanced** and select **iSCSI CHAP**.

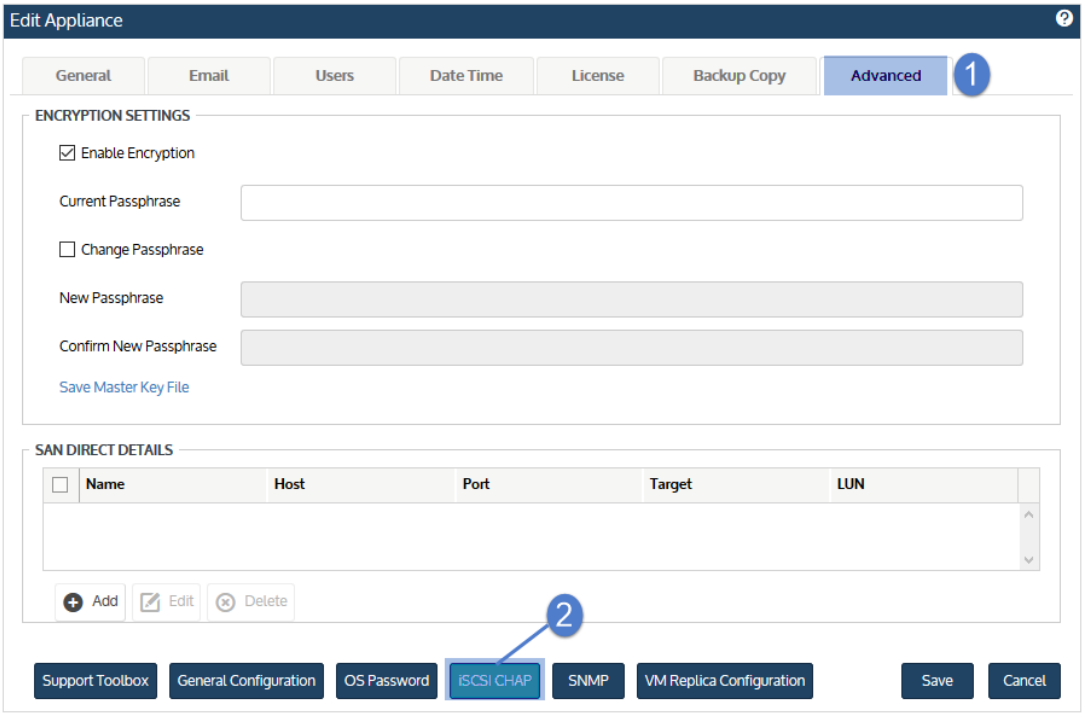

4 Verify that the **Use System CHAP Credentials** box is checked.

5 Enter credentials in the **Username**, **CHAP Password**, and **Confirm CHAP Password** fields, then click **Save**. One set of credentials is used to access all iSCSI targets that have been configured to use CHAP authentication.

- By default, **Username** contains the appliance's iSCSI qualified name (IQN). It is required that the username and password on the initiator (backup appliance) match those defined on the targets. Modify the Username entry if necessary.

- The password must be 12-16 characters in length.

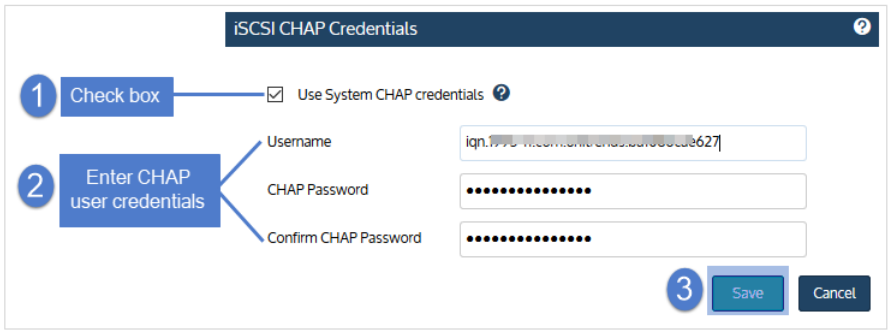

Support Toolbox advanced administration tasks

In most cases, you use the main UI pages (Dashboard, Protect, Recover, Jobs, Reports, and Configure) for appliance administration tasks. In some cases, you may need to access additional information, such as log files, lists of running processes and services, or disk status. The Support Toolbox provides an easy way to access this lower-level appliance information and perform related tasks.

To use the Support Toolbox

1 Log in to the appliance UI.

You must log in directly to the appliance. You cannot access the Support Toolbox of a managed appliance.

2 On the **Configure > Appliances** page, select the appliance and click **Edit**.

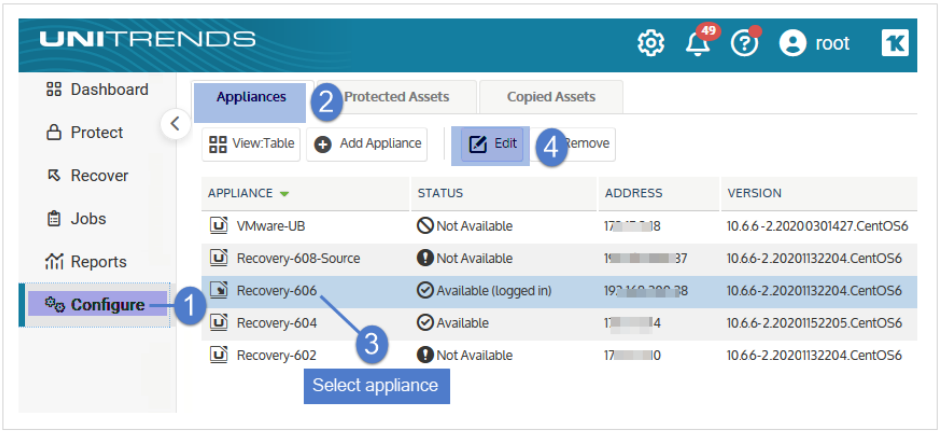

3 Click **Advanced** and select **Support Toolbox**.

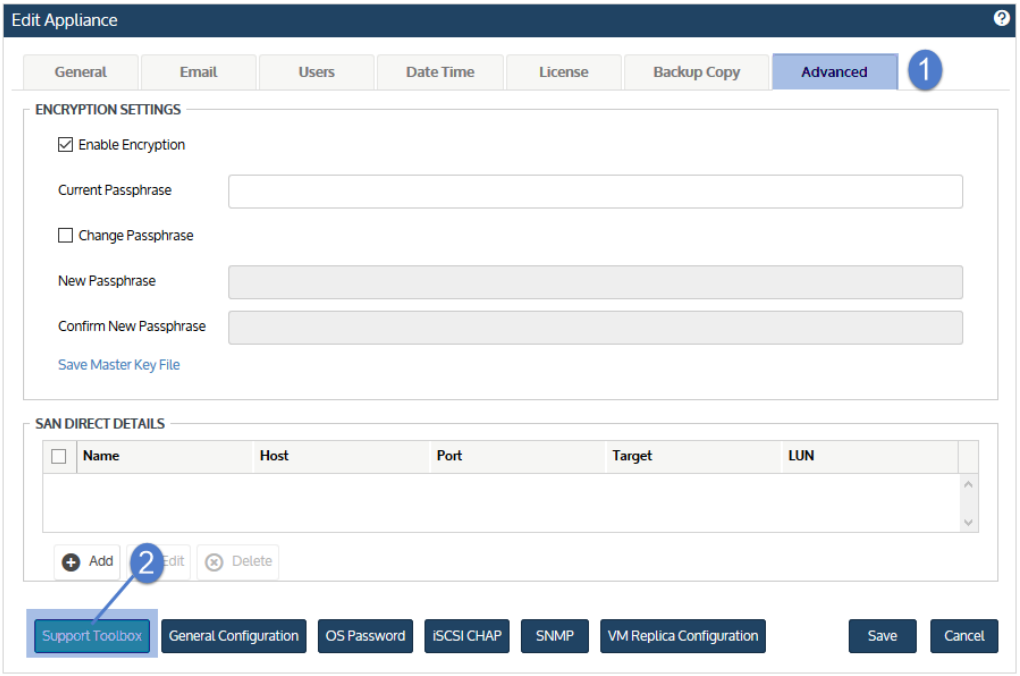

4 The Support Toolbox (Advanced) dialog displays.

- Scroll through the toolbox to find the information or task you are interested in.
- Hover over the options to see descriptions and helpful tips.

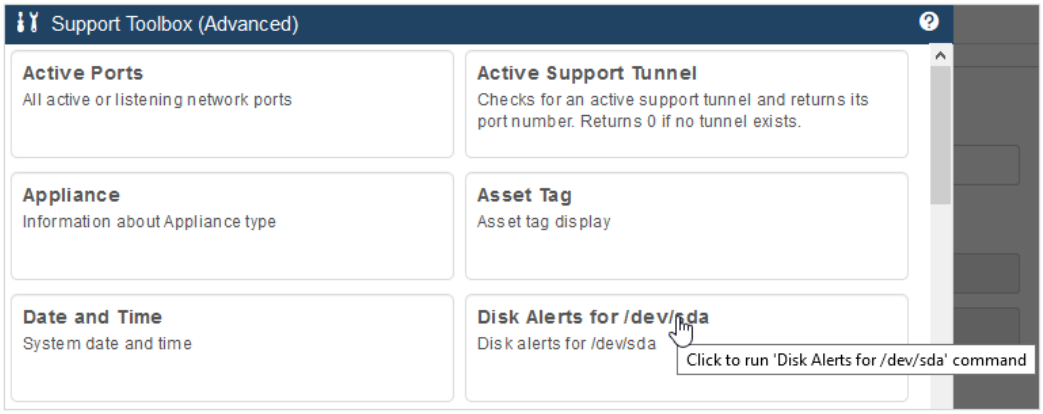

- Click an item to run a command. Click **Close** to exit.

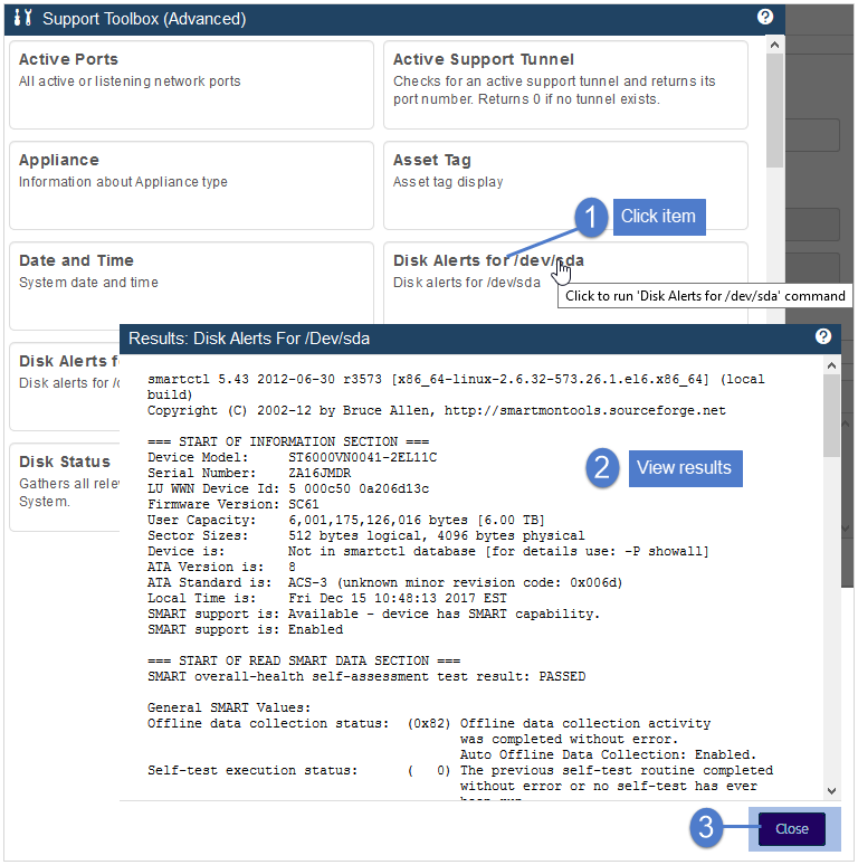

5 Click **Close** to exit.

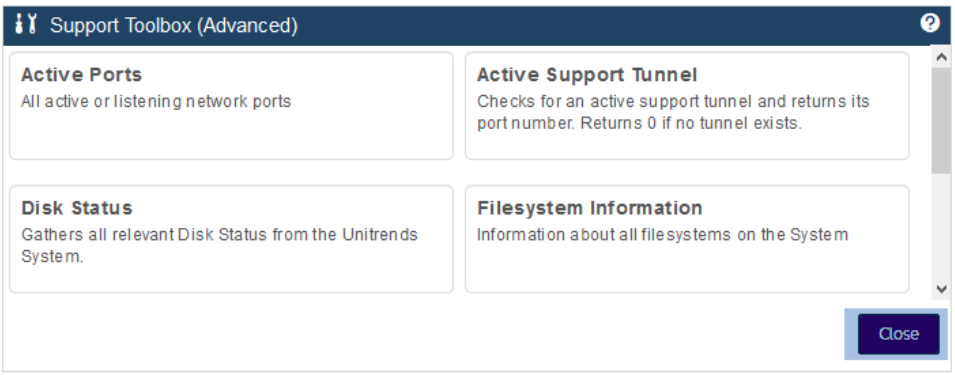

Additional appliance settings

The appliance is automatically configured to use the best settings for the appliance model and other factors in your environment. In most cases you will never need to modify these settings. If you are an advanced user and want to adjust deep configuration settings, such as MaxBlockSize and QuickSeek, you can edit these settings. Do not modify these settings if you are not familiar with how the change will impact appliance performance and on-appliance

retention.

To configure advanced settings

1 Log in to the appliance UI.

You must log in directly to the appliance. You cannot change the configuration settings of a managed appliance.

2 On the **Configure > Appliances** page, select the appliance and click **Edit**.

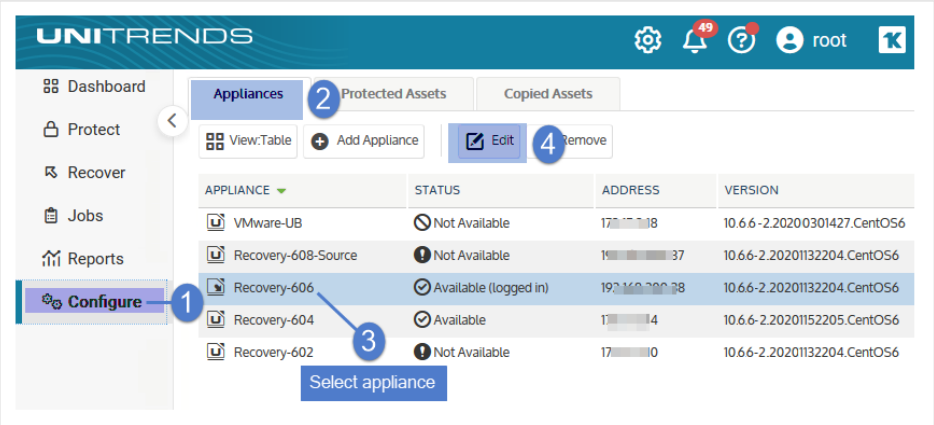

3 Click **Advanced** and select **General Configuration**.

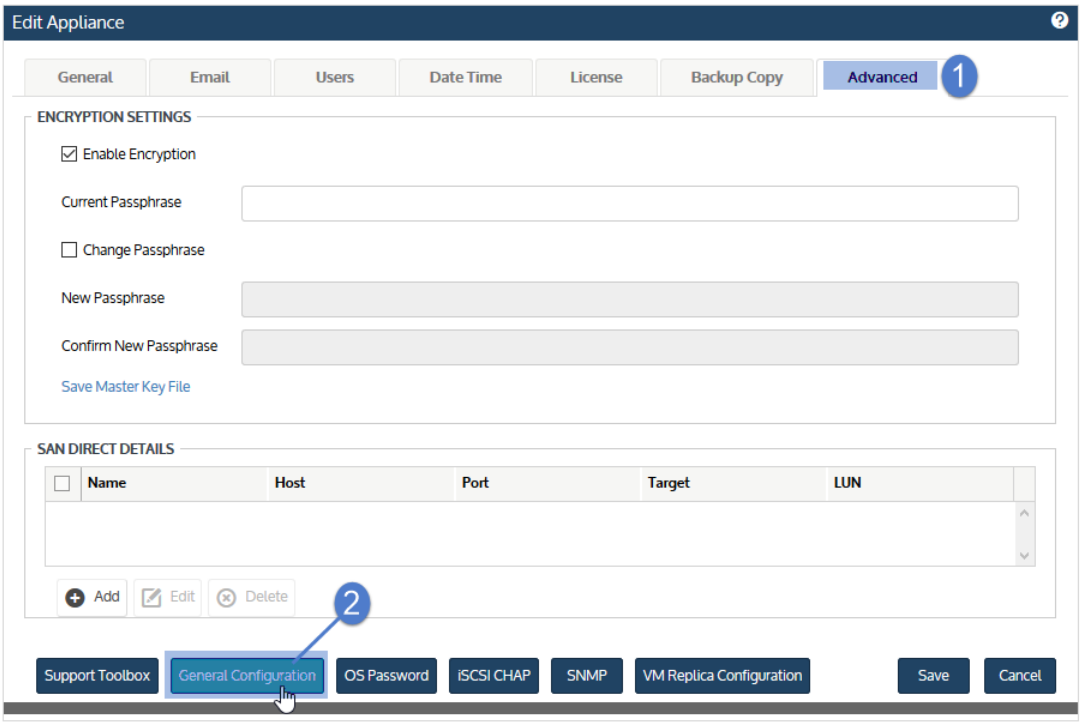

The General Configuration (Advanced) dialog displays.

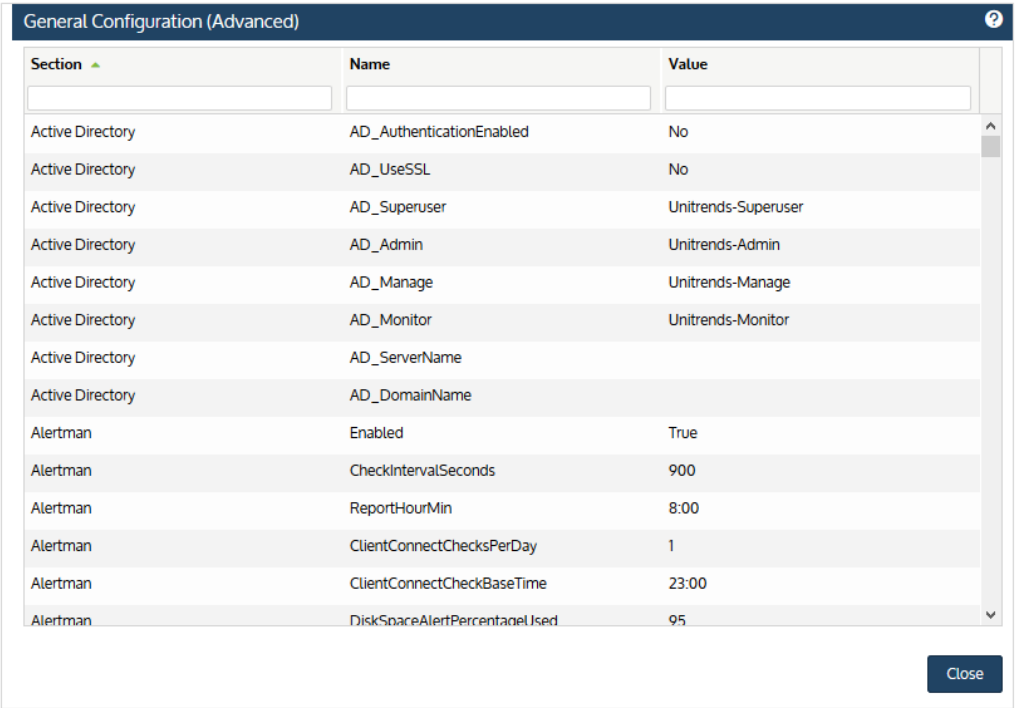

4 Scroll or use the column sort and filter fields to view appliance settings.

- Click on any column to sort alphabetically (a to z). To reverse the order (z to a), click again.

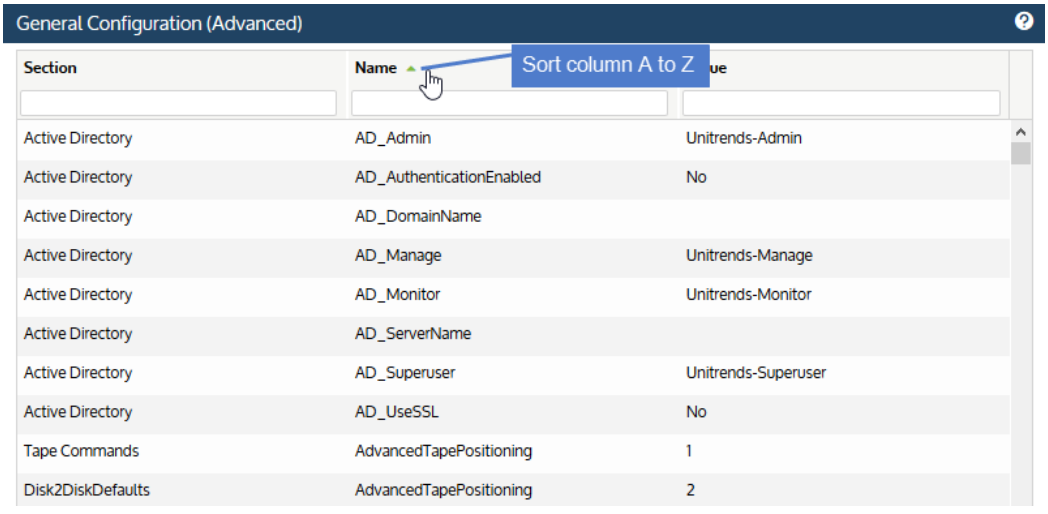

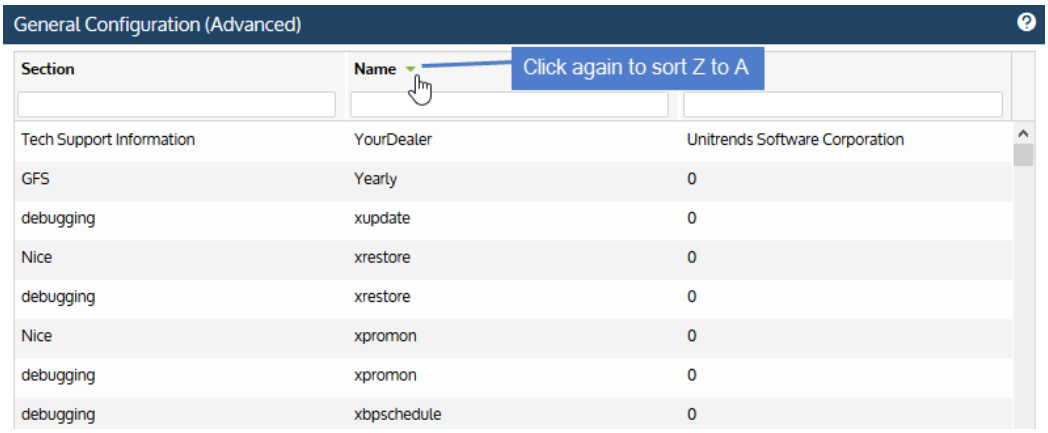

- Filter the display by entering text in a column's filter field.

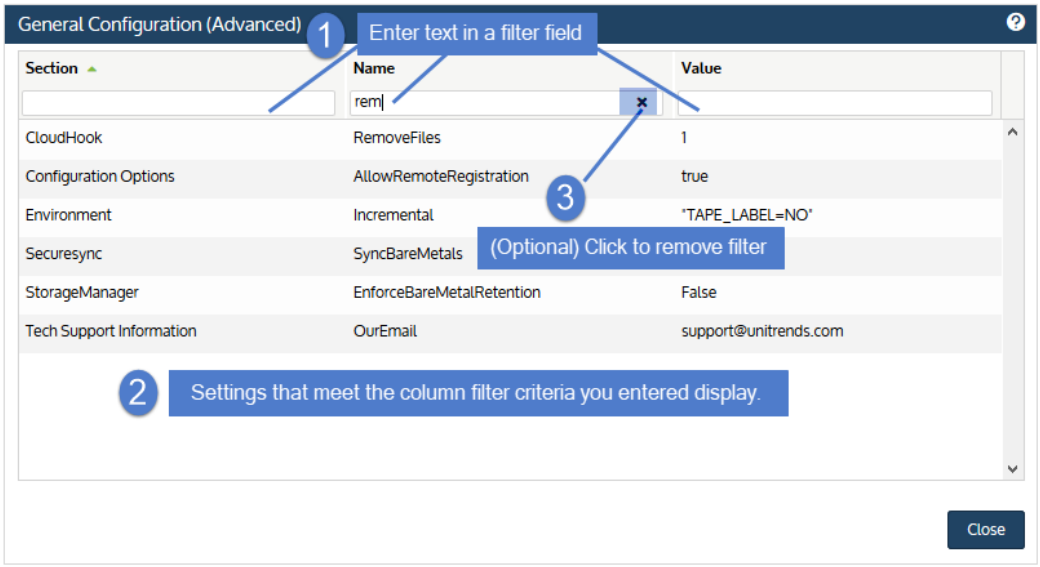

5 (Optional) Modify a setting. Select the setting, enter a new Value, then **Save**. Repeat this step as needed to modify additional settings.

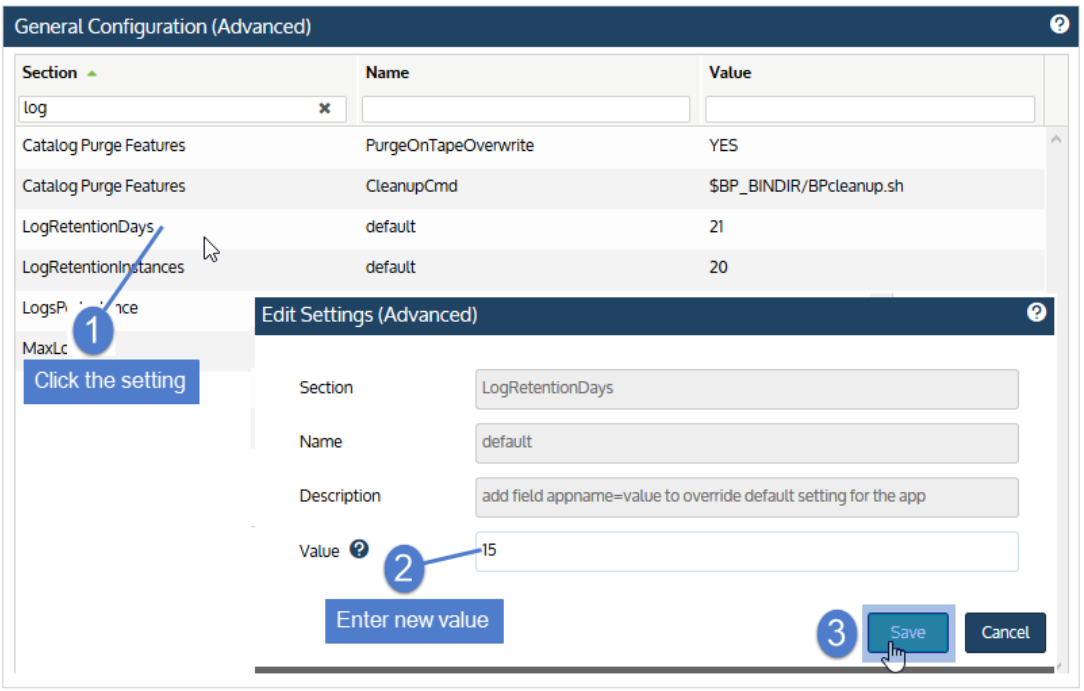

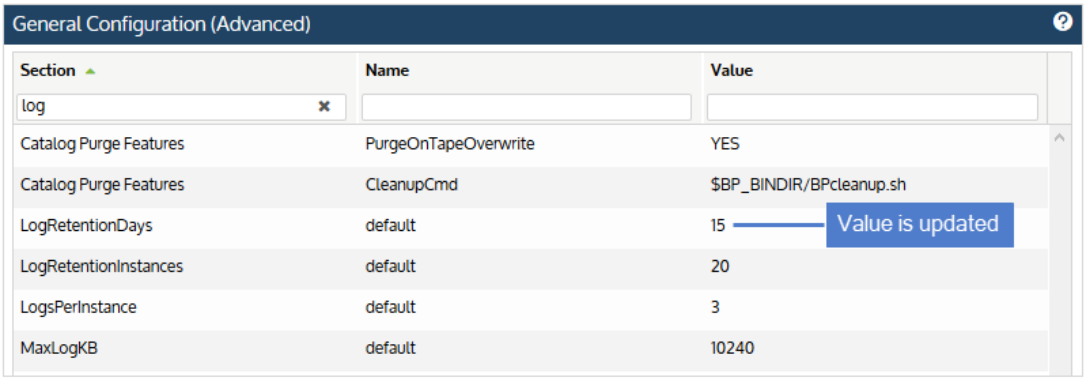

6 Click **Close** to exit.

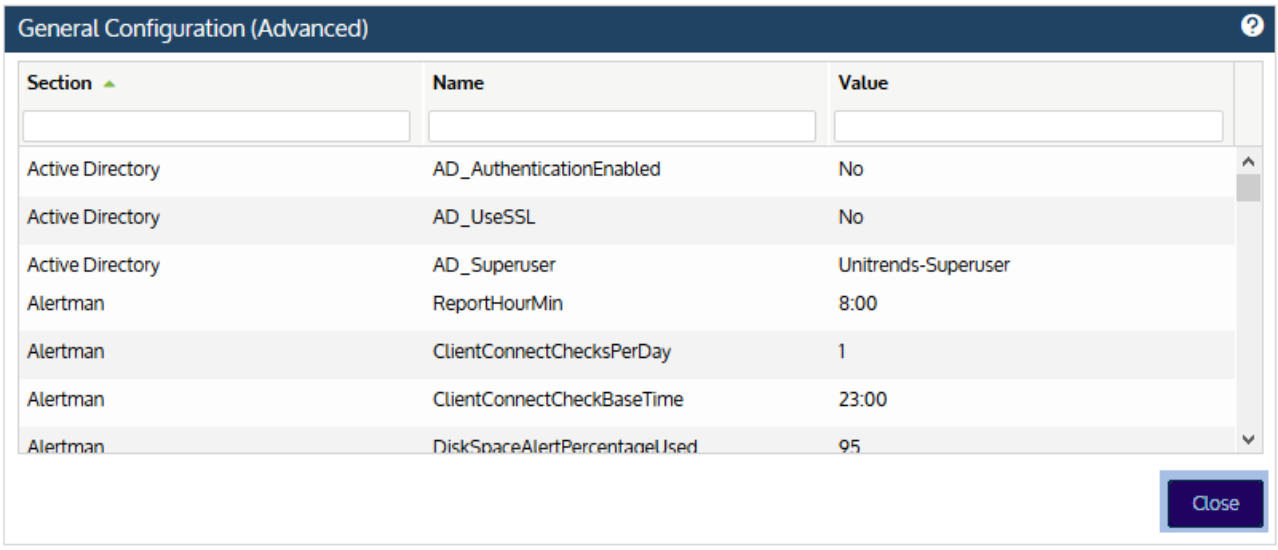

Create a separate database partition on your Unitrends Backup appliance

Note: This feature applies to Unitrends Backup appliances only.

When deployed, the Unitrends Backup database is located in the same partition as stored backups. You now have the option to configure a separate partition to house the Unitrends Backup database. With this configuration, the database resides in its own partition on a separate logical volume than the backups themselves. Use this option to increase backup performance and stability by using faster performing storage for the database and lower-tier storage for the backup data itself. This is great when you are using slower backup storage that communicates with Unitrends Backup over NFS or CIFS protocols.

Requirements and considerations for creating a separate database partition on your Unitrends Backup appliance

By moving the database to a different location than the stored backups, you add hardware. This introduces additional potential points of failure. Be sure to:

- Implement proper measures to ensure hardware reliability of all storage.
- Store copies of your backups on secondary storage to avoid losing backup data in the event of a hardware failure. For details, see "Backup copies" on page 101.

These requirements must be met to create a separate database partition:

- The appliance must be a Unitrends Backup appliances.
- To create a separate database partition, you must first add a disk to the Unitrends Backup appliance.
- The disk you add must be at least 100GB or twice the running database size, whichever is greater.

- To create the database disk, Unitrends recommends that you add a disk in the same way you added the initial backup storage during Unitrends Backup deployment. For example, if you created the initial backup storage using direct attached storage (DAS), use DAS for the database partition. For details, see the applicable Unitrends Backup deployment guide:
  - Deployment Guide for Unitrends Backup on VMware
  - Deployment Guide for Unitrends Backup on Hyper-V
  - Deployment Guide for Unitrends Backup on Citrix XenServer
  - Deployment Guide for Unitrends Backup on Nutanix AHV
  - Deployment Guide for Unitrends Backup in Microsoft Azure
  - Deployment Guide for Unitrends Backup in Amazon Web Services
  - Deployment Guide for Unitrends Free on VMware
  - Deployment Guide for Unitrends Free on Hyper-V

  WARNING! Unitrends strongly recommends that all Unitrends Backup storage is either direct attached storage (DAS, internal to the hypervisor) or resides on one external storage array. If you configure storage across multiple storage arrays and one becomes unavailable, all backup data ends up corrupted, resulting in total data loss.

Once the above requirements have been met and you have added the database disk to your Unitrends Backup appliance, proceed to "To create a separate database partition on your Unitrends Backup appliance" to set up the partition and migrate the database.

To create a separate database partition on your Unitrends Backup appliance

This procedure assumes you have added a disk to your Unitrends Backup appliance that meets the requirements above.

1. Using a terminal emulator, such as PuTTY, connect to the appliance using the following:
   - Appliance IP address
   - Port 22
   - SSH connection type
2. Log in as user *root*.
3. Issue this command to download the create database partition script:

```
# wget https://sftp.kaseya.com/utilities/newdisk
```

4. Issue this command to add execute permissions:

```
# chmod +x newdisk
```

5. Issue this command to run the script:

```
# ./newdisk
```

6 In the script output, you see one of the following:

- Output similar to the following, indicating that an eligible database disk is available. Note the device partition (*/dev/sdd* in the example). You will need it in the next step. Proceed to step 7.

  ```
  Available added disk = /dev/sdd (size 107.2 GB)
  ```

- The following message, indicating that an eligible disk is not available. Add an eligible disk and rerun this procedure.

  ```
  No available added disks found
  ```

- A message similar to the following, indicating that the disk you added is less than 100GB in size. Add an eligible disk and rerun this procedure.

  ```
  Disk /dev/sdd size X GB is less than minimum 100 GB
  ```

- A message similar to the following, indicating that the disk you added is less than twice the current database size. Add an eligible disk and rerun this procedure.

  ```
  Disk /dev/sdd size X GB is less than twice database X GB
  ```

7 Issue the **cmc_stateless dbpart** command followed by the device partition you noted above to format the disk and migrate the database. The following example uses */dev/sdd* as the device partition:

```
# cmc_stateless dbpart /dev/sdd
```

Configure deduplication settings on your Unitrends Backup appliance

Notes:

- This feature applies to Unitrends Backup virtual appliances only.
- For Unitrends Backup on Citrix XenServer, the appliance is configured to optimize retention (Level 3) because only full backups are supported.
- Level 3 is required for ransomware detection functionality.

Deduplication is a data compression technique that eliminates duplicate data blocks. To yield fastest performance, the Unitrends Backup appliance is configured to use the Level 1 deduplication setting. You can opt to modify this setting to increase on-appliance retention. Keep in mind that increasing the deduplication level decreases job speed.

To configure deduplication settings:

1 Log in to the appliance UI.

2 Click the gear icon and select **Deduplication Settings**.

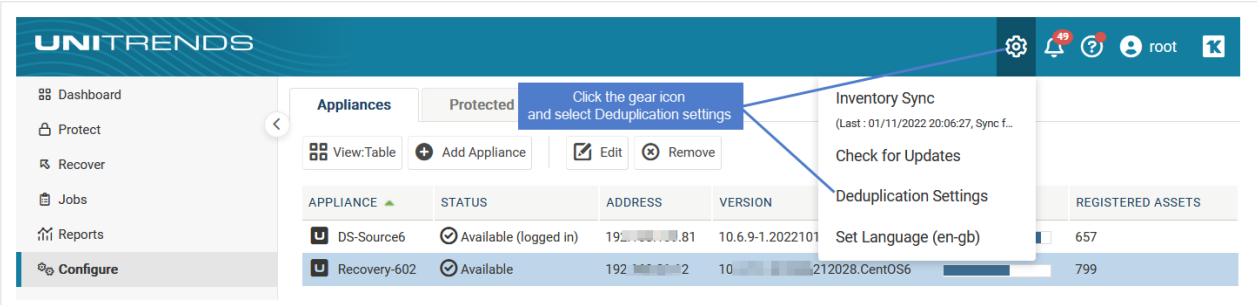

3 Select one of the following deduplication settings:

- Level 1 — Use this setting to optimize performance.
- Level 2 — Use this setting to balance performance and on-appliance retention.
- Level 3 — Use this setting to optimize retention.

4 Click **Apply Settings**.

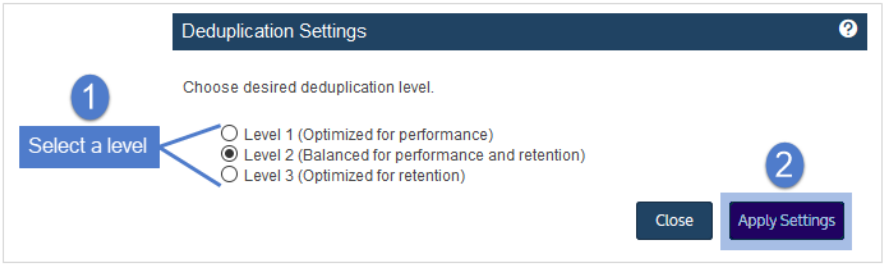

Set appliance language

Use this procedure to select a language for the appliance.

1 Log in to the appliance UI.

2 Click the gear icon and select **Set Language**.

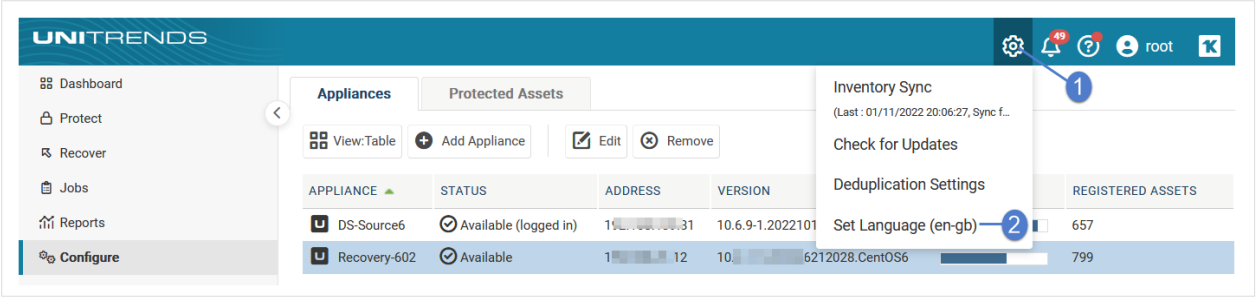

3 Select one of the following language options:

- US English (en)
- French (fr)
- Europe English (en-gb)

4 Click **Apply**.

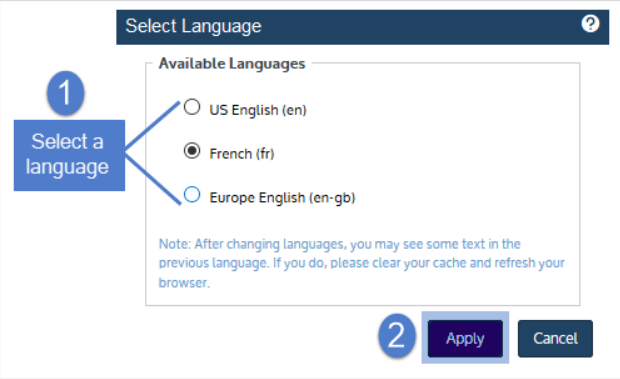

Appliance Samba share

The Unitrends agent needs access to the appliance's Samba share for backup and recovery operations. The SMB 2.0 security option is enabled by default on Unitrends appliances that were originally imaged or deployed with version 10.4.8 or higher. (The SMB 1.0 security option is enabled by default on appliances that were originally imaged or deployed with a pre-10.4.8 version. Upgrading the appliance does not change the SMB 1.0 setting.)

The following requirements apply to appliances configured with SMB 2.0:

Note: See How Unitrends supports SMB2 for SMB 2.0 configuration procedures.

Feature	SMB 2.0 Requirements
Oracle on Solaris	The Unitrends agent must have access to the appliance's SMB 2.0 Samba share to perform backup and recovery operations. These requirements apply: • A Samba client must be enabled. See How Unitrends supports SMB2 for details. • A Samba key must be added for the backup appliance. To add the key, issue this command (the default password is *samba*): `smbadm add-key -u samba@<applianceIP>`
Oracle on Windows	SMB 2.0 must be enabled on the Windows server so that the Unitrends agent can access the appliance's SMB 2.0 Samba share when performing backup and recovery operations.

Feature	SMB 2.0 Requirements
Recover files from host-level backups of Windows VMs	To use a CIFS share for the recovery, SMB 2.0 must be enabled on the target Windows asset.
SharePoint	To perform backup and recovery operations, SMB 2.0 must be enabled on each node in the farm. **Notes:** • SharePoint 2007 on Windows 2003 and prior is not supported on SMB 2.0 appliances. (To configure your appliance to use SMB 1.0, see How Unitrends supports SMB2.) • SharePoint may require custom client configuration for use with SMB 2.0. If SharePoint backups do not run successfully, see this Microsoft article for client configuration details: SharePoint Ports, Proxies and Protocols...An overview of farm communications.
Windows agent push	To push install the Windows agent, SMB 2.0 must be enabled on the Windows asset.
Windows replica on Hyper-V	To run a Windows replica on Hyper-V, SMB 2.0 must be enabled on the Hyper-V server.

Backup storage

Use the **Configure > Appliances > Storage** tab to manage your appliance's backup storage. Supported storage procedures vary by Unitrends appliance type:

- Recovery Series, Recovery MAX, and ION/ION+ physical appliances come with a set amount of backup storage. You cannot add backup storage to the appliance.
- Unitrends Backup virtual appliances deploy as virtual machines. During deployment, the initial backup storage was created using either a virtual attached disk, a SAN LUN, or a NAS share. After initial deployment, you can add more backup storage. See "About adding backup storage to a Unitrends Backup appliance" on page 199 for details.

Instant recovery write space

Appliance storage can be used to store backups, for instant recovery of VMs and Windows image-level backups, and for Windows replicas. To use the instant recovery and Windows replica features, you must allocate a portion of the appliance storage to Instant Recovery in the Edit Storage dialog (as described in "To edit storage allocation" on page 198). For more on these features, see "Virtual machine instant recovery" on page 904, "Instant recovery of Windows

image-level backups" on page 1055, and "Windows file-level replicas" on page 993.

Once storage is allocated to instant recovery, it can be used only for Windows replicas and instant recovery. The storage is reserved and cannot be used for other purposes, such as backups or deduplication (but you can modify your storage allocation at any time).

Because the appliance is designed to retain as many local backups as possible, it is best to reserve instant recovery space soon after initial deployment. If necessary, you can allocate instant recovery space later, but doing so may require the appliance to purge local backups to make room for the newly allocated instant recovery space.

Instant recovery storage requirements

Before allocating storage, determine the percentage to use for backups and the percentage to reserve for Windows replicas and instant recovery. Instant recovery storage requirements are given here:

- VM instant recovery and Windows image-level instant recovery – The disks for the recovered asset will reside on the appliance until storage migration completes. Allocate at least 20% of the space used on the original Windows asset or virtual machine to instant recovery storage. Additionally, ensure that there is enough extra storage to account for anticipated growth during the recovery process, especially for long-running recoveries. Once disks have been migrated to another location, or BMR is complete and the IR is torn down, the appliance's instant recovery storage is no longer needed.
- Windows replicas – Requirements vary by recovery target:
  - Virtual host target – If recovering to an ESXi or Hyper-V server, no appliance instant recovery storage is needed.
  - Recovery Series, Recovery MAX, or ION/ION+ appliance target – If recovering to the Unitrends appliance, the disks for the recovered asset reside on the appliance. Ensure that instant recovery storage space equal to the amount of space used on the Windows asset's original disks is available on the appliance.

For Recovery Series, Recovery MAX, ION/ION+, and Unitrends Backup appliances, you can view backup storage and modify the amount of storage allocated for backups and instant recovery. See these procedures for details:

- "To view backup storage"
- "To edit storage allocation" on page 198

To view backup storage

1. On the **Configure > Appliances** page, select the appliance.
2. Select the **Storage** tab below.

   This information displays for each storage device: name, status, type, capacity, total size, and free space.

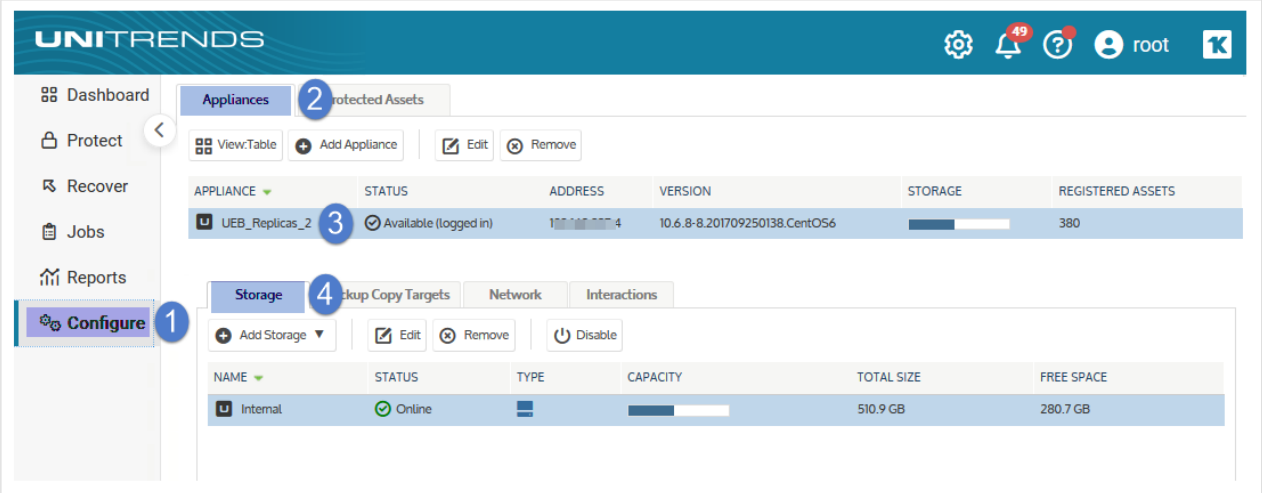

To edit storage allocation

1 On the **Configure > Appliances** page, select the appliance.

2 On the **Storage** tab, select the **Internal** storage and click **Edit**.

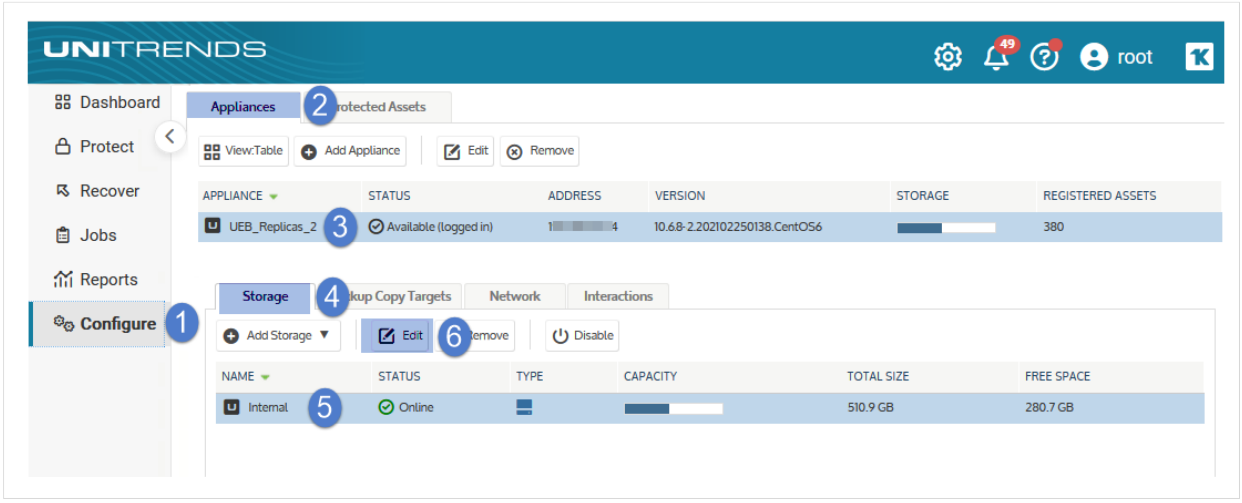

3 Modify the percentages used for backups versus instant recovery, and click **Save**.

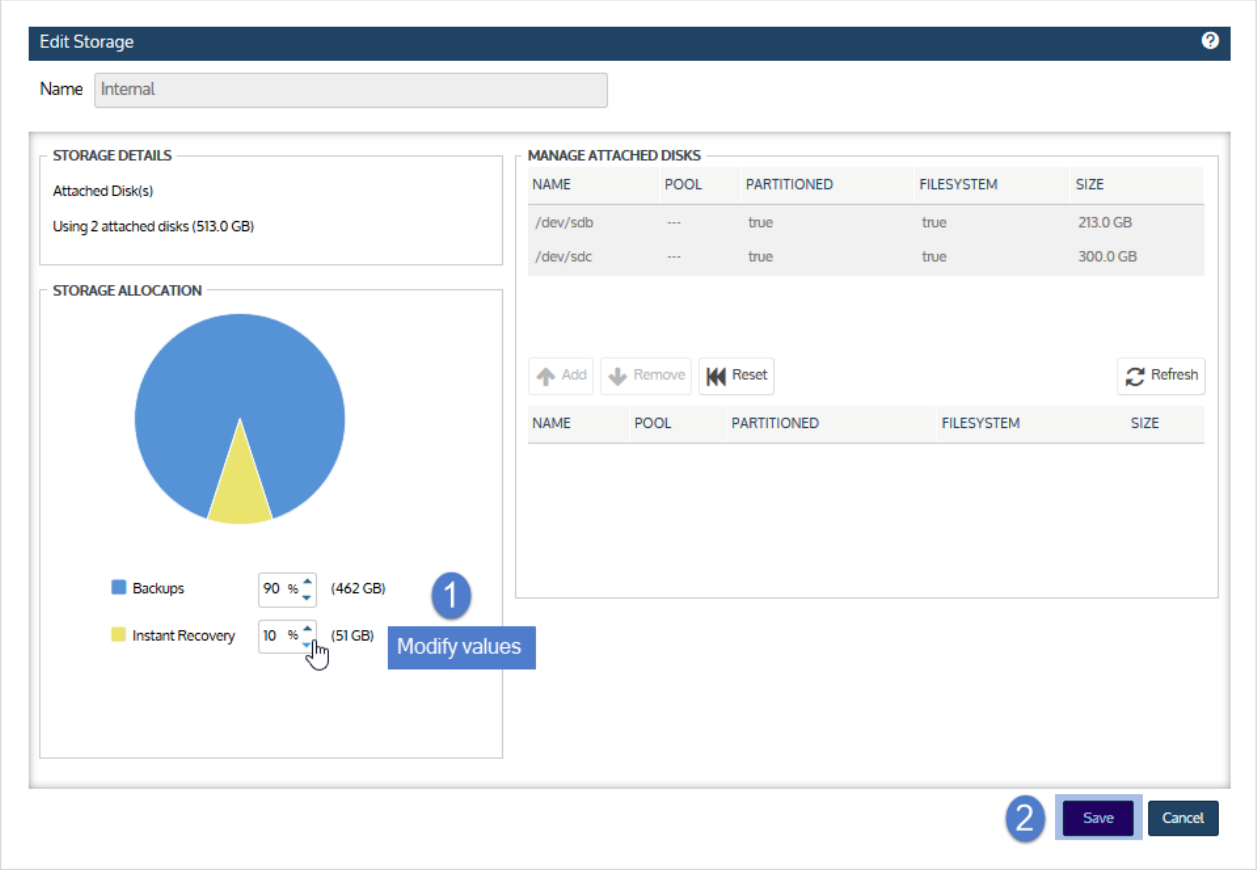

About adding backup storage to a Unitrends Backup appliance

To add more backup storage to your Unitrends Backup appliance, Unitrends recommends adding storage in the same manner as the initial backup storage:

Note: These procedures apply to adding add more storage to your Unitrends Backup appliance. Additional requirements and considerations apply to creating the initial backup storage. If you have not yet deployed the appliance and created the initial backup storage, follow the procedures in the applicable deployment guide instead. For links to the deployment guides, see "About this Guide" on page 18.

- If you used virtual attached disk storage, Unitrends recommends using the host to add virtual disks to the Unitrends Backup VM. Once the disks are added to the VM, use the Unitrends Backup UI to expand the initial backup storage to include these new disks. See "Procedures for adding attached disk backup storage" on page 200 for details.
- If a NAS share was attached to the host , we recommend adding another share in the same manner. Then go to the host and add new virtual disks to the Unitrends Backup VM using storage on the share you added. Once the disks are added to the VM, use the Unitrends Backup UI to expand the initial backup storage to include these new disks. See "Procedures for adding attached disk backup storage" on page 200 for details.

- If a SAN LUN was connected to the host, we recommend adding another LUN in the same manner. Then go to the host and add new virtual disks to the Unitrends Backup VM using storage on the LUN you added. Once the disks are added to the VM, use the Unitrends Backup UI to expand the initial backup storage to include these new disks. See "Procedures for adding attached disk backup storage" on page 200 for details.
- If a SAN or NAS was directly attached to the Unitrends Backup appliance, expanding the initial backup storage is not supported. Instead, LUNs or shares can be added to the appliance as separate storage areas. See "Procedures for adding external storage" on page 210 for details.

  Although you can attach external storage directly to the Unitrends Backup appliance, Unitrends does not recommend this approach. If you must connect external storage to the Unitrends Backup VM directly through network protocols (CIFS, NFS, or iSCSI), make sure to use a supported vendor from the list in Supported external storage vendors for use with Unitrends Backup appliances.

Additional recommendations

Review the following recommendations before adding storage:

WARNING! Unitrends strongly recommends that all storage is either direct attached storage (DAS, internal to the hypervisor) or resides on one external storage array. If you configure storage across multiple storage arrays and one becomes unavailable, all backup data is corrupted, resulting in total data loss.

- As you add storage, be sure to add resources to the Unitrends Backup virtual machine, such as CPU and memory.
- Storage should be dedicated to the Unitrends Backup VM and not shared by other virtual machines, applications, etc.
  - If you are using external SAN or NAS storage, the shares or LUNs used by the Unitrends Backup VM should be dedicated to that Unitrends Backup VM.
  - The Unitrends Backup VM can be deployed on a host in a cluster configuration and can use shared storage. However, in this configuration, the Unitrends Backup VM should use a dedicated NAS share or SAN LUN.
- For best performance with SAN storage, use a thickly provisioned LUN and either a thick provisioned eager zeroed VMDK or a fixed size VHD(X). (For XenServer, VHDs are always fixed size.)
- For VMware environments, do not use Storage vMotion. Storage must remain in a fixed location.
- For Hyper-V environments, do not use Storage Migration. Storage must remain in a fixed location.

Procedures for adding attached disk backup storage

Use these procedures to add attached disk storage to the Unitrends Backup appliance. You can view the appliance's disk storage in the Edit Storage dialog. Disks that have been added to the Unitrends Backup VM by using the hypervisor display as */dev/sdx/*, where *x* indicates alphabetically the order in which the disks were added to the VM. For example, the initial disk is always */dev/sda/*, the first disk that was added as the initial backup storage device is */dev/sdb/*, the next would be */dev/sdc/*, etc.

See the following topics for details:

- "To view or edit attached disk backup storage"

- "To add a disk and expand the storage device"
- "To add a disk as a separate storage area" on page 207

To view or edit attached disk backup storage

1. On the **Configure > Appliances** page, select your appliance.
2. Click the **Storage** tab.
3. Select the *Internal* storage and click **Edit**.

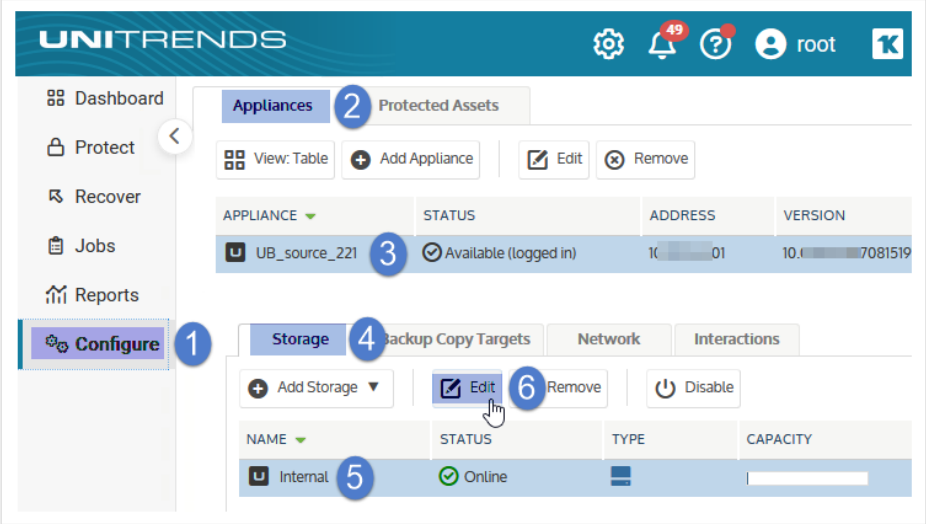

4. Manage the attached disks as desired. If a virtual disk you would like to add does not display, click **Refresh**.
5. Once finished editing the storage, click **Save**.

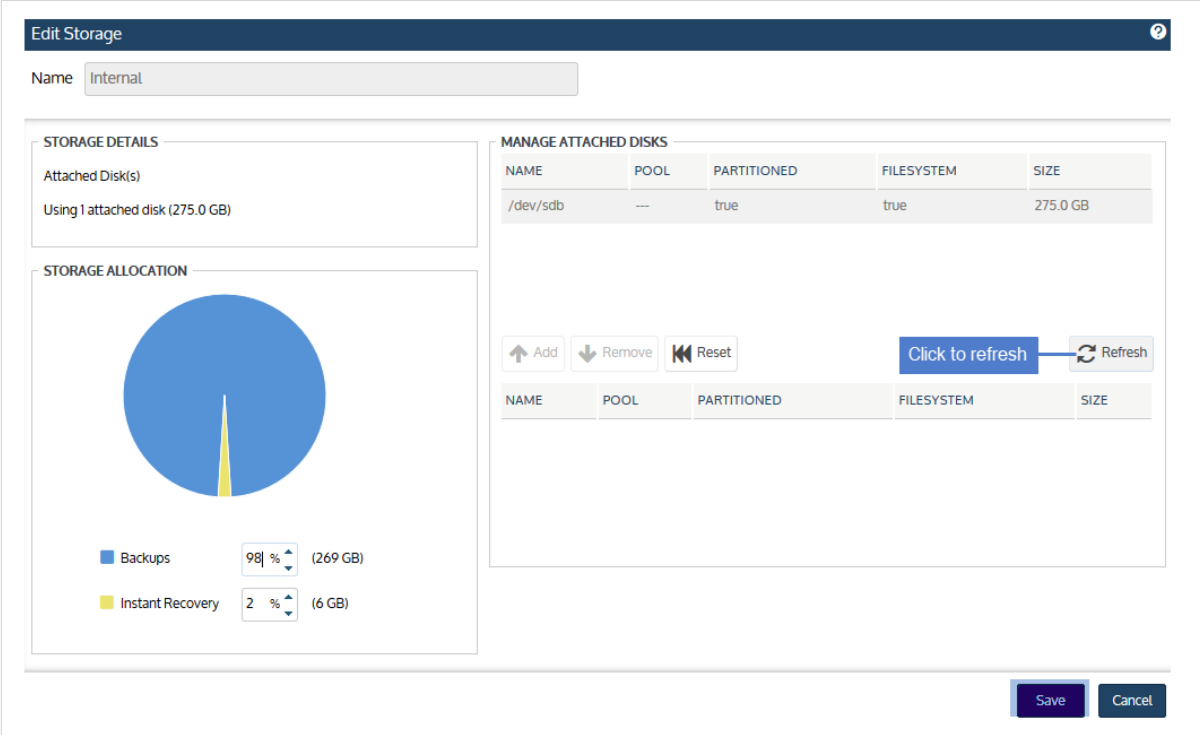

To add a disk and expand the storage device

Note: Once a disk has been added, it cannot be removed.

1 Go to the hypervisor and add a virtual disk to the Unitrends Backup VM.

- For an appliance whose initial backup storage is DAS internal to an ESXi host, see one of the following VMware documents:
  - vSphere 5.1: Create a Virtual Disk in vSphere Client 5.1
  - vSphere 5.5: Create a Virtual Disk in vSphere Client 5.5
  - vSphere 6.0: Create a Virtual Disk in vSphere Client 6.0
  - vSphere 6.5: Add a Hard Disk to a Virtual Machine
- For an appliance whose initial backup storage is DAS internal to a Hyper-V host, Unitrends recommends that you use a VHD(X) disk and that you add the disk to the SCSI controller. See the following Microsoft documents:
  - To create a virtual hard disk
  - To add a hard disk to a virtual machine

- For an appliance whose initial backup storage is DAS internal to an AHV host, edit the appliance VM to add a vDisk to the ISCSI controller (by selecting **ISCSI** in the Bus Type list). For details, see Managing a VM (AHV) in the Prism Web Console Guide.
- For an appliance whose initial backup storage is DAS internal to a XenServer host, Unitrends recommends that you add the VHD disk to the SCSI controller. VM disks cannot be attached as *Read Only*. Be sure to use the *Read Only* = *No* setting when attaching disks.
- For an appliance whose initial backup storage is on an external SAN connected to the hypervisor, add another LUN and expose it to the Unitrends Backup VM. Then go to the hypervisor and add new virtual disks to the VM using storage on the LUN you added.
- For an appliance whose initial backup storage is on an external NAS share connected to the hypervisor, add another NAS share. Then go to the hypervisor and add new virtual disks to the Unitrends Backup VM using storage on the share you added.

2 On the **Configure > Appliances** page, select your appliance.

3 Click the **Storage** tab.

4 Select the *Internal* storage and click **Edit**.

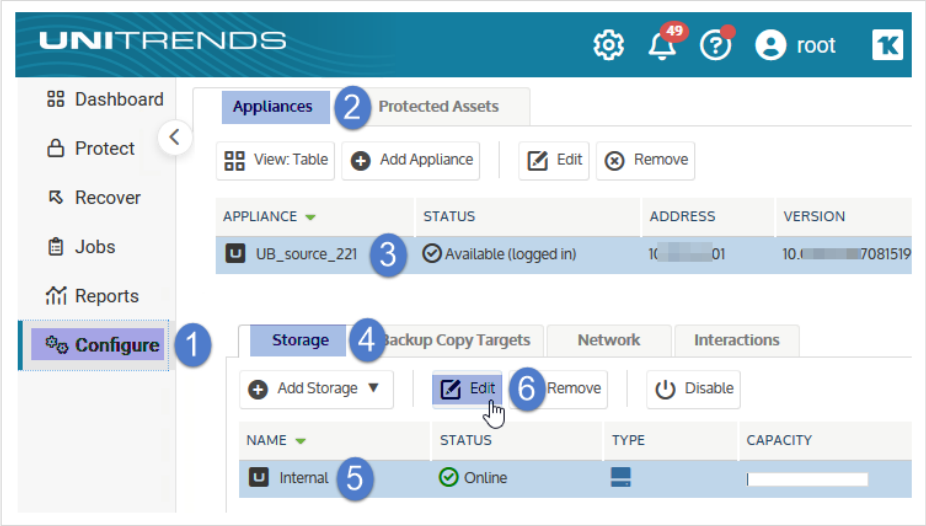

5 In the Manage Attached Disks area:

- To add an attached disk, select a disk from the list of available attached disks and click **Add**.

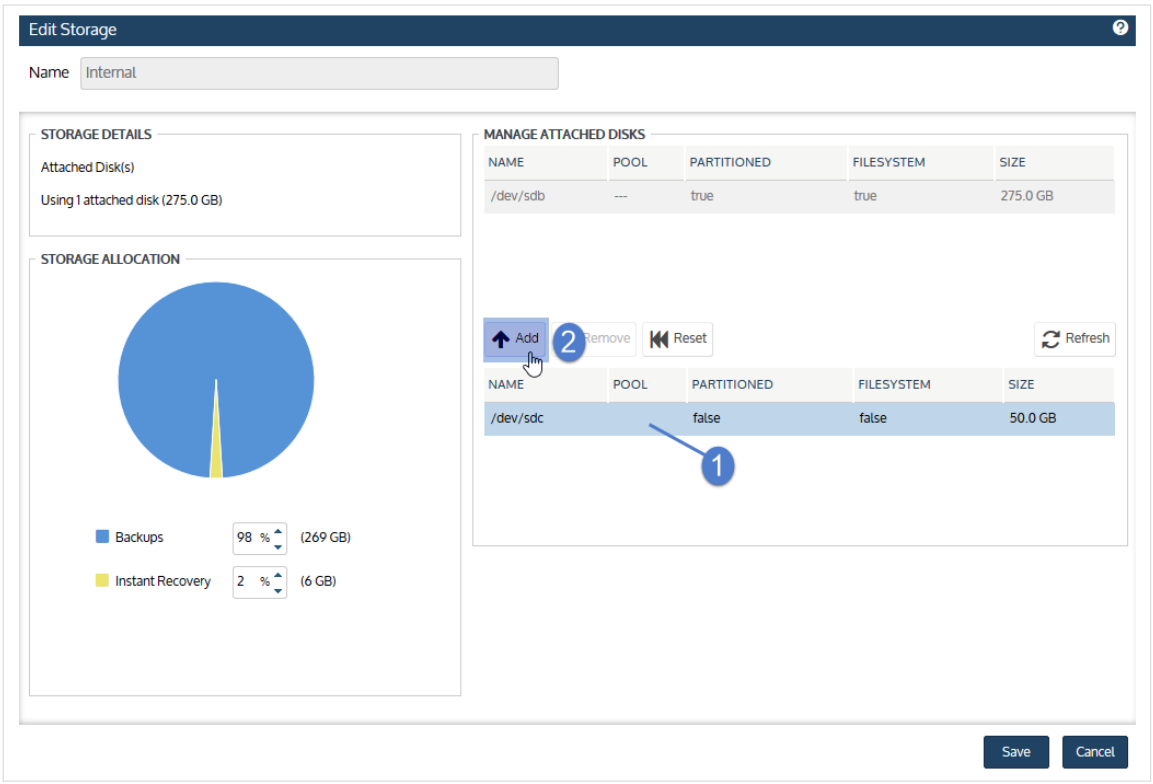

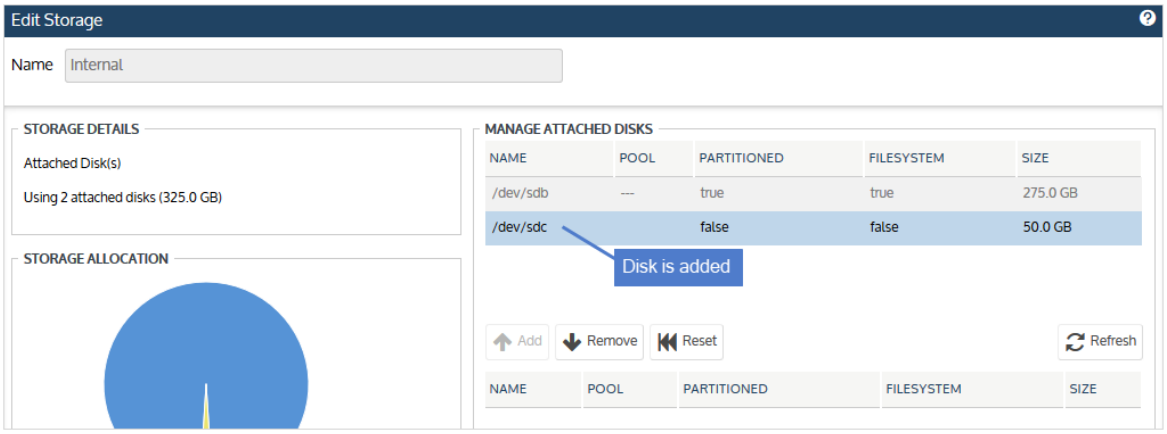

- To remove a disk, select it in the list and click **Remove**.

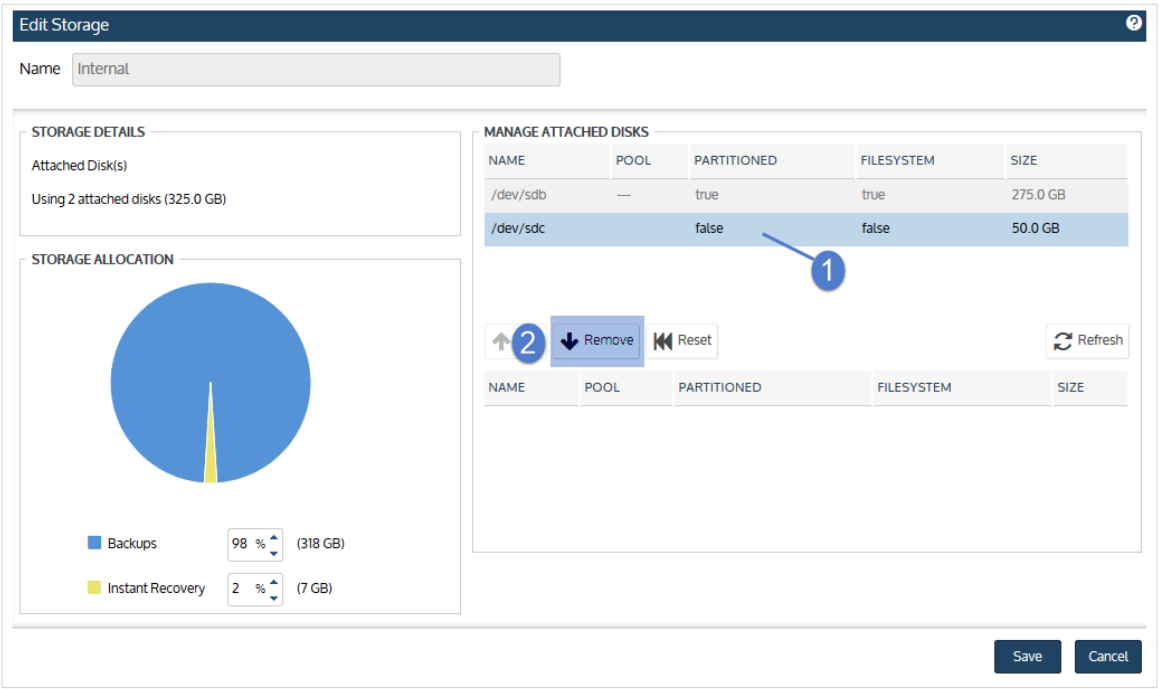

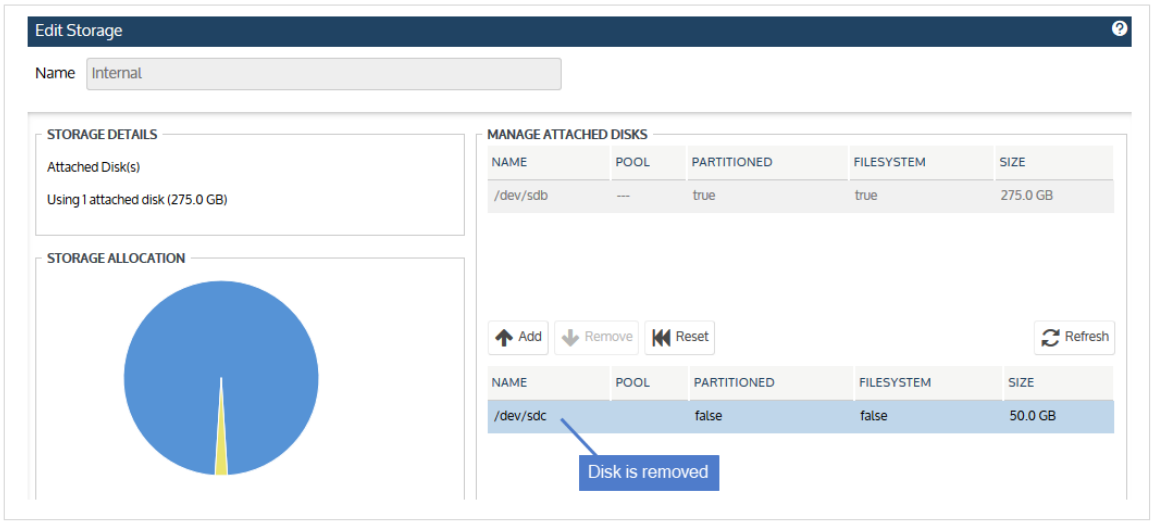

- Click **Reset** to revert all disk settings to the original settings.

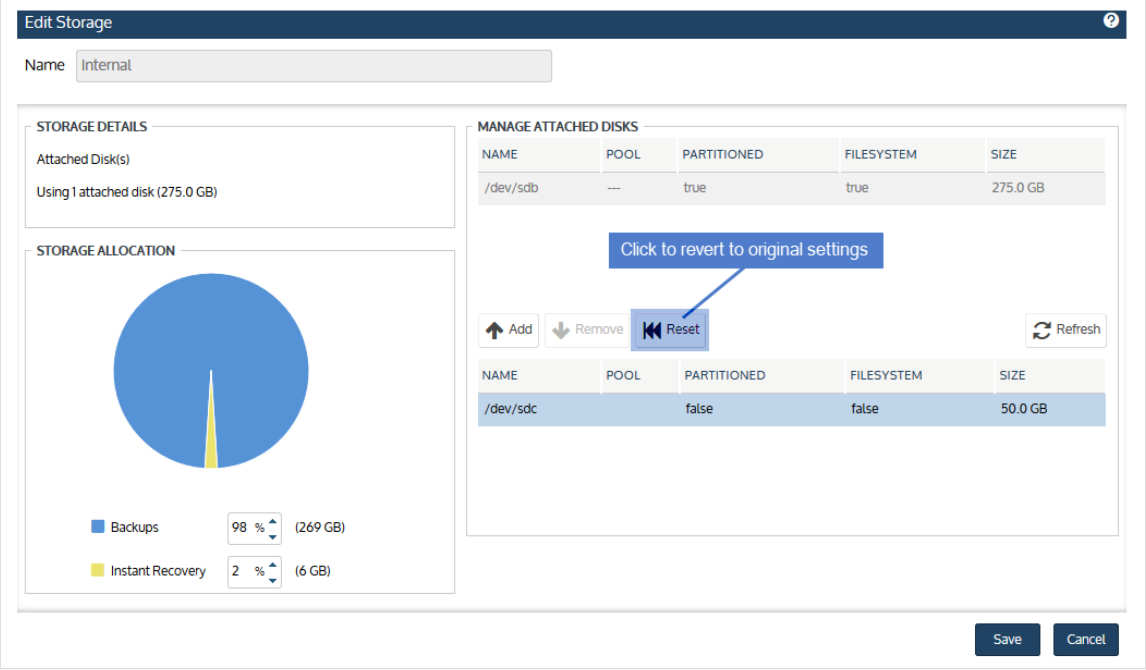

- Click **Refresh** to refresh the list of available disks.

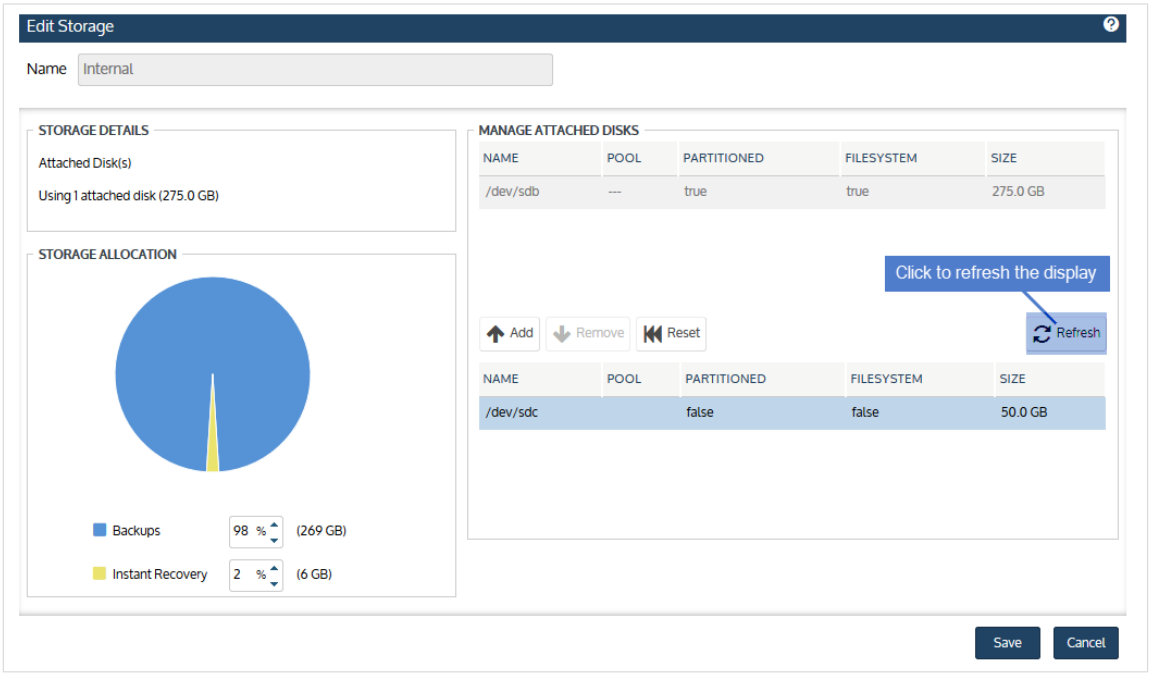

6 (Optional) Adjust the storage allocation for Backups versus Instant Recovery.

7 Click **Save**.

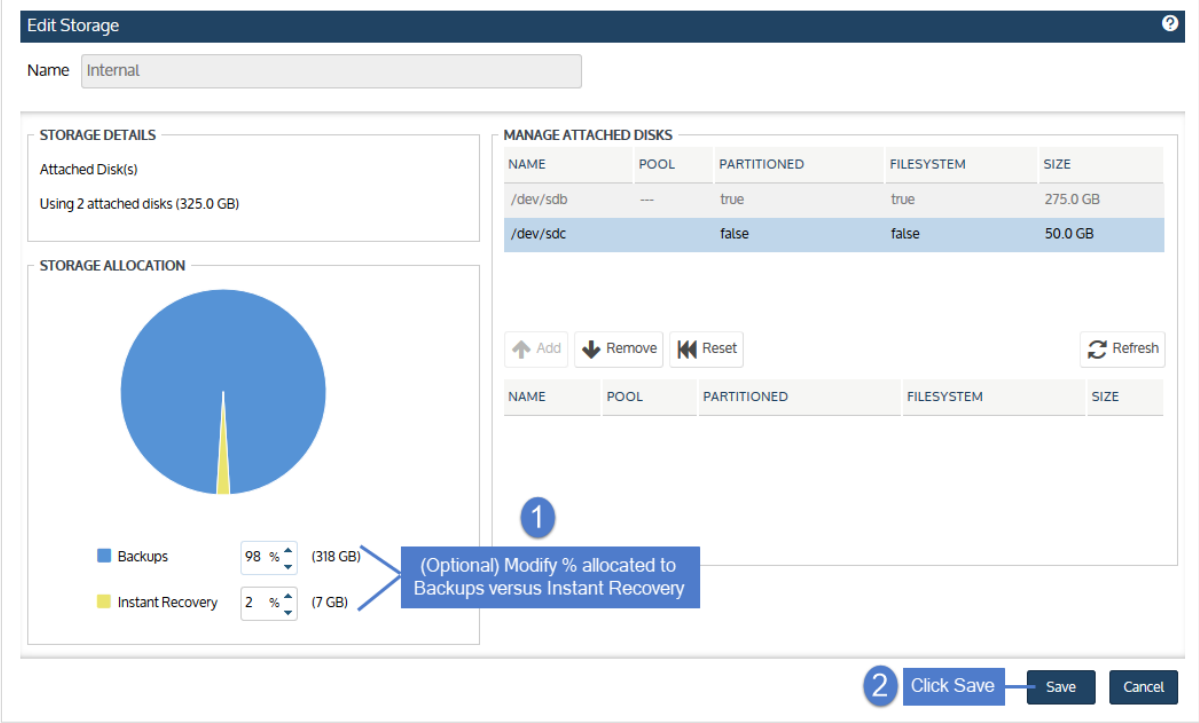

On the **Configure > Appliances** page, the status initially displays as *Pending*. When the storage has expanded to include the new disk, its status changes to *Online* and the storage can be used.

8 (Optional) repeat this procedure to add another disk.

To add a disk as a separate storage area

1 Go to the hypervisor and add a virtual disk to the Unitrends Backup VM:

- For an appliance whose initial backup storage is DAS internal to an ESXi host, see one of the following VMware documents:
  - vSphere 5.1: Create a Virtual Disk in vSphere Client 5.1
  - vSphere 5.5: Create a Virtual Disk in vSphere Client 5.5
  - vSphere 6.0: Create a Virtual Disk in vSphere Client 6.0
  - vSphere 6.5: Add a Hard Disk to a Virtual Machine
- For an appliance whose initial backup storage is DAS internal to a Hyper-V host, Unitrends recommends that you use a VHD(X) disk and that you add the disk to the SCSI controller. See the following Microsoft documents:
  - To create a virtual hard disk
  - To add a hard disk to a virtual machine

- For an appliance whose initial backup storage is DAS internal to an AHV host, edit the appliance VM to add a vDisk to the ISCSI controller (by selecting **ISCSI** in the Bus Type list). For details, see Managing a VM (AHV) in the Prism Web Console Guide.
- For an appliance whose initial backup storage is DAS internal to a XenServer host, Unitrends recommends that you add the VHD disk to the SCSI controller. VM disks cannot be attached as *Read Only*. Be sure to use the *Read Only* = *No* setting when attaching disks.
- For an appliance whose initial backup storage is on an external SAN connected to the hypervisor, add another LUN and expose it to the Unitrends Backup VM. Then go to the hypervisor and add new virtual disks to the VM using storage on the LUN you added.
- For an appliance whose initial backup storage is on an external NAS share connected to the hypervisor, add another NAS share. Then go to the hypervisor and add new virtual disks to the Unitrends Backup VM using storage on the share you added.

2 On the **Configure > Appliances** page, select your appliance.

3 Click the **Storage** tab.

4 Select **Add Storage > Attached.**

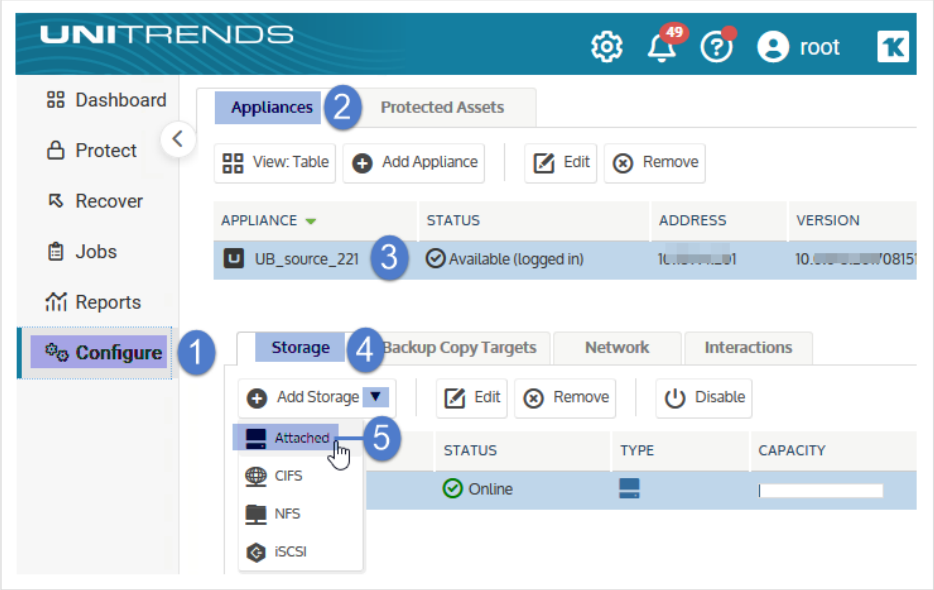

5 Enter a unique **Name** for the storage. This name cannot contain spaces.

6 In the Manage Attached Disks area, select a disk from the list of available attached disks and click **Add**.

- To refresh the list of available disks, click **Refresh**.
- To remove an attached disk, select it from the list of available attached disks and click **Remove**.
- To revert to the original disk settings, click **Reset**.

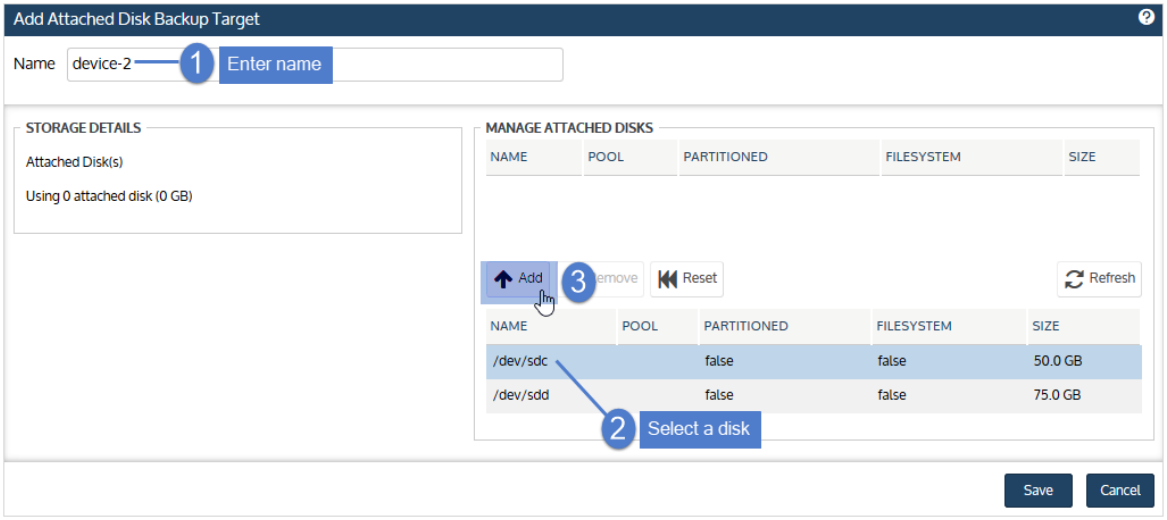

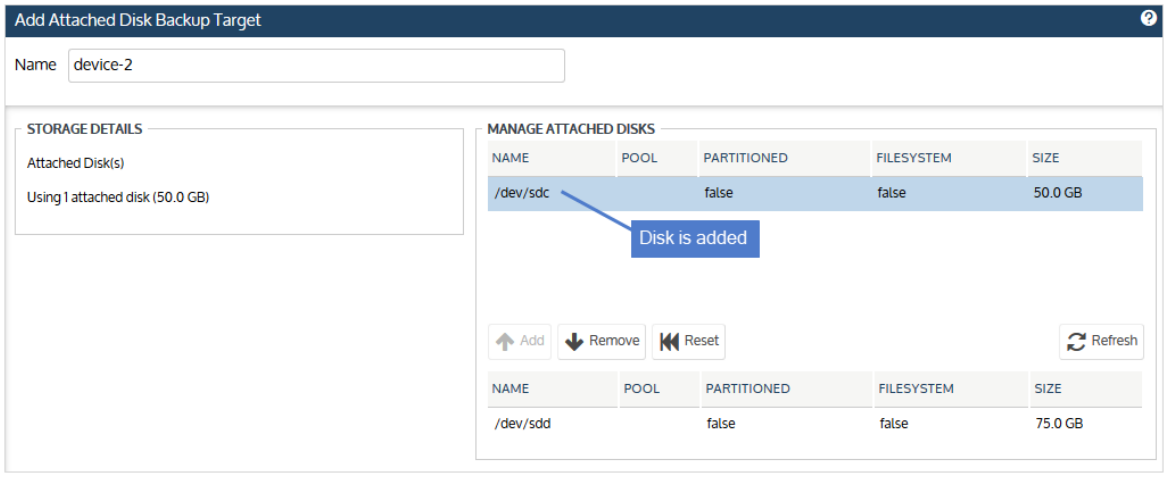

7 Click **Save** to create the storage device.

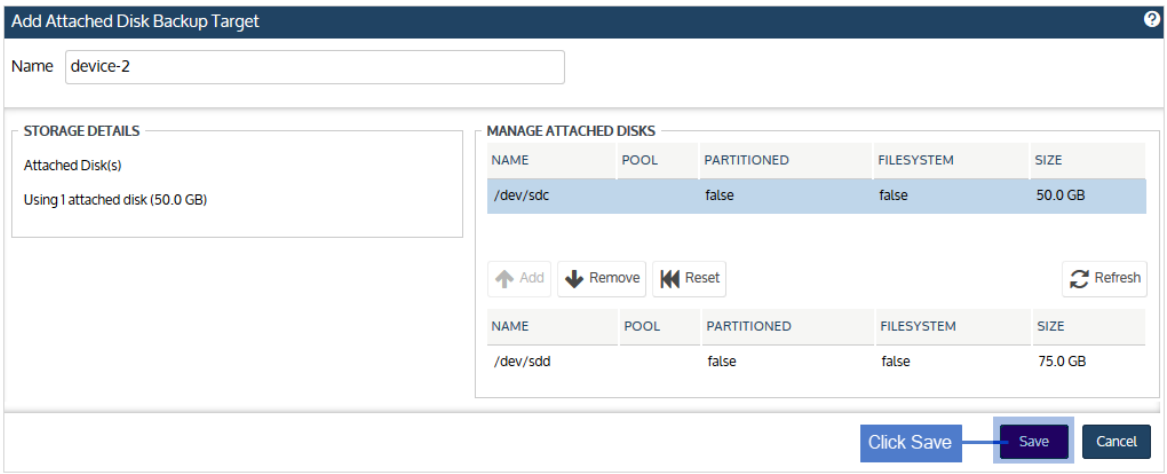

The new device is added and displays on the Storage tab.

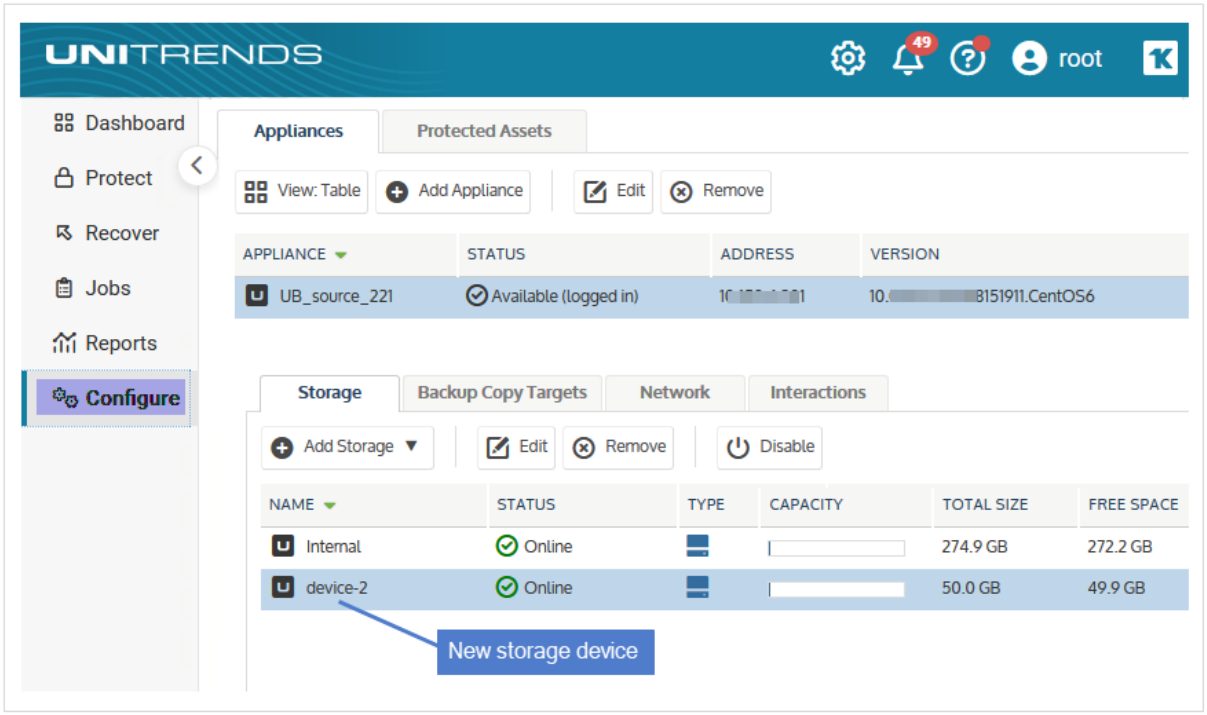

8 (Optional) repeat this procedure to add another disk.

Procedures for adding external storage

Use these procedures to attach external storage directly to the Unitrends Backup VM and configure this storage on the Unitrends Backup appliance. Each SAN LUN or NAS share is added as a separate storage area.

Notes:

- Unitrends recommends that you connect external storage to the hypervisor, rather than directly to the Unitrends Backup VM. Once you connect the SAN LUN or NAS share to the hypervisor, use the hypervisor to create virtual disks from this external storage, then use the procedures in "Procedures for adding attached disk backup storage" on page 200 to add the disks to the Unitrends Backup VM.
- Attaching external storage directly to the appliance VM is not supported for Unitrends Backup on Nutanix AHV.

Before adding storage, be sure to review the recommendations in "About adding backup storage to a Unitrends Backup appliance" on page 199.

Procedures for adding external storage:

- "To add iSCSI storage" on page 211
- "To add Fibre Channel storage" on page 211
- "To add CIFS storage" on page 212

- "To add NFS storage" on page 213

To add iSCSI storage

Use this procedure to connect a SAN LUN directly to the Unitrends Backup VM using the iSCSI protocol, and then configure this storage so it can be used by the Unitrends Backup appliance.

Note: The appliance must be running release 9.0.0-12 or higher to use CHAP authentication.

1. Allocate a LUN on the SAN and expose it to the Unitrends Backup VM.
2. Log in to the appliance UI.
3. On the **Configure > Appliances** page, select your appliance.
4. Click the **Storage** tab.
5. Select **Add Storage > iSCSI**.
6. Enter a unique **Name** for the storage device. This name cannot contain spaces.
7. Enter the IP address of the SAN storage array in the **Host** field.
8. The default port used for iSCSI communication is 3260. If the LUN is configured to use a different port, enter it in the **Port** field.
9. Click **Scan for targets** to retrieve a list of targets on the remote storage array, then choose one from the list.

   Notes: If you do not see the LUN in the list, go to your SAN manager and check the following:
   - Verify that you can see the Unitrends appliance in your SAN manager.
   - Verify that you have a LUN assigned to the Unitrends appliance with the correct permissions.
   - Check with your Storage Administrator for more information.
10. Click **Scan for LUNs** and select one from the list.

    Note: If you receive an error indicating CHAP authentication has failed, CHAP has been configured on the target and either CHAP has not been enabled on the Unitrends appliance, or the Unitrends CHAP credentials do not match those of the target. To configure the appliance to use CHAP, see "To configure iSCSI CHAP authentication" on page 182.
11. Click **Save**.

    The storage is added as a separate device. To store backups on this device, select it when creating backup jobs. See "Creating backup jobs" on page 433 for details.

To add Fibre Channel storage

Use this procedure to connect a SAN LUN directly to the Unitrends Backup VM using Fibre Channel, and then configure this storage so it can be used by the Unitrends Backup appliance.

1. Allocate a LUN on the SAN and expose it to the Unitrends Backup VM.
2. Log in to the appliance UI.

3 On the **Configure > Appliances** page, select your appliance.

4 Click the **Storage** tab.

5 Select **Add Storage > FC**.

6 Enter a unique **Name** for the storage device. This name cannot contain spaces.

7 Enter the IP address of the SAN storage array in the **Host** field.

8 Click **Scan for targets** to retrieve a list of targets on the remote storage array, then select one in the list.

9 Click **Scan for LUNs** and select one in the list.

Notes: If you do not see the LUN in the list, go to your SAN manager and check the following:

- Verify that you can see the Unitrends appliance in your SAN manager.
- Verify that you have a LUN assigned to the Unitrends appliance with the correct permissions.
- You may need to reboot the Unitrends appliance to enable it to discover the storage device.
- Check with your Storage Administrator for more information.

10 Click **Save**.

The storage is added as a separate device. To store backups on this device, select it when creating backup jobs. See "Creating backup jobs" on page 433 for details.

Note: To remove the LUN from Fibre Channel storage on the Unitrends appliance, you must go to the SAN manager and indicate that the SAN should no longer use the LUN.

To add CIFS storage

Use this procedure to connect a NAS share directly to the Unitrends Backup VM using the CIFS protocol, and then configure this storage so it can be used by the Unitrends Backup appliance.

1 Allocate a share on the NAS.

2 Log in to the appliance UI.

3 On the **Configure > Appliances** page, select your appliance.

4 Click the **Storage** tab.

5 Select **Add Storage > CIFS**.

6 Enter the required CIFS share information and click **Save**.

The storage is added as a separate device. To store backups on this device, select it when creating backup jobs. See "Creating backup jobs" on page 433 for details.

CIFS configuration details

Field	Description
Name	Name of the storage. Cannot contain spaces.
Host	IP address or hostname of the NAS share.
Port	Contains the default CIFS port. To use a custom port, enter that port number.
Share Name	Enter the full directory pathname of the NAS share.
Username (optional)	If the share is configured for authentication, enter the domain username as user@domain.com.
Password (optional)	If the share is configured for authentication, enter the password.

To add NFS storage

Use this procedure to connect a NAS share directly to the Unitrends Backup VM using the NFS protocol, and then configure this storage so it can be used by the Unitrends Backup appliance.

1. Allocate a share on the NAS.
2. Log in to the appliance UI.
3. On the **Configure > Appliances** page, select your appliance.
4. Click the **Storage** tab.
5. Select **Add Storage > NFS.**
6. Enter the required NFS share information and click **Save**.

The storage is added as a separate device. To store backups on this device, select it when creating backup jobs. See "Creating backup jobs" on page 433 for details.

NFS configuration details

Field	Description
Name	Name of the storage. Cannot contain spaces.
Host	IP address or hostname of the NAS share.

Field	Description
Port	Contains the default NFS port. To use a custom port, enter that port number.
Share Name	Enter the full directory pathname of the NAS share.
Username (optional)	If the share is configured for authentication, enter the domain username as user@domain.com.
Password (optional)	If the share is configured for authentication, enter the password.

Backup copy targets

Backup copy targets retain copies of your backups on storage external to the appliance. Backup copies provide another layer of protection for your data and should be used for longer-term retention and disaster recovery. Retention options vary by backup copy target, but all targets enable you to define settings for long-term retention.

Backup copy management and administration procedures vary by target type. Once you have added a backup copy target:

- See the applicable topic in "Managing backup copy targets" on page 260 for details on working with the target you added.
- See "Creating backup copy jobs" on page 491 to start copying backups to the target.

Notes:

- The types of backup copy targets supported vary by Unitrends appliance model. For example, attached disk is not supported on Recovery Series, Recovery MAX, and ION/ION+ appliances.
- When adding backup copy targets to an appliance, only supported types display in the **Configure > Appliances > Backup Copy Targets > Add Target** list.
- To determine which backup copy targets are supported for your appliance, see *Archiving Adapters* in Unitrends Recovery Series and Recovery Series MAX or ION and ION+ Appliance Models, or *Local Backup Copies* in Unitrends Backup Editions.

Use these procedures to add backup copy targets to your Unitrends appliance:

- "Adding a Unitrends Cloud backup copy target" on page 215
- "Adding a Unitrends appliance backup copy target" on page 215
- "Adding an eSATA or USB backup copy target" on page 233
- "Adding a tape backup copy target" on page 234
- "Adding a third-party cloud backup copy target" on page 244

- "Adding an attached disk backup copy target" on page 251
- "Adding a NAS backup copy target" on page 254
- "Adding a SAN backup copy target" on page 258

Adding a Unitrends Cloud backup copy target

You can use Unitrends Cloud Backup to store copies of your backups. To purchase this offering, go to http://www.unitrends.com/products/cloud-solutions/unitrends-cloud-backup.

To add a Unitrends Cloud backup copy target

1. On the **Configure > Appliances** page, select your appliance.
2. Click the **Backup Copy Targets** tab.
3. Select **Add Target > Unitrends > Unitrends Cloud**.
4. Enter your Unitrends Cloud license string in the text box and click **Submit**.
5. Final configuration steps take place in the background and can take up to 15 minutes to complete. When complete:
   - A Unitrends Cloud device has been added as a backup copy target.
   - The Unitrends appliance name has been modified to include an additional 12 characters, required to uniquely identify this appliance in the Unitrends Cloud.
6. Fine-tune settings as described in "Step 4: Return to the source backup appliance and fine-tune settings by adjusting connection options" on page 225.
7. Create a job to start copying backups to the Unitrends Cloud. For details, see "To create a backup copy job for a Unitrends Cloud target" on page 501.

Adding a Unitrends appliance backup copy target

To add an appliance target, review the requirements in "Preparing to add a Unitrends appliance backup copy target", then set up the target as described in "To add a Unitrends backup copy target appliance" on page 217.

Preparing to add a Unitrends appliance backup copy target

Before adding the backup copy target, verify that you have met appliance requirements and seed the target appliance (optional).

For use as a backup copy *target*, an appliance must meet these requirements:

- Must be a Unitrends Backup appliance, an ION/ION+ appliance, a Recovery MAX appliance, or a Recovery Series model listed in the Recovery Series Appliance Family Datasheet.

  Notes:

  - Unitrends Backup appliances deployed to the following environments can be used either as a backup appliance or as a backup copy target (performing both roles is not supported):

  - VMware
  - Hyper-V
  - Nutanix AHV
  - Citrix XenServer
  - Microsoft Azure
  - Amazon Web Services
- Recovery Series, Recovery MAX, and ION/ION+ appliances can be used as both a backup appliance and a backup copy target, if desired.

- For Unitrends Backup, you must license and register the appliance.
- The appliance must have a static IP address assigned. An appliance that is configured to use DHCP cannot be used as a backup copy target.
- The appliance must have at least 200GB of available backup storage space.
- You must be able to log into the appliance by entering ***https://<applianceIPaddress>/ui/*** in a browser on the local network. If you enter ***https://<applianceIPaddress>/ui/*** and receive the message *Managed by UniView*, local network access has been disabled on the target appliance. You must log in to UniView and enable local access before running the "To add a Unitrends backup copy target appliance" procedure. Once the backup copy target has been added, return to UniView and disable local network access on the target appliance. For details on enabling and disabling local network access, see "Disable or enable local network access to an appliance" on page 170.
- The configuration process creates a secure VPN tunnel. The VPN tunnel must be on a dedicated subnet. Be sure to use a network that is not in use in your environment. The VPN tunnel cannot use the following subnets as these are reserved by Unitrends: 172.17.0.0 through 172.17.10.0.
- If adding multiple appliance targets to one source, a separate secure VPN tunnel is created for each target you configure. Use a separate dedicated subnet for each (for example, 172.17.13.0, 172.17.14.0, and 172.17.15.0). To view the IPs that are already in use by other targets, view the hosts file on the source appliance (**Configure > Appliances > Network > Edit Hosts File**):

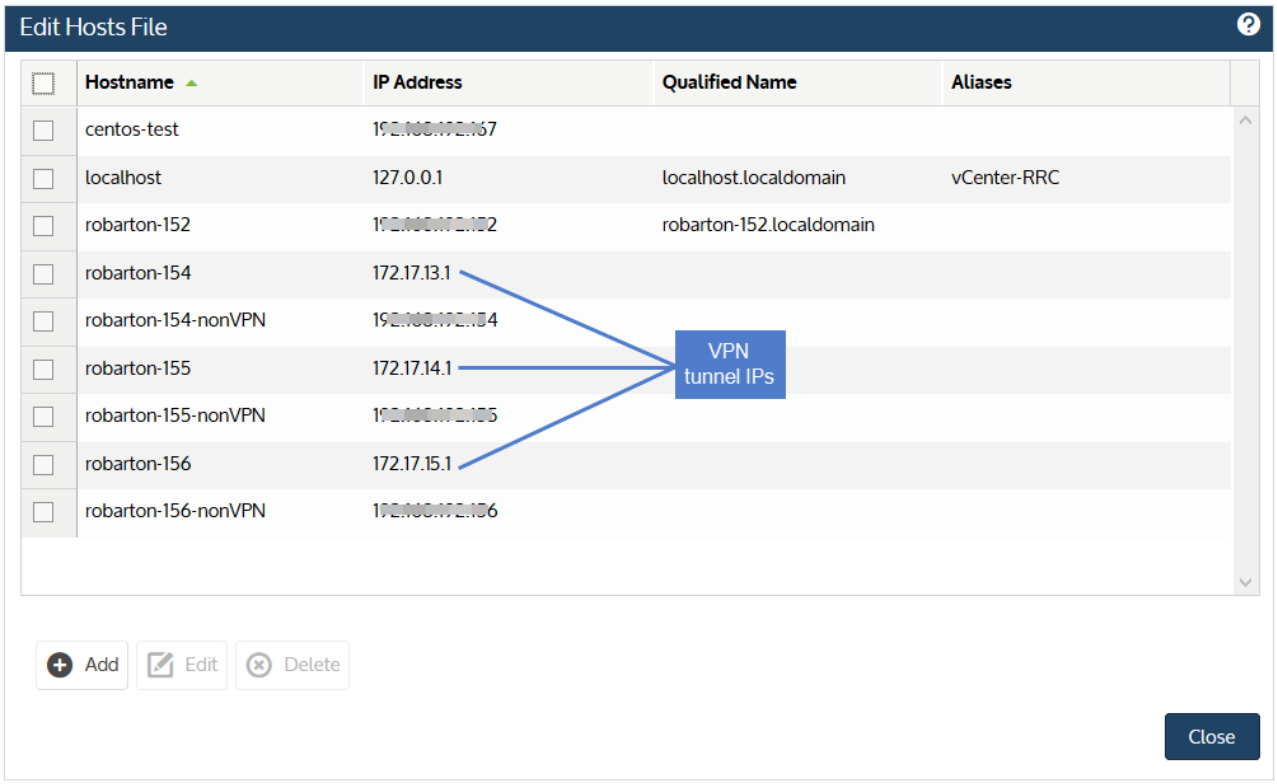

- If one appliance is already managing the other (you are accessing both from one UI), then the manager appliance can be configured as the backup copy target only. To use the manager appliance as the source backup appliance, log in to the managing appliance and remove the managed appliance (as described in "To remove a managed appliance" on page 420) before configuring the backup copy target.

Be aware that the appliance's backup storage is used to store the backup copies, so on-appliance retention of local backups will be impacted if the appliance is also being used to run and store local backups.

For large data sets, Unitrends recommends that you seed the initial data set to the backup copy target by using removable media (disk or NAS). This seeding greatly reduces the time required to copy the first backups. For details, see the Rapid Seed for Backup Copy Target Appliances document.

To add a Unitrends backup copy target appliance

This procedure uses the following terms:

- *Backup copy target* is the appliance that will be receiving and storing backup copies.
- *Source backup appliance* is the appliance running the local backups that will be copied to the target appliance.

Step 1: Do these steps on the backup copy target appliance

1. Log in to the backup copy target appliance.
2. On the **Configure > Appliances** page, select the target appliance and click **Edit**.

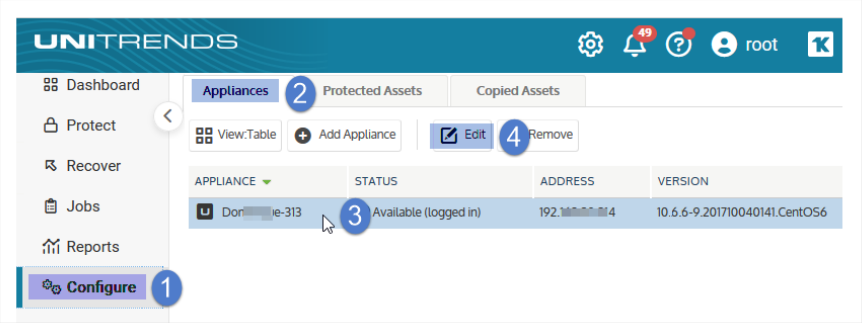

3 On the **Backup Copy** tab, select **Enable this appliance as a Backup Copy Target** and enter the following information to configure a secure tunnel connection:

- Secure Network IP – This must be an unused network in your environment. The VPN tunnel cannot use the following subnets as these are reserved by Unitrends: 172.17.0.0 through 172.17.10.0.
- Secure Netmask IP – The netmask associated with the secure network above.
- Secure Port – Port number to use for OpenVPN.

Notes:

- If adding multiple appliance targets to one source, a separate, secure VPN tunnel is created for each target you configure. Use a separate dedicated subnet for each (for example, 172.17.13.0, 172.17.14.0, and 172.17.15.0).
- To view the IPs that are already in use by other targets, view the hosts file on the source appliance (**Configure > Appliances > Network > Edit Hosts File**):

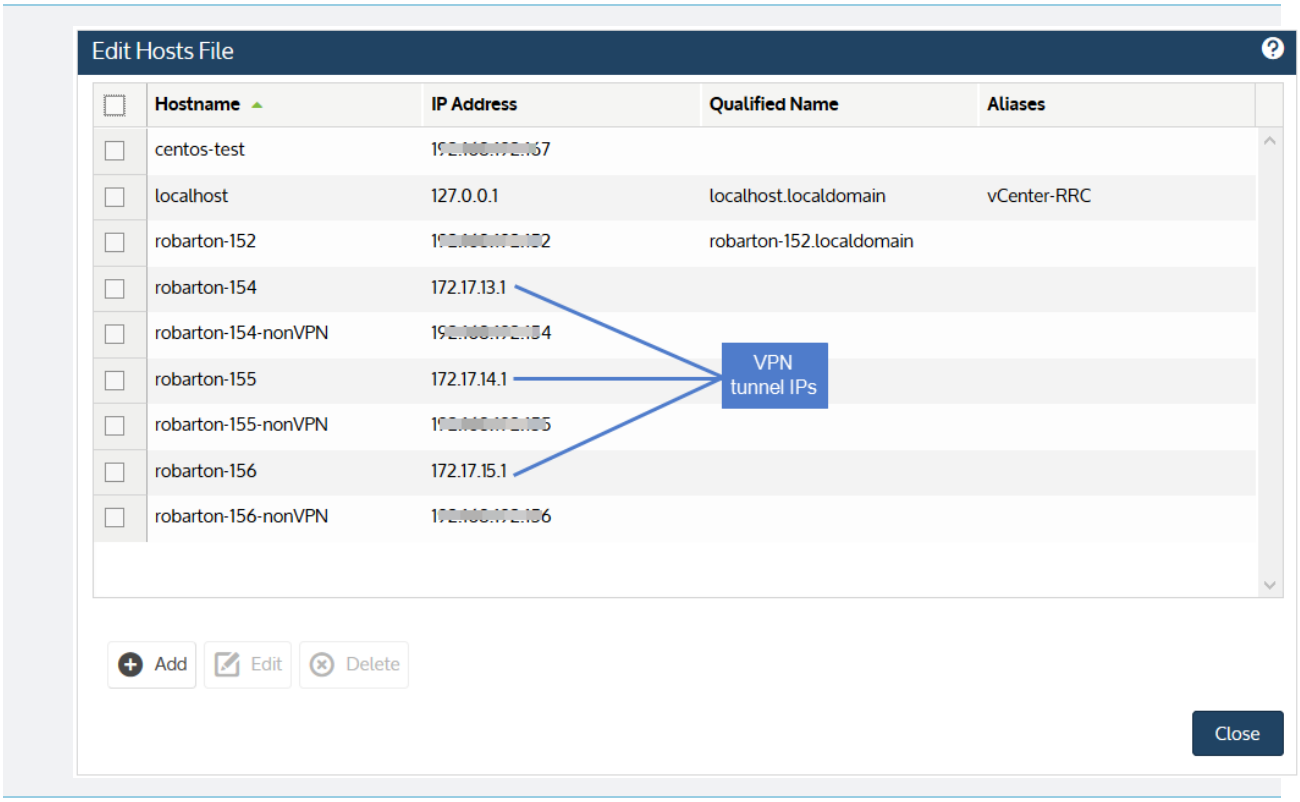
Edit Hosts File
Hostname
IP Address
Qualified Name
Aliases
centos-test
localhost
127.0.0.1
localhost.localdomain
vCenter-RRC
robarton-152
robarton-152.localdomain
robarton-154
172.17.13.1
robarton-154-nonVPN
robarton-155
172.17.14.1
robarton-155-nonVPN
robarton-156
172.17.15.1
robarton-156-nonVPN
VPN tunnel IPs
Add
Edit
Delete
Close

4 Click **Save**.

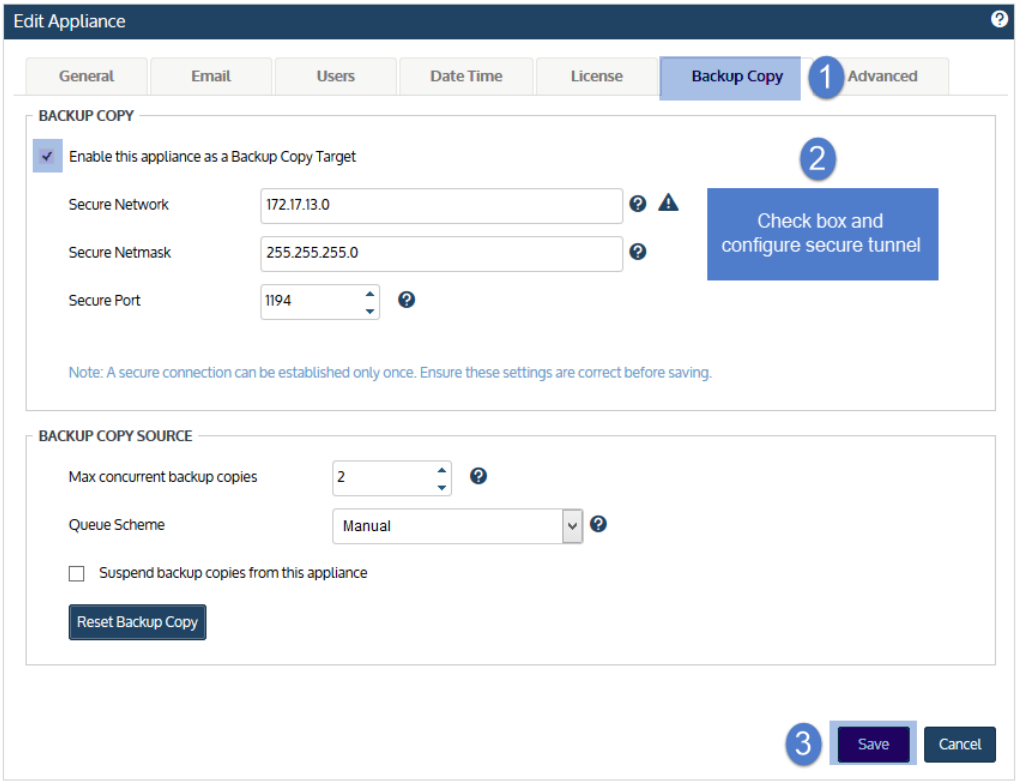

Step 2: Do these steps on the source backup appliance

1. Log in to the source backup appliance.
2. On the **Configure > Appliances** page, select the source backup appliance.
3. Click the **Backup Copy Targets** tab.
4. Select **Add Target > Unitrends > Unitrends appliance**.

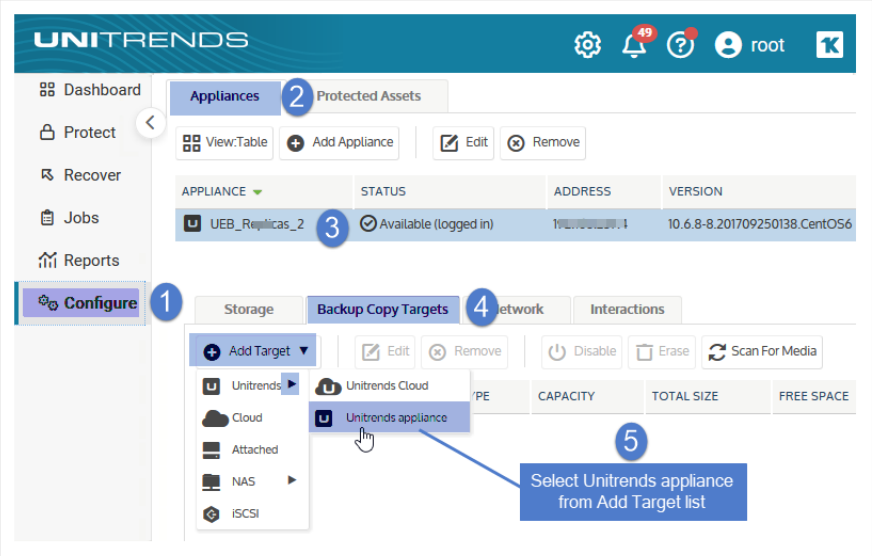

5 Enter the Host Name, Host IP address, User Name (optional), and Password (optional) of the backup copy target appliance, and click **Save**.

> **Note:** If you receive the message *Access to the target is restricted*, local network access has been disabled on the target appliance. You must log in to UniView and enable local access before continuing with this procedure. Once the backup copy target has been added, return to UniView and disable local network access on the target appliance. For details on enabling and disabling local network access, see this KB article: Replication setup when a target is only accessible via UniView.

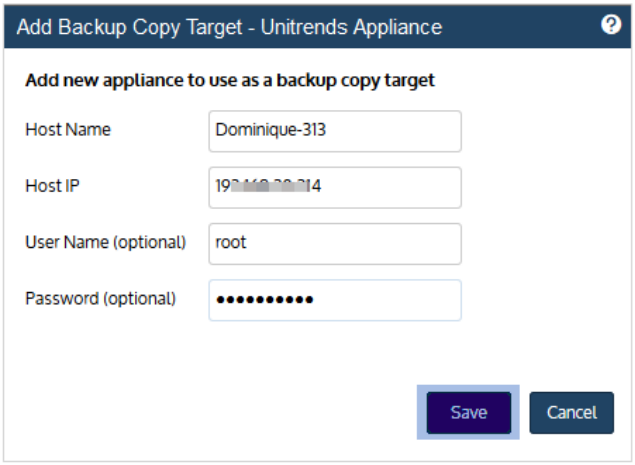

6 You receive a *Connection could not be verified* warning (unless you have previously installed your own custom security certificates on the target appliance.) As long as you trust your network configuration, you can safely click **Continue**.

7 After clicking **Continue**, it takes several minutes to configure the OpenVPN tunnel. Before continuing, wait for the confirmation message that the appliance has been added as a backup copy target.

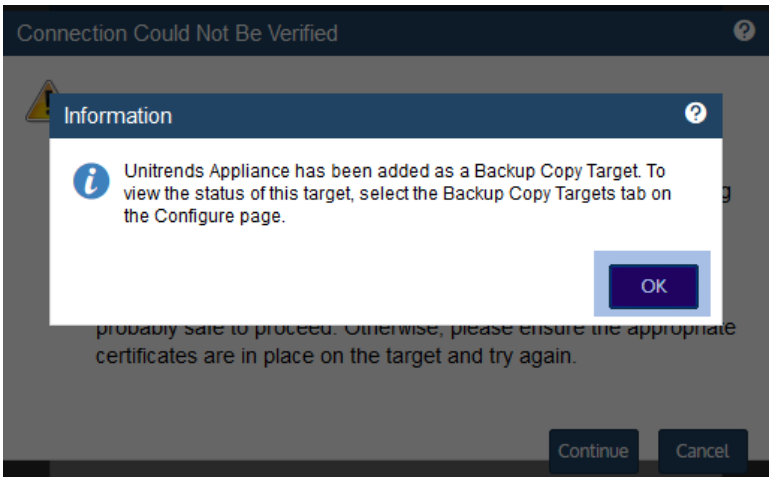

Step 3: Return to the backup copy target appliance and do these steps

1. Log in to the backup copy target appliance.
2. On the **Configure > Appliances** page, select the source appliance.

   The source has an **!** icon and a Pending status. You may need to refresh to page to display the newly added source appliance.
3. Click the **!** icon and enter the following:
   - Select **Allow backup copies to be sent from *sourceAppliance***.
   - Select the storage device where backup copies will be stored on the target appliance.
4. Click **Save**.

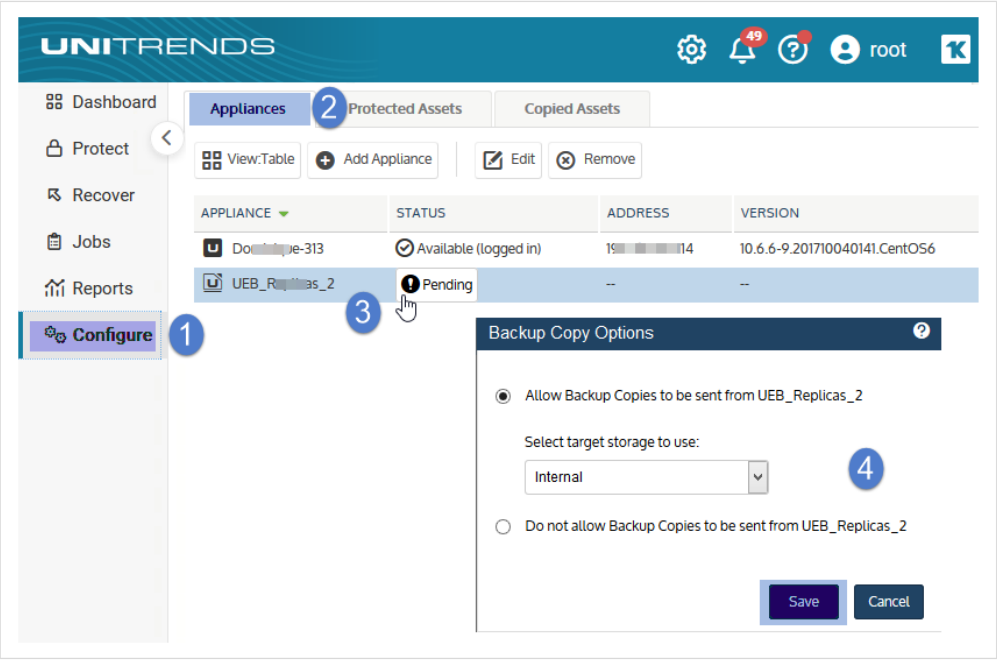

5 Final configuration steps take place in the background and can take up to 20 minutes to complete. Configuration completes when you see the following:

- On the target's Appliances tab, the source appliance status has changed from Pending to Not Available:

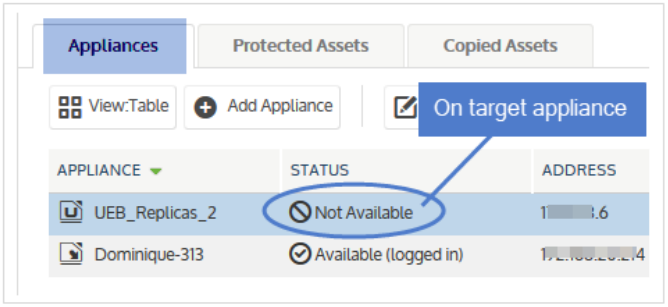

- On the source backup appliance Backup Copy Targets tab, the status of the backup copy target has changed to Online:

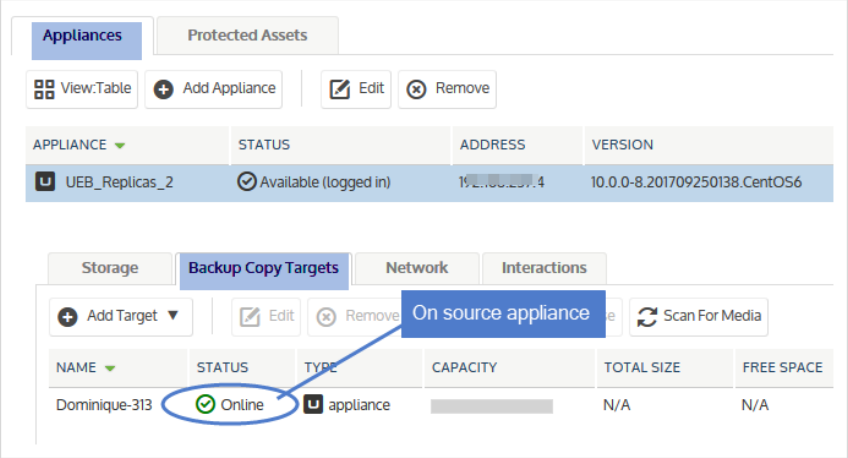

Once the above statuses have changed, you can use the target.

Notes:

- To check the appliance status, reload the page to refresh the display.
- The Not Available status simply means the target appliance does not have credentials to manage the source beyond receiving its backup copies.

6 (Optional) Add management credentials to the source appliance:

Notes: By default, the backup copy target appliance does not manage backup copy source appliances.

When the source appliance is NOT managed by the target:

- Backup copy jobs do not display on the **Jobs > Active Jobs** tab or on the **Dashboard > Active Jobs** tile on the target. To view backup copy jobs from the target appliance, you must add management credentials for the source appliance.
- You cannot edit settings of copied assets. The Edit button does not display on the Copied Assets tab. To enable edits, you must add management credentials for the source appliance. For more on editing these settings, see "Copied Assets" on page 403.
- You can not manage retention policies for copied assets. To manage retention policies, you must add management credentials for the source appliance. For details, see "Copied Assets" on page 403.

- Select the source appliance and click **Edit**.
- Check **Enable Management of this appliance**, supply **User Name** and **Password** credentials, and click **Save**. Click **OK**.

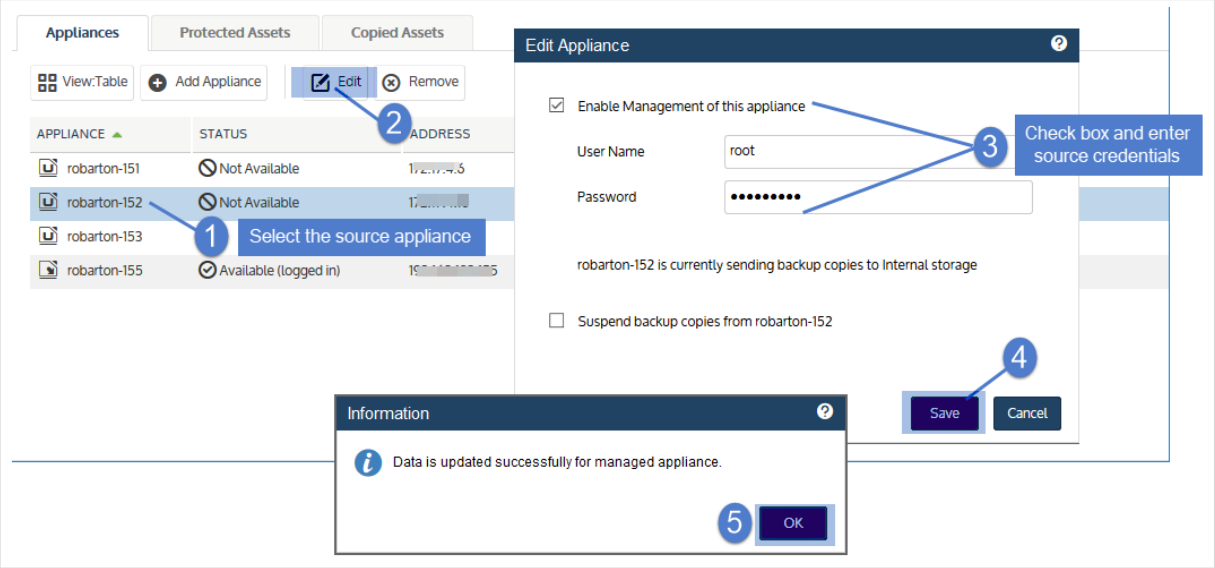

- The source appliance status changes to Available, indicating that the target can now manage the source.

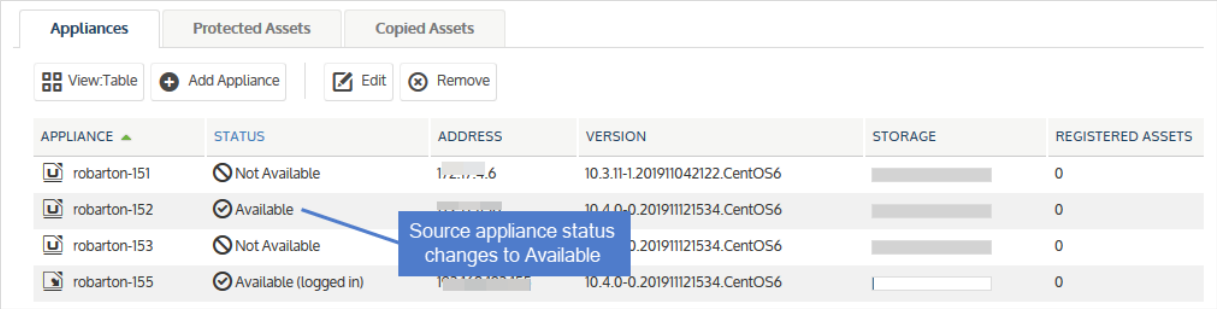

Step 4: Return to the source backup appliance and fine-tune settings by adjusting connection options

For Unitrends Cloud and Unitrends appliance hot backup copy targets, you can adjust the following settings for optimal performance in your environment:

Setting	Description	Procedure
Max concurrent backup copies	Determines how many backups can be copied concurrently. While the default setting of two is adequate for most deployments, you may wish to increase the concurrency when you have enough WAN bandwidth to support more concurrent copies.	See "To configure connection options" on page 231.
Queue scheme	Determines the order in which the source appliance copies backups to the target: • Recency – By default, the source backup appliance sends	See "To configure connection options" on

Setting	Description	Procedure
	copies to the target using the *recency* queue scheme, where the most recent backups are copied first. Unitrends recommends this approach because it supports recovering from availability issues with the target appliance (or the WAN connecting to the target appliance) by skipping over older backups when a newer backup arrives. This is particularly important if the connection to the target appliance is unreliable. • Maximize retention – If it is important to you to ensure that every backup on the source is copied to the target, choose the *Maximize retention* queue scheme. • Manual – To copy backups to the target manually, choose the *Manual* queue scheme. With the Manual scheme, the appliance does not add backups to the queue. Instead, you must copy backups by using the procedure "To copy a full backup to a hot backup copy target on-demand" on page 499.	page 231.
Suspend backup copies from this appliance	Check this box to stop sending backup copies from the source appliance. This option may become necessary when either your target appliance or the connection to your target becomes unavailable for an extended period of time.	See "To configure connection options" on page 231.
Reset Backup Copy	Use to stop active copy jobs, reset the backup copy processes, then restart active jobs that were stopped. Use only when working with Support or following troubleshooting instructions in a Unitrends KB article.	See "To configure connection options" on page 231.
Backup copy bandwidth throttling	If the WAN connection to your backup copy target is shared with general purpose Internet use during normal business hours, you may wish to throttle the amount of bandwidth that backup copies can use during these hours.	See "To set bandwidth throttling" on page 227.

Setting	Description	Procedure
	Note: The bandwidth throttling setting limits the maximum amount of bandwidth that backup copy jobs can use. The amount of bandwidth available for the jobs is also effected by environmental factors, such as actual network throughput (which may be constrained by intermediate nodes between the source and target) and other backup copy job tasks, such as encryption and compression.	

To set bandwidth throttling

Note: This procedure sets bandwidth throttling for all hot backup copy targets that have been added to the source appliance. Be sure to consider the impact to all hot copy targets when configuring these settings.

1. Select **Configure > Appliances**.
2. Select the source backup appliance and click the **Backup Copy Targets** tab below.
3. Click **Configure Bandwidth**.

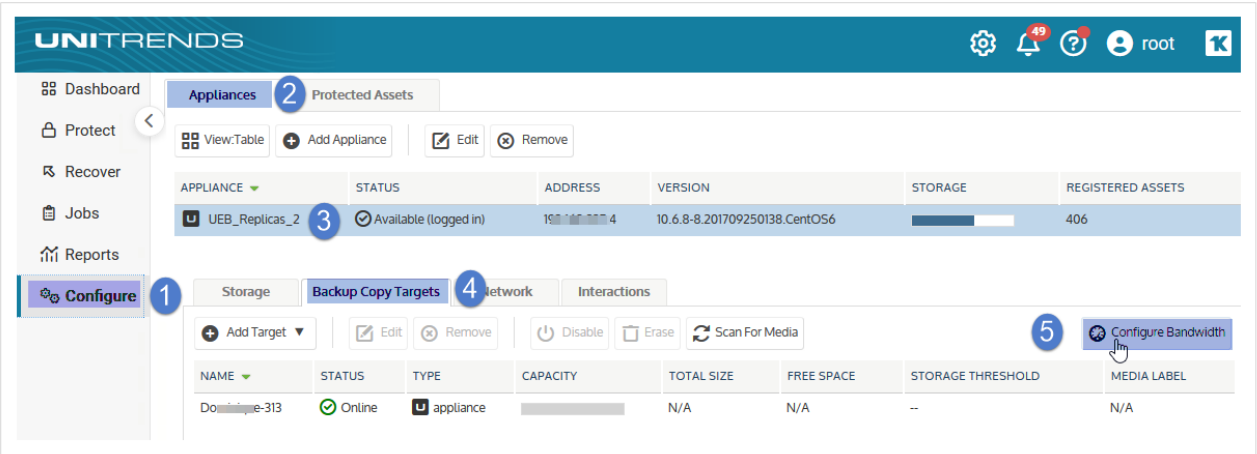

4. Choose the **Connection Type** that most closely matches your WAN bandwidth.

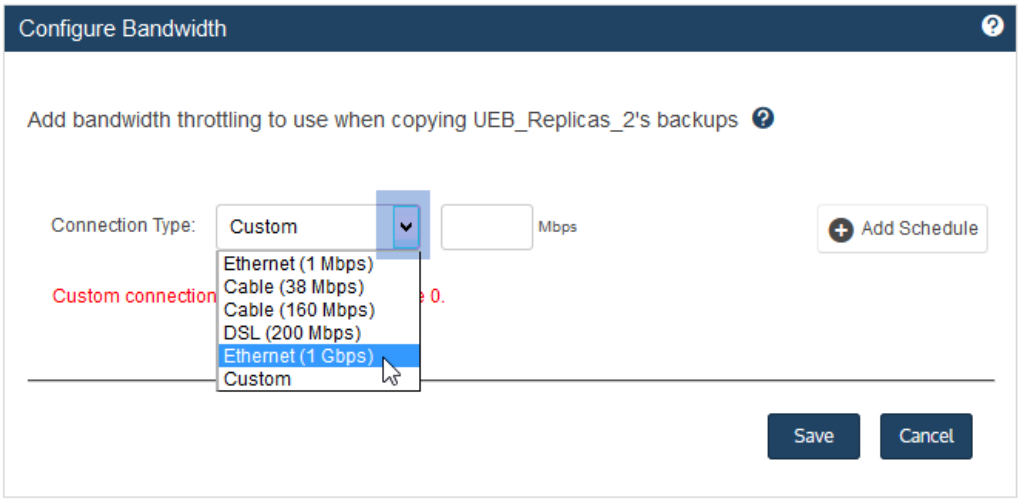

5 Click **Add Schedule**.

6 Choose a Throttle percentage and define the days of the week and times when this percentage will be used.

- The maximum bandwidth that backup copy jobs can use during the scheduled times is *X* percent of the Connection Type you chose in step 4 on page 227 above.
- For Hours values, From begins at *:00* minutes and To ends at *:59* minutes. In this example, 50% throttling applies each weekday from 6:00 AM through 6:59 PM:

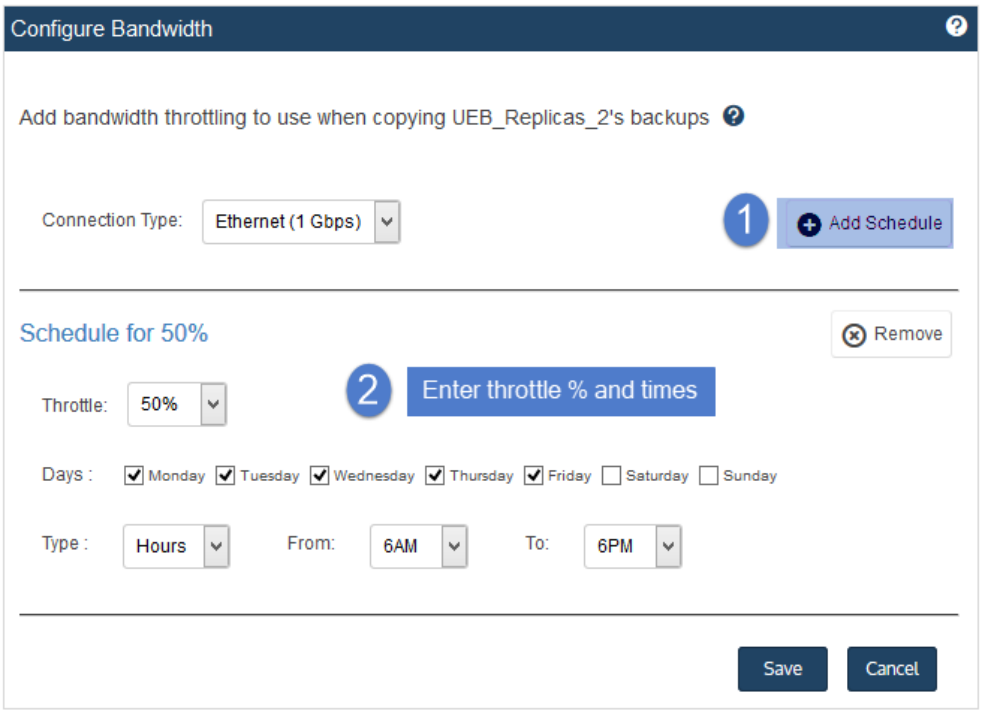

7 Repeat as necessary to create additional throttling schedules:

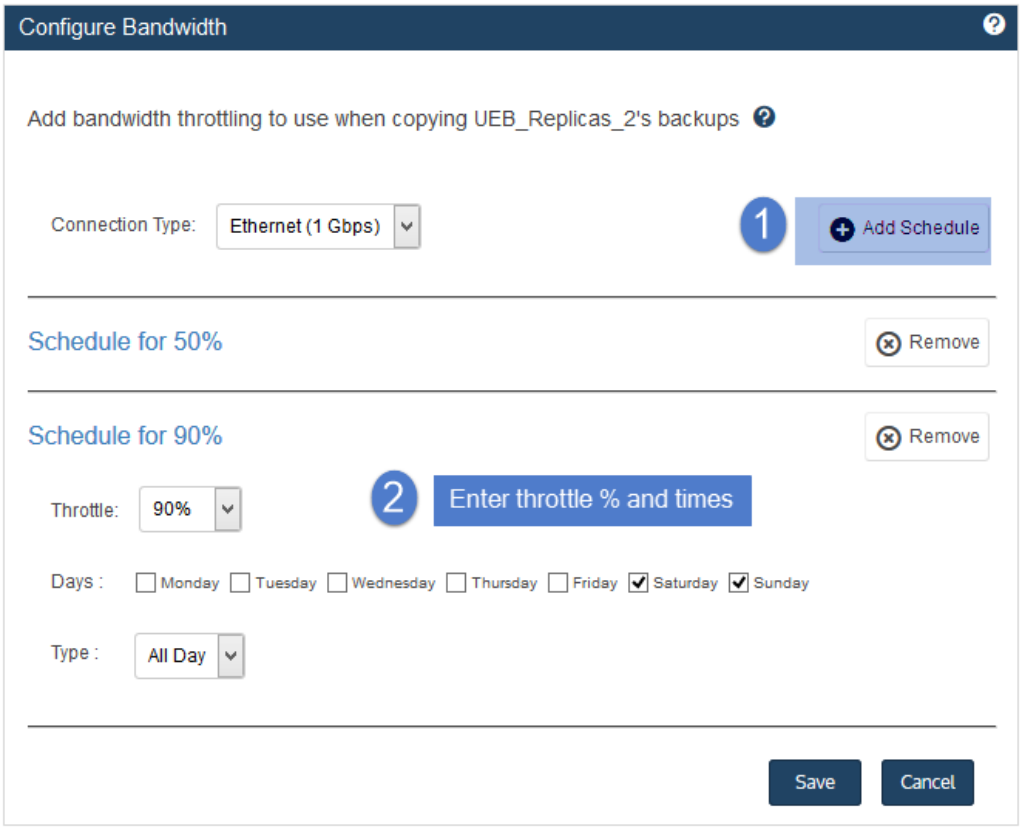
Configure Bandwidth
Add bandwidth throttling to use when copying UEB_Replicas_2's backups
Connection Type:
Ethernet (1 Gbps)
1
Add Schedule
Schedule for 50%
Remove
Schedule for 90%
Remove
Throttle:
90%
2
Enter throttle % and times
Days :
Monday
Tuesday
Wednesday
Thursday
Friday
Saturday
Sunday
Type :
All Day
Save
Cancel

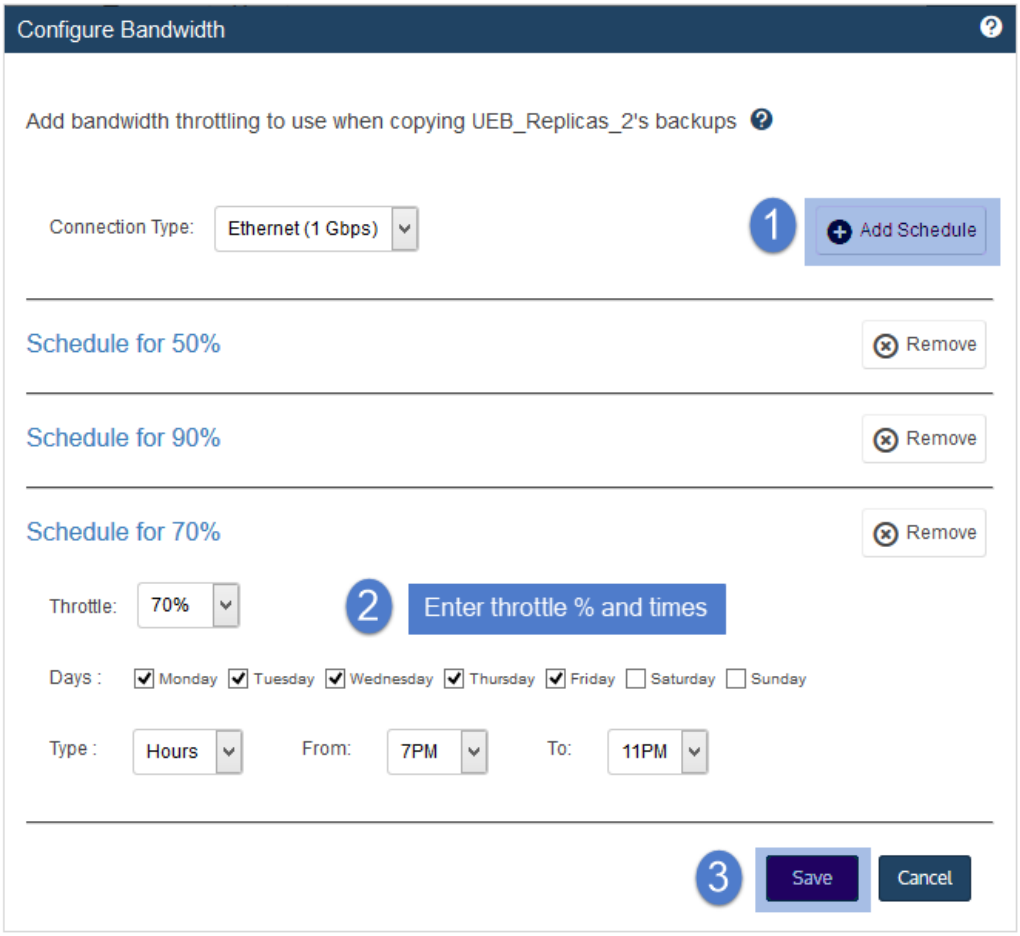

8 Click **Save**, then click **OK** to close the Information message.

9 (Optional) Review and edit bandwidth schedules as needed.

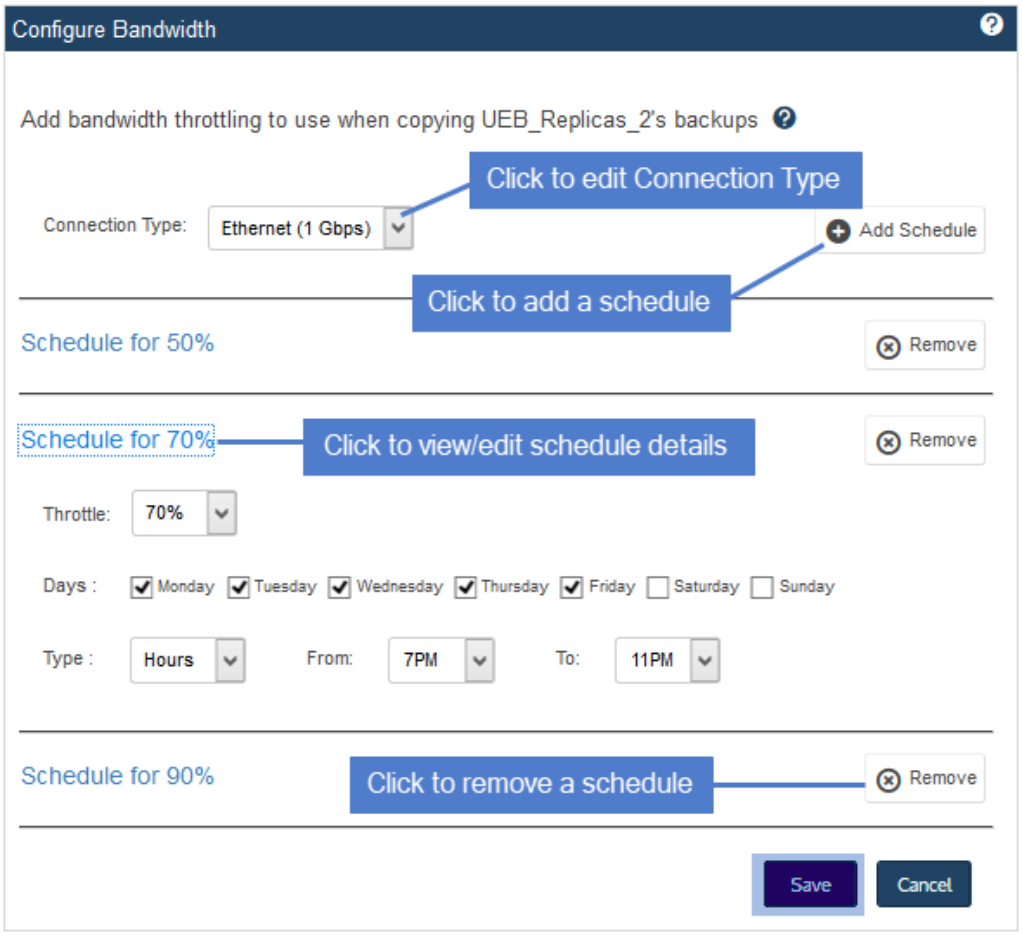

To configure connection options

1. Select **Configure > Appliances**.
2. Select the source backup appliance and click **Edit**.

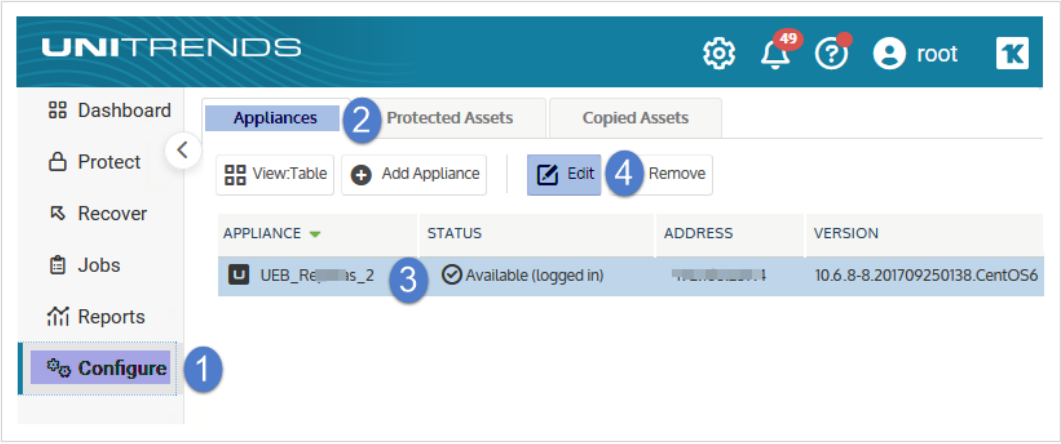

3 On the **Backup Copy** tab, modify Backup Copy Source settings for this hot backup copy target.

Note: If the appliance has multiple hot backup copy targets, a scroll bar displays. If needed, scroll to locate the target to modify.

4 Click **Save**.

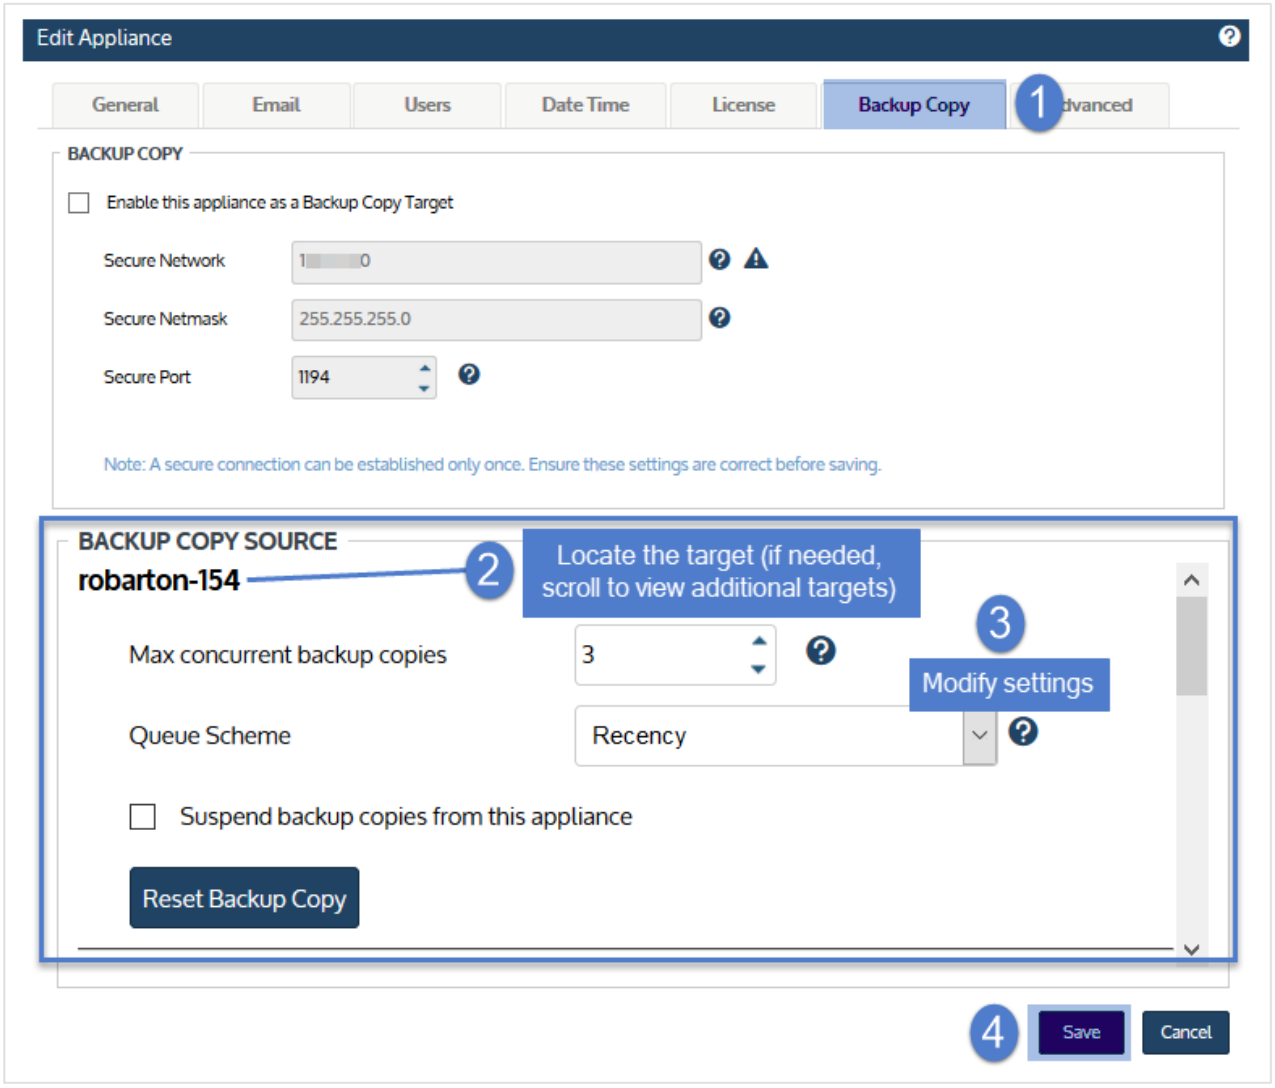

Step 5: On the source backup appliance, create a job to start copying your backups.

For details, see "To create a backup copy job for a Unitrends appliance target" on page 496.

Adding an eSATA or USB backup copy target

Use this procedure to add an eSATA or USB backup copy target. For eSATA and USB devices, you must initialize new drives before using them for the first time. This permanently deletes any existing data and formats the drives.

Additional considerations for using a device with multiple drives:

- All drives attached to a single appliance must have equal capacity. Drives may be of varying capacity if they are attached to different appliances.
- Within the backup copy target, all drives attached to a single appliance are treated as one logical device. When you add a multi-drive device and initialize the drives, the appliance formats them as a single unit. Data is then written across all drives as if they are one larger drive. Once you copy backups to the device, these drives must be treated as a single entity. You must remove them as a set and, to recover data, you must insert all drives in the set back into the eSATA or USB device so the appliance can read and import backup data. If you separate a drive from the set, all data is lost.

To add an eSATA or USB backup copy target

1. Connect the eSATA or USB device to the Unitrends appliance and power it on. For details, see the instructions you received with the device.
2. Log in to the appliance UI.
3. On the **Configure > Appliances** page, select the source backup appliance.
4. Click the **Backup Copy Targets** tab below.
5. Click **Scan For Media**.
   - The appliance discovers the device and it displays on the Backup Copy Targets tab.
   - The device displays as Type *eSATA* or *USB* and its status is *Offline*.

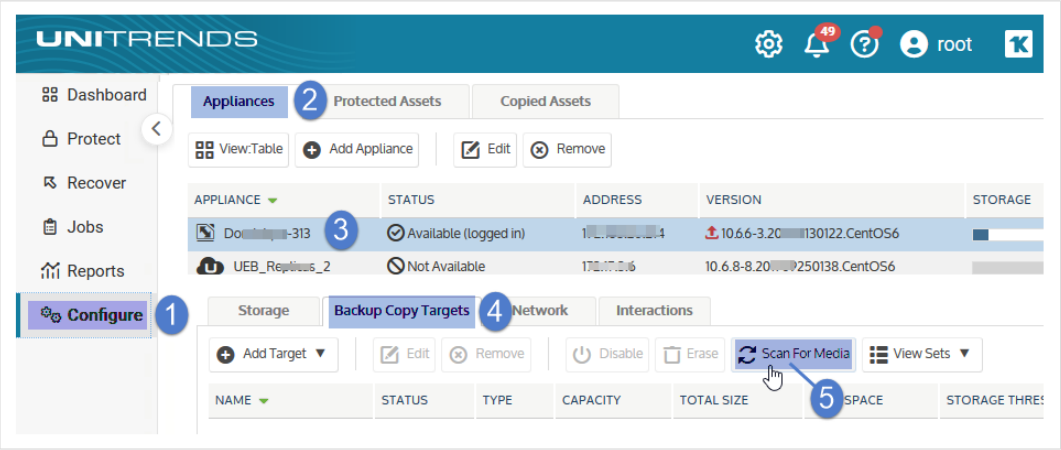

6. Do one of the following:
   - If your drive(s) contains backup copies from another Unitrends appliance, select the device on the Backup Copy Targets tab and click **Enable**. The appliance brings the device online and imports reference information

about those backup copies.

Note: If you receive a message indicating the drive has not been initialized, you must Erase the drive instead (described below).

- If your drive(s) does NOT contain backup copies from another Unitrends appliance, select the device on the Backup Copy Targets tab and click **Erase**. In the Confirm Erase dialog, enter a Media Label (optional) and click **Erase Backup Copies**.

  The appliance erases any existing data and formats the disk(s). If you did not enter a Media Label, the appliance creates one.

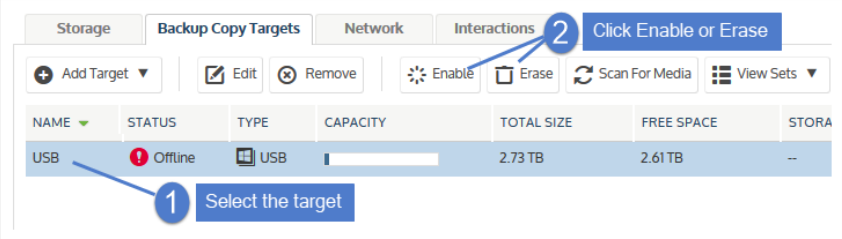

7 The target is ready for use. Create a job to start copying your backups, as described in "To create a backup copy job for an eSATA or USB target" on page 523.

Notes:

- The backup copy job mounts the media, writes the copy, then unmounts the media. On the Backup Copy Targets tab, the device remains *Offline* when the media is not mounted. The device is automatically brought *Online* when a copy job runs.
- If using an eSATA or USB dock, be sure to use these steps when swapping the drive:
  - Power down the dock.
  - Swap the drive.
  - Power on the dock.

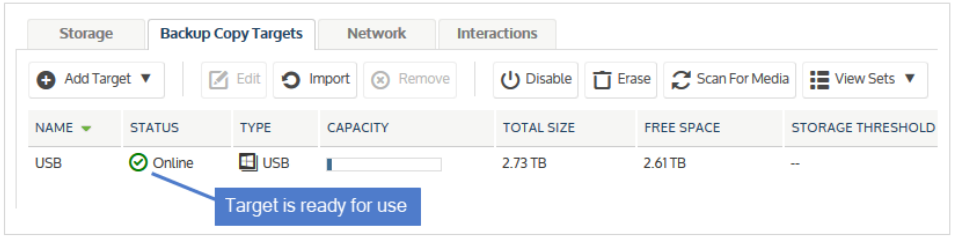

Adding a tape backup copy target

Following is a summary of the steps required to add a tape backup copy target. The links provide detailed instructions for each procedure.

Step 1: Review the requirements in "Preparing to add a tape backup copy target".

Step 2: Connect the tape device as described in "To connect the tape device" on page 236.

Step 3: Add the tape backup copy target to the appliance and configure it as described in "Configuring the tape device on the Unitrends appliance" on page 237.

Preparing to add a tape backup copy target

Before adding the tape backup copy target, review these requirements and considerations.

Unitrends supports copying to tape from Recovery MAX appliances and from select Recovery Series appliance models. Tape is not supported with ION/ION+ and Unitrends Backup virtual appliances. For supported Recovery Series models, see the Recovery Series Appliance Family Data Sheet.

The following requirements must be met before setting up a tape target on a Recovery Series or Recovery MAX appliance:

System	Requirement
Recovery Series or Recovery MAX appliance	The appliance must be licensed with the advanced archiving (ADX) feature. Check for *ADX* in the Feature String under **Configure > Appliances > Edit > License**.
Tape device	• The tape device must be either SCSI, SAS, or Fibre Channel. • The tape device must be configured as described in "Configuring the tape device on the Unitrends appliance" on page 237. • The tape or set of tapes must have adequate space to store the data being copied. If the copy does not fit, the job fails. • If using tape barcodes, your tape device must have a barcode reader and tapes must have valid barcode labels.

Before copying to tape, note these additional considerations:

Tape feature	Description
Multi-tape devices	The following considerations apply to devices with multiple tapes: • All tapes configured for the backup copy job must be rotated as a set as the job writes across all tapes. • All tapes must be available to recover data from the backup copy. • To recover data from tapes that do NOT have barcode labels, the tapes must be loaded into the same slot position as when the backup copy was written.

Tape feature	Description
	• For tapes that do NOT have barcode labels, it is strongly recommended that you develop a labeling system to help you manage the tapes. Prior to pulling a tape without a barcode, note its slot number. When you pull a set of tapes, be sure to *physically label* each tape with the slot number and other identifying information for speedy recovery. • If your tapes have barcodes, there are no special procedures when recovering from a backup copy. The appliance automatically uses the barcode during the recovery process. If you have moved tapes with barcodes to different slots, the appliance reads the barcodes to determine the location of the tapes.
Multiple tape drives	The following considerations apply to devices with multiple tape drives: • Even though you can connect multiple tape drives, the appliance only uses one tape drive for backup copies. • You can connect an autochanger that has multiple tape drives; however, only one drive can be enabled for use by the appliance.
Tape devices with barcodes	On tape devices that support barcodes, the appliance recognizes the barcode as soon as you insert the tape into the library.

To connect the tape device

1. Connect the tape drive or autochanger to the Unitrends appliance using a SAS or LVD SCSI cable. If using a LVD SCSI cable, ensure that a SCSI bus terminator is installed on the tape device according to the vendor's documentation.
2. Once connected, power on the tape device.
3. Once the tape device initializes, log in to the Unitrends appliance and reboot by selecting **Configure > Appliance > Edit > Shutdown/Restart** (on the General tab). This enables the appliance to discover the tape device.

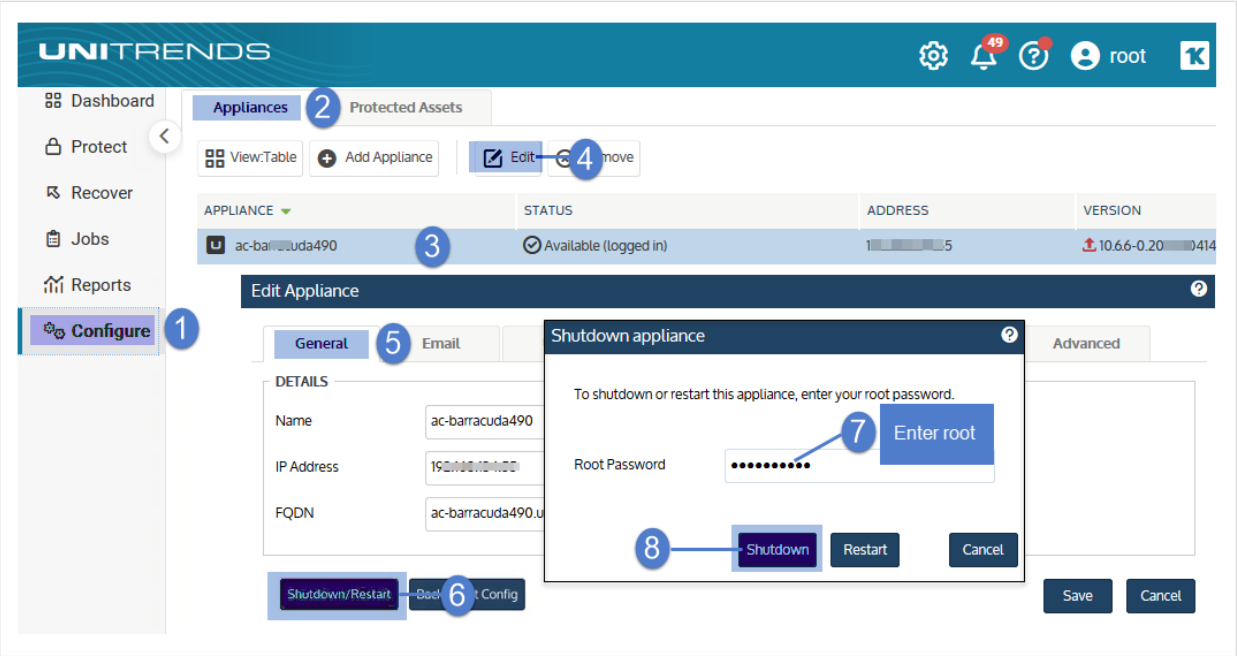

Configuring the tape device on the Unitrends appliance

This section discusses how to configure your tape device after it has been connected to the appliance. Tape drives and autochangers must be configured before you can begin copying backups. After configuring the tape drive and autochanger, you must prepare the tape media. The media can be prepared manually or automatically by the Unitrends appliance.

Tape configuration procedures are given below. Use the procedure that applies to your tape device. (Use only one procedure. The autochanger procedure configures both the changer and the tape drive.)

- "To configure an autochanger" below
- "To configure a tape drive" on page 242 (use this procedure if your tape device does not have an autochanger)

To configure an autochanger

1. Log in to the appliance UI.
2. On the **Configure > Appliances** page, select the source backup appliance.
3. Click the **Backup Copy Targets** tab below.
4. Click **Scan For Media**.

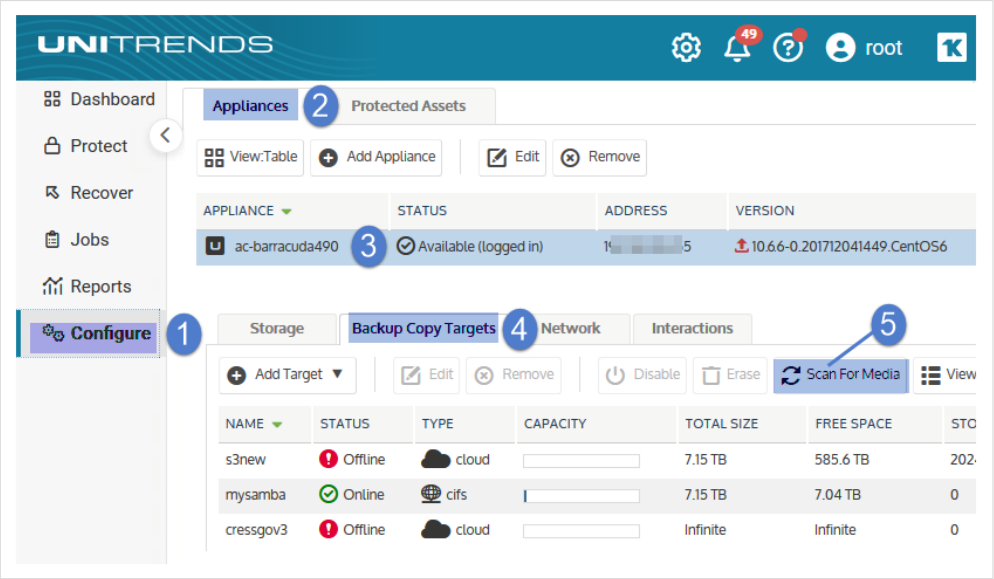

5 You receive a message indicating that you must associate the changer with a tape drive. Select a **Drive** from the list and click **Yes**. (If your changer has only one drive, just click **Yes**.)

> **Note:** Only one drive can be used by the appliance.

The autochanger and tape drive display on the Backup Copy Targets tab in *Offline* status.

- The tape drive displays as Type *tape*.
- The autochanger displays as Type *changer*.

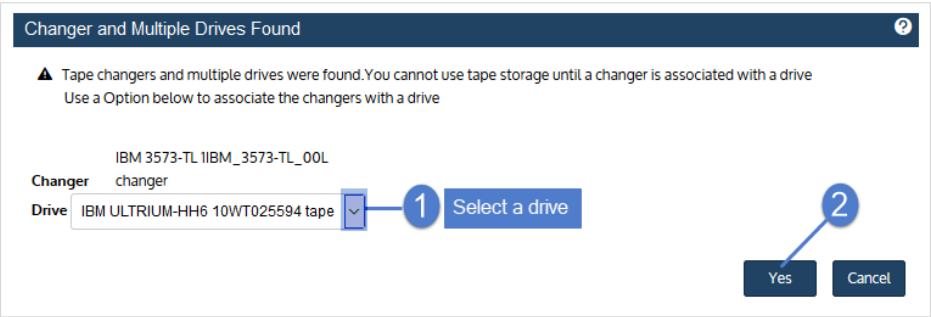

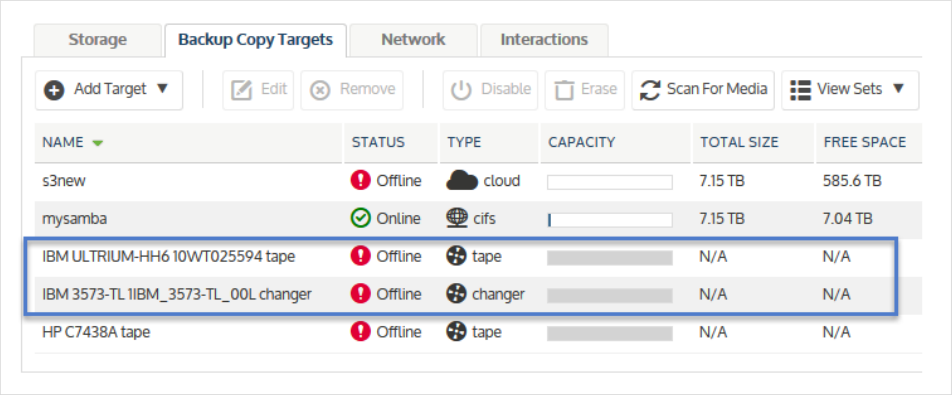

6 Enable the tape drive and autochanger:

- Select the tape drive and click **Enable**. Its status changes to *Online*.

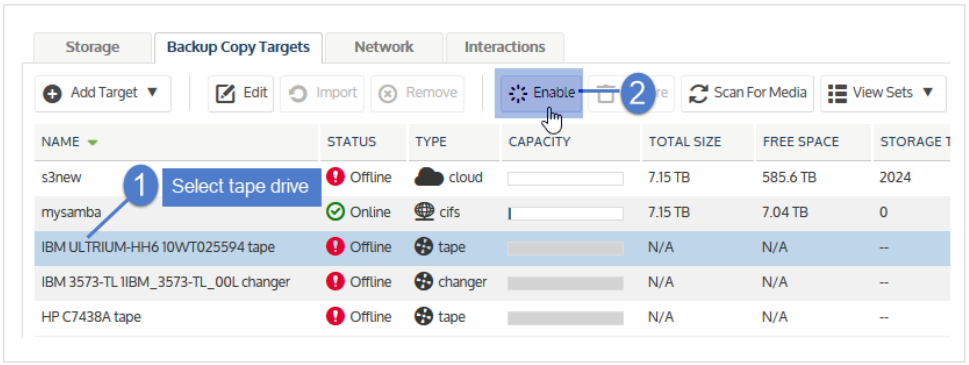

- Select the autochanger and click **Enable**. Its status changes to *Online*.

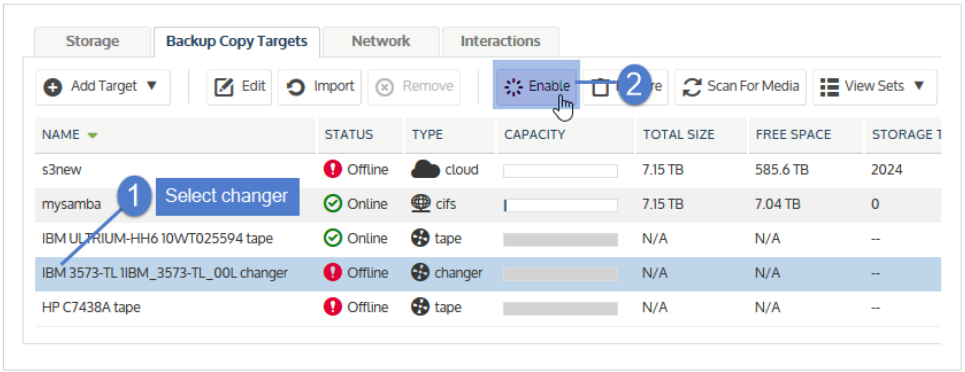

7 If a tape contains Unitrends backup copies, you are asked if you would like to import the data. Selecting **Yes** imports reference information about the backup copies. You must import this data to be able to recover those copies. Do one of the following:

- Check the **Force** option and click **Yes** to import all data, regardless of whether it was written from this Unitrends appliance.
- Leave the Force option unchecked and click **Yes** to import only the copies that were written from this Unitrends appliance.
- Click **No** to continue without importing reference information about these copies.

8 (Optional) Configure the appliance to automatically prepare tapes for first use.

IMPORTANT! Preparing formats the tapes, permanently deleting any existing data. To preserve data on a tape, do not configure this option. If the tape contains Unitrends backup copies, it has already been formatted and can be used for subsequent backup copy jobs. For tapes that do not contain Unitrends backup copies, you can manually format the tapes later in this procedure.

To configure the appliance to automatically prepare tapes:

- Select the tape drive.
- Click **Edit**.
- Check the Use Unlabeled Tapes box.
- Click **Save**.

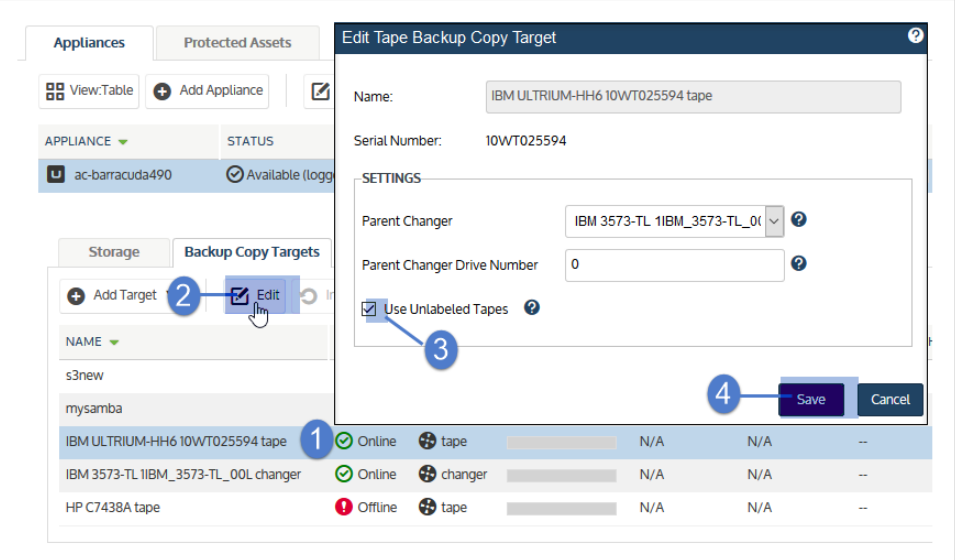

9 (Optional) Modify slots that can be used when writing backups to tape.

Each slot that contains a tape is automatically enabled for use (but you can copy to a subset of these enabled tapes by entering slot numbers when you create the backup copy job.) If you have tapes that cannot be used for backup copies, you can specify the slots that are enabled for use.

To specify the slots that are enabled for use:

- Select the autochanger.
- Click **Edit** and enter the Slots that can be used when writing backups to tape:
  - Enter slot numbers using a comma-separated list and/or ranges. Example: 1,2,3,5-8
  - The slots you enter must contain a tape.
  - To see which slots contain tapes, look at the Tape Library Information area below. A green check indicates the slot contains an available tape. A red X indicates the slot does not contain a tape. To refresh this list, click the arrows in the upper-right corner.
- Click **Save**.

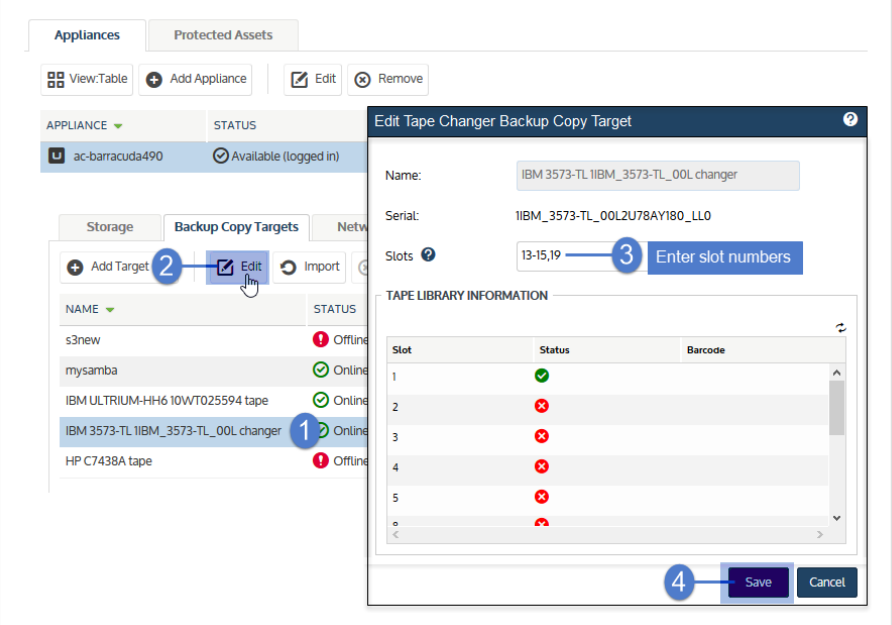

10 Proceed to one of following:

- If you did check the *Use Unlabeled Tapes* box above, you are ready to start copying to tape. Create a job to start copying your backups, as described in "To create a backup copy job for a tape target" on page 528.
- If you did NOT check the *Use Unlabeled Tapes* box above, proceed to the next step to prepare the tapes.

11 Select the autochanger and click **Prepare**.

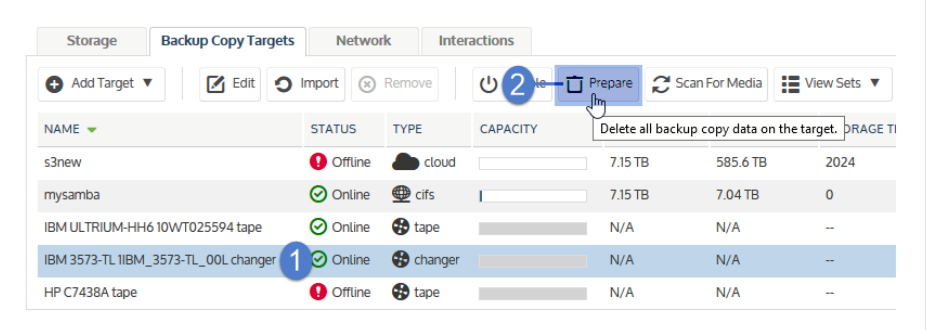

12 In the Confirm Prepare dialog:

- (Optional) Enter a Media Label. The label can contain up to 12 alphanumeric characters or underscores. If you do not enter a label, the appliance generates one for you.
- Enter the Slots whose tapes will be prepared. Enter slot numbers using a comma-separated list and/or ranges. Example: 1,2,3,5-8

13 Click **Prepare** to prepare the tapes in the Slots you entered.

WARNING! Clicking **Prepare** permanently deletes any existing data and formats the tapes.

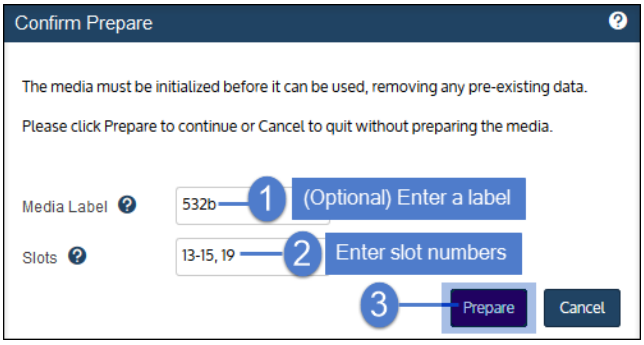

14 Create a job to start copying your backups, as described in "To create a backup copy job for a tape target" on page 528.

Note: The backup copy job mounts the media, writes the copy, then unmounts the media. On the Backup Copy Targets tab, the device remains *Offline* when the media is not mounted. The device is automatically brought *Online* when a copy job runs.

To configure a tape drive

1 Log in to the appliance UI.

2 On the **Configure > Appliances** page, select the source backup appliance.

3 Click the **Backup Copy Targets** tab below.

4 Click **Scan For Media**.

- The appliance discovers the drive and it displays on the Backup Copy Targets tab.
- The tape drive displays as Type *tape* and its status is *Offline*.

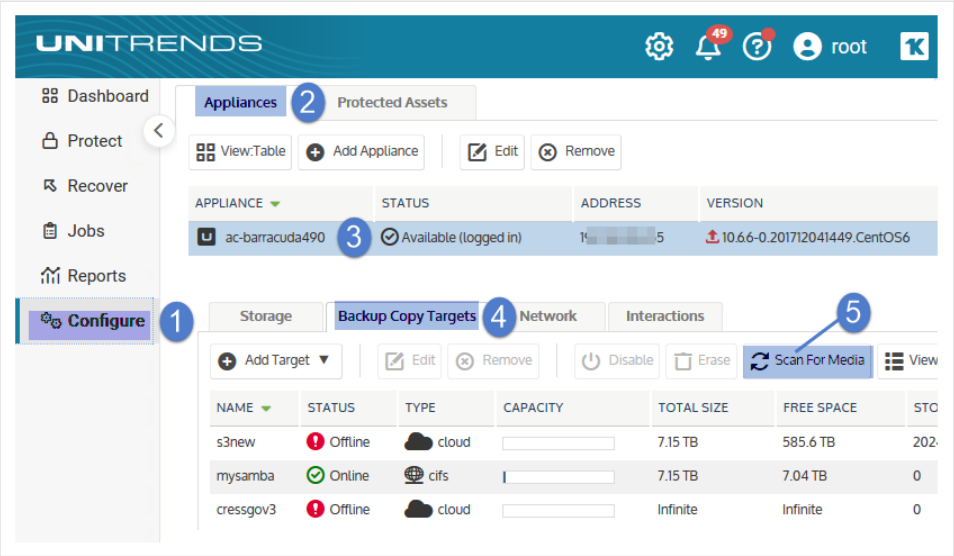

5 Do one of the following:

- If your tape contains backup copies from another Unitrends appliance, select the device and click **Enable**. The appliance brings the device online and imports reference information about those backup copies.

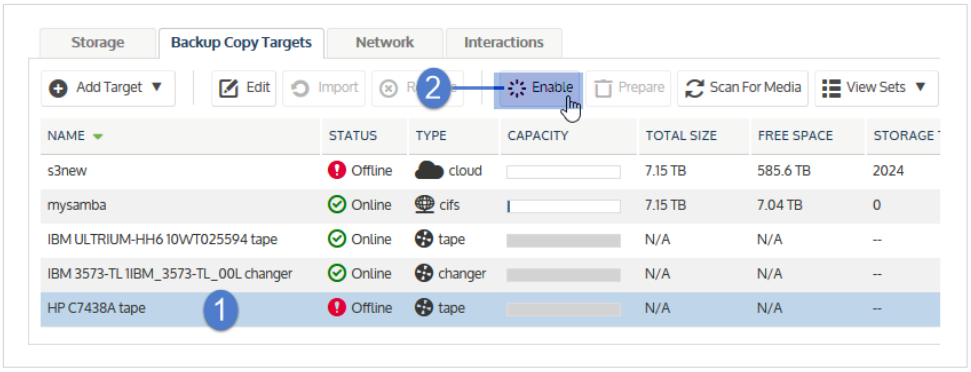

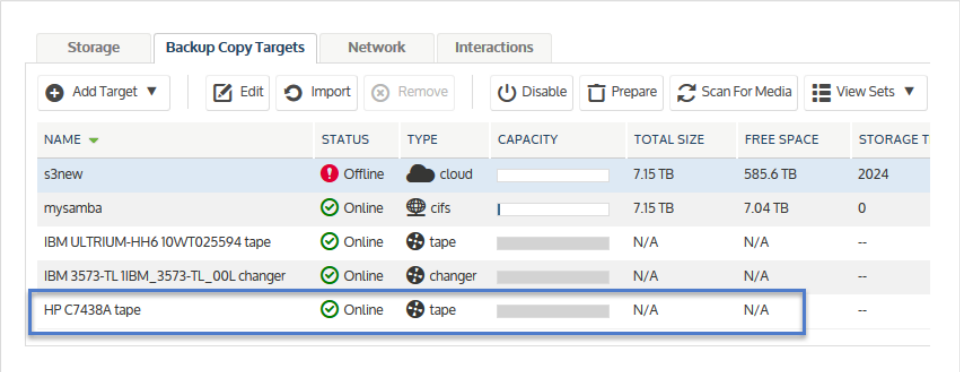

- If your tape does NOT contain backup copies from another Unitrends appliance, select the device and click **Prepare**. Enter a Media Label (optional) and click **Prepare**.

  The appliance erases any existing data and formats the tape. If you did not enter a Media Label, the appliance creates one.

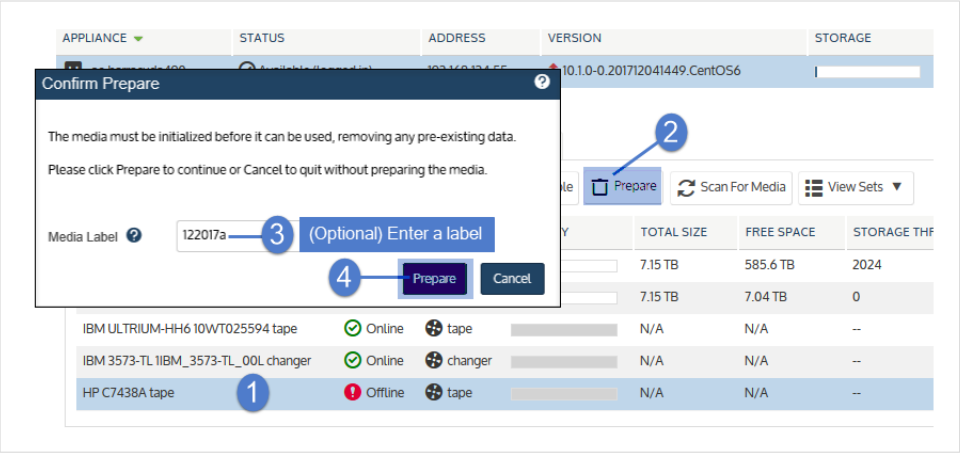

6 Create a job to start copying your backups, as described in "To create a backup copy job for a tape target" on page 528.

Note: The backup copy job mounts the media, writes the copy, then unmounts the media. On the Backup Copy Targets tab, the device remains *Offline* when the media is not mounted. The device is automatically brought *Online* when a copy job runs.

Adding a third-party cloud backup copy target

You can use cloud storage managed by select providers to store copies of your backups. Use this procedure to add cloud storage to your appliance as a backup copy target.

Preparing to add a third-party cloud backup copy target

Cloud storage must meet these requirements to be used as a backup copy target:

- You must have an account with one of the following cloud storage providers:
  - Google Cloud Storage (Standard or Nearline)
  - Amazon S3
  - Amazon S3-IA
  - Rackspace Cloud Files
  - Wasabi (Standard or S3)
  - Azure Blob

  Note: Azure Blob storage is not supported for appliances running on CentOS 6.

  For details about creating an Amazon, Google, or Azure account and purchasing storage, see How do I create a cloud storage account for Unitrends?.
- For bucket or container names, only the following characters are supported: upper and lowercase letters, numbers, dots, and dashes. Buckets with names containing other characters cannot be added to a Unitrends appliance.
- You can use existing buckets or containers that follow the supported naming conventions identified above. However, we recommend that you create unique folders for your Unitrends data.
- The cloud storage bucket or container must not be configured to use any lock features, WORM modes, or other immutability features that prevent overwrite and deletion of data by the Unitrends appliance. (For example, a Google Cloud Storage bucket must not be configured to use the Cloud Storage Bucket Lock feature.)
- Amazon's Reduced Redundancy Storage (RRS) option is not supported.

You must provide the following account information when adding a cloud backup copy target to the appliance:

- Credentials for the storage bucket or container that you are adding to the appliance.

  Note: These are the credentials you use to access the particular bucket or container that you are adding to the appliance, and not the username and password you use to log in to your storage provider account. If you do not know these credentials, contact your storage provider. Unitrends does not have access to this information.
- Name of the cloud storage type: Google Cloud Storage (Standard or Nearline), Amazon S3, Amazon S3-IA, Rackspace Cloud Files, Wasabi (Standard or S3), or Azure Blob.
- Name of the bucket or container that you are adding to the appliance.
- For Amazon S3 and Amazon S3-IA, you must enter an S3 region. For example, *us-east-1*. See this Amazon AWS article for a list of valid regions: Regions and Availability Zones.
- For Wasabi S3 , you must enter an S3 region. For example, *s3.us-east-1.wasabisys.com*. See this article article for a list of valid regions: What are the service URLs for Wasabi's different regions?

Additional considerations:

- Accounting and billing management for your cloud storage occur between you and the storage provider. You cannot manage your cloud storage account from the appliance's UI, and Unitrends cannot answer questions about this account. You must contact your provider with any questions you have about your cloud storage account.
- It is extremely important that you understand the amount of data you are copying and the related charges from your cloud storage provider. To manage the amount of space you are using in cloud storage, you should specify a storage threshold.
- Sending backup copies to the cloud does not require any special network configurations. Just add a backup copy cloud target to the appliance and create backup copy jobs.
- The speed at which backup copies can be sent to the cloud depends on a number of factors, including memory and network bandwidth. We recommend that you test a small backup copy to determine the speed prior to sending larger backup copies to the cloud.
- When recovering backup copies from the cloud, the cloud backup copy target can be attached to any Unitrends appliance. You do not have to recover to the original appliance.
- If you add a backup copy target and the appliance recognizes backups on it, those backups are added to the Backup Catalog and are available for import.

Managing the amount of data copied to a third-party cloud target

It is extremely important that you monitor the amount of backup data that your Unitrends appliance is copying to the cloud because cloud storage providers charge based on the amount of storage you use. When adding cloud backup copy target to the appliance, you can specify a storage threshold equal to the maximum amount of space the appliance can use to store backup copies in the cloud storage bucket or container.

You can monitor the amount of data in your cloud storage from the **Configure > Appliance > Backup Copy Targets** page:

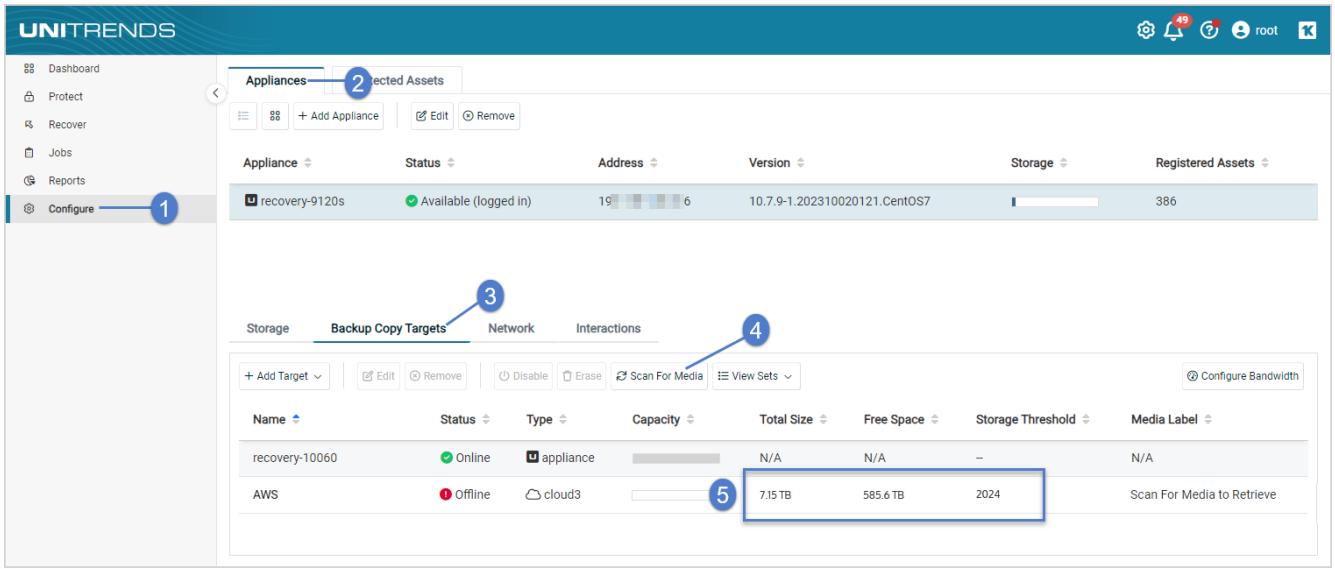

About the storage threshold for cloud backup copy targets

The storage threshold setting for cloud backup copy targets is intended to aid in managing the total amount of data you are copying to the cloud. The user-defined storage threshold functions as a maximum amount of data the appliance can write to the cloud bucket or container.

Note that there are instances when backup copy jobs can use slightly more storage space than the storage threshold you specified. When you initiate a backup copy job, the appliance estimates the amount of space needed for the job, and certain factors can cause it to underestimate. To avoid unexpected charges from your cloud storage provider, it is highly recommended that you develop a policy for managing the amount of backup data that you copy to the cloud in addition to setting a storage threshold.

The storage threshold can be increased or decreased at any time, and changes are applied to all subsequent jobs copying backups to the bucket or container. If you increase the storage threshold, your storage provider will bill you for the additional storage space you are using.

Note that decreasing the threshold to a value that is less than the amount of space currently used to store backup copies can result in the deletion of backup copies. The next time the job runs, the appliance recognizes that the amount of data on the backup copy target is greater than the storage threshold and invokes your selected behavior: either deleting older backups to free space or failing the job. If you select to delete older backup data to free space for the job, backup copies exceeding the new storage threshold are purged the next time the job runs. Backup groups are still recognized in backup copies, and backups are deleted as a group beginning with the oldest backup copies. If you chose to fail the backup copy job and send an alert, your new backup copies are not written to the cloud and an alert notifies you of this failure.

To add a third-party cloud backup copy target

Use this procedure to add a Google, Amazon, AWS, Rackspace, Wasabi, or Azure Blob cloud target.

1. On the **Configure > Appliances** page, select your appliance.
2. Select the **Backup Copy Targets** tab.
3. Click **Add Target > Cloud**.

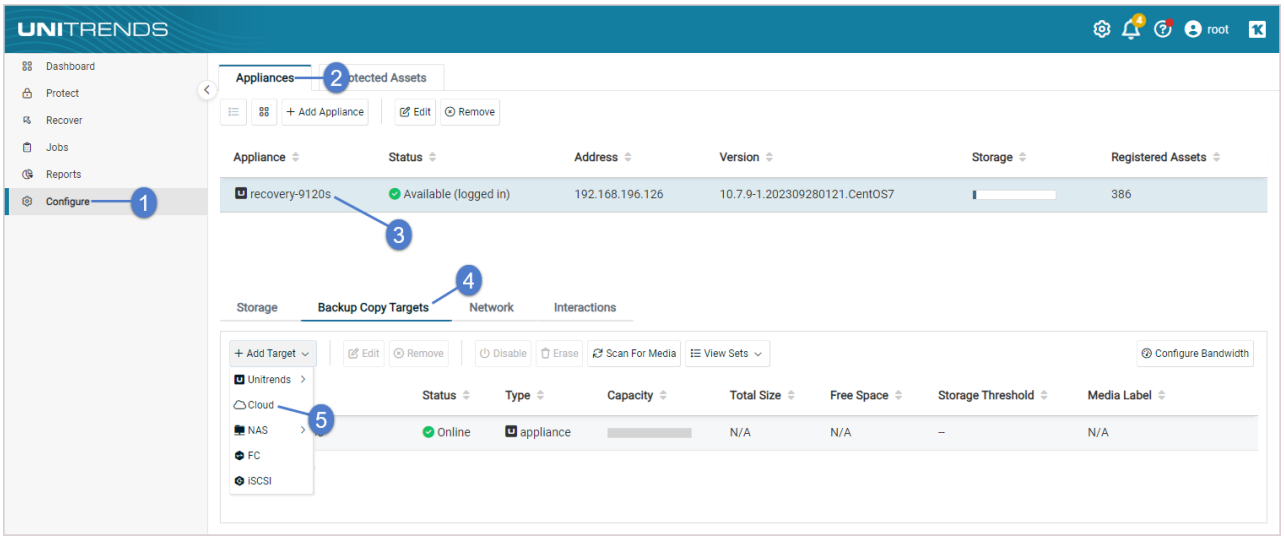

4 Enter a **Name** for the cloud storage.

5 Select the storage type in the **Cloud Storage** drop-down.

> **Note:** If you do not see the *Azure Blob* storage type in the list, it is not supported for your appliance model.

6 Enter the bucket or container name in the **Storage Path** field.

To create a sub-folder within the bucket or container, add a forward slash followed by the name of the new folder and another forward slash. For example: If your bucket is named *mybackup* and you want to create a new sub-folder within that bucket called *myfolder*, enter: *mybackups/myfolder/*

7 (S3 only) For Amazon S3, Amazon S3-IA, or Wasabi S3 you must enter an S3 region. For example, *us-east-1*. (Leave the S3 Region field empty for Wasabi Standard.)

- See this Amazon AWS article for a list of valid Amazon S3 regions: Regions and Availability Zones.
- See this article for a list of valid Wasabi S3 regions: What are the service URLs for Wasabi's different regions?

8 Enter the required credentials. These vary by storage provider:

> **Note:** Be sure to enter the credentials you use to access the particular bucket or container that you are adding to the appliance. These credentials are not the same as the username and password you use to access your storage provider account.

Provider	Required Credentials
Google Cloud Storage (Nearline or Standard)	Access Key Secret
Amazon (S3, S3-IA, or Wasabi)	Access Key ID Secret Access Key
Rackspace	Username API Key
Azure Blob	Storage Account Shared Key

9 (Optional) We recommend enabling and setting a storage threshold. This functions as the maximum amount of data the appliance can copy to the cloud. For more information, see "Managing the amount of data copied to a third-party cloud target" on page 246.

10 Click **Save**.

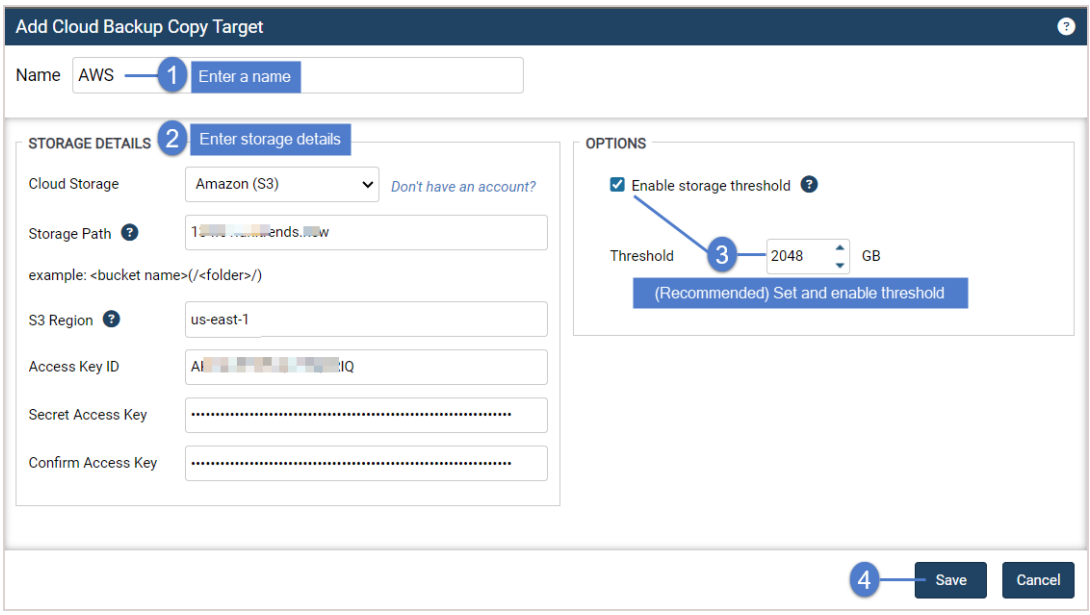

11 The cloud target is added to the appliance.

Cloud storage is mounted only while a backup copy job is running. At other times, cloud storage is unmounted. This is the recommended approach to prevent issues that may occur if the connection to cloud storage is interrupted. If desired, you can configure a persistent connection as described in "To disable automatic cloud mount/unmount" below.

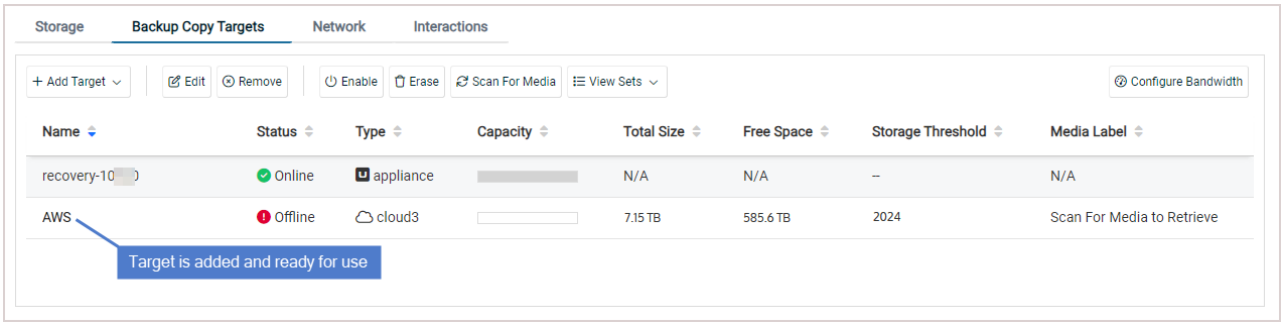

12 Create a job to start copying your backups, as described in "To create a backup copy job for a third-party cloud target" on page 503.

To disable automatic cloud mount/unmount

Various external network events can interrupt the connection to cloud storage. To prevent issues that may occur if this connection is interrupted, cloud storage remains unmounted until a backup copy job runs. The job mounts the cloud storage, copies data to the target, and unmounts the storage. If desired, you can configure the appliance to maintain a persistent connection to the cloud storage target, as described here:

1 On the **Configure > Appliances** page, select the appliance and click **Edit**.

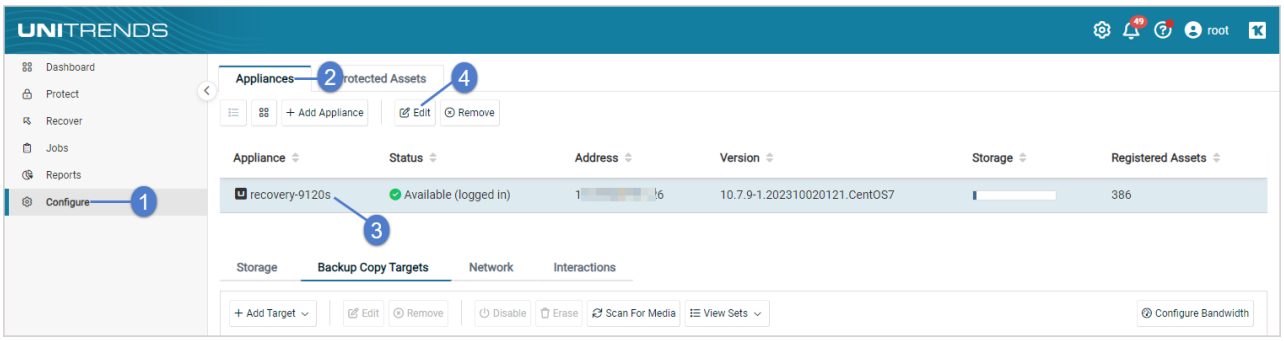

2 Click **Advanced** and select **General Configuration**.

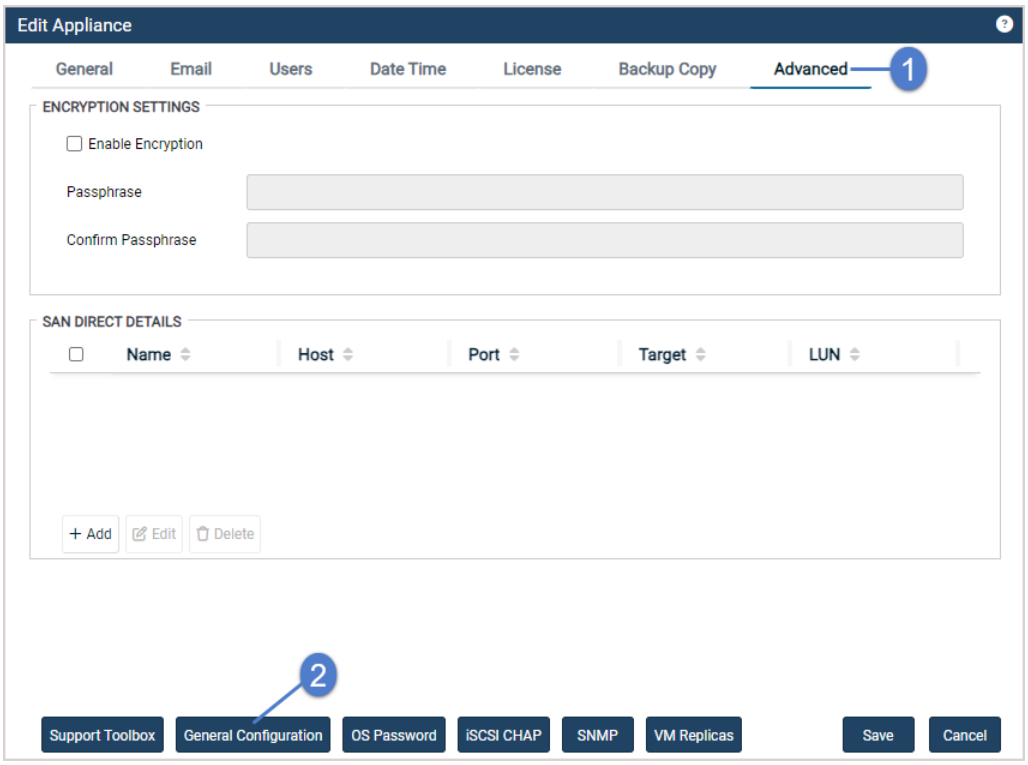

3 Change the CloudHook *UnmountAfter* setting to **0** and click **Save**.

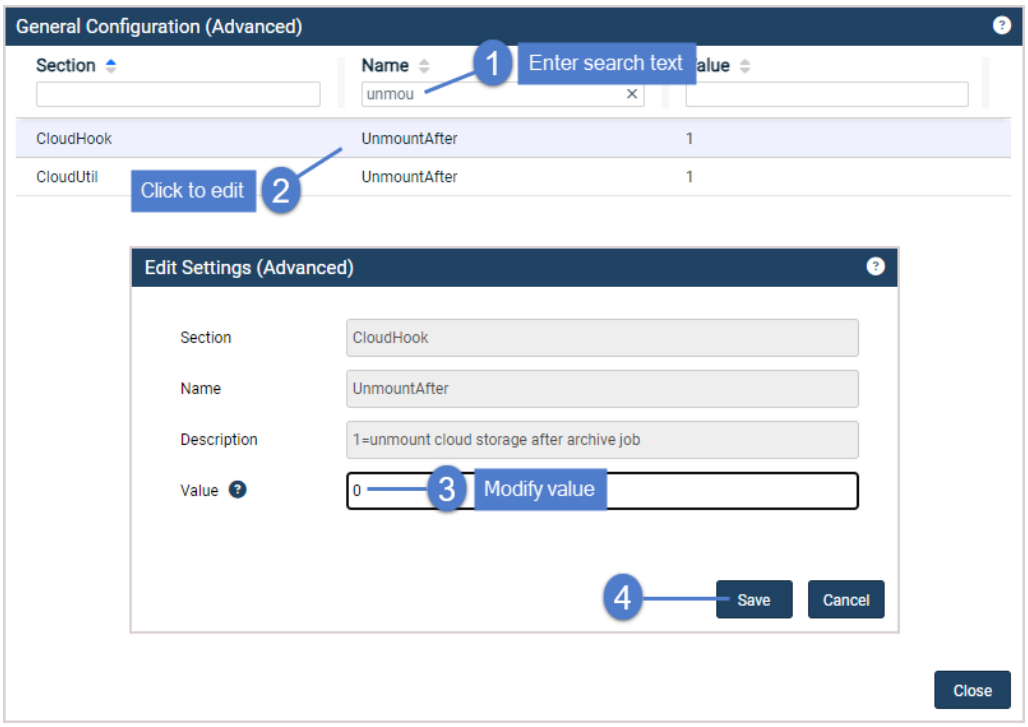

4 Click **Close** to exit.

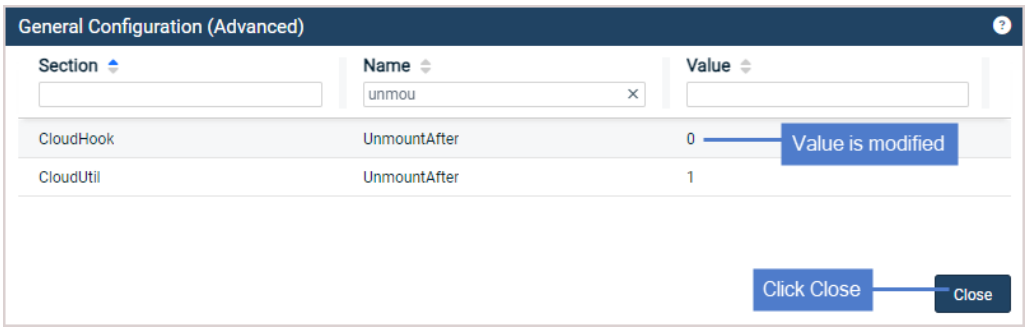

Adding an attached disk backup copy target

Use this procedure to copy backups to attached disk storage. This target type is supported only for Unitrends Backup appliances deployed on Hyper-V, VMware, Nutanix AHV, or XenServer.

Preparing to add an attached disk backup copy target

Before adding the backup copy target, add an attached disk to the Unitrends Backup VM or to external storage attached to the Unitrends Backup hypervisor. To provide redundancy, Unitrends recommends that backup copy storage be on a different datastore or different type of storage than the backup storage. Choose from the following methods to add a disk:

- For Unitrends Backup on VMware, add a VMDK disk to the Unitrends Backup VM by using the ESXi host. See one of the following VMware documents for instructions:
  - vSphere 5.1: Create a Virtual Disk in vSphere Client 5.1
  - vSphere 5.5: Create a Virtual Disk in vSphere Client 5.5
  - vSphere 6.0: Create a Virtual Disk in vSphere Client 6.0
  - vSphere 6.5: Add a Hard Disk to a Virtual Machine
- For Unitrends Backup on Hyper-V, add a virtual disk to the Unitrends Backup VM by using the Hyper-V host. Unitrends recommends that you use a VHD(X) disk and that you add the disk to the SCSI controller. See the following Microsoft documents for instructions:
  - To create a virtual hard disk
  - To add a hard disk to a virtual machine
- For Unitrends Backup on Nutanix AHV, add a virtual disk to the Unitrends Backup VM by using the Prism Web Console. Edit the appliance VM to add a vDisk to the ISCSI controller (by selecting **ISCSI** in the Bus Type list). For details, see Managing a VM (AHV) in the Prism Web Console Guide.
- For Unitrends Backup on XenServer, add a VHD disk to the Unitrends Backup VM by using the XenServer host. Unitrends recommends that you add the disk to the SCSI controller.
- Add a LUN to an external SAN and expose it to the Unitrends Backup VM. Then go to the host and add a new virtual disk to the VM using storage on the LUN you added.
- Connect a NAS share to the Unitrends Backup VM using the NFS or CIFS protocol. Then go to the host and add a new virtual disk to the VM using storage on the share you added.

To add the attached disk backup copy target

1. Attach a disk to your Unitrends Backup VM as described above in "Preparing to add an attached disk backup copy target" on page 251.
2. Log in to the Unitrends Backup UI.
3. On the **Configure > Appliances** page, select your appliance.
4. Select the **Backup Copy Targets** tab.
5. Click **Add Target > Attached**.

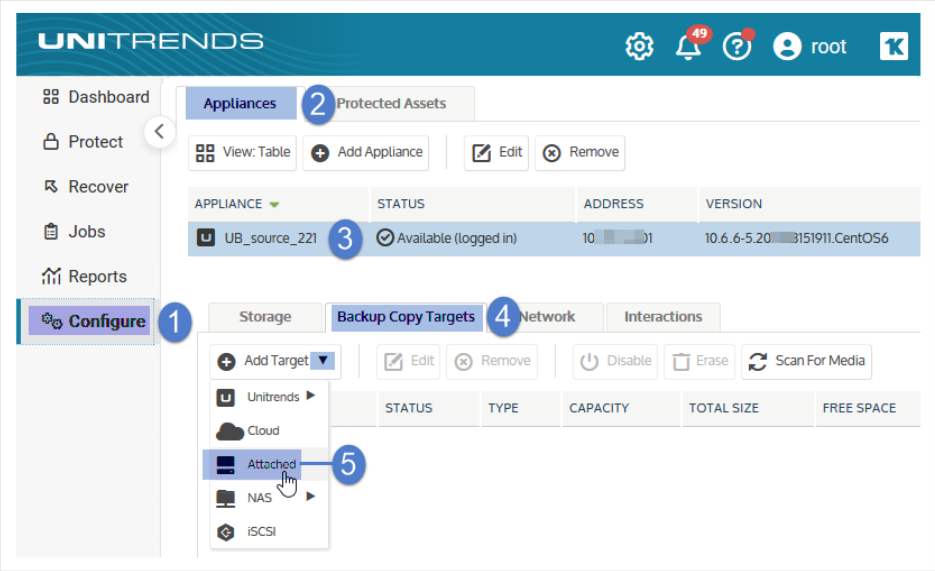

6 Enter a **Name** for the storage.

7 In the Manage Attached Disks area, select the disk from the list of available attached disks and click **Add**.

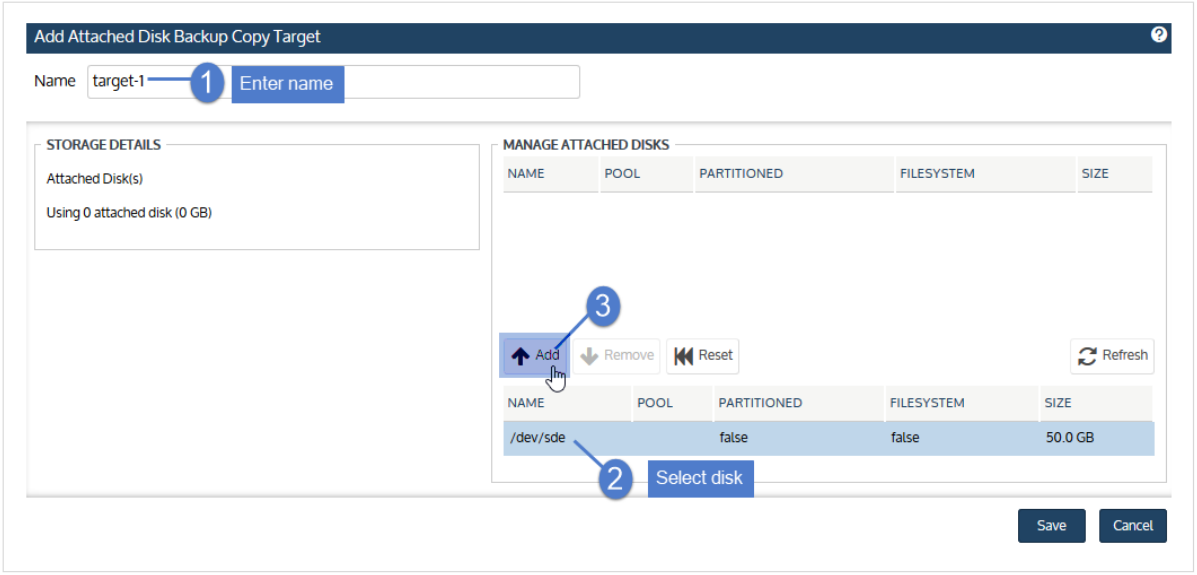

8 Click **Save**.

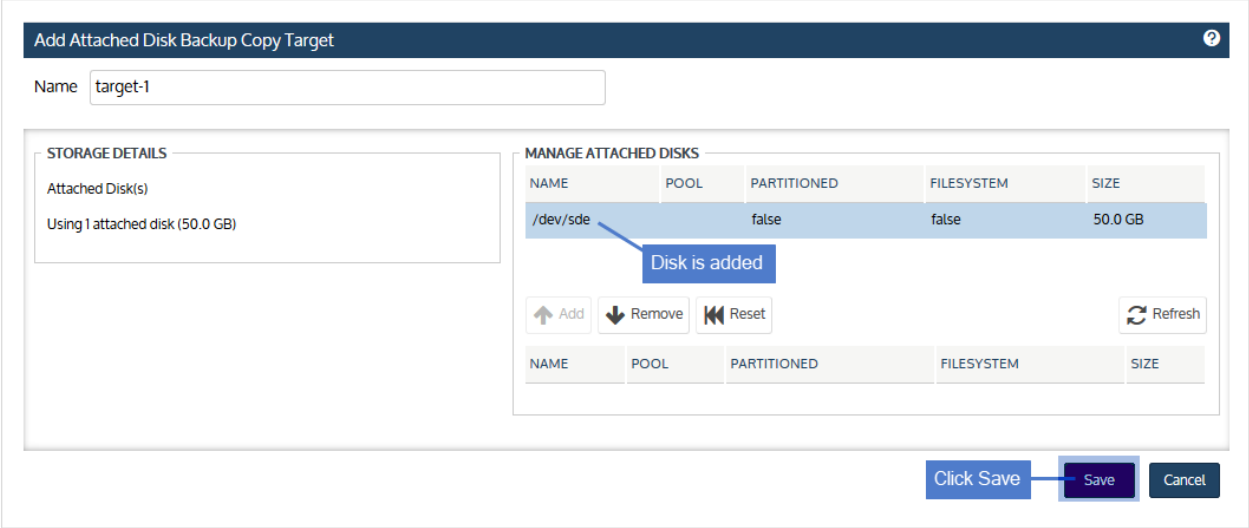

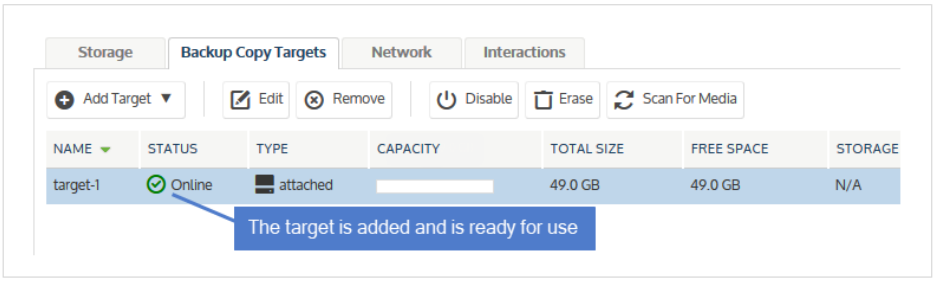

9 Create a job to start copying your backups, as described in "To create a backup copy job for an attached disk target" on page 508.

Adding a NAS backup copy target

Use these procedures to store backup copies on a NAS share:

- "To add a NAS backup copy target that uses the CIFS protocol" below
- "To add a NAS backup copy target that uses the NFS protocol" on page 256

To add a NAS backup copy target that uses the CIFS protocol

1 Allocate a share on the NAS.

2 Log in to the appliance UI.

3 On the **Configure > Appliances** page, select your appliance.

4 Select the **Backup Copy Targets** tab.

5 Click **Add Target > NAS > CIFS**.

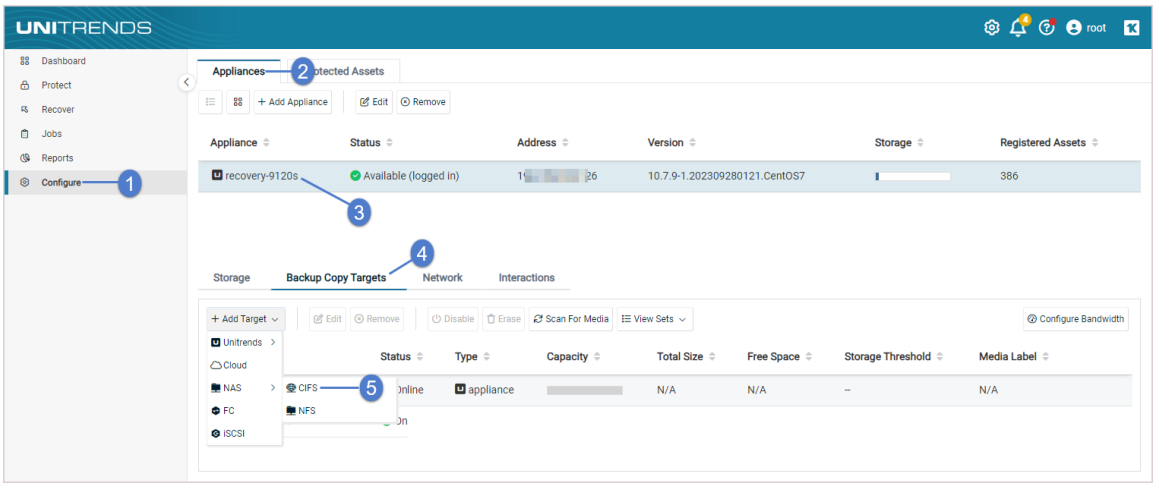

6 Enter the required CIFS share information and click **Save**. For a description of each field, see the table below.

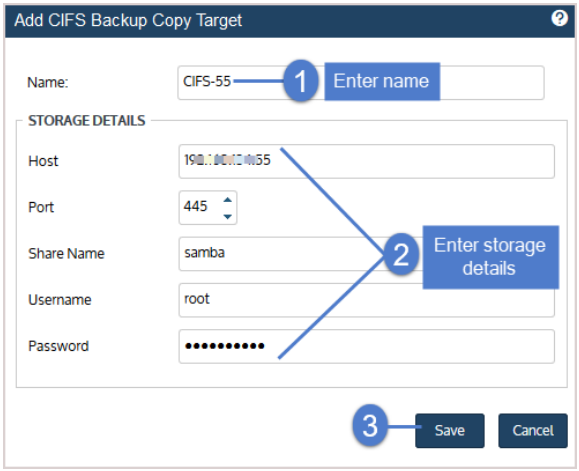

Field	Description
Name	Name of the storage. Cannot contain spaces.
Host	IP address or hostname of the NAS share.
Port	Contains the default CIFS port. To use a custom port, enter that port number.
Share Name	Enter the full directory pathname of the NAS share.
Username (optional)	If the share is configured for authentication, enter the domain username as user@domain.com.

Field	Description
Password (optional)	If the share is configured for authentication, enter the password.

7 The appliance adds the target and imports any existing cold backup copies. Click **OK**.

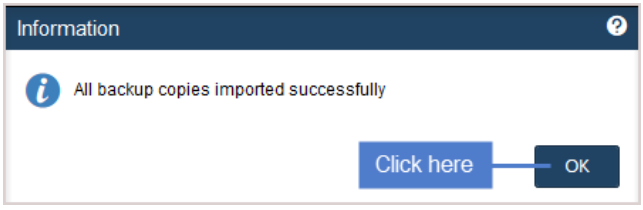

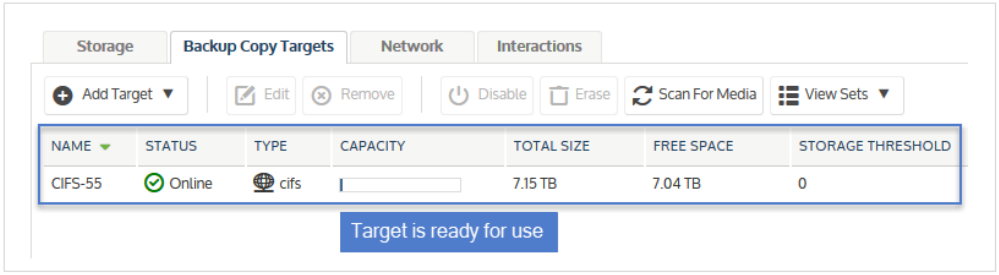

The target is ready for use. Create a job to start copying your backups, as described in "To create a backup copy job for a NAS target" on page 513.

To add a NAS backup copy target that uses the NFS protocol

1 Allocate a share on the NAS.

2 Log in to the appliance UI.

3 On the **Configure > Appliances** page, select your appliance.

4 Select the **Backup Copy Targets** tab.

5 Click **Add Target > NAS > NFS**.

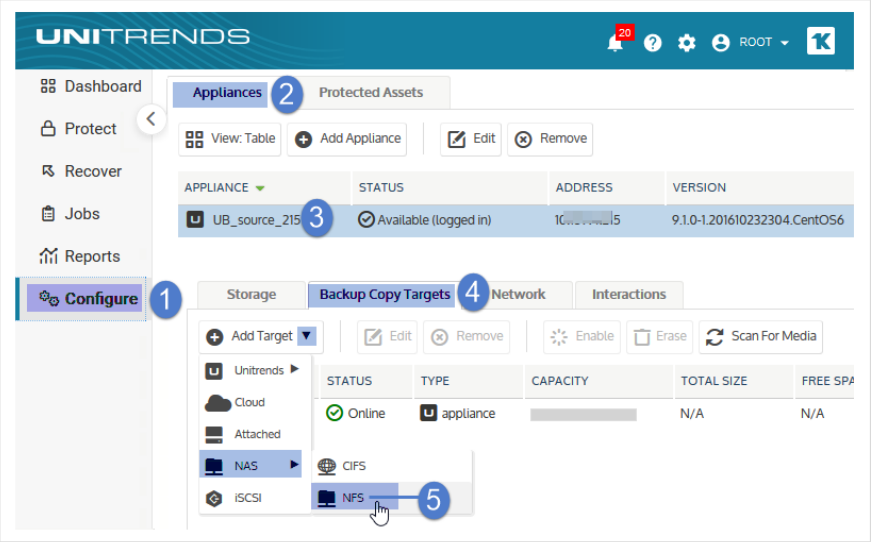

6 Enter the required NFS share information and click **Save**. For a description of each field, see the table below.

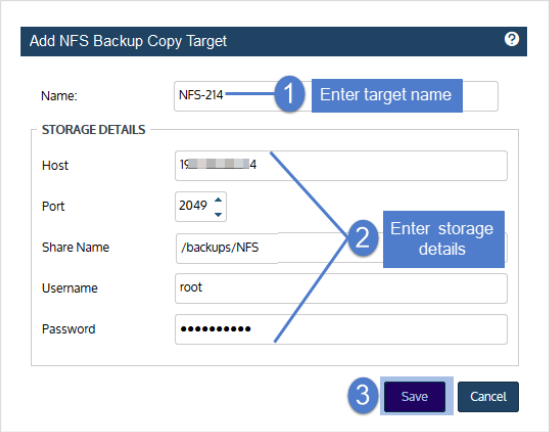

Field	Description
Name	Name of the storage. Cannot contain spaces.
Host	IP address or hostname of the NAS share.
Port	Contains the default NFS port. To use a custom port, enter that port number.
Share Name	Enter the full directory pathname of the NAS share.
Username	If the share is configured for authentication, enter the domain username as

Field	Description
(optional)	user@domain.com.
Password (optional)	If the share is configured for authentication, enter the password.

7 The appliance adds the target and imports any existing cold backup copies. Click **OK**.

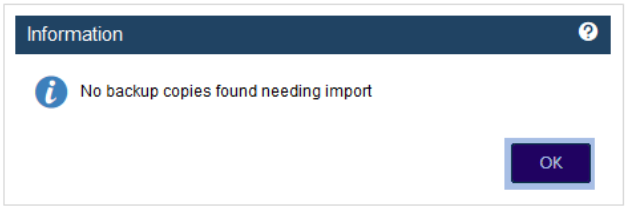

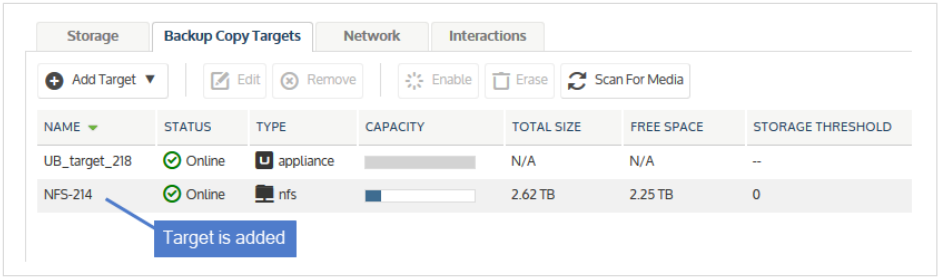

The target is ready for use. Create a job to start copying your backups, as described in "To create a backup copy job for a NAS target" on page 513.

Adding a SAN backup copy target

Use these procedures to store backup copies on a SAN LUN:

- "To add a SAN backup copy target that uses the iSCSI protocol" below
- "To add a SAN backup copy target that uses Fibre Channel" on page 259

To add a SAN backup copy target that uses the iSCSI protocol

1 Allocate a LUN on the SAN.

2 Log in to the appliance UI.

3 On the **Configure > Appliances** page, select your appliance.

4 Select the **Backup Copy Targets** tab.

5 Click **Add Target > iSCSI**.

6 Enter a unique **Name** for the storage device. This name cannot contain spaces.

7 Enter the IP address of the SAN storage array in the **Host** field.

8 The default port used for iSCSI communication is 3260. If the LUN is configured to use a different port, enter it in the **Port** field.

9 Click **Scan for targets** to retrieve a list of targets on the remote storage array, then choose one from the list.

> **Notes:** If you do not see the LUN in the list, go to your SAN manager and check the following:
>
> - Verify that you can see the Unitrends appliance in your SAN manager.
> - Verify that you have a LUN assigned to the Unitrends appliance with the correct permissions.
> - Check with your Storage Administrator for more information.

10 Click **Scan for LUNs** and select one from the list.

> **Note:** If you receive an error indicating CHAP authentication has failed, CHAP has been configured on the target and either CHAP has not been enabled on the Unitrends appliance, or the Unitrends CHAP credentials do not match those of the target. To configure the appliance to use CHAP, see "To configure iSCSI CHAP authentication" on page 182.

11 Click **Save**.

12 The appliance adds the target and checks for any existing cold backup copies. Do one of the following:

- If no copies were found, click **OK**. The target is ready for use. Create a job to start copying your backups, as described in "To create a backup copy job for a SAN target" on page 518.
- If existing copies were found, you are asked if you would like to import the data. Selecting **Yes** imports reference information about the backup copies. You must import this data to be able to recover the copies. Do one of the following:
  - Check the **Force** option and click **Yes** to import all data, regardless of whether it was written from this Unitrends appliance.
  - Leave the Force option unchecked and click **Yes** to import only the copies that were written from this Unitrends appliance.
  - Click **No** to continue without importing reference information about these copies.

Once copies have been imported, the target is ready for use. Create a job to start copying your backups, as described in "To create a backup copy job for a SAN target" on page 518.

To add a SAN backup copy target that uses Fibre Channel

1 Allocate a LUN on the SAN.

2 Log in to the appliance UI.

3 On the **Configure > Appliances** page, select your appliance.

4 Select the **Backup Copy Targets** tab.

5 Click **Add Target > FC**.

6 Enter a unique **Name** for the storage device. This name cannot contain spaces.

7 Enter the IP address of the SAN storage array in the **Host** field.

8 Click **Scan for targets** to retrieve a list of targets on the remote storage array, then select one in the list.

9 Click **Scan for LUNs** and select one in the list.

Notes: If you do not see the LUN in the list, go to your SAN manager and check the following:

- Verify that you can see the Unitrends appliance in your SAN manager.
- Verify that you have a LUN assigned to the Unitrends appliance with the correct permissions.
- You may need to reboot the Unitrends appliance to enable it to discover the storage device.
- Check with your Storage Administrator for more information.

10 Click **Save**.

11 The appliance adds the target and checks for any existing cold backup copies. Do one of the following:

- If no copies were found, click **OK**. The target is ready for use. Create a job to start copying your backups, as described in "To create a backup copy job for a SAN target" on page 518.
- If existing copies were found, you are asked if you would like to import the data. Selecting **Yes** imports reference information about the backup copies. You must import this data to be able to recover the copies. Do one of the following:
  - Check the **Force** option and click **Yes** to import all data, regardless of whether it was written from this Unitrends appliance.
  - Leave the Force option unchecked and click **Yes** to import only the copies that were written from this Unitrends appliance.
  - Click **No** to continue without importing reference information about these copies.

Once copies have been imported, the target is ready for use. Create a job to start copying your backups, as described in "To create a backup copy job for a SAN target" on page 518.

Managing backup copy targets

Once a backup copy target has been added to the appliance, you can monitor its status, such as amount of space used, from the Storage tile on the dashboard or from the **Configure > Appliance > Backup Copy Targets** page.

To view or modify a backup copy target, use the following procedures:

Note: Options vary by backup copy target type. Options not supported for your target type are disabled in the UI.

- "To view or edit a backup copy target" on page 261
- "To reduce the amount of space used on a third-party cloud backup copy target" on page 262

To view or edit a backup copy target

1 On the **Configure > Appliances** page, select your appliance.

2 Click the **Backup Copy Targets** tab, and select a backup copy target.

3 Click **Edit**.

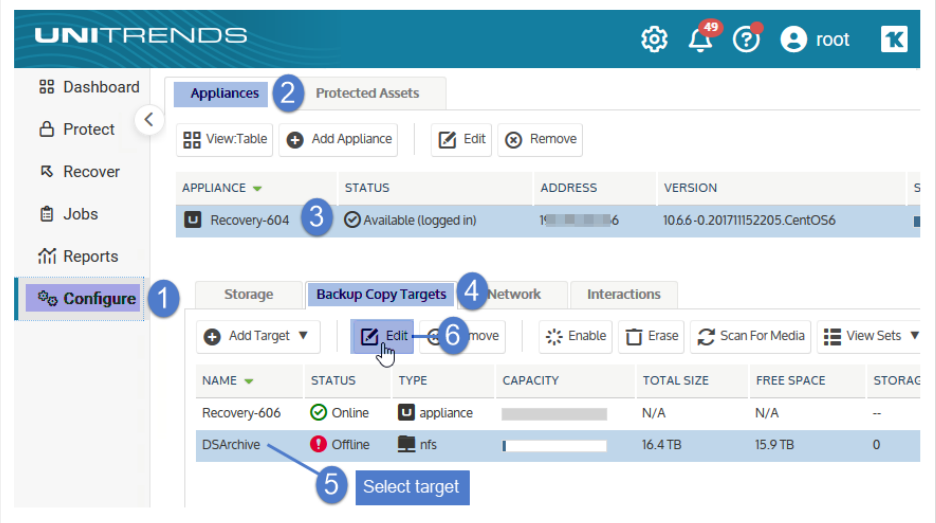

4 (Optional) Modify information.

Note: For Google, Amazon, AWS, and Rackspace cloud targets, you can increase or decrease the storage threshold from this dialog. These changes are applied to all subsequent backup copy jobs that write to the bucket. For more information about the storage threshold, see "Managing the amount of data copied to a third-party cloud target" on page 246.

5 Click **Save** to retain changes or **Cancel** to exit without saving.

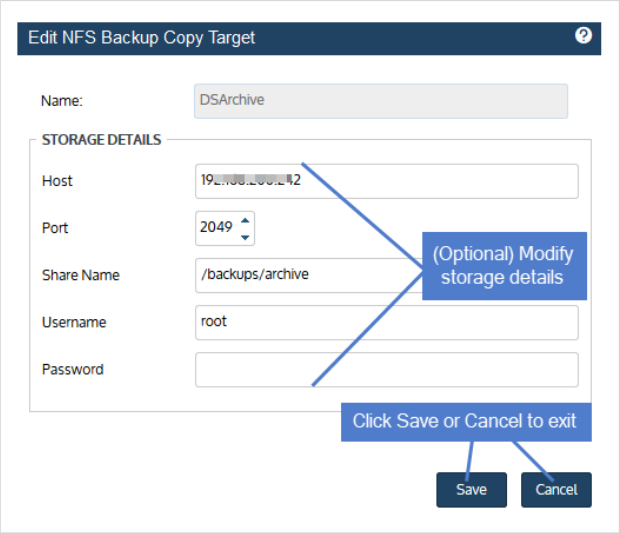

To reduce the amount of space used on a third-party cloud backup copy target

Use this procedure to remove older backup copies and to reduce the amount of data that can be written to a Google, Amazon, AWS, or Rackspace cloud backup copy target. The new threshold does not result in immediate data reduction on the backup copy target. Instead, the next time the backup copy job runs, the appliance purges older backup copies to meet the new storage threshold and to make space for the new backup copy.

1 On the **Configure > Appliances** page, select your appliance.

2 Click the **Backup Copy Targets** tab, select the cloud target on which you want to reduce the space used, and click **Edit**.

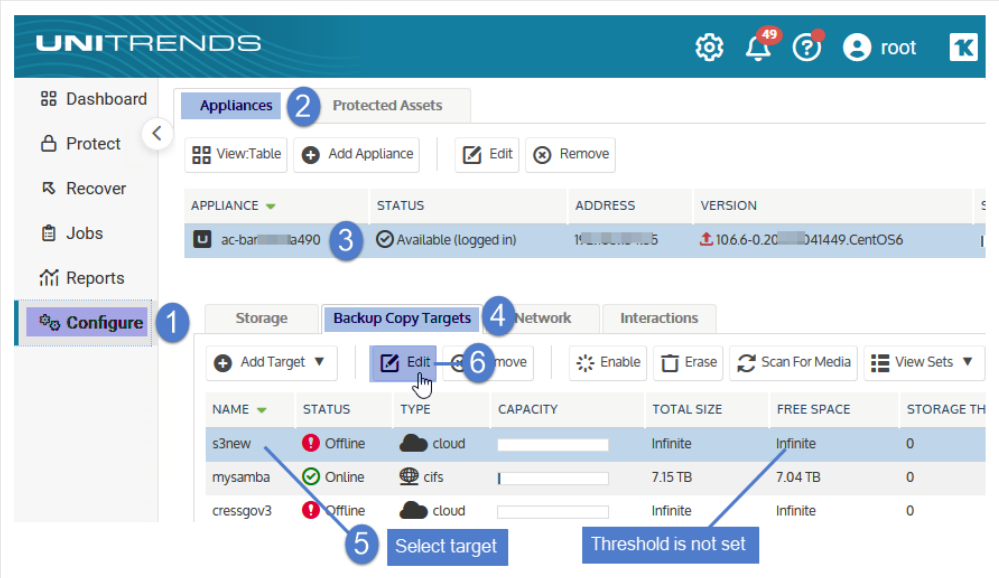

3 Adjust the storage threshold as desired and click **Save**.

- If you do not have a threshold, check the **Enable storage threshold** box, and adjust the **Threshold** value as desired. The threshold functions as a maximum amount of data the appliance can store on the backup copy target.
- If you need to adjust your existing threshold, decrease the **Threshold** value to the new maximum amount of data you want the appliance to store on the backup copy target.

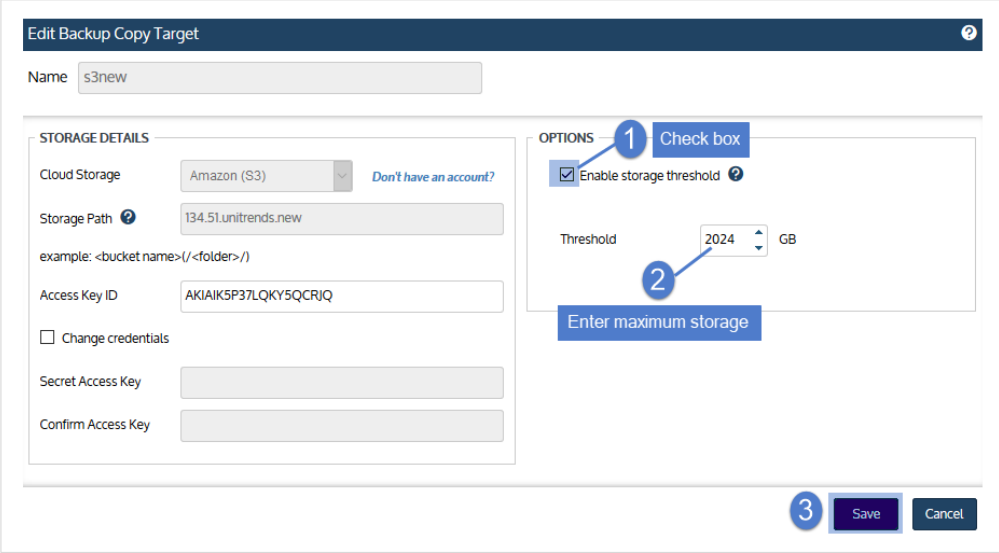

4 On the **Jobs > Job Manager** page, select the job that writes to the backup copy target and click **Edit**.

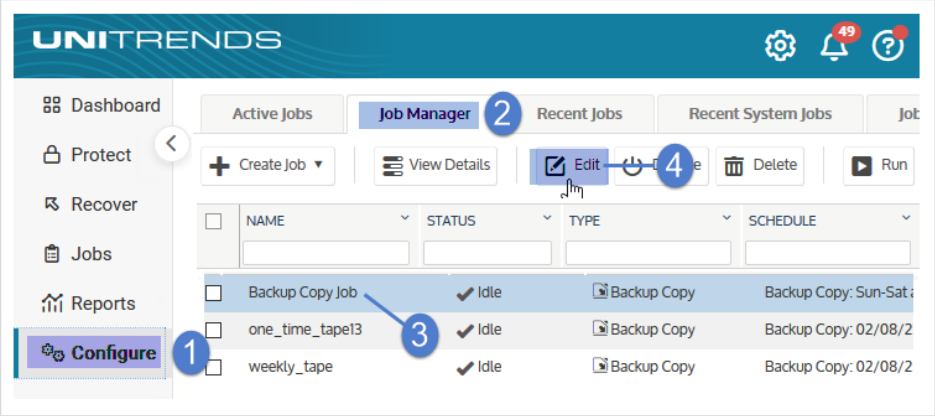

5 Click **Next**.

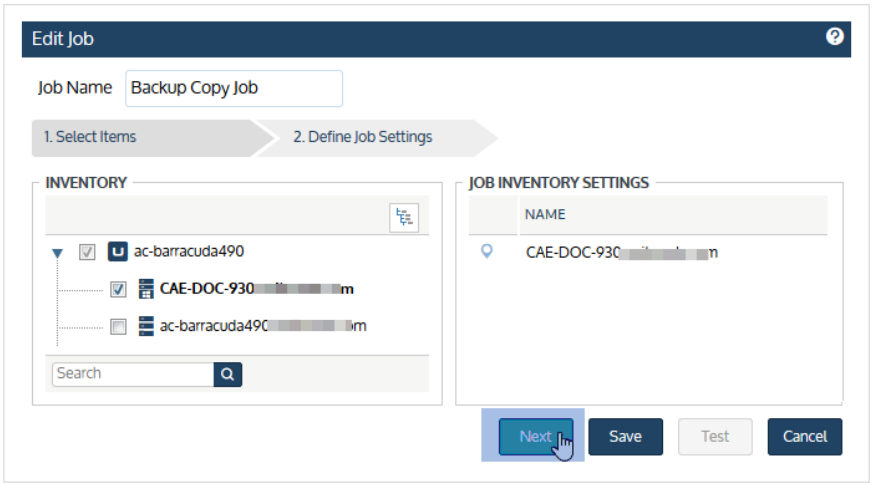

6 Select **Delete older backup data to free space** if the storage threshold is reached and click **Save**.

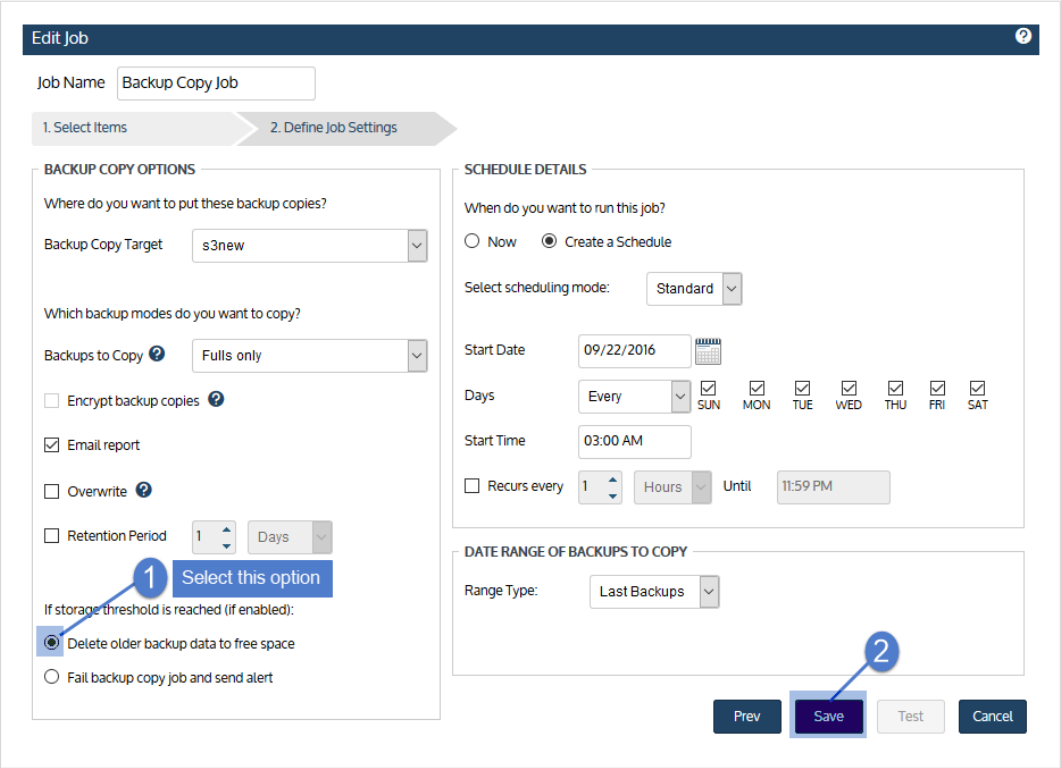

To tune connection options for a Unitrends Cloud or Unitrends appliance target

For Unitrends Cloud and Unitrends appliance hot backup copy targets, you can adjust the following settings for optimal performance in your environment:

Setting	Description	Procedure
Max concurrent backup copies	Determines how many backups can be copied concurrently. While the default setting of two is adequate for most deployments, you may wish to increase the concurrency when you have enough WAN bandwidth to support more concurrent copies.	See "To configure connection options" on page 231.
Queue scheme	Determines the order in which the source appliance copies backups to the target: • Recency – By default, the source backup appliance sends copies to the target using the *recency* queue scheme, where the most recent backups are copied first. Unitrends recommends this approach because it supports recovering from availability issues with the target appliance (or the WAN connecting to the target appliance) by skipping over older backups when a newer backup	See "To configure connection options" on page 231.

Setting	Description	Procedure
	arrives. This is particularly important if the connection to the target appliance is unreliable. • Maximize retention – If it is important to you to ensure that every backup on the source is copied to the target, choose the *Maximize retention* queue scheme. • Manual – To copy backups to the target manually, choose the *Manual* queue scheme. With the Manual scheme, the appliance does not add backups to the queue. Instead, you must copy backups by using the procedure "To copy a full backup to a hot backup copy target on-demand" on page 499.	
Suspend backup copies from this appliance	Check this box to stop sending backup copies from the source appliance. This option may become necessary when either your target appliance or the connection to your target becomes unavailable for an extended period of time.	See "To configure connection options" on page 231.
Reset Backup Copy	Use to stop active copy jobs, reset the backup copy processes, then restart active jobs that were stopped. Use only when working with Support or following troubleshooting instructions in a Unitrends KB article.	See "To configure connection options" on page 231.
Backup copy bandwidth throttling	If the WAN connection to your backup copy target is shared with general purpose Internet use during normal business hours, you may wish to throttle the amount of bandwidth that backup copies can use during these hours. **Note:** The bandwidth throttling setting limits the maximum amount of bandwidth that backup copy jobs can use. The amount of bandwidth available for the jobs is also effected by environmental factors, such as actual network throughput (which may be constrained by intermediate nodes between the source and target) and other backup copy job tasks, such as encryption and compression.	See "To set bandwidth throttling" on page 227.

To suspend hot backup copies

Run this procedure to stop sending backup copies to a Unitrends appliance target or to the Unitrends Cloud. After suspending hot copies, you can resume copies at any time by using the "To resume hot backup copies" procedure.

1 Log in to the source backup appliance.

2 Select **Configure > Appliances**.

3 Select the source backup appliance and click **Edit**.

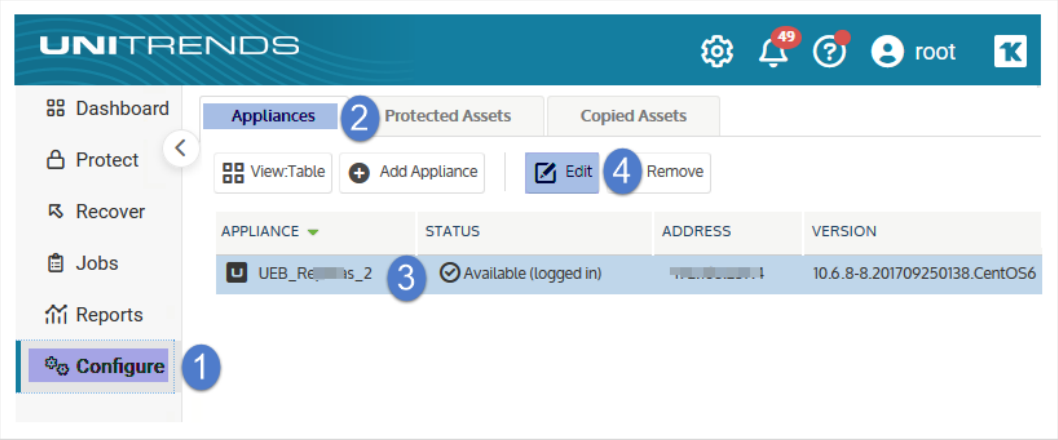

4 On the **Backup Copy** tab, locate the source appliance in the Backup Copy Source area and check its **Suspend backup copies from this appliance** box.

> **Note:** If the appliance has multiple hot backup copy targets, a scroll bar displays. If needed, scroll to locate the target whose copies will be suspended.

5 Click **Save**.

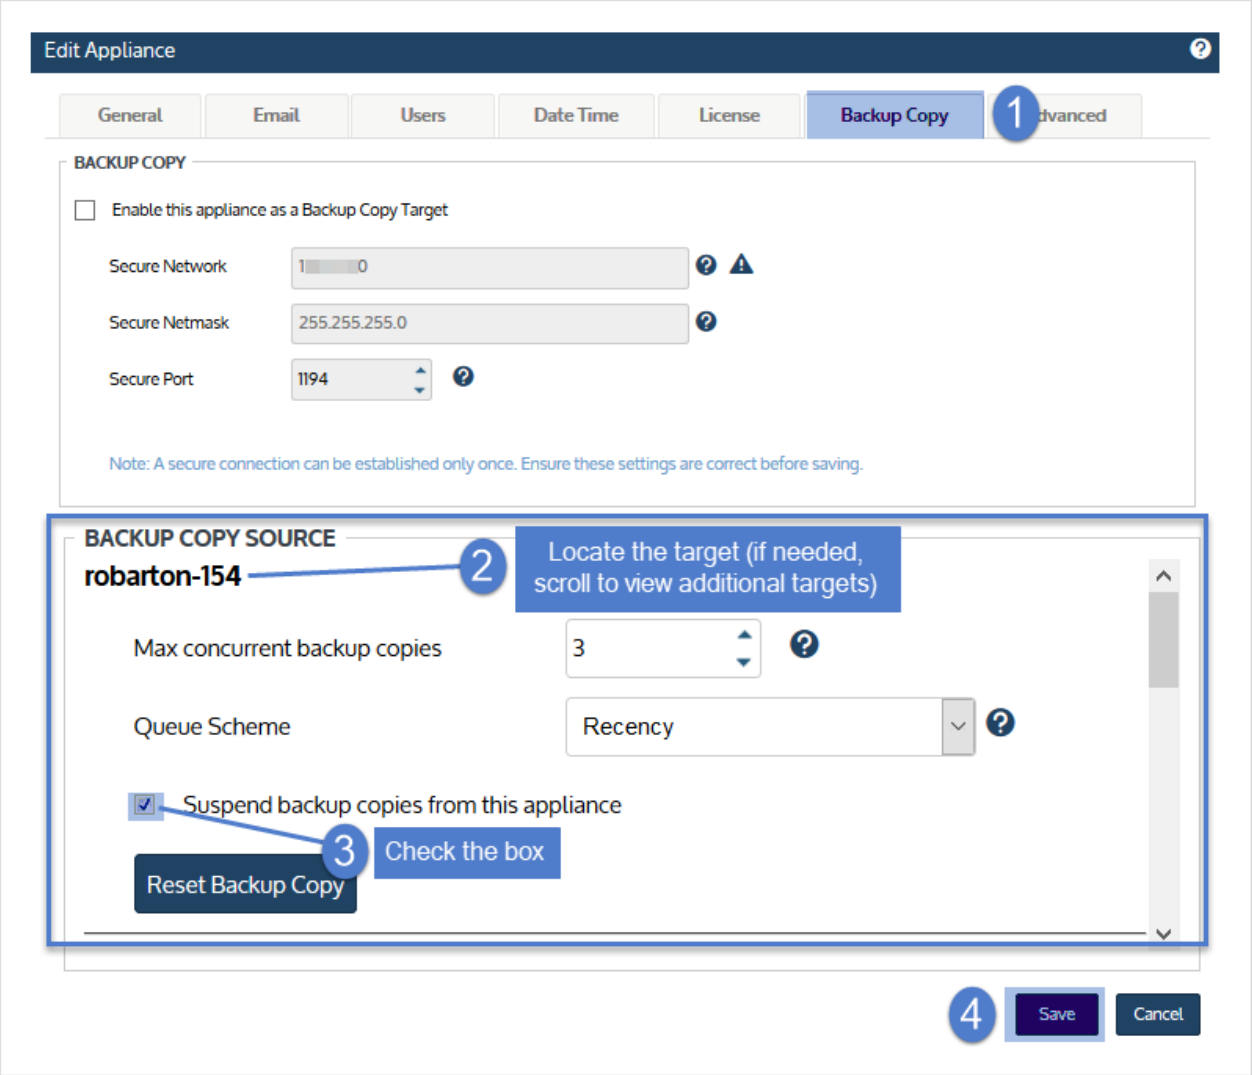

To resume hot backup copies

After suspending hot copies, you can resume copies at any time by using this procedure.

1. Log in to the source backup appliance.
2. Select **Configure > Appliances**.
3. Select the source backup appliance and click **Edit**.

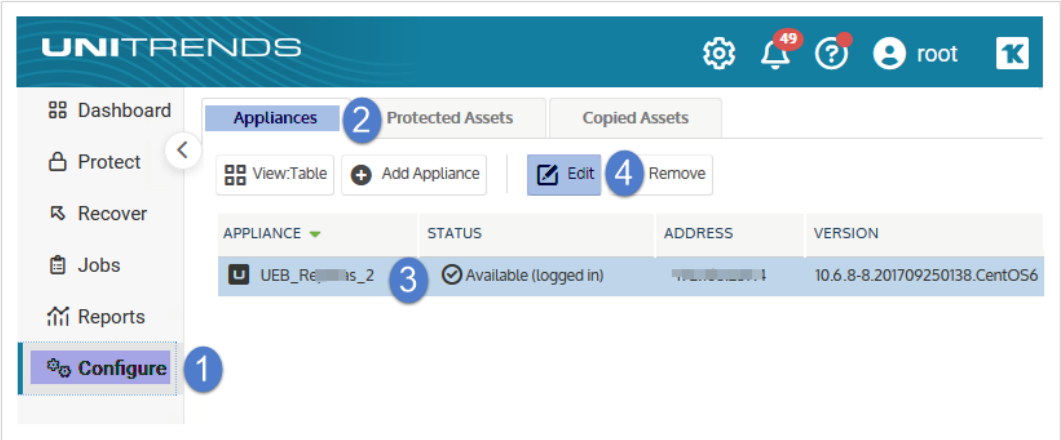

4 On the **Backup Copy** tab, locate the source appliance in the Backup Copy Source area and clear its **Suspend backup copies from this appliance** box.

> **Note:** If the appliance has multiple hot backup copy targets, a scroll bar displays. If needed, scroll to locate the target whose copies will be resumed.

5 Click **Save**.

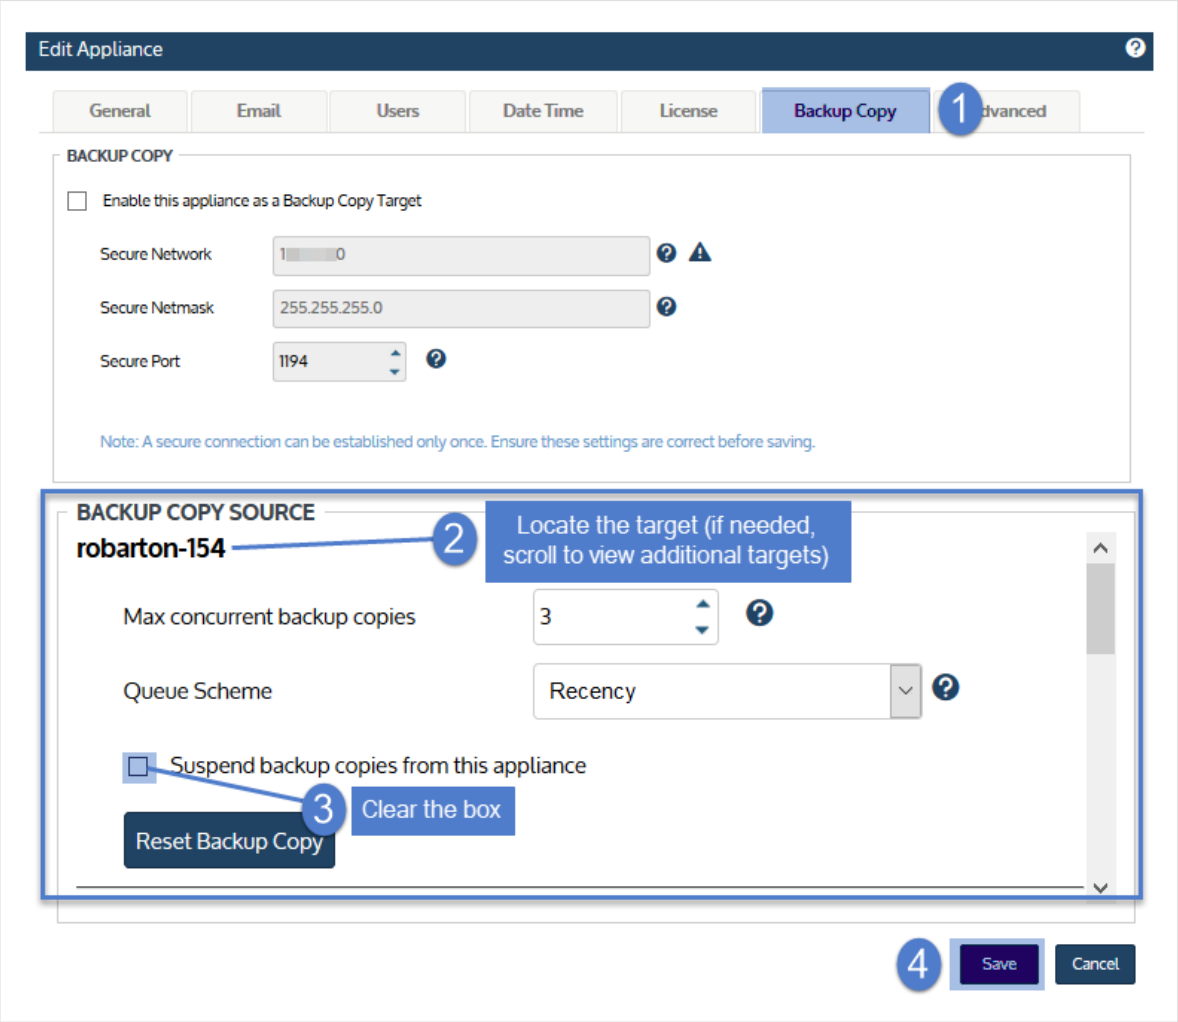

Do not cancel an active hot backup copy job

Do not cancel an active hot backup copy job. Instead, suspend hot backup copies.

Each successful backup of the original asset is copied to the hot backup copy target as soon as the backup completes. (The target can be the Unitrends Cloud or another Unitrends appliance.) To copy a backup, the appliance runs a hot backup copy job, which displays on the Active Jobs tab as shown here:

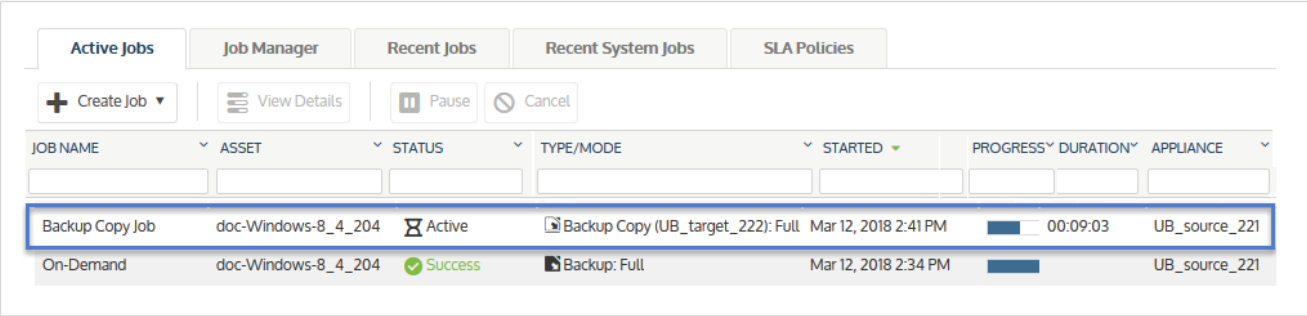

If you cancel a hot copy job by using the Cancel button on the Active Jobs page, the appliance automatically creates a new copy job to replace the one you canceled. To temporarily stop copying backups, suspend backup copies instead (as described in "To suspend hot backup copies" on page 266). Use the procedure "To resume hot backup copies" on page 268 to start sending hot copies again. Note that copies of all backups that ran while copies were suspended will be sent to the hot backup copy target once you resume. You cannot skip copying a specific backup.

To initialize and erase cold backup copy media

For tape drive, USB, and eSATA devices, you must initialize new tapes or drives before they can be used for the first time. This removes any existing data and formats the media. After you have written data to the media, you can also use this procedure to erase all backup copies.

Notes:

- Use this procedure only if you wish to remove *all backup data* from the media.
- This procedure is not used for tape devices with autochangers. If your tape device has an autochanger, see "To prepare tapes for use with an autochanger device".

To initialize and erase cold backup copy media:

1. Load the tape or drive(s) you want to initialize and erase.

   Notes:

   If you need to remove a drive from an eSATA or USB dock, be sure to:

   - Power down the dock.
   - Swap the drive.
   - Power on the dock.

2. On the **Configure > Appliances** page, select your appliance.
3. Click the **Backup Copy Targets** tab below.
4. Click **Scan for Media**.

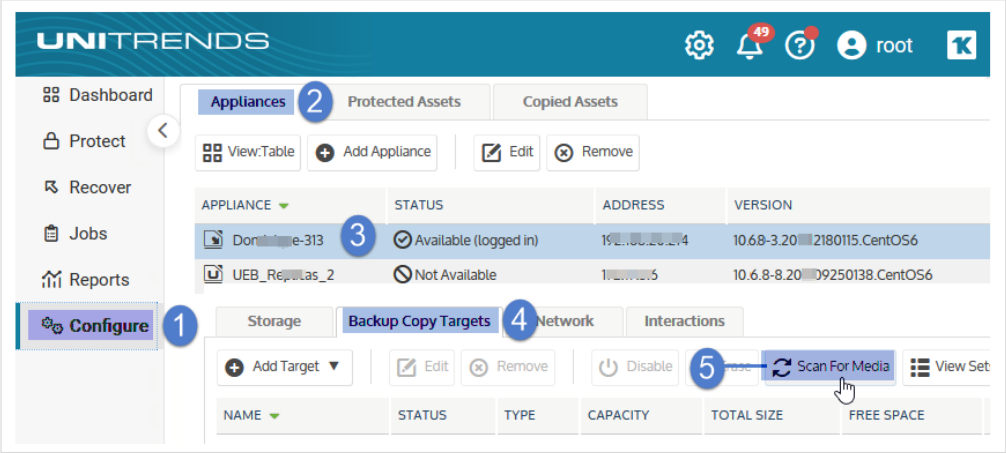

5 Select the target you wish to erase.

6 (Optional) If you want to check whether there are copies on the media, mount the media by clicking **Enable**. The target's status changes to Online.

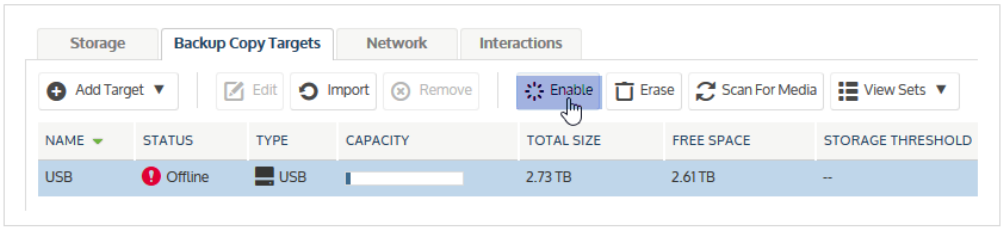

7 Click **Erase**.

8 (Optional) Check for existing backup copies.

9 (Optional) Enter a Media Label.

- The label can contain up to 12 alphanumeric characters or underscores.
- If you do not enter a label, the appliances generates one for you.

10 Click **Erase Backup Copies**. The appliance permanently deletes any existing data and formats the media with the Unitrends file system.

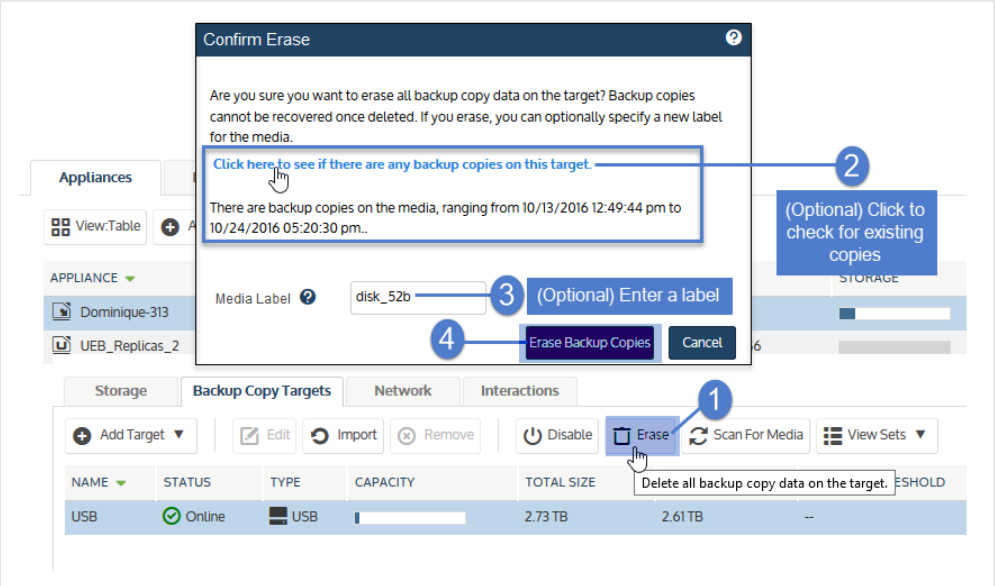

To prepare tapes for use with an autochanger device

For tape autochangers, you must prepare new tapes before they can be used for the first time. The prepare option removes any existing data and formats the media. After you have written data to the media, you can also use this procedure to erase all backup copies.

WARNING! Use this option with caution. Any existing data is permanently deleted from the media.

To prepare tape media:

1. Load the tape(s) you want to prepare.
2. On the **Configure > Appliances** page, select your appliance.
3. Click the **Backup Copy Targets** tab below.
4. Click **Scan for Media**
5. Select the autochanger in the list.
6. Click **Prepare**.
7. (Optional) Enter a Media Label.
   - The label can contain up to 12 alphanumeric characters or underscores.
   - If you do not enter a label, the appliance generates one for you.
8. Enter the Slots whose tapes will be prepared. Enter slot numbers using a comma-separated list and/or ranges. Example: 1,2,3,5-8
9. Click **Prepare**. The appliance permanently deletes any existing data and formats the media with the Unitrends file system.

To import a cold backup copy that was run by a different appliance

If you have run cold backup copies on one appliance, you can import reference information about those copies to a second appliance. Once reference data has been imported, you are able to recover those copies by using the second appliance.

Use this procedure to view the cold copy sets that reside on media that is currently connected to the backup copy target and to import a set's reference information to the cold copy catalog on the appliance.

Notes:

- Backup copy jobs mount and unmount the target automatically. The target remains offline unless a copy job is running. To view sets, you must enable the target, which also imports reference information for any backup copies that were not found on the appliance.
- A new cold copy format was introduced beginning in Unitrends release 10.1. If you attempt to import a cold copy and receive the message *Failed to read archive file*, verify that the appliance is running version 10.1 or higher. If not, install appliance updates, then try the import again.
- To view additional cold copy job details, use the "Managing backup copy targets" on page 260 procedure instead.

1. Connect the media containing the cold copies to the target appliance that will import the reference information.

   Notes:

   If you need to remove a drive from an eSATA or USB dock, be sure to:

   - Power down the dock.
   - Swap the drive.
   - Power on the dock.

2. Log in to the target appliance.
3. On the **Configure > Appliances** page, select the appliance.
4. Click the **Backup Copy Targets** tab below.
5. Click **Scan for Media**.

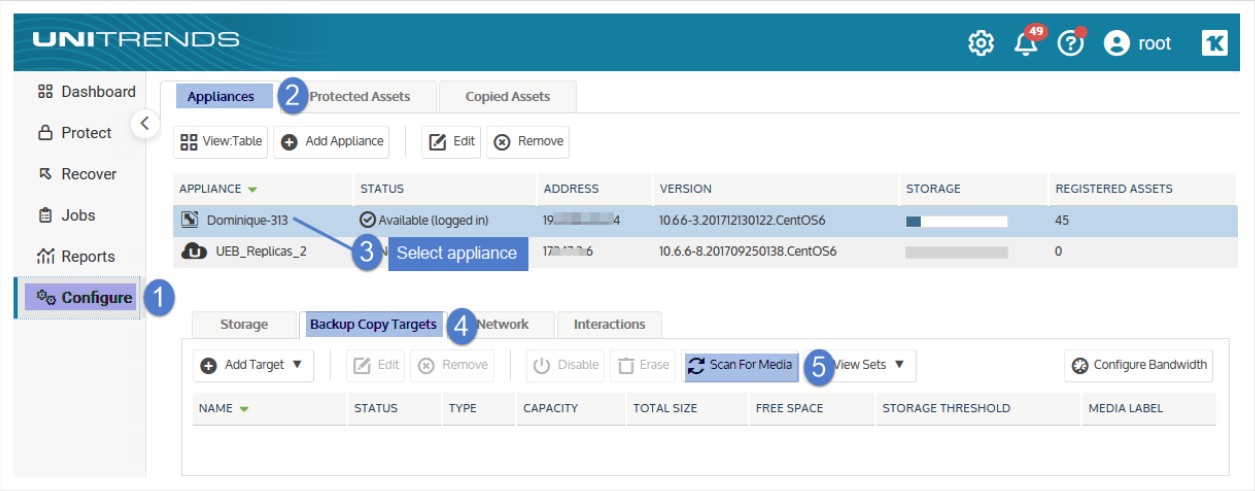

6 Select the cold backup copy target.

7 (If needed) Click **Enable** to bring the target online. (If the target is already enabled, you see a Disable button instead and you can skip this step.)

Notes:

If you receive a message that the media has not been initialized, you must erase or prepare the media before you can enable the target. Erasing or preparing the media removes all backup copies stored on the media. For details, see one of the following:

- "To prepare tapes for use with an autochanger device" on page 273 for tape autochanger targets.
- "To initialize and erase cold backup copy media" on page 271 for all other cold backup copy targets, including single tape drive devices.

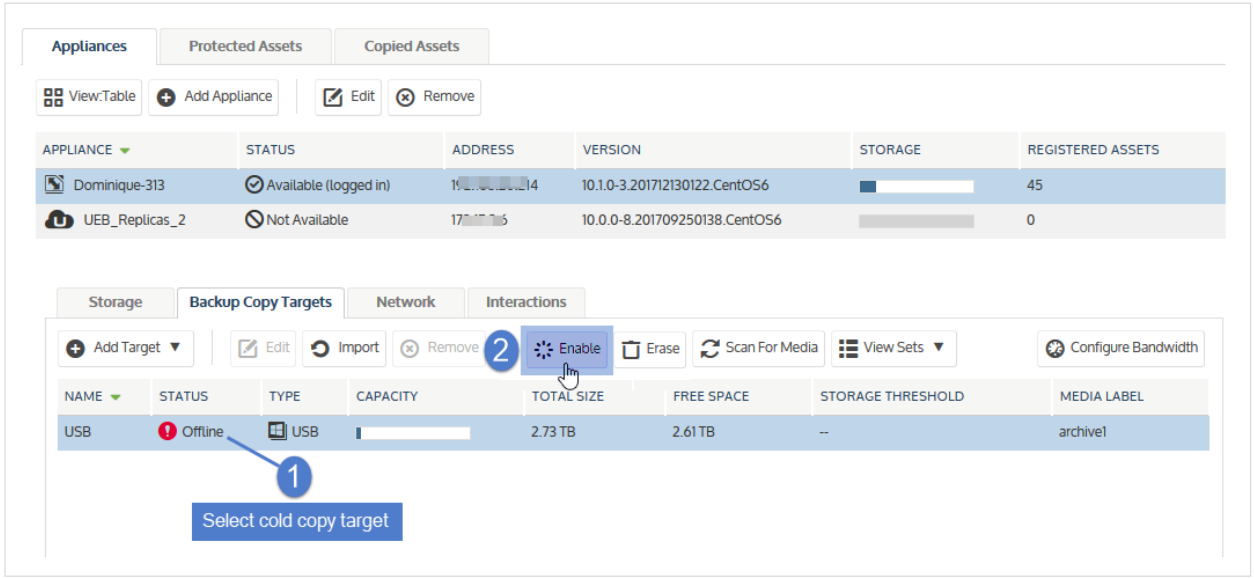

8 Select **On Media** from the **View Sets** list:

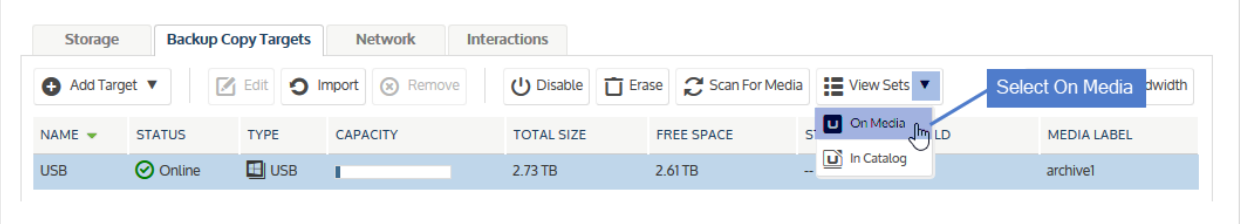

9 Sets that are stored on the media display in a list. The following is given for each set:

- Description – Name of the backup copy job that created the set.
- Date – Date and time that the set was copied to the media.
- Imported – Indicates whether the set's reference information has been imported to the cold target's catalog on this backup appliance. (Reference information must be imported so the set can be discovered in the Backup Catalog, where you can then opt to import the set's backup copies to the appliance. For details, see "To import a cold backup copy" on page 786.)

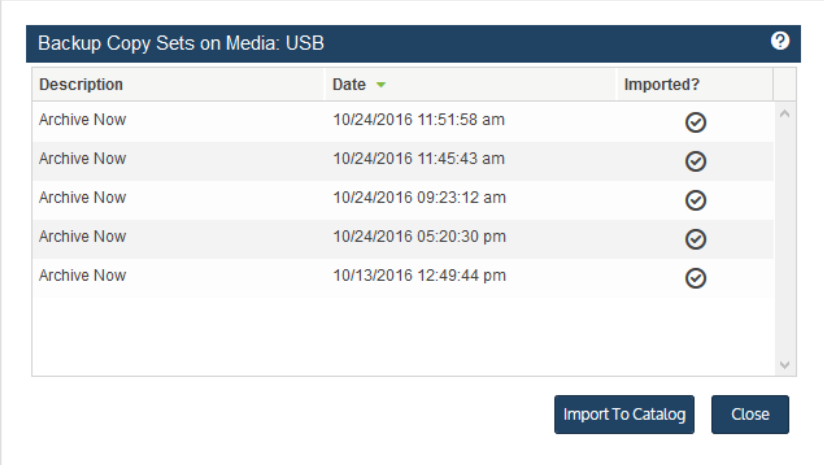

10 Do one of the following:

- To import a set's reference information, select it in the list and click **Import to Catalog**.
- To exit, click **Close**.

Reference information is imported. You can now import the cold copies to the appliance. See "To import a cold backup copy" on page 786 for details.

To enable a backup copy target

Use this procedure to mount the backup copy target.

Notes:

- Backup copy jobs mount and unmount the target automatically. The target remains offline unless a copy job is running, but you can use this procedure as needed to mount the target manually.
- This procedure also imports reference information for any backup copies that were not found on the appliance.

1 On the **Configure > Appliances** page, select your appliance.

2 Click the **Backup Copy Targets** tab below.

3 Click **Scan for Media**.

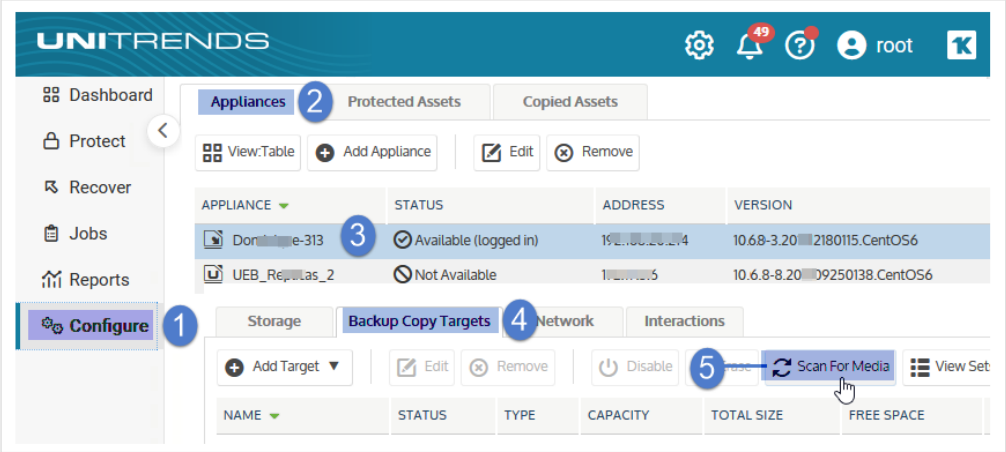

4 Select the offline target.

5 Click **Enable** to bring the target online.

Notes:

If you receive a message that the media has not been initialized, you must erase or prepare the media before you can enable the target. Erasing or preparing the media removes all backup copies stored on the media. For details, see one of the following:

- "To prepare tapes for use with an autochanger device" on page 273 for tape autochanger targets.
- "To initialize and erase cold backup copy media" on page 271 for all other cold backup copy targets, including single tape drive devices.

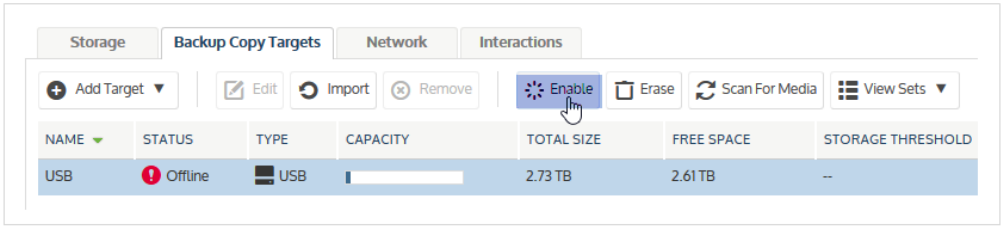

The target is enabled.

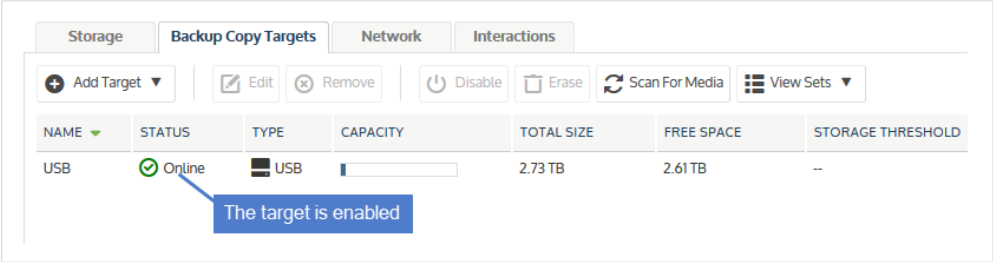

To swap out drives in a multi-drive Recovery Archive unit

You can swap out drives as needed, without powering down your Recovery Archive unit. Before swapping drives, note these considerations for using a device with multiple drives:

- All drives attached to a single appliance must have equal capacity. Drives may be of varying capacity if they are attached to different appliances.
- Within the backup copy target, all drives attached to a single appliance are treated as one logical device. When you add a multi-drive device and initialize the drives, the appliance formats them as a single unit. Data is then written across all drives as if they are one larger drive. Once you copy backups to the device, these drives must be treated as a single entity. You must remove them as a set and, to recover data, you must insert all drives in the set back into the eSATA or USB device so the appliance can read and import backup data. If you separate a drive from the set, all data is lost.

To swap out a drive in an eSATA or USB dock

1. Power down the dock.
2. Swap the drive.
3. Power on the dock.

To remove a backup copy target

Use this procedure to remove a target that is not directly connected physically to the appliance. Applies to these backup copy target types: Unitrends Cloud, Unitrends appliance, third-party cloud, NAS, and SAN.

Notes:

- To remove an iSCSI LUN from the Unitrends appliance, you must go to the SAN manager and indicate that the SAN should no longer use the LUN.
- For targets that are directly connected to the appliance, such as eSATA, USB, and tape devices, you do not remove the target from the UI. Instead, physically disconnect the target.

1. On the **Configure > Appliances** page, select your appliance.
2. Click the **Backup Copy Targets** tab, and select the target you wish to remove.
3. Click **Remove**.

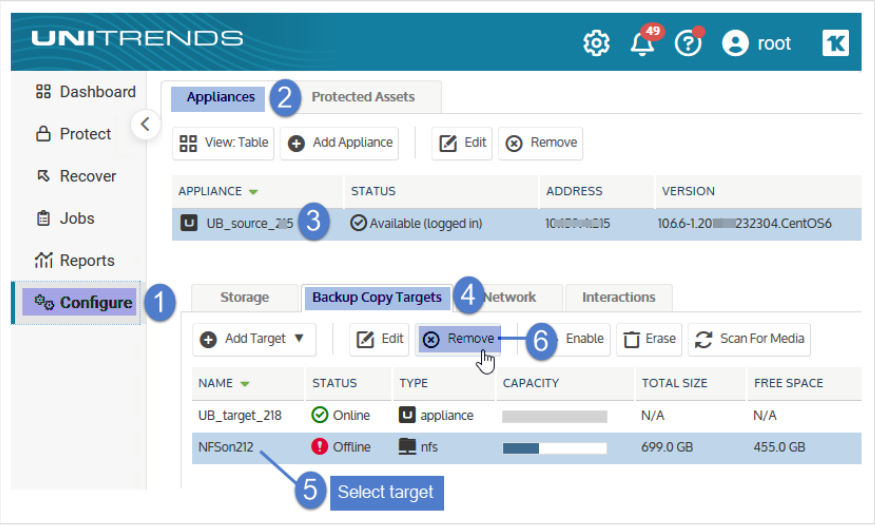

4 Click **Confirm**.

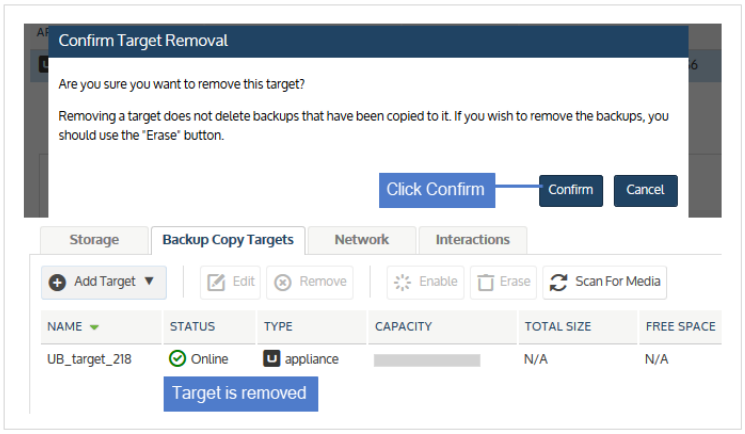

Protected assets

Any physical machine, virtual machine, or application you wish to protect is an *asset*.

Preparing to manage assets

The first step in protecting an asset is adding it to the appliance. Before you begin, determine which features you will configure for your assets and perform any required setup procedures. You can edit an asset at any time to implement a feature. If you are not sure which features you want to use, add the asset without optional features and configure these features later as desired.

Installing the Unitrends agent

Before you can protect a physical asset, you must install the Unitrends agent. (You can also opt to install the agent on virtual machines if you prefer to use file-level protection.) For most Windows assets, the appliance can push-install the agent when you add the asset. For other physical assets, you must install the Unitrends agent manually before you add the asset.

Note: A Unitrends agent is not used to protect iSeries assets, For details on iSeries, see "iSeries Backups Overview and Procedures" on page 767.

Agent installation procedures vary by operating system. See the following topics for details:

Operating system	Agent install procedure
Microsoft Windows	"Installing the Windows agent" on page 362
Linux	"Installing and updating the Linux agent" on page 387
CentOS	"Installing and updating the Linux agent" on page 387
Debian	"Installing and updating the Linux agent" on page 387
Red Hat	"Installing and updating the Linux agent" on page 387
SUSE	"Installing and updating the Linux agent" on page 387
Ubuntu	"Installing and updating the Linux agent" on page 387
Solaris	"Installing and updating the Solaris agent" on page 401
Novell Netware	"Installing and updating the Novell Netware agent" on page 399
Mac	"Installing and updating the Mac agent" on page 398
AIX	"Installing and updating the AIX agent" on page 396
UnixWare	"Installing and updating the UnixWare agent" on page 402
HP-UX	"Installing and updating the HP-UX agent" on page 397

Configurable features for protected assets

The following table describes the features that can be configured when adding or editing a protected asset. A description of each feature follows. For procedures used to add or edit an asset, see "Managing protected assets" on page 286.

	Supported for protected asset type?			
Feature	**Physical**	**Virtual**	**Application**	**Configured where?**
"Asset credentials" on page 282	Yes	Yes	Yes	Create the credential, then apply using: • Add Asset • Add Virtual Host • Add NAS • Edit Asset • Edit Virtual Host
"Retention settings" on page 283	Yes	Yes	Yes	Apply using Edit Asset.
"Encrypt backups" on page 283	Yes	Yes	Yes	Configure on Edit Appliance, then apply using: • Add Asset • Edit Asset
Index settings	Yes, Windows image-level only	Yes, VMware Windows VMs only	No	Enables quick file recovery using filename search across all backups of the asset. Apply using Edit Asset (see "To edit an agent-based asset" on page 293 or "To edit a virtual machine asset" on page 317).
"Quiesce settings for host-level backups" on page 283	No	Yes, VMware, AHV, and XenServer only	No	Configure globally, by host, or by VM:

Feature	Supported for protected asset type? Physical	Virtual	Application	Configured where?
				Note: The application aware quiesce setting must be set at the VM level. Global and host-level settings do not overwrite any application aware setting. • Set up globally using Manage Global VM Settings. (Applies to all VMs on the selected appliance.) • Set up by host using Edit Virtual Host. (Applies to one host's VMs.) • Set up by VM using Edit Asset. (Applies to selected VMs .)

Asset credentials

Credentials are used to establish a trust relationship between the Unitrends appliance and its assets. Once you apply a credential to an asset, the appliance can only access the asset using the associated administrative username and password. If the username and password are not valid, access is denied.

You can apply credentials when adding or editing the following asset types:

- Virtual hosts - Credentials are required for each vCenter, ESXi, Hyper-V, Nutanix AHV, or XenServer host asset you add to the appliance. These credentials are required for the Unitrends appliance to run host-level backup and recovery jobs for hosted virtual machines (VMs). You must enter credentials when you add the virtual host.
- Agent-based assets - Credentials are optional for assets (typically physical machines) that you add individually. Credentials are recommended for Windows assets to enable push-installation of the Unitrends agent and agent updates. When you add an asset, you can enter new credentials or apply existing credentials.

- NAS assets - Credentials are required only if the NAS share is configured for authentication. Enter NAS credentials while you add the NAS asset.
- Hosted virtual machines - After you add a virtual host, any hosted VMs are discovered and display in the inventory tree under their virtual host. Credentials are required to enable application-aware protection of VMware Windows VMs and are optional for other VMs. Edit a VM asset to add new credentials or apply existing credentials.
- Hosted applications - When you add an asset, any hosted applications are discovered and display in the inventory tree under their host machines. Credentials are required for these applications: Cisco UCS, NDMP, Oracle, and SharePoint full farm installations. Credentials are optional for other application types. For considerations and requirements, see "Application Backups Overview" on page 733.

Retention settings

The Unitrends appliance ingests new backups and retains them until additional backup storage space is needed. The oldest backups are then purged to make room for newer ones. However, the Unitrends appliance will not delete the latest backups of any type for a given asset, or any backups that are held by a retention policy.

Retention policies assure that the necessary recovery points are available on your appliance. Furthermore, backups can be copied to an off-site target as described in "Backup copies" on page 101.

Appliances are configured with a default backup retention policy of 30 days. This 30-day policy is applied to each protected asset. To apply custom retention policies, see "Managing retention with long-term data management" on page 328.

Notes:

- The 30-day default retention policy applies to appliances imaged with release 10.7.8 or higher. This default policy does not apply to appliances that were originally imaged with an earlier release. Upgrading an appliance that was imaged with a pre-10.7.8 release does not modify its retention policies in any way.
- The 30-day default retention policy ensures that 7 daily backups and 4 weekly backups are retained for each protected asset.

Encrypt backups

Use the this option to encrypt an asset's backups using an AES-256 bit algorithm. Before an asset's backups can be encrypted, you must set up encryption on the appliance as described in "Encryption" on page 155. To encrypt an asset's backups, Use the Edit Asset dialog. (Go to **Configure > Protected Assets**, select the asset in the list, click **Edit Asset**, then select an encryption setting in the **Encrypt Backups** drop-down.

Quiesce settings for host-level backups

For host-level backups of VMware, AHV, and XenServer VMs, quiesce settings determine how the VM is brought to a consistent state in preparation for backup. (Quiesce settings do not apply to Hyper-V VMs.) Unitrends provides these quiesce settings: *crash consistent*, *application consistent*, and *application aware.* Detailed descriptions of each setting are described in the table below.

Consider the following when working with quiesce settings:

- For appliances that were deployed with release 9.1 or later, *crash consistent* is the default quiesce setting for newly added virtual hosts and VMs.

- For appliances that were deployed with a pre-9.1 release, *application consistent* is the default quiesce setting for newly added virtual hosts and VMs. The application consistent default persists upon upgrading to later releases.
- The *application aware* setting must be applied to VMs individually. Applying a quiesce setting globally or to one host's VMs does not overwrite a VM's application aware quiesce setting. A VM's quiesce setting is overwritten only if it was set to *crash consistent* or *application consistent*.
- Backups are run using a cascading fall-back approach. If a backup attempt fails, the appliance tries again with a less stringent quiesce setting:
  - *Application aware* falls back to *application consistent*.
  - *Application consistent* falls back to *crash consistent*.
  - If *crash consistent* fails, the backup fails. (There is no fall back.)
  - For VMware and AHV, a backup that was run with a lesser quiesce setting is marked with a Warning status.
  - For XenServer, a backup that was run with a lesser quiesce setting is NOT marked with a Warning status.
  - To determine which quiesce setting was used, go to the Backup History report and select the backup. Detailed messages display in the Backup Status window. Look at the *Snapshot for this Backup was created with...* entry in the *vProtect Messages*, *aProtect Messages*,or *xProtect Messages* section.

Detailed descriptions of each quiesce setting are given in the following table:

Quiesce setting	Description	Apply globally to all VMs	Apply to one host's VMs	Apply to selected VMs
Crash consistent	The VM is not quiesced before the backup runs. The backup takes a snapshot of the VM disks in their current state. This is the fastest quiesce setting.	Apply to newly discovered VMs by selecting **Crash Consistent** in the Global Virtual Machine Settings dialog. Optionally, use **Apply to all current VMs** to apply to existing VMs. For details, see "To manage global quiesce settings" on page 314.	To apply to one host's VMs, select **Overwrite this hypervisor's VMs to Crash Consistent** in the Edit Virtual Host dialog. For details, see "To apply a quiesce setting to one host's VMs" on page 315.	To apply to one or more selected VMs, select **Crash Consistent** in the Edit Assets dialog. For details, see "To edit a virtual machine asset" on page 317.
Application consistent	The VM guest operating system invokes processes to flush application and filesystem transactions and place the VM into an idle state while a	Apply to newly discovered VMs by selecting **Application Consistent** in the Global Virtual Machine Settings	To apply to hosted VMs, select **Overwrite this hypervisor's VMs to Application**	To apply to one or more selected VMs, select **Application Consistent** in the Edit Assets dialog. For details, see "To edit a virtual

Quiesce setting	Description	Apply globally to all VMs	Apply to one host's VMs	Apply to selected VMs
	VM disk snapshot is taken.	dialog. Optionally, use **Apply to all current VMs** to apply to existing VMs. For details, see "To manage global quiesce settings" on page 314.	**Consistent** in the Edit Virtual Host dialog. For details, see "To apply a quiesce setting to one host's VMs" on page 315.	machine asset" on page 317.
Application aware (VMware Windows VMs only)	Use this option for application-aware protection of hosted Exchange or SQL simple recovery model applications. Leverages VSS writers to provide application consistent quiesce and additional post-backup processing. Exchange logs are truncated with VMware full and incremental backups. SQL logs are not truncated. See "Recommendations for protecting SQL databases hosted on VMware virtual machines" on page 285 for best practices.	Not applicable. Must be applied to VMs individually.	Not applicable. Must be applied to VMs individually.	Apply to one or more selected VMs. To set up application-aware protection, use the Edit Assets dialog to supply administrative credentials and to select the **Application Aware** quiesce setting. For details, see "Using application aware quiesce" on page 319.

Recommendations for protecting SQL databases hosted on VMware virtual machines

Application aware quiesce does not truncate SQL logs. Follow these recommendations to protect SQL databases that are hosted on VMware virtual machines:

- Simple recovery model - No logs are created. Run host-level backups with the application aware quiesce setting.
- Full recovery model - Do one of the following:
  - Use agent backups.
  - Use host-level backups with the application aware quiesce setting along with separate transaction log backups to truncate logs. (Schedule periodic transaction log backups using a SQL Maintenance Plan. Do not use SQL Maintenance Plan with agent-based backups.)

- Bulk-logged recovery model - Use agent backups. See "Recommendations for bulk-logged recovery model" on page 748 for details.

Managing protected assets

Use these procedures to view, add, edit, and remove protected assets. These procedures include options to configure or modify various features. We recommend reviewing "Preparing to manage assets" on page 279 for details on these features before running the following procedures.

See the following topics to manage assets:

Viewing all protected assets

The **Configure > Protected Assets** tab displays assets in an inventory tree where:

- Each physical asset and virtual host displays as a top-level node.

  Note: If you have opted to install the Unitrends agent on a VM, it is treated as a physical asset and displays as a top-level node. Use the "Managing agent-based assets" on page 288 procedures to protect the VM.

- Each application displays as a sub-node under its host asset.
- Each VM displays as a sub-node under its virtual host.

To view the assets that have been added to an appliance

1. Select **Configure > Protected Assets**.
2. Use these options to customize your view:
   - View options:
     - To view assets in a list, click  . Each row in the list describes a single asset.

  - To view assets in a table, click . Assets display in tiles on the left. Click an asset to view its details.

- Display options for list view:
  - Click the icon to add or remove columns. Click **Reset column defaults** to restore the default display. Click **Clear all filters** to clear any column filters you have applied.
  - Click the arrow next to a column name to sort values in ascending or descending order.
  - Hover over a column border and drag to resize a column.
  - Enter text in a column's filter field to display only rows that contain the string you entered.

See the table below for a description of each Protected Assets column.

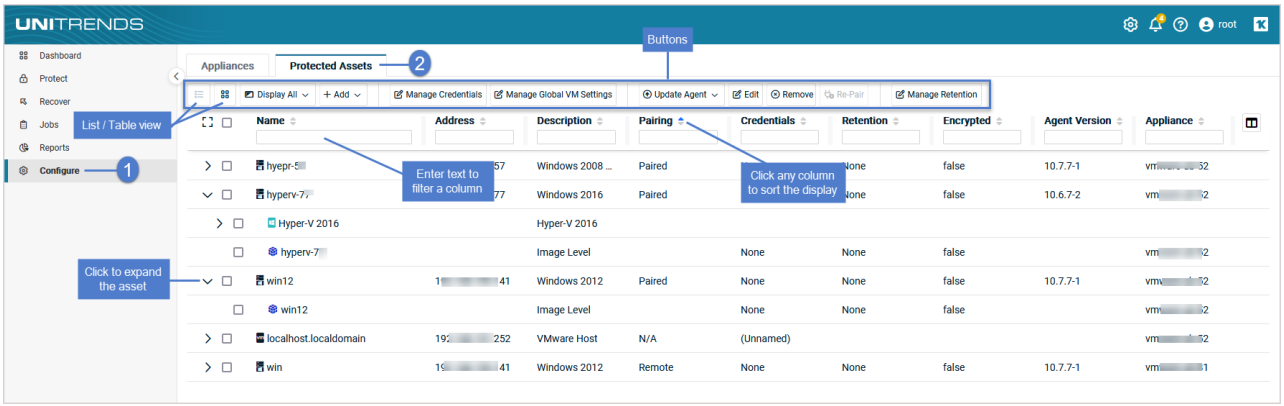

Column descriptions

Column	Description
Name	Name of the protected asset. If needed, click the arrow to view hosted VMs or applications.
Address	Asset's IP address, if applicable.
Description	Description of the asset.
Pairing	Asset's secure agent pairing status. The asset's pairing status is updated when a backup runs, during an inventory sync, or any time the asset is re-saved. • To update the pairing status of all applicable assets, click and select **Inventory Sync**. • To update the pairing status of one asset, select the asset, click **Edit**, then click **Save** in the Edit Asset dialog. The asset's last pairing status displays on the Protected Assets tab. Statuses include:

Column	Description

	• Paired – The agent has been paired.
	• Failed – The agent is not paired. Hover to see the error message. One of these errors has occurred:
	  – The agent pairing time window has expired.
	  – The agent can't save pairing keys.
	  – The agent pairing request failed.
	• Unsupported – Pairing is disabled or the agent pairing version is not compatible.
	• N/A – Not applicable for this asset.
	• Remote – This asset resides on a managed appliance. Log in to its appliance directly to see the asset's pairing status.
	For more on agent pairing, see "Secure agent pairing for Windows and Linux agents" on page 338, "Windows agent requirements" on page 362, and "Requirements for secure pairing of Unitrends Linux agents" on page 391.

Credentials	Name of the credentials that have been applied to this asset. *None* if no credentials have been applied. *Unnamed* if credentials have been applied but cannot be applied to any other asset. For details, see "Managing asset credentials" on page 322.
Retention	Asset's retention policy. *None* if no policy has been applied. For details, see "Managing retention with long-term data management" on page 328.
Encrypted	Indicates whether the asset's backups are encrypted: *True* if encrypted, *False* if not encrypted. For details, see "Encryption" on page 155.
Agent Version	Unitrends agent version running on the asset, if applicable.
Appliance	Name of the Unitrends appliance that is protecting the asset.

Managing agent-based assets

An *agent-based asset* is a machine that is protected by installing a Unitrends agent and running file-level or Windows image-level backups. Agent-based assets include:

- All physical machines other than iSeries servers. (For iSeries assets, see "iSeries Backups Overview and Procedures" on page 767 instead.)

- Any virtual machine that is protected by installing a Unitrends agent and running file-level backups. (In most cases, VMs are protected at the host-level. For host-level procedures, see "Managing virtual machine assets" on page 317 procedures instead.)

Use these procedures to manage agent-based assets:

- "To add an agent-based asset"
- "Creating aliases for agent-based assets" on page 291
- "To edit an agent-based asset" on page 293
- "Removing an agent-based asset" on page 295
- "Secure agent pairing for Windows and Linux agents" on page 338 for a description of this feature and related procedures.

To add an agent-based asset

The Unitrends agent must be installed for file-level or Windows image-level protection. For most Windows assets, the appliance can push-install the agent when you add the asset. For other physical assets, you must install the Unitrends agent manually before you add the asset. For procedures, see "Installing the Unitrends agent" on page 280.

1 Select **Configure > Protected Assets**.

2 Click **Add > Asset**.

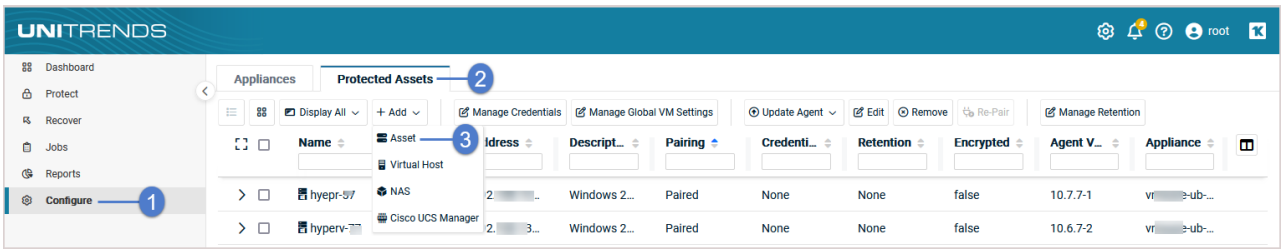

3 Enter the asset's hostname.

4 Enter the asset's IP address. This is optional in some cases, as described here:

- For Hyper-V hosts and Windows, Linux, or Mac assets, you can use DNS rather than entering a static IP address.
- DNS registration should be used for assets that obtain their network settings through DHCP. It is optional for assets with static IP addresses.
- If you do not enter a static IP address, make sure that both the asset and the appliance have DNS entries and that reverse lookup is configured.
- If you enter a static IP address, the appliance attempts to connect using this address, but if the attempt fails, it will try to add the asset using DNS.

5 Do one of the following:

- To push install the Windows agent, check the **Install Agent** box and enter credentials. (For push install requirements, see "Windows agent requirements" on page 362.)
- In all other cases, clear the **Install Agent** box.

6 Click **Save**.

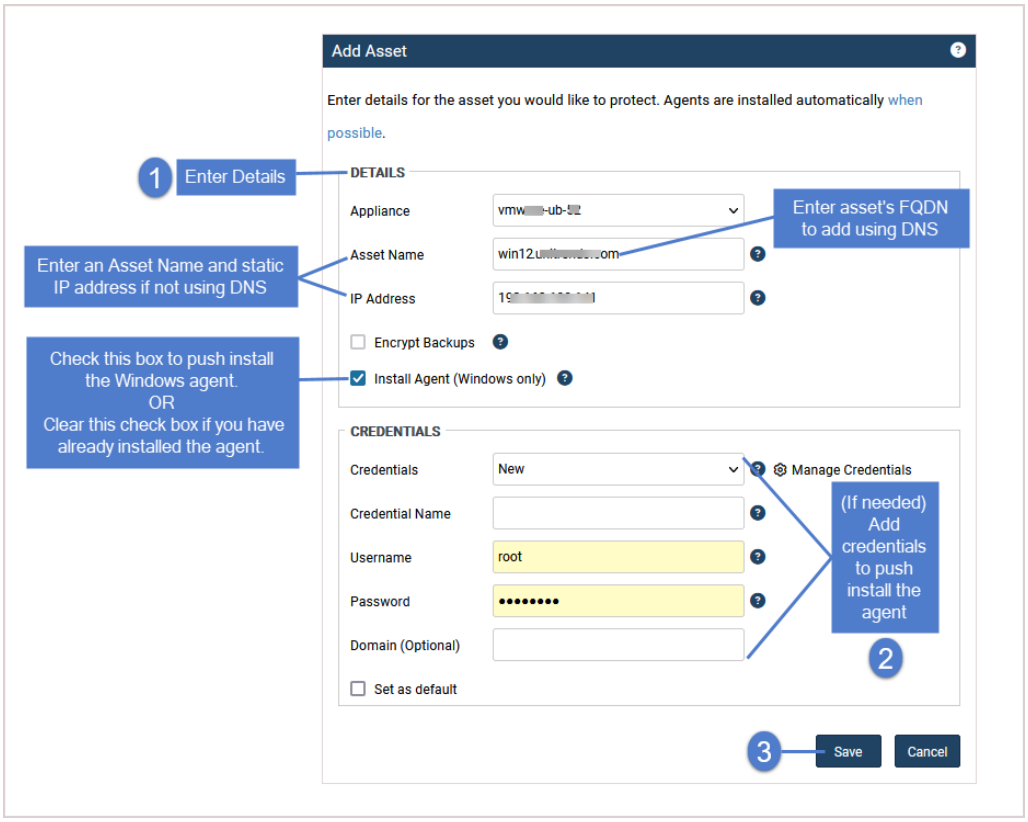

The asset is added to the appliance:

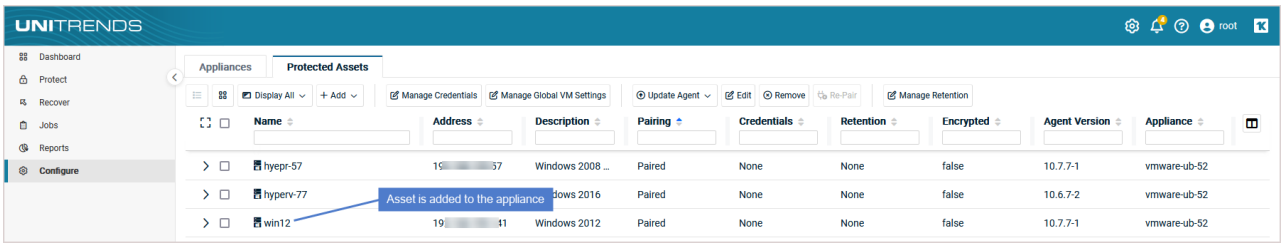

Note: If you receive the following error, enable pairing mode as described in "To enable pairing mode on a Windows asset" or "To enable pairing mode on a Linux asset": *Failed to save client...The agent pairing time window is expired*. Once pairing mode is enabled, run this procedure to add the asset.

7 (Optional) Exchange and SQL servers – To protect hosted applications with application aware image-level backups, use this procedure to enable the application aware setting: "To edit an agent-based asset" on page 293.

Creating aliases for agent-based assets

For large file servers, file-level backups may take a long time to run. To reduce overall runtime, you can create aliases for a single asset and run separate, radically different backup schedules for each.

Note: For Windows assets, you can opt to run image-level backups instead of creating aliases. This is a simpler approach. Because image-level backups are taken at the disk and volume level, they run much faster than file-level backups of servers that house many small files.

For example:

- Using aliases, you can break apart large data stores. Because backups are smaller, they run more quickly. Because less data is copied in the job, network traffic is reduced.
- You can have multiple fulls running at different times. Because a full cannot be purged until a new one is created, separating a large full into smaller ones can increase the space available on the appliance by enabling separate purging.
- When scheduling backup jobs for aliased assets, you must include the system state on the asset whose backups contain the boot and critical OS volumes (this is typically the *C:* volume) and exclude the system state on the other aliased assets.
- You can exclude the system state from an asset's backups when creating or editing the backup schedule. For details, see "To create a file-level backup job" on page 437.

To create an alias for an agent-based asset

Note: Add the asset to the appliance, as described in "To add an agent-based asset" on page 289, before running this procedure.

1. On the **Configure > Appliances** page, select the appliance and click the **Network** tab below.
2. Select the adapter and click **Edit Hosts File**.

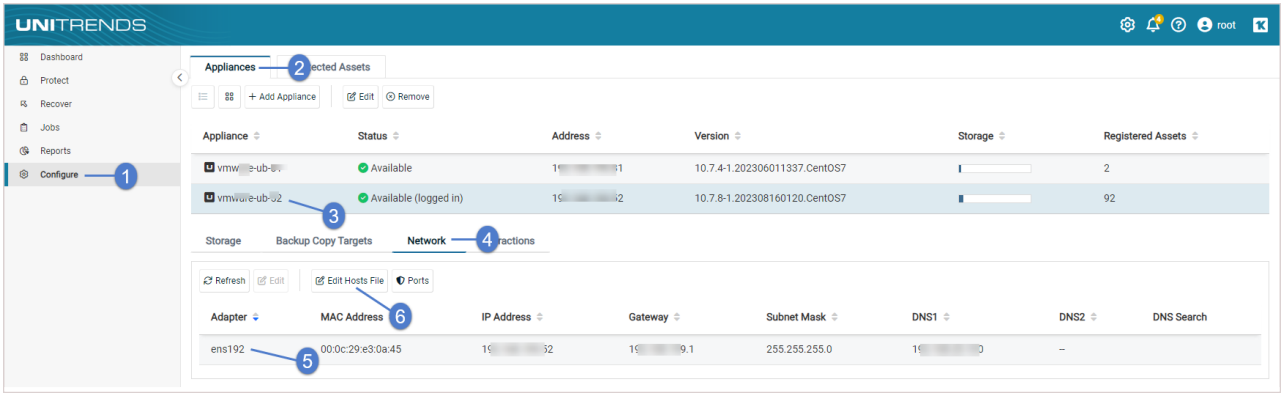

3. Select the asset in the list and click **Edit**.
4. Enter a second host name in the Aliases field, and click **Save**.

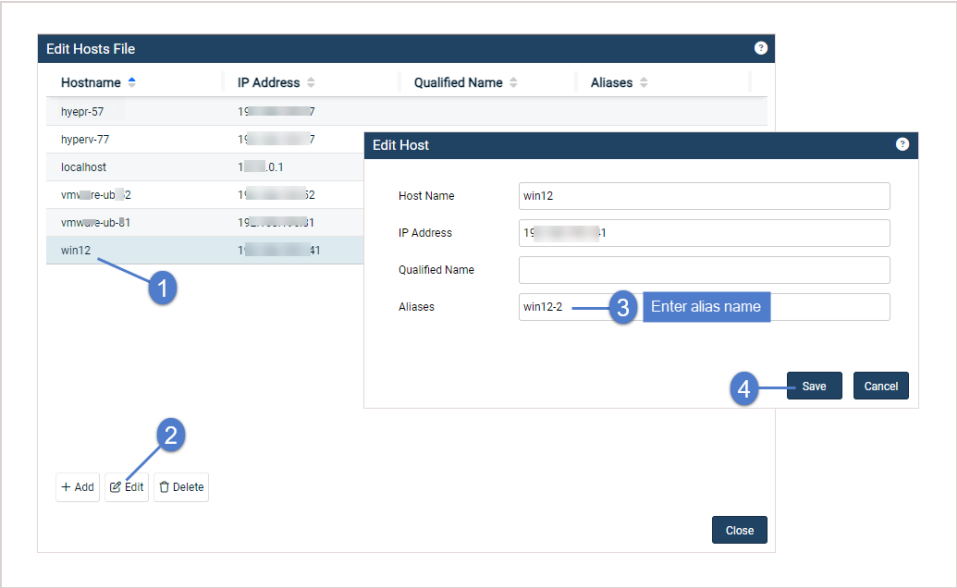

5 Click **Close** to exit.

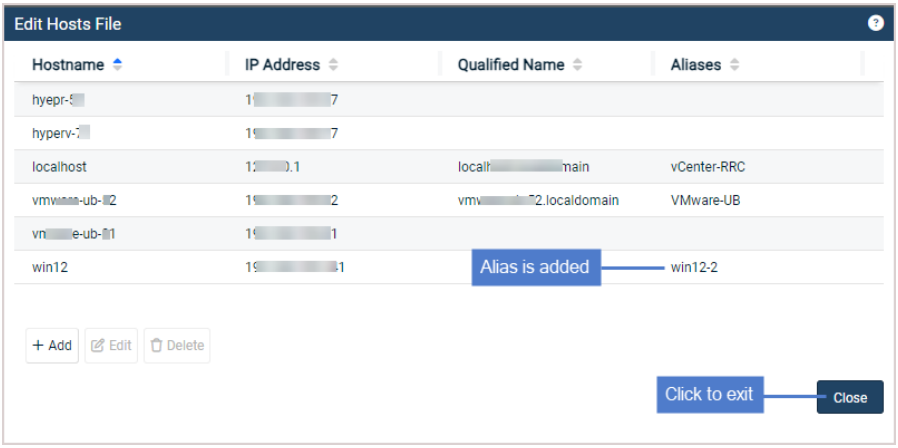

6 Add the alias as a separate protected asset:

- For Asset Name, enter the alias you entered in the hosts file.
- Do not enter anything in the IP Address field (the field must be empty).

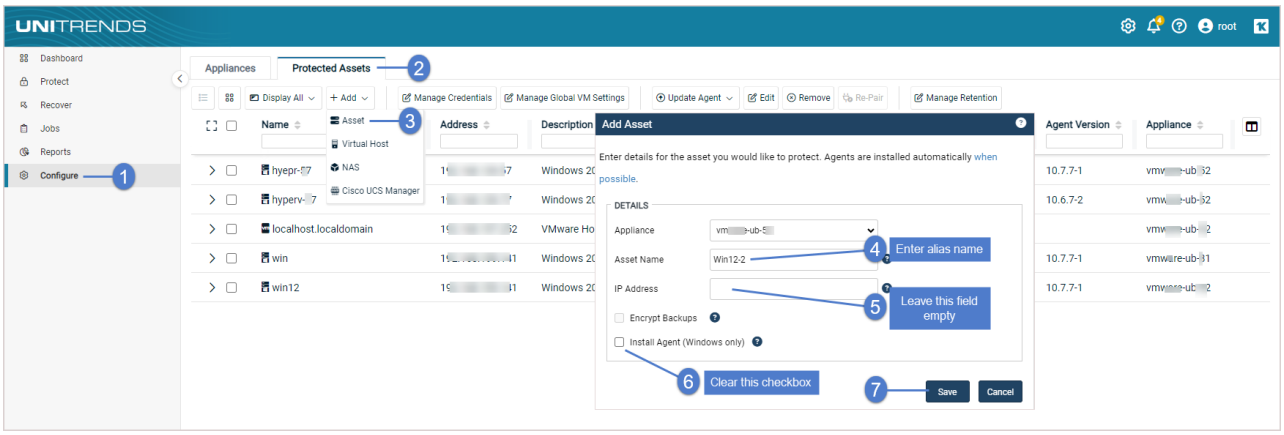

The aliased asset is added to the appliance:

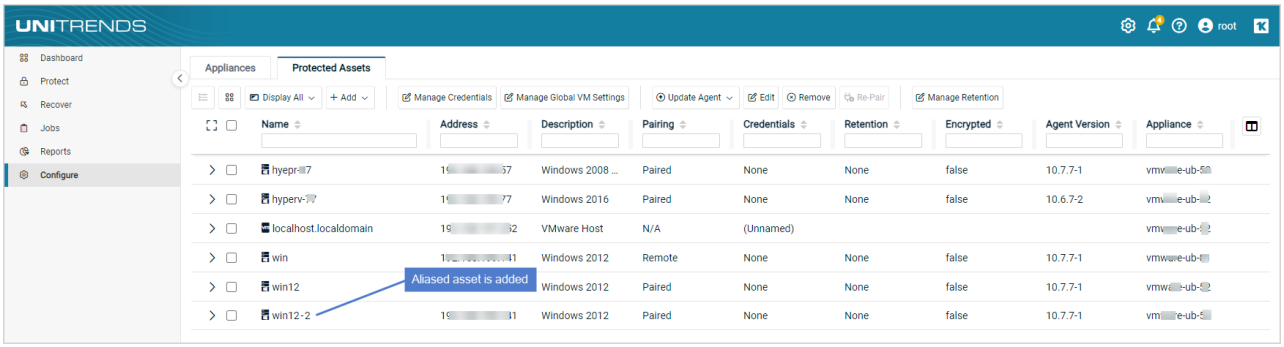

7 Create backup schedules to begin protecting the asset and its alias. Be sure to include the system state on the asset whose backups contain the boot and critical OS volumes (this is typically the *C:* volume) and exclude the system state on the other aliased assets. For details, see "To create a file-level backup job" on page 437.

To edit an agent-based asset

Note: Because each asset can have only one retention policy, you cannot edit an asset's retention settings if the asset has been added to an SLA policy. For more on SLA policies, see "Creating SLA policies" on page 536.

1 Select **Configure > Protected Assets**.

2 Select the asset and click **Edit**.

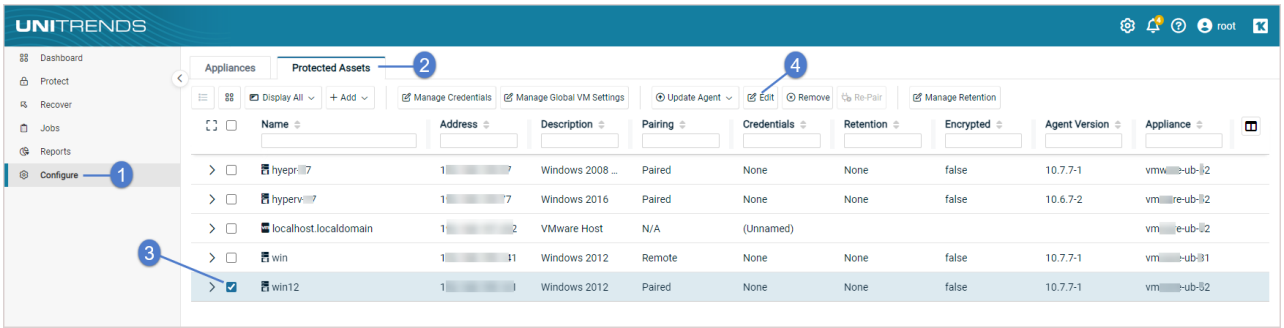

3 Modify settings and click **Save**.

- For more on the Details group settings, see "To add an agent-based asset" on page 289.
- For more on Credentials settings, see "Managing asset credentials" on page 322.
- For more on Retention settings, see "Managing retention with long-term data management" on page 328.
- Application aware setting for Windows image-level backups –
  - To configure a Windows asset to use the application aware feature when running image-level backups, check these boxes: **Show Image Level Backup Settings** and **Allow application aware**. For more on protecting hosted applications with image-level backups, see "Windows Image-level Backups Overview" on page 709.
  - To stop using the application aware feature when running image-level backups, check the **Show Image Level Backup Settings** box and clear the **Allow application aware** box. After saving your change, re-enable any Exchange and SQL backup schedules for this asset. For details, see "To enable or disable a job" on page 583.

  Note: Upon modifying the *Allow application aware* setting, any running Exchange or SQL jobs continue. Once these jobs complete, the new application aware setting is applied to subsequent image-level and application backups for this asset.

- Index setting for Windows image-level backups – You can index Windows image-level backups so you can quickly search for and recover individual files.

  Notes:

  - Assets with high-frequency backups or with very large file counts can add considerable load to the appliance. Consider appliance load when enabling the index option for these types of assets.
  - To index the backup, the job creates and mounts an object. If a file recovery object is already mounted for the asset, the backup runs but no index is created (as only one object per asset can be mounted at any given time). The resulting backup completes in warning status, with a message indicating that no index was created.

  - To index this asset's image-level backups, check these boxes: **Show Image Level Backup Settings** and **Index Image-Level Backups**.
  - To stop indexing this asset's image-level backups, check the **Show Image Level Backup Settings** box and clear the **Index Image-Level Backups** box.

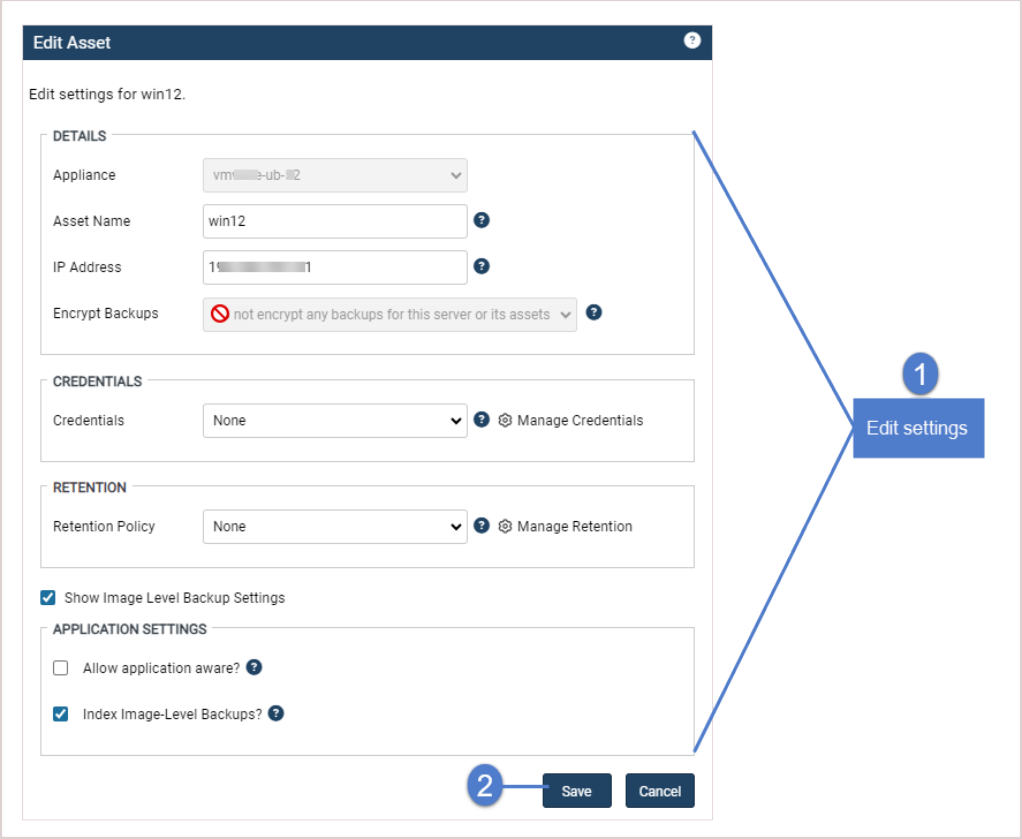

Removing an agent-based asset

CAUTION! When an asset is removed, all associated backups of that asset are also deleted. Please use caution when removing an asset.

Preparing to remove an asset

Before removing an asset, you must remove the asset from all job schedules.

Notes:

- Unitrends asset configuration settings are saved in the *master.ini* file. Note that deleting the asset from the Unitrends appliance also removes this file from the asset itself and any customized settings you have added are lost. Be sure to save the asset's *master.ini* file before deleting if you think you may want to add the asset to this or another Unitrends appliance and want to use these settings. After adding the asset back to an appliance, replace the standard *master.ini* file with the one you have saved.
- If you are using Windows replicas and you remove the Windows asset while a virtual recovery is in progress, the deletion may not be instantaneous. The clean up takes time because the recovery is shut down and the virtual replica asset is removed.

To remove an agent-based asset

1 Select **Configure > Protected Assets**.

2 Select the asset you want to remove.

3 Click **Remove**.

4 Check the **I understand...** box and click **Remove Asset and Delete Backups**.

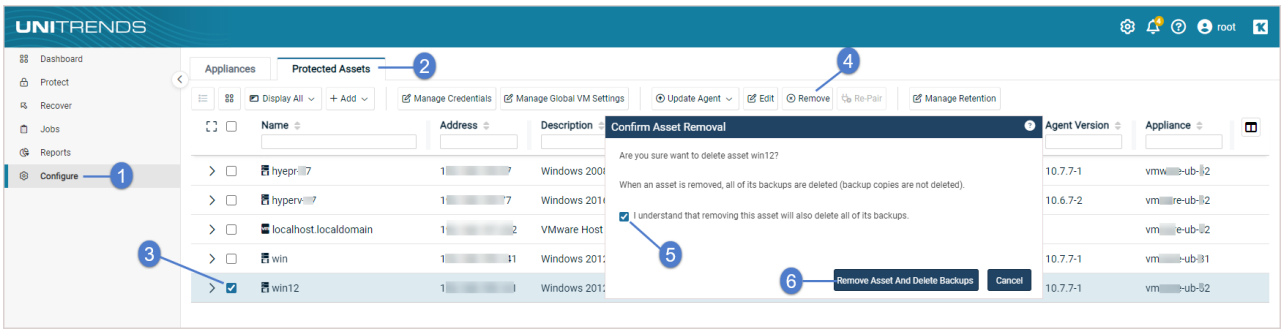

5 Click **OK** to close the Information message. The asset is removed.

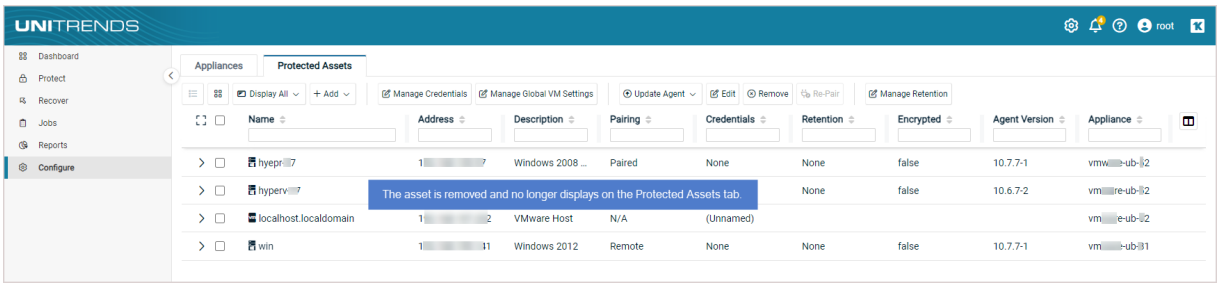

Managing NAS assets

Use these procedures to protect data stored on a NAS share using the CIFS, NFS, or NDMP protocol:

- "To add a NAS CIFS or NFS asset" on page 296
- "To add a NAS NDMP asset" on page 298
- "To edit a NAS asset" on page 299
- "Removing a NAS asset" on page 300

To add a NAS CIFS or NFS asset

1 Select **Configure > Protected Assets**.

2 Click **Add > NAS**.

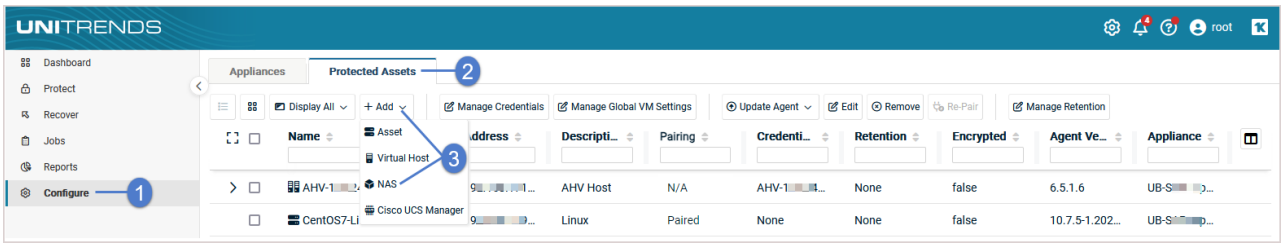

3 Enter the NAS Name. The name cannot contain spaces.

4 Select the Appliance that will protect this asset.

5 Enter the NAS IP address or resolvable hostname.

6 Select the CIFS or NFS protocol.

7 The Port field contains the default for the protocol selected. If the protocol uses a custom port, enter that port number.

8 Enter the full directory pathname of the NAS share in the Share Name field. Do not use leading or ending slashes.

Example pathname: *parentShare/subDirectory1/subDirectory2*.

- To protect only the *subDirectory2* share and its subdirectories, enter *parentShare/subDirectory1/subDirectory2*.
- To protect *parentShare* and all of its subdirectories, enter *parentShare*.
- If credentials are required to access the share and these credentials enable access to a parent directory only, enter the full path to the parent directory. You can specify desired folders and files to include in the backup when you create the job.

9 If credentials are required to access the NAS share, enter the Username and Password.

> **Note:** For the CIFS protocol, the password must not contain spaces. (Spaces can be used in passwords for the NFS protocol only.)

10 Click **Save**.

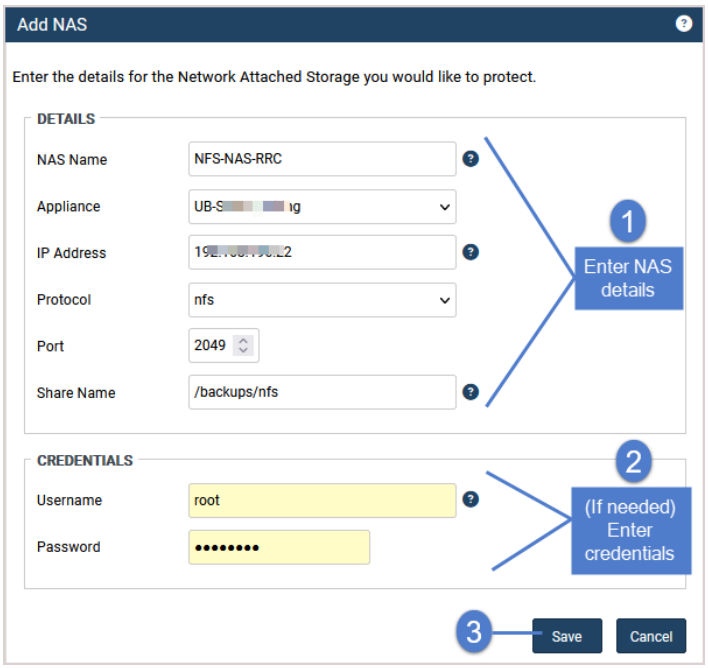

The NAS asset is added. To start protecting the NAS, see "To create a NAS CIFS or NFS backup job" on page 470.

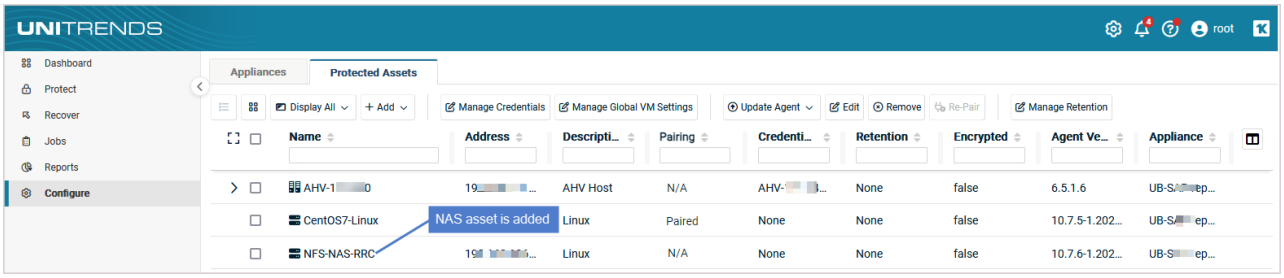

To add a NAS NDMP asset

1. Select **Configure > Protected Assets**.
2. Click **Add > NAS**.

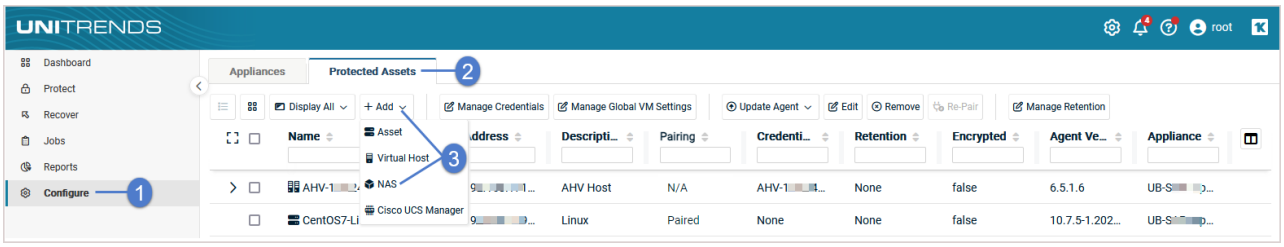

3. Enter the NAS Name. The name cannot contain spaces.
4. Select the Appliance that will protect this asset.

5 Enter the NAS IP address or resolvable hostname.

6 Select the NDMP protocol.

7 The Port field contains the default for the protocol selected. If the protocol uses a custom port, enter that port number.

8 Enter NDMP credentials:

- To use existing credentials, select one from the Credentials list.
- To add credentials, click **Manage Credentials > Add**, enter details, then **Save**.

9 Click **Save**.

The NAS asset is added. To start protecting the NAS, see "To create a NAS NDMP backup job" on page 475.

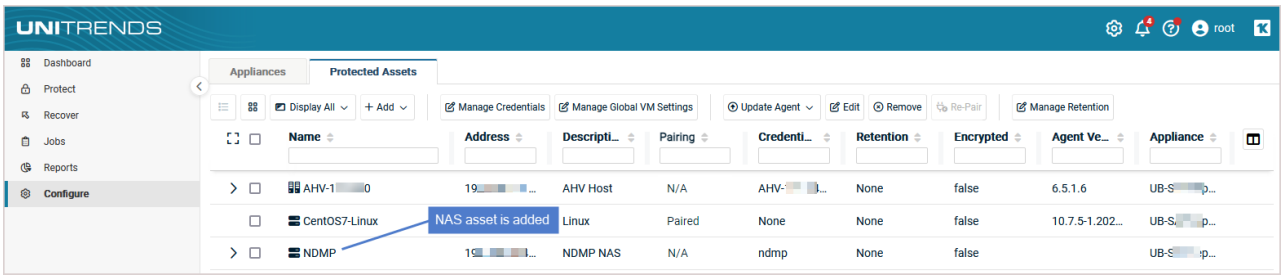

To edit a NAS asset

1 Select **Configure > Protected Assets**.

2 Select the NAS and click **Edit**.

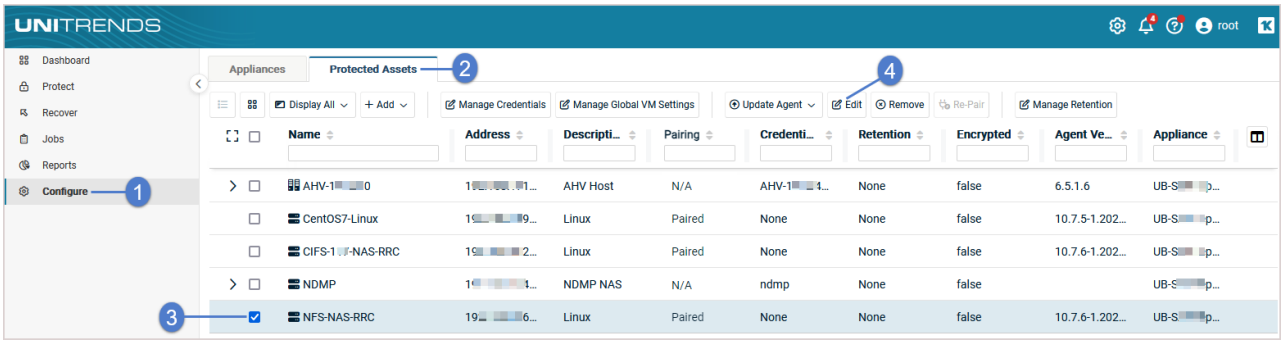

3 Modify settings and click **Save**. (For more on retention, see "Managing retention with long-term data management" on page 328.)

Note: For the CIFS protocol, the password must not contain spaces. (Spaces can be used in passwords for the NFS protocol only.)

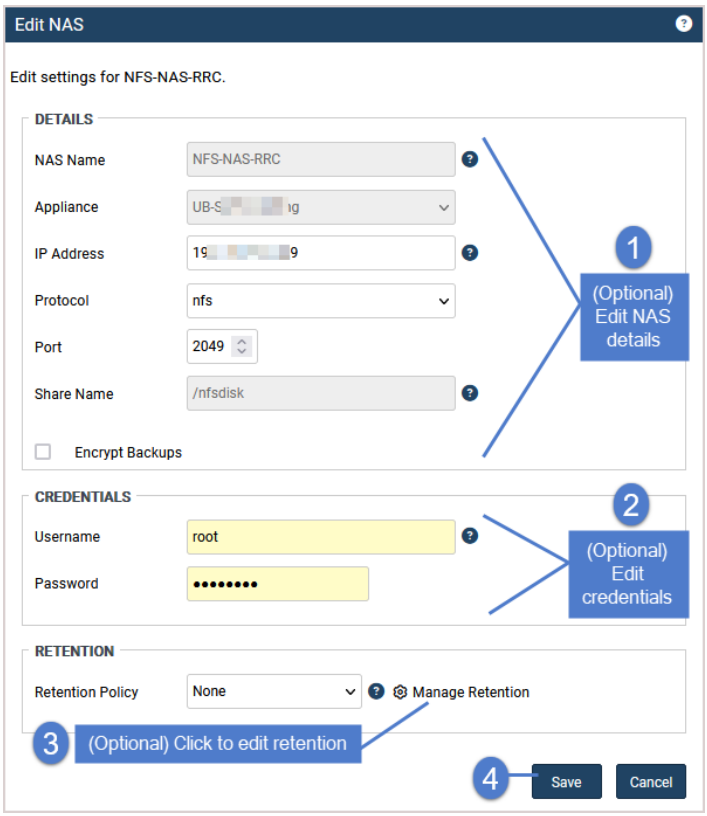

Removing a NAS asset

CAUTION! When an asset is removed, all associated backups of that asset are also deleted. Please use caution when removing an asset.

Preparing to remove an asset

Before removing an asset, you must remove the asset from all job schedules.

To remove a NAS asset

1 Select **Configure > Protected Assets**.

2 Select the NAS asset you want to remove.

3 Click **Remove**.

4 Check the **I understand...** box and click **Remove Asset and Delete Backups**.

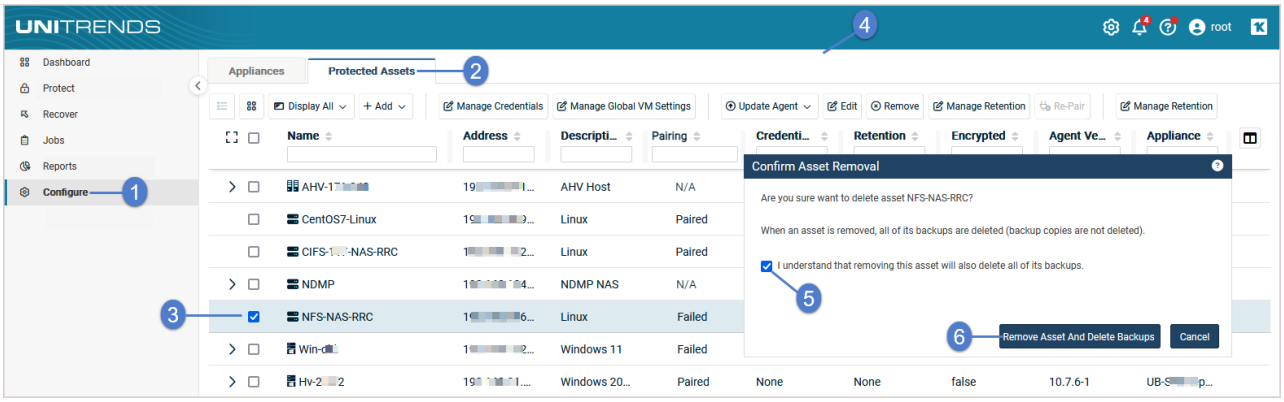

5 Click **OK** to close the Information message. The asset is removed.

Managing application assets

Use these procedures to manage application assets:

- "To add an application" on page 301
- "To add a UCS manager asset" on page 302
- "To edit a UCS manager asset" on page 303
- "To remove a UCS manager asset" on page 303
- "To edit an application" on page 303
- "To remove an application" on page 307

To add an application

To protect an application, you do not add the application itself. Instead, add its host server to the Unitrends appliance using the applicable procedure in the table below. Once you've added the host asset, run backups as described in "Creating backup jobs" on page 433. Before adding and protecting an application, be sure to review the applicable requirements and considerations, described in these topics:

- "Exchange backup requirements and considerations" on page 733
- "SQL backup requirements and considerations" on page 737

- "SharePoint backup requirements and considerations" on page 755
- "Oracle backup requirements and considerations" on page 759
- "Cisco UCS service profile backup requirements and considerations" on page 764

Application	Add host procedure
Exchange	Add the Exchange server using the "To add an agent-based asset" on page 289 procedure. All hosted databases or storage groups display under the asset you have added.
SQL non-clustered instance	Add the SQL server using the "To add an agent-based asset" on page 289 procedure. All hosted, non-clustered instances and databases display under the asset you have added.
SQL failover clusters and availability groups	Because clustered instances and availability groups can move between nodes in the cluster, you must add each server node, clustered instance, and availability group to the backup appliance as a separate asset, as described in "Protecting SQL clusters and availability groups" on page 748.
SharePoint	Configure and add the SharePoint server using the instructions in "SharePoint configuration prerequisites" on page 757. The SharePoint application displays under the asset you have added.
Oracle	Add the Oracle server using the "To add an agent-based asset" on page 289 procedure. All hosted Oracle instances display under the asset you have added.
Cisco UCS service profiles	Add the UCS manager using the "To add a UCS manager asset" procedure. The hosted service profile application displays under the asset you have added.

To add a UCS manager asset

To protect service profiles, add the UCS manager to the appliance using this procedure.

Note: To protect servers in your UCS environment, add the server to the Unitrends appliance using the applicable add asset procedure (see "To add an agent-based asset" on page 289 for physical servers or "Adding a virtual host" on page 308 for virtual machines).

1. Select **Configure > Protected Assets**.
2. Click **Add > Cisco UCS Manager**.
3. Select an Appliance.
4. Enter the asset's hostname as follows:
   - If your UCS is in a stand-alone configuration that consists of one physical UCS fabric interconnect that runs a single UCS manager, enter the hostname of the physical UCS node.

- If your UCS is configured in a cluster comprised of two physical Cisco UCS fabric interconnects (one active and one standby) with a UCS manager running on each, enter the cluster node name. Be sure to use the cluster name. Do not use the name of either fabric interconnect. With this approach, Unitrends can connect to the UCS manager regardless of which fabric interconnect is currently active.

5 Enter the asset's IP address, if required. (You do not need to enter an IP address if DNS is setup in your environment.)

- If your UCS is in a stand-alone configuration, enter the IP of the physical UCS node.
- If your UCS is configured in a cluster, enter the cluster IP address. Be sure to use the cluster IP. Do not use the IP of either fabric interconnect.
- Use DNS registration for assets that obtain their network settings through DHCP. This is optional for assets with static IP addresses.
- If you do not enter a static IP address, make sure that both the asset and the appliance have DNS entries and that reverse lookup is configured.
- If you enter a static IP address, the appliance attempts to connect using this address. If connecting by IP fails, the appliance attempts to add the asset using DNS.

6 Click **Manage Credentials > Add**, supply required credential information, and click **Save**.

The credentials you supply must support native backup and restore of UCS service profiles. To ensure sufficient privilege, the user must have Cisco UCS administrator privileges.

7 Click **Save** to add the asset.

To edit a UCS manager asset

1 Select **Configure > Protected Assets**.

2 Select the UCS manager asset (do not select the service profile application below).

3 Select the desired application and click **Edit**.

4 Modify settings as desired, and click **Save**.

To remove a UCS manager asset

CAUTION! When an asset is removed, all associated backups are also deleted. Please use caution when removing an asset.

1 Select **Configure > Protected Assets**.

2 Select the UCS manager asset.

3 Click **Remove > Confirm**.

To edit an application

Use this procedure to edit the encryption and retention settings of hosted applications. To edit other settings, such as credentials and IP address, edit the host server instead (for details, see "To edit an agent-based asset" on page 293).

1 Select **Configure > Protected Assets**.

2 Click to expand the application's host.

3 Click to expand the application instance.

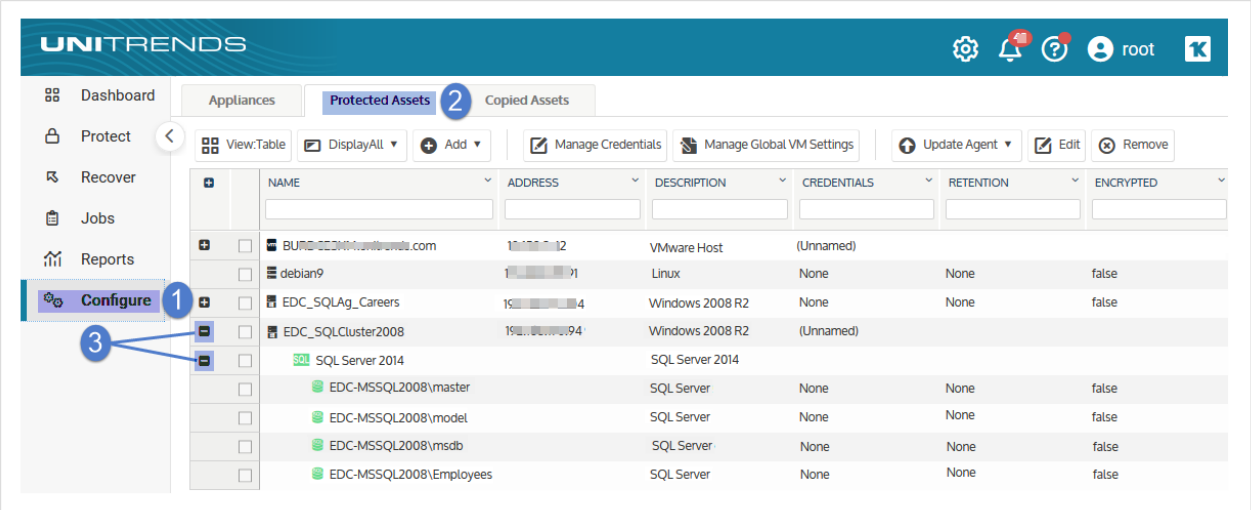

4 Select the hosted databases to which settings will be applied:

- Select the application instance to apply settings to all of its databases.

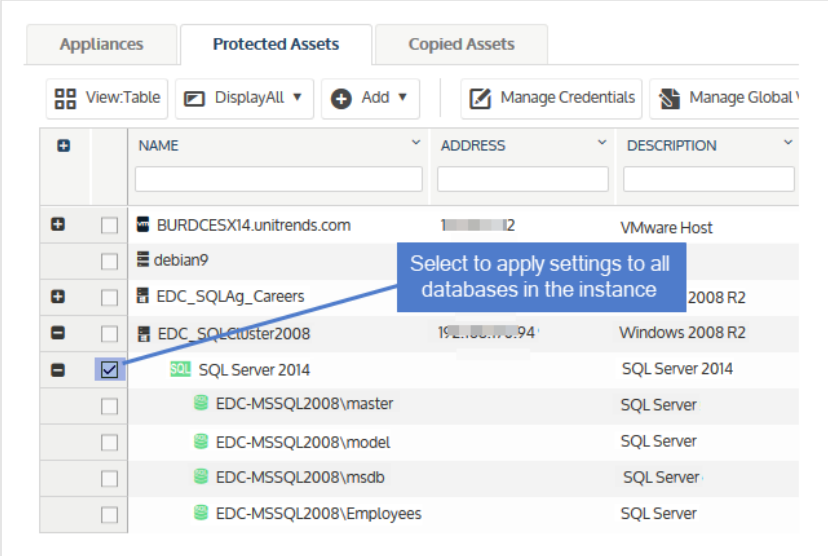

- Select individual databases to apply settings to a subset of databases.

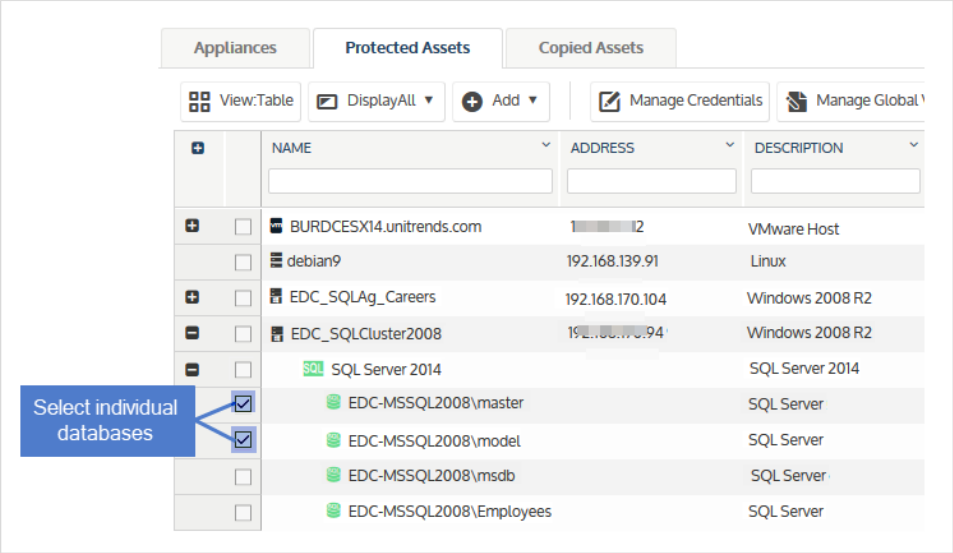

5 Click **Edit**.

6 Modify settings as needed:

- Encryption – Check **Encrypt Backups** to encrypt backups of the selected databases or clear the **Encrypt Backups** box to no longer encrypt backups of the selected databases. Click **Save**.

  Notes:

  - Encryption must be set up on the appliance before backups can be encrypted. (For details, see "Encryption" on page 155.)
  - Once encryption has been set up on the appliance, checking **Encrypt Backups** causes subsequent backups to be encrypted. Any existing unencrypted backups remain unencrypted.
  - Clearing the **Encrypt Backups** box causes subsequent backups to be unencrypted. Any existing encrypted backups remain encrypted.

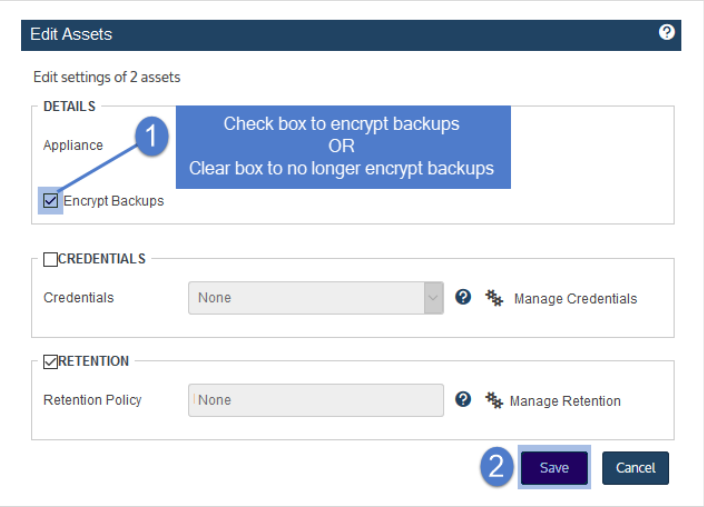

- Retention – Click **Manage Retention** to modify retention settings for backups of the selected databases. Enter settings and click **Save**.

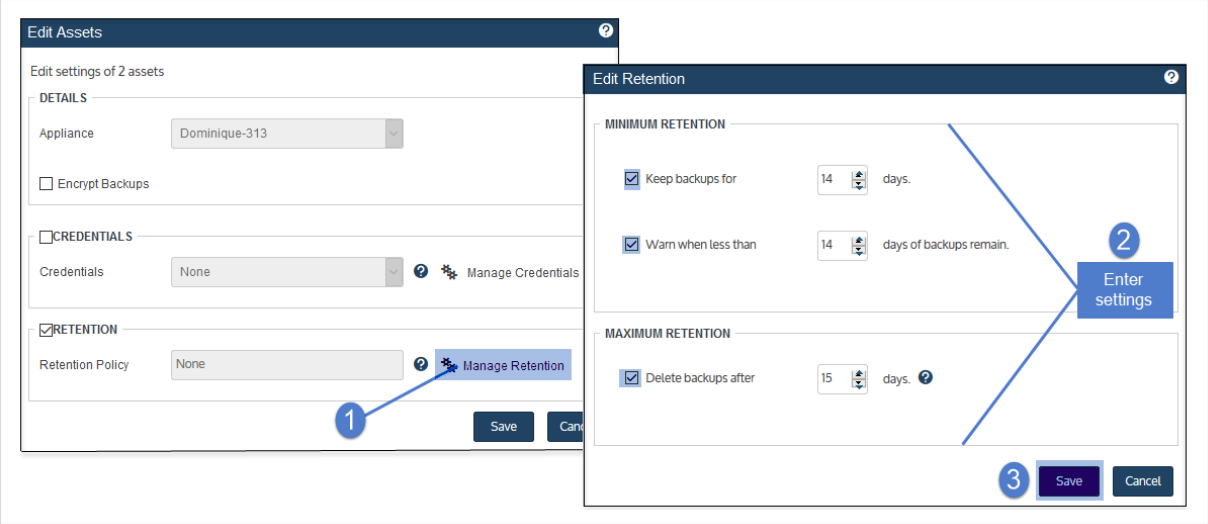

7 Settings are applied to the databases you selected and display on the Protected Assets tab.

- Example with settings applied to all databases (application instance was selected):

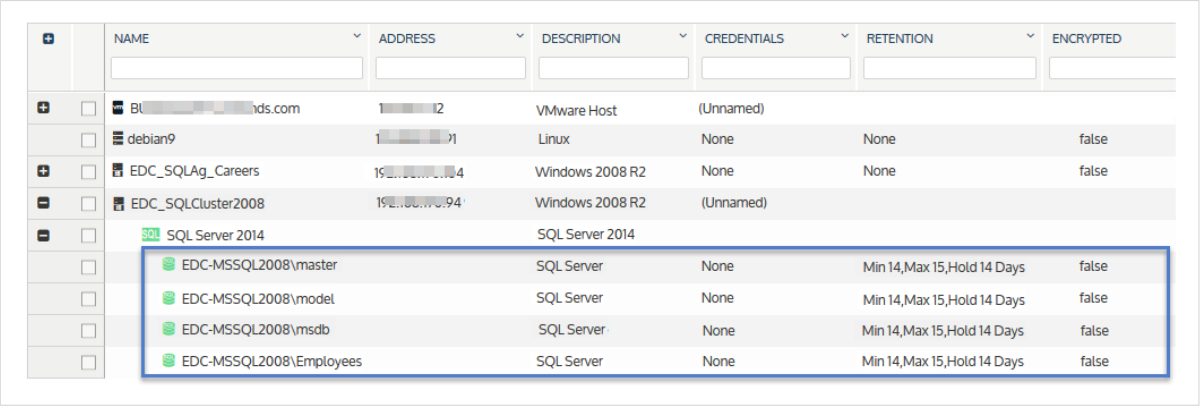

- Example with settings applied to individual databases:

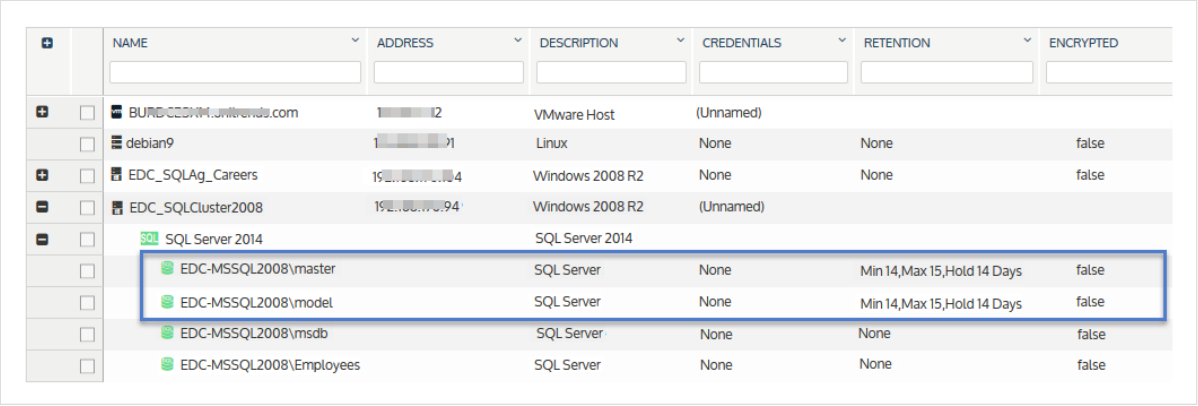

To remove an application

When a host asset is added, its applications are discovered and display in the UI. You cannot remove individual applications.

Managing virtual hosts

To protect VMs at the host level, you must add the VM's virtual host to the Unitrends appliance. The host can be an ESXi server, a Hyper-V server, a Citrix XenServer, or a Nutanix AHV cluster. Note the following:

- XenServer hosts are supported on Unitrends Backup on Citrix XenServer appliances only.
- Servers running free ESXi versions cannot be added as virtual hosts. To protect VMs hosted on free ESXi, you must install the Unitrends agent on each VM and protect them with file-level backups.
- If a vCenter is managing your ESXi servers, Unitrends recommends that you add to the appliance each ESXi server and the vCenter server itself. Some features, such as VM instant recovery, require a vCenter server. To enable these features, you must add both the ESXi host and the vCenter server to the appliance.
- Supported virtual hosts vary by Unitrends appliance, as described in the following table:

Unitrends appliance	Supported virtual hosts
Any of these appliance types: • Recovery Series • Recovery MAX • ION/ION+ • Unitrends Backup on VMware • Unitrends Backup on Hyper-V • Unitrends Backup on Nutanix AHV	• VMware vCenter • VMware ESXi server • Hyper-V server • Nutanix AHV cluster
Unitrends Backup on Citrix XenServer	• VMware vCenter • VMware ESXi server • Hyper-V server • Nutanix AHV cluster • XenServer

Use these procedures to manage virtual hosts:

- "Adding a virtual host" on page 308
- "To edit a virtual host asset" on page 312
- "To manage global quiesce settings" on page 314
- "To apply a quiesce setting to one host's VMs" on page 315
- "To upgrade a virtual host" on page 316
- "Removing a virtual host asset" on page 316

Adding a virtual host

Review the applicable considerations in the following table, then add each host as described in "To add a virtual host asset" on page 311. Once a virtual host is added, all VMs on that host are automatically discovered and can be selected for protection.

Virtual environment	Considerations
VMware	To protect hosted virtual machines, use the "To add a virtual host asset" on page 311

Virtual environment	Considerations
	procedure to add the following VMware servers to the Unitrends appliance: • vCenter and managed ESXi servers – If ESXi servers belong to a vCenter and both are accessible on the network, Unitrends recommends that you add the vCenter and its ESXi servers to the appliance. This enables the appliance to contact the vCenter for management operations (including vMotion support) and to directly contact the ESXi servers for backup and recovery, potentially improving performance by reducing network traffic around the vCenter server. • vCenter only – If the ESXi servers are accessible through a vCenter, adding the vCenter to the Unitrends appliance automatically detects all of the associated ESXi servers and their hosted virtual machines. This also enables the Unitrends appliance to be compatible with vMotion, a process through which VMs can migrate among the vCenter's ESXi servers. In this case, the appliance detects when VMs move between ESXi servers in a cluster and contacts the appropriate server to perform backups. • ESXi server only – If ESXi servers are not accessible through a vCenter, or if only a subset of the VMs hosted on the vCenter's ESXi servers are to be protected, you can add individual ESXi servers. In this case, the appliance contacts the ESXi servers directly for backup and recovery. **Notes:** • Servers running free ESXi versions cannot be added as virtual hosts. To protect VMs hosted on free ESX, you must install the Unitrends agent on each VM and protect them with file-level backups. • To protect templates or clustered VMs, you must add a vCenter server. • For additional requirements, see "Best practices and requirements for VMware protection" on page 666.
Hyper-V	Before adding the virtual host to the appliance, you must install the Unitrends Windows agent on the Hyper-V server. (See "Installing the Windows agent" on page 362 for details.) Once the agent has been installed, use the "To add a virtual host asset" procedure to add the Hyper-V host to the appliance. For additional requirements, see "Best practices and requirements for Hyper-V protection" on page 654.
Nutanix AHV	To protect VMs hosted in a Nutanix Acropolis Hypervisor (AHV) environment, you must add the AHV host cluster to the Unitrends appliance as a virtual host asset. Use the "To add a virtual host asset" procedure to add the AHV host. In the Add Virtual Host dialog, you must enter the following:

Virtual environment	Considerations
	• Hostname – Enter a unique name to identify the AHV cluster. This is the display name used by the appliance UI and does not need to match the actual hostname of the AHV cluster. • IP address – Enter the Nutanix cluster virtual IP address. This is a highly-available IP address used to reach the management services running on the Nutanix AHV cluster. • Username and Password: – For Nutanix AHV clusters that are NOT configured to use directory services authentication and are running a pre-5.5 AOS release, enter the credentials of the Nutanix **cluster admin** user account. You must use the **cluster admin** account. Other users with administrative privileges are not supported. – For Nutanix AHV clusters that are NOT configured to use directory services authentication and are running AOS release 5.5, enter the credentials of any local Nutanix cluster account that has been assigned the **cluster admin** or **user admin** role. – For Nutanix AHV clusters that are configured to use directory services authentication, enter the credentials of an LDAP user that has the **cluster admin** role. You must specify a domain in addition to the username. For AOS 5.1, the username and domain are case sensitive and you must match the case that was entered in the self service portal (SSP). In the Username field, enter the username and domain in this format: *user@domain*. For example, **jalvarez@unitrends.com**. (For configuration requirements, see "Requirements for directory services authentication in AOS 5.1" on page 696 or "Requirements for directory services authentication in AOS 5.5" on page 696.)
XenServer	Host-level protection of XenServer VMs is supported on Unitrends Backup on Citrix XenServer appliances only. You can add only one XenServer host to the appliance. The host must be one of the following: • A XenServer pool master host meeting both of these criteria: – The Unitrends Backup VM resides either on the pool master host itself or on one of the pool master's slave hosts. – The Unitrends Backup VM has been granted access to the shared storage used by the pool master host.

Virtual environment	Considerations
	• A stand-alone XenServer host where the Unitrends Backup VM resides. Use the "To add a virtual host asset" procedure to add the XenServer host. For additional requirements, see "Best practices and requirements for XenServer protection" on page 689.

To add a virtual host asset

Use this procedure to add any of the following to the appliance: an ESXi server, a vCenter server, a Hyper-V server, a Citrix XenServer, or an AHV cluster.

1 Select **Configure > Protected Assets**.

2 Click **Add > Virtual Host**.

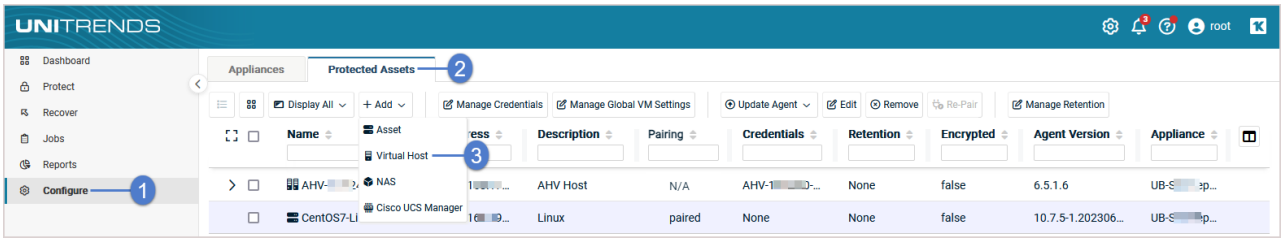

3 Complete all fields on the Add Virtual Host page.

Notes:

- The Hypervisor dropdown contains only the hypervisor types that can be protected on the Unitrends appliance. For supported hypervisors by appliance type see "Managing virtual hosts" on page 307.
- For AHV, user account requirements vary by cluster configuration. For details, see "Nutanix AHV" on page 309.

4 Click **Save**.

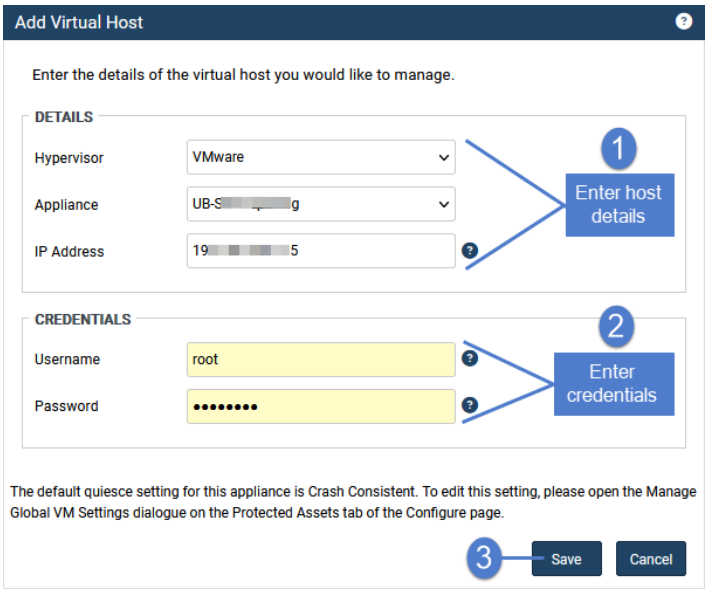

The host asset is added to the appliance. To start protecting the hosted VMs, see "Backup Administration and Procedures" on page 425.

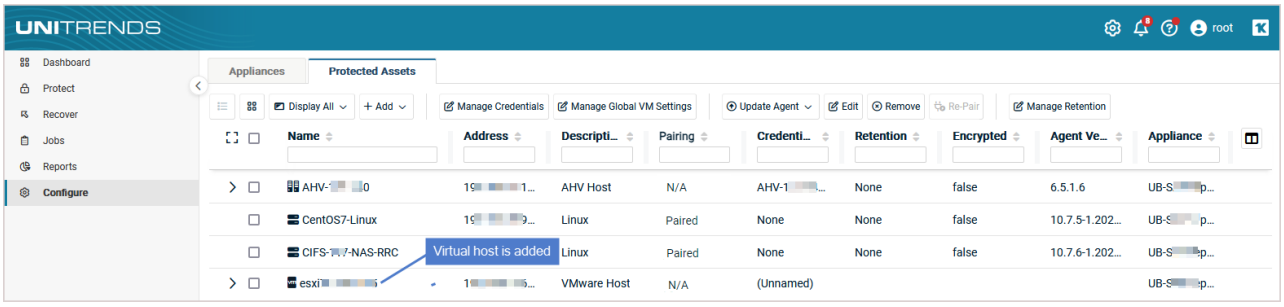

To edit a virtual host asset

Note: Because each asset can have only one retention policy, you cannot edit an asset's retention settings if the asset has been added to an SLA policy. For more on SLA policies, see "About creating backup and backup copy jobs" on page 426.

1. Select **Configure > Protected Assets**.
2. Select the virtual host asset.
3. Click **Edit**.

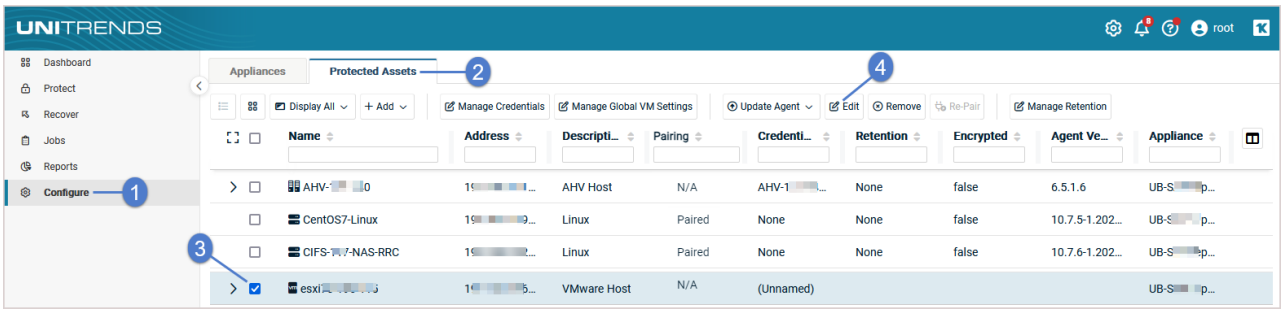

4 Modify settings and click **Save**.

Note: Quiesce settings display for VMware, AHV, and XenServer hosts only.

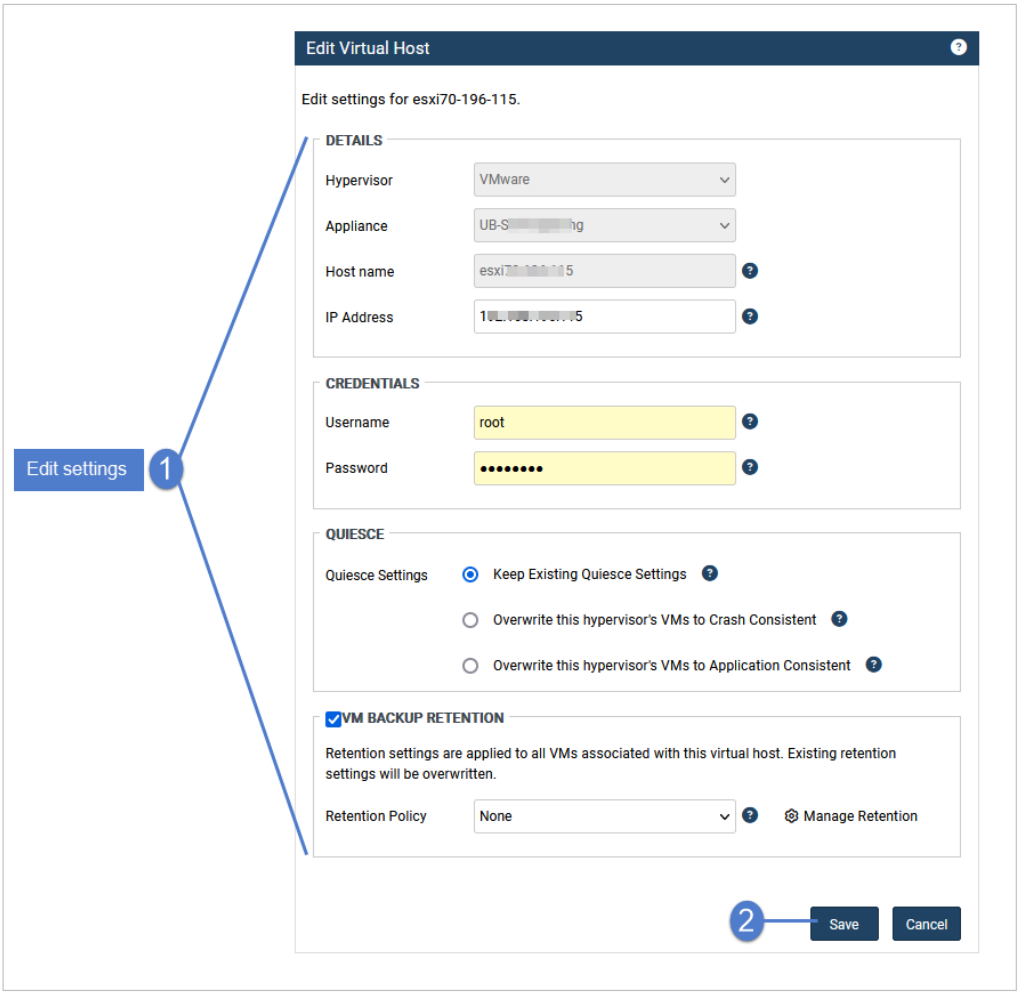

For details on these settings, see the following topics:

- "Managing asset credentials" on page 322

- "Managing retention with long-term data management" on page 328
- "Quiesce settings for host-level backups" on page 283 (applies to VMware, AHV, and XenServer VMs only)
- "To manage global quiesce settings" (applies to VMware, AHV, and XenServer VMs only)
- "To apply a quiesce setting to one host's VMs" on page 315 (applies to VMware, AHV, and XenServer hosts only)

To manage global quiesce settings

The global quiesce setting applies to all VMware, AHV, and XenServer VMs on the selected appliance. (Quiesce is not used for Hyper-V.) The setting controls how newly discovered VMs are quiesced in preparation for backup. There is also an option to overwrite the quiesce setting of existing VMs.

1 Select **Configure > Protected Assets**.

2 Click **Manage Global VM Settings**.

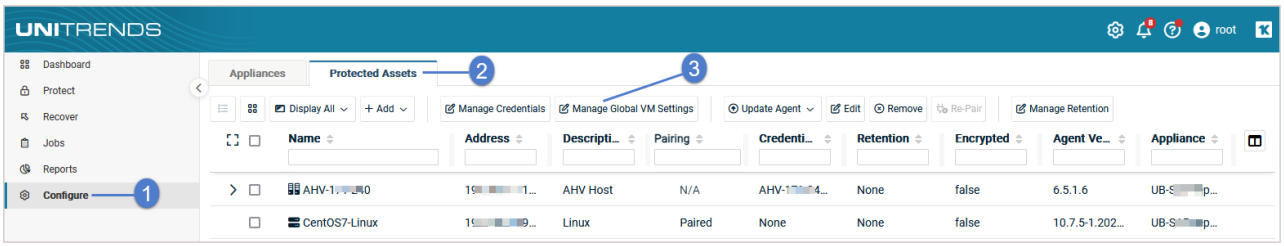

3 Select an **Appliance**.

4 Select **Crash Consistent** or **Application Consistent**.

5 (Optional) Select **Apply to all current VMs** if you want to overwrite the quiesce setting of existing VMs. Do not select this option if you want to apply the Quiesce Setting to newly discovered VMs only.

> **Note:** A VM's application aware quiesce setting is not overwritten by this procedure. A VM's quiesce setting is overwritten only if it was set to crash consistent or application consistent. For details, see "Quiesce settings for host-level backups" on page 283.

6 Click **Save**.

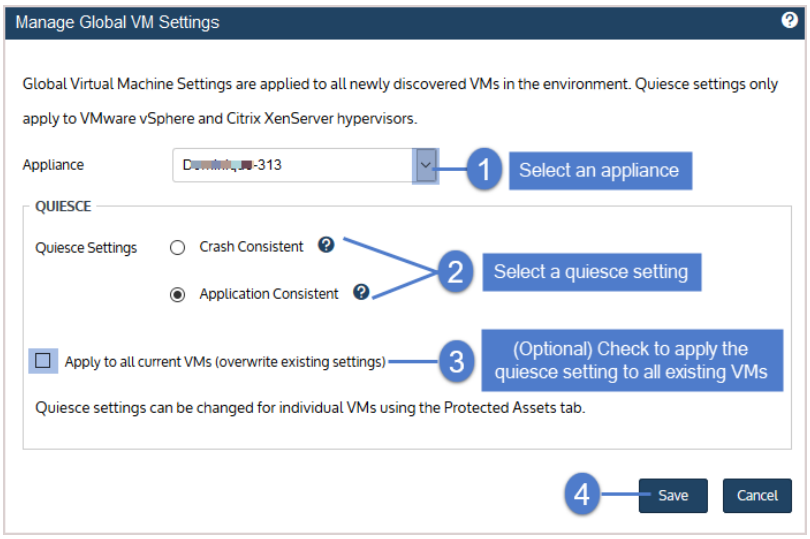

To apply a quiesce setting to one host's VMs

A virtual host's quiesce setting controls how newly discovered VMs are quiesced in preparation for backup. There is also an option to overwrite the quiesce setting of the host's existing VMs. This setting applies to VMware, AHV, and XenServer virtual hosts. (Quiesce is not used for Hyper-V.)

Use this procedure to apply the selected quiesce setting to all hosted VMs.

1 Select **Configure > Protected Assets**.

2 Select the virtual host and click **Edit**.

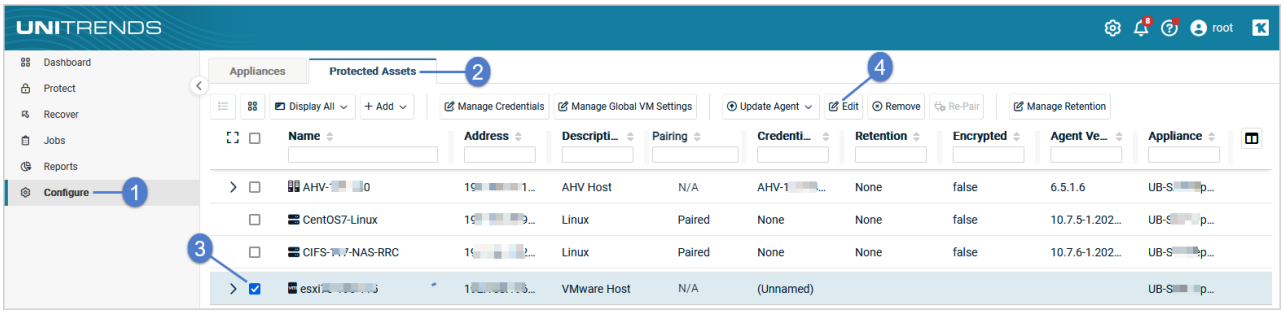

3 In the Quiesce area, select one of the following:

- Overwrite this hypervisor's VMs to Crash Consistent.
- Overwrite this hypervisor's VMs to Application Consistent.

> **Note:** A VM's application aware quiesce setting is not overwritten by this procedure. A VM's quiesce setting is overwritten only if it was set to crash consistent or application consistent. For details, see "Quiesce settings for host-level backups" on page 283.

4 Click **Save**.

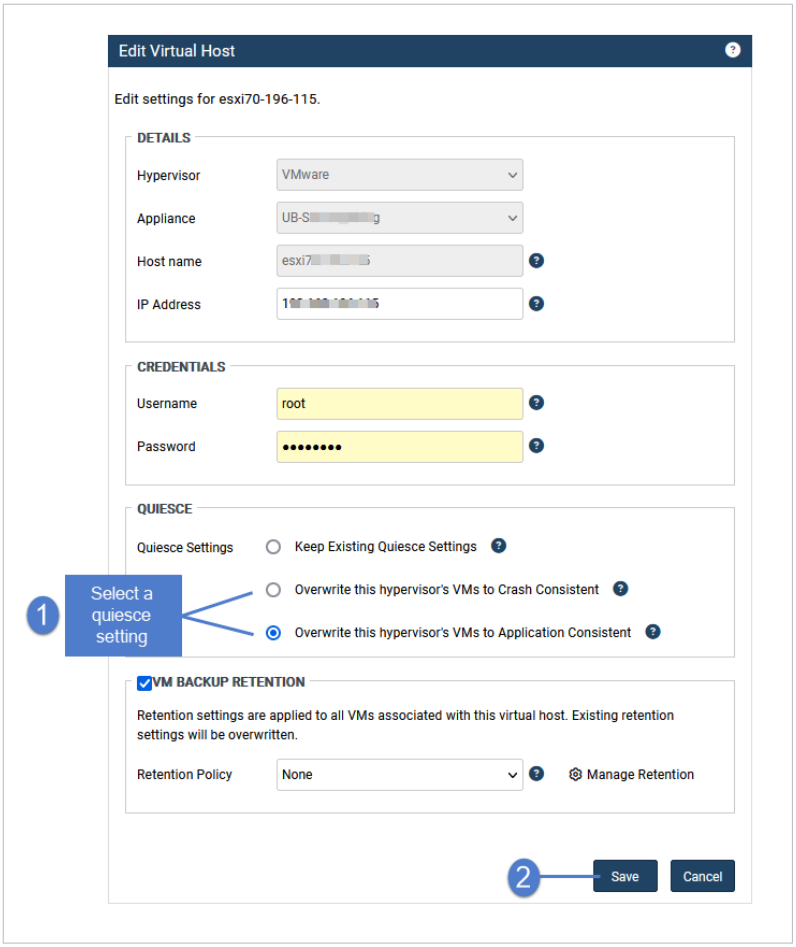

To upgrade a virtual host

Unitrends recommends upgrading virtual hosts to the latest supported version. Refer to the appropriate vendor documentation for instructions on upgrading. Note the following when upgrading:

- Your Unitrends appliance continues to protect the host with existing schedules as long as the hostname and IP address remain unchanged.
- If you change the hostname or IP address during the upgrade, update these settings in the appliance UI as described in "To edit a virtual host asset" on page 312. Existing schedules can then continue to protect the host's VMs.
- If VMs are added or removed on the host during the upgrade, refresh the VMs on the appliance to reflect the changes by selecting the **Options** icon in the top-right and clicking **Inventory Sync**.

Removing a virtual host asset

CAUTION! When a virtual host is removed, all backups of its VMs are also deleted. Please use caution when removing a virtual host asset.

Use this procedure to remove a vCenter, an ESXi host, a Hyper-V host, a XenServer host, or a Nutanix AHV cluster from the Unitrends appliance. When you remove a virtual host, all backups of its VMs are also deleted. However, if you have added a vCenter server and the ESXi hosts it's managing, the VM backups are not deleted from the appliance if you remove only the vCenter server. The backups are not deleted unless you also remove the ESXi host servers.

To remove a virtual host asset

1. Select **Configure > Protected Assets**.
2. Select the virtual host you want to remove.
3. Click **Remove**.
4. Check the **I understand...** box and click **Remove Asset and Delete Backups**.

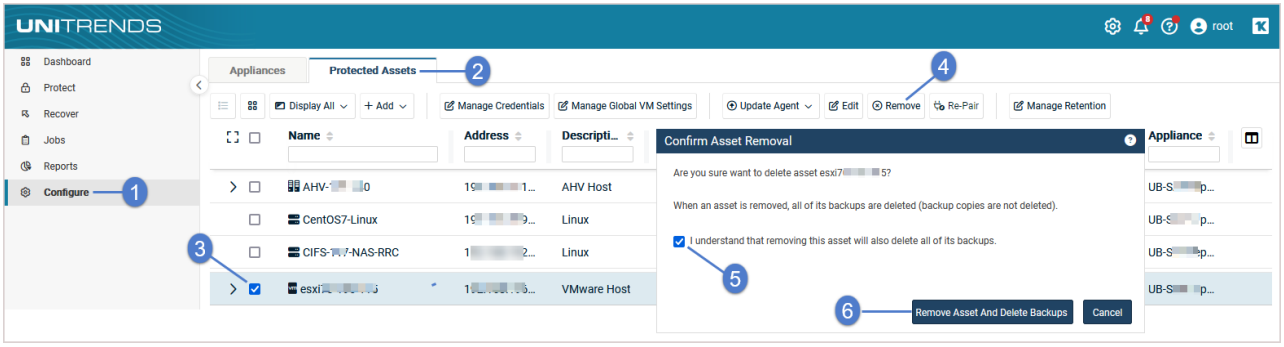

Managing virtual machine assets

Use these procedures to manage VMs you are protecting at the host level.

To add a virtual machine asset

For host-level protection of a VM, you do not add the VM itself. Instead, add its virtual host as described in "To add a virtual host asset" on page 311. All hosted VMs display under the host you have added.

To edit a virtual machine asset

Note: Because each asset can have only one retention policy, you cannot edit an asset's retention settings if the asset has been added to an SLA policy. For more on SLA policies, see "Creating SLA policies" on page 536

1. Select **Configure > Protected Assets**.
2. Click to expand the VM's virtual host to display its VMs.
3. Select the VM and click **Edit**.

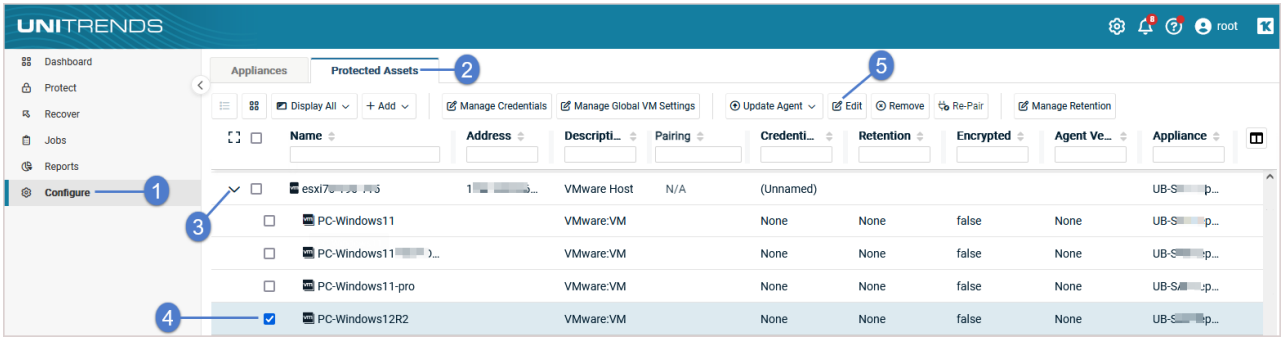

4 Modify settings and click **Save**.

Notes:

- Quiesce settings display for VMware, AHV, and XenServer VMs only.
- The Index Backups setting is supported for VMware Windows VMs only. With indexed backups you can use wild card searches to find files and folders across all backups of the virtual machine, rather than mounting and browsing each backup individually.

For details on these settings, see the following topics:

- "Managing asset credentials" on page 322
- "Managing retention with long-term data management" on page 328
- "Quiesce settings for host-level backups" on page 283 (applies to VMware, AHV, and XenServer VMs only)
- "Using application aware quiesce" on page 319 (applies to VMware Windows VMs only)
- "Indexed VMware Windows backups" on page 669 (applies to VMware Windows VMs only)

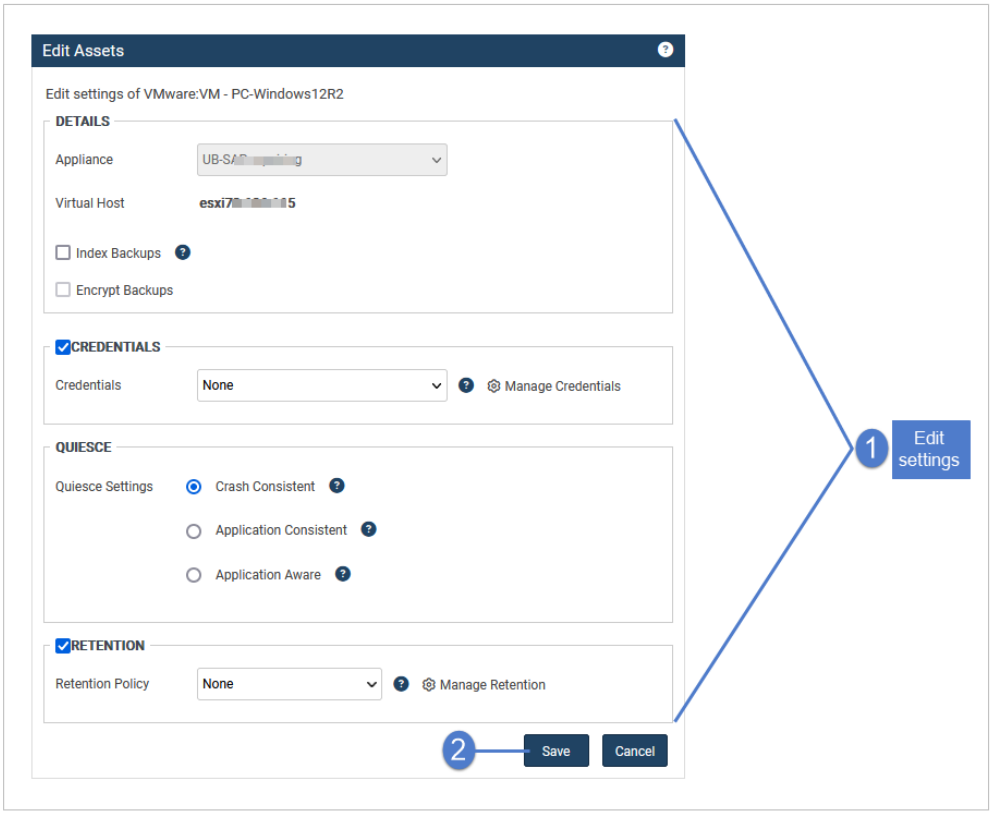

Using application aware quiesce

For VMware Windows VMs, you can opt to use the application aware quiesce setting to protect hosted Exchange and SQL simple recovery model applications. For more on how this quiesce setting works, see "Quiesce settings for host-level backups" on page 283.

Preparing for application-aware protection

Before you start, create administrative credentials for the Windows VM, as described in "To add a credential" on page 322.

To apply the application aware quiesce setting

1. Select **Configure > Protected Assets**.
2. Click to expand the VM's virtual host.
3. Select the VM and click **Edit**.

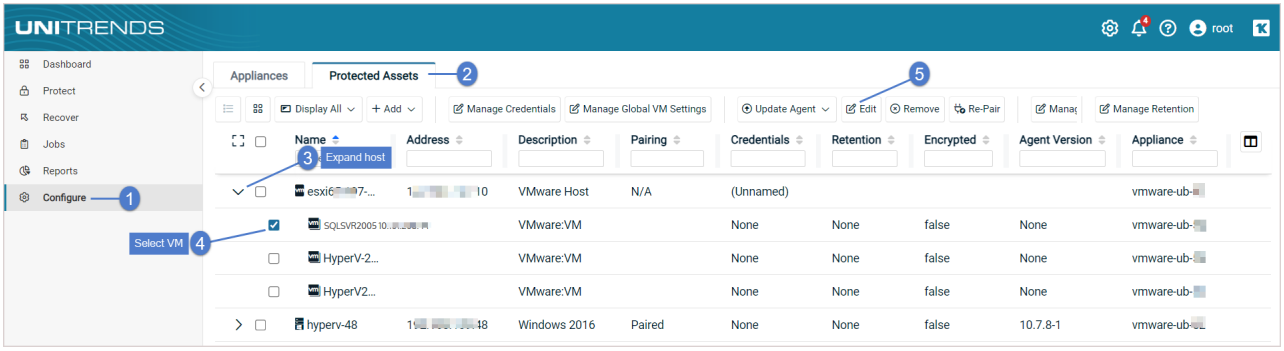

4 In the Credentials list, select the credential you created for this VM.

5 In the Quiesce area, select **Application Aware**.

6 Click **Save**.

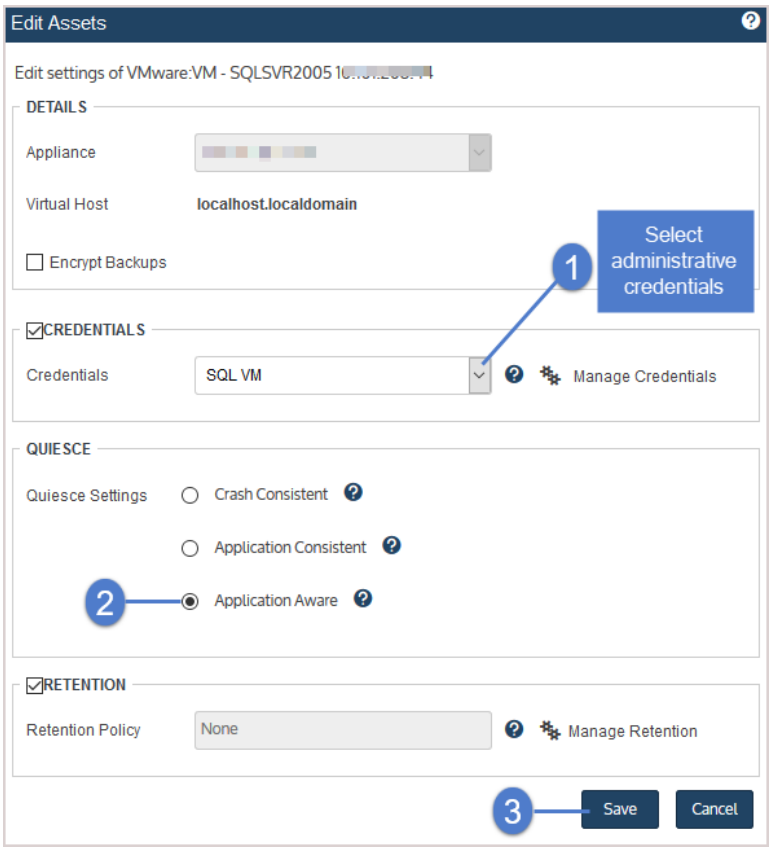

Encrypting backups

Consider the following before encrypting an asset's backups:

- Encryption must be set up on the appliance before backups can be encrypted.

- Once encryption has been set up on the appliance, selecting an **Encrypt...** setting in the procedure below causes subsequent backups to be encrypted. Any existing unencrypted backups remain unencrypted.
- Selecting a **Do not encrypt...** setting in the procedure below causes subsequent backups to be unencrypted. Any existing encrypted backups remain encrypted.

To encrypt an asset's backups

1. Set up encryption on the appliance, as described in .
2. Select **Configure > Protected Assets**.
3. Select the asset and click **Edit**.
4. Select a setting from the **Encrypt Backups** list and click **Save**.

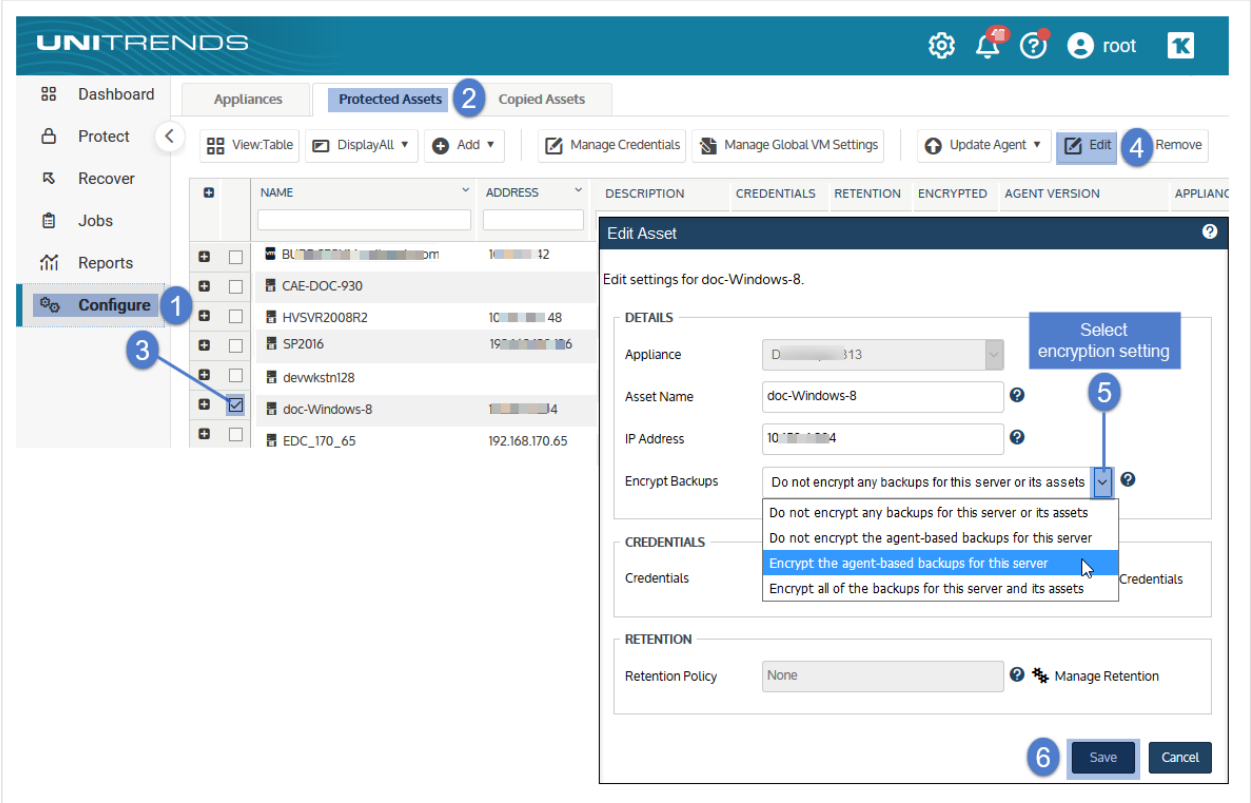

The encryption setting is applied to the asset:

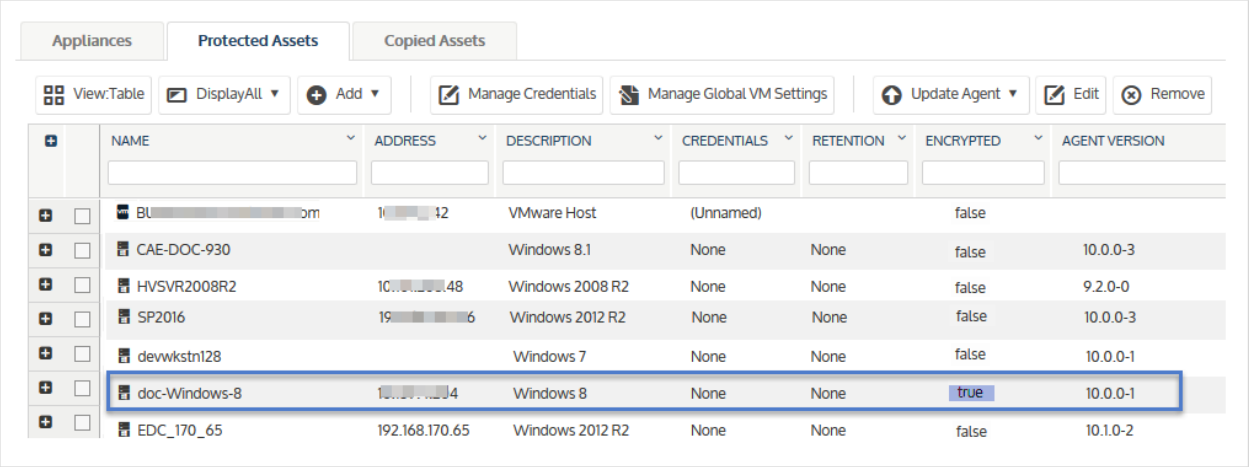

Managing asset credentials

Credentials are used to establish a trust relationship between the Unitrends appliance and its assets. For an overview of how credentials are used, see "Asset credentials" on page 282. For credential requirements for the asset you are protecting, see the applicable Backups Overview chapter in this guide.

Use these procedures to add, edit, and delete credentials, and to apply credentials to your assets (physical assets, applications, virtual hosts, and virtual machines):

- "To add a credential" on page 322
- "To view all credentials" on page 324
- "To edit a credential" on page 324
- "To remove a credential" on page 325
- "To apply a credential to an asset" on page 326
- "Managing protected assets" on page 286

To add a credential

Use this procedure to add a credential. Once added, you can apply it to any protected asset.

Note: You can also add a credential while editing an asset. See the applicable procedure for details: "To edit an agent-based asset" on page 293, "To edit a NAS asset" on page 299, "To edit a UCS manager asset" on page 303, or "To edit a virtual machine asset" on page 317.

1 Click **Configure > Protected Assets > Manage Credentials**.

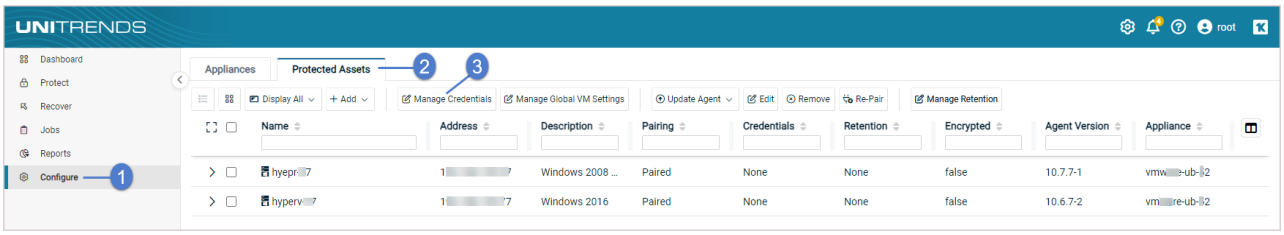

2 Click **Add**.

3 Enter credential information and click **Save**.

Note: Checking the **Set as default** box adds the text *(Default)* to the credential name. The default credential is not automatically applied to any asset.

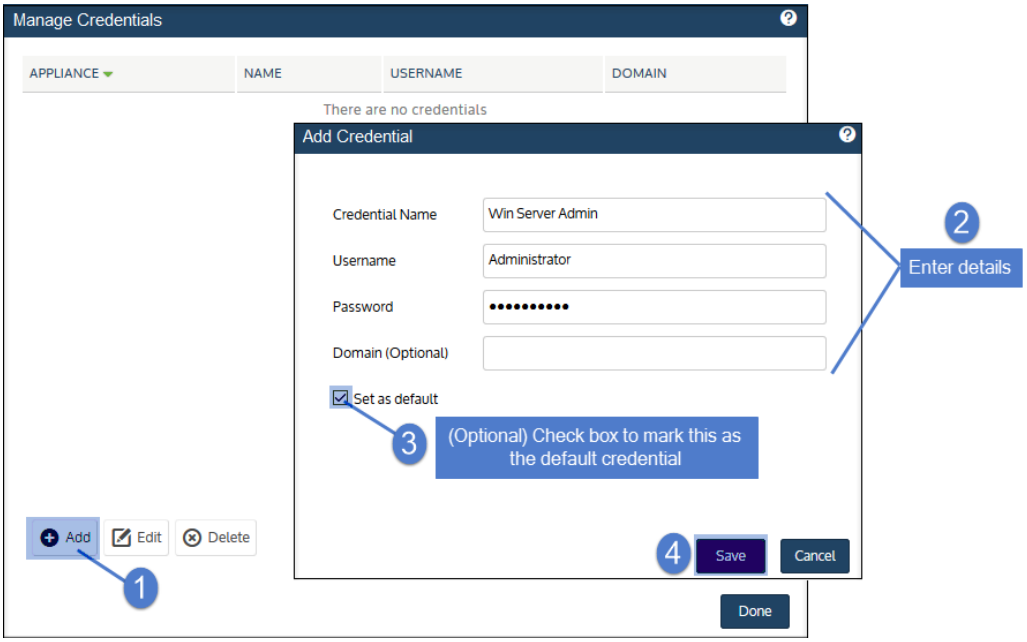

4 Click **Done** to exit.

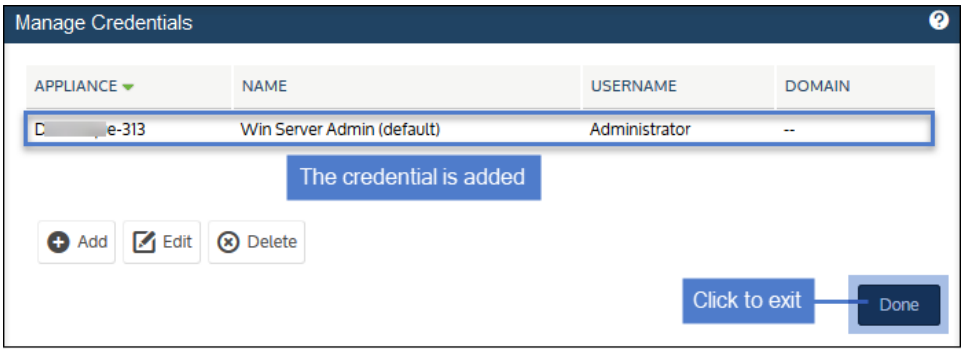

This credential can now be applied to your assets. (For details, see "To apply a credential to an asset" on page 326.)

To view all credentials

1 Click **Configure > Protected Assets > Manage Credentials**.

All credentials defined for this appliance display.

2 Click **Done** to exit.

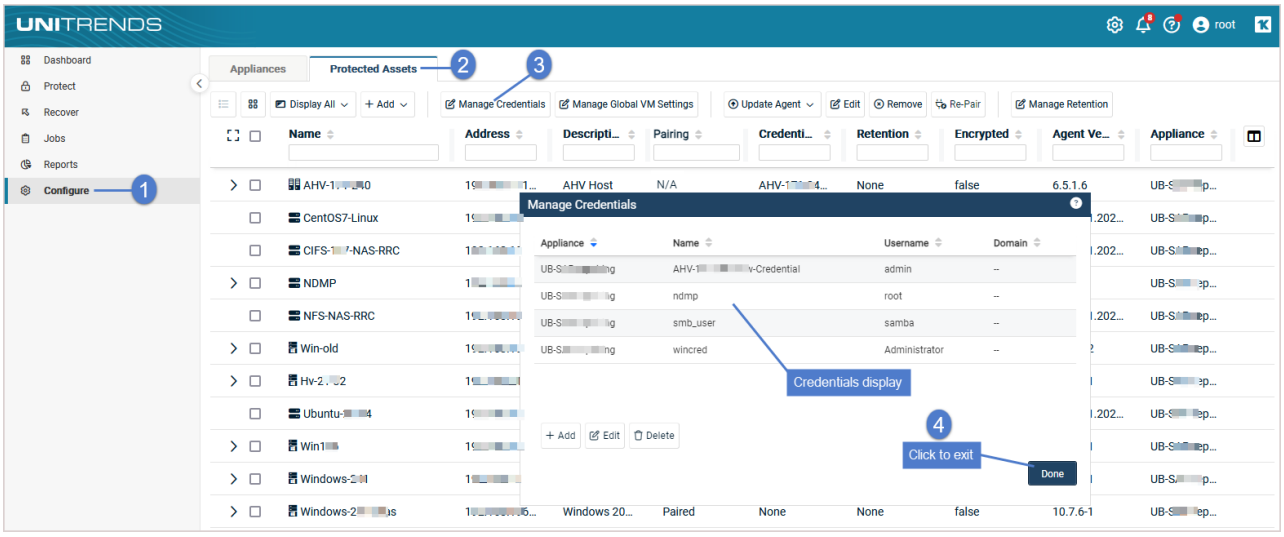

To edit a credential

1 Click **Configure > Protected Assets > Manage Credentials**.

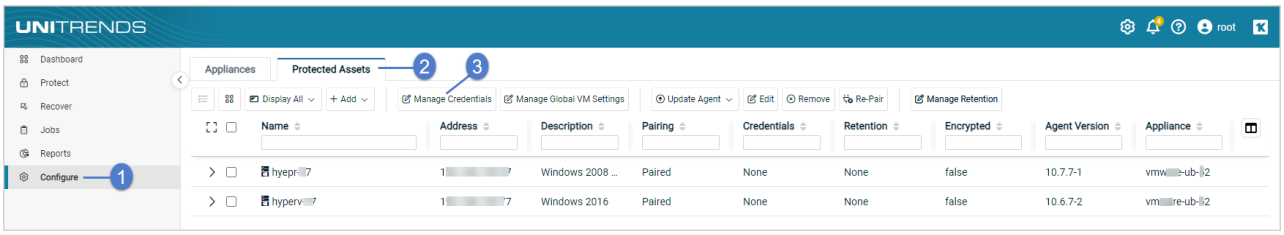

2 Select a credential and click **Edit**.

3 Modify information and click **Save**.

Note: Checking the **Set as default** box adds the text *(Default)* to the credential name. The default credential is not automatically applied to any asset.

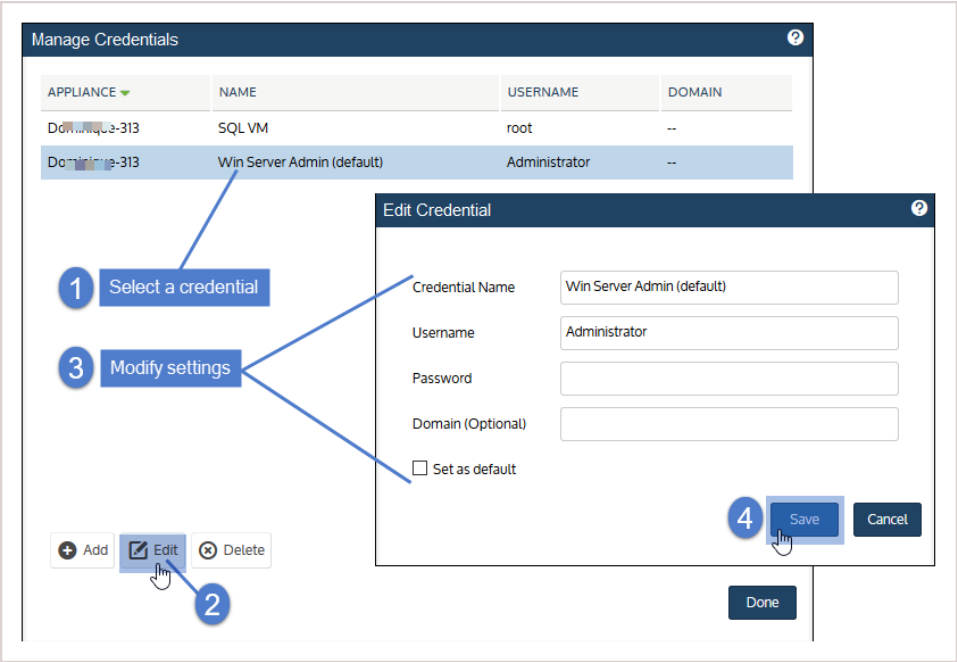

4 Click **Done** to exit.

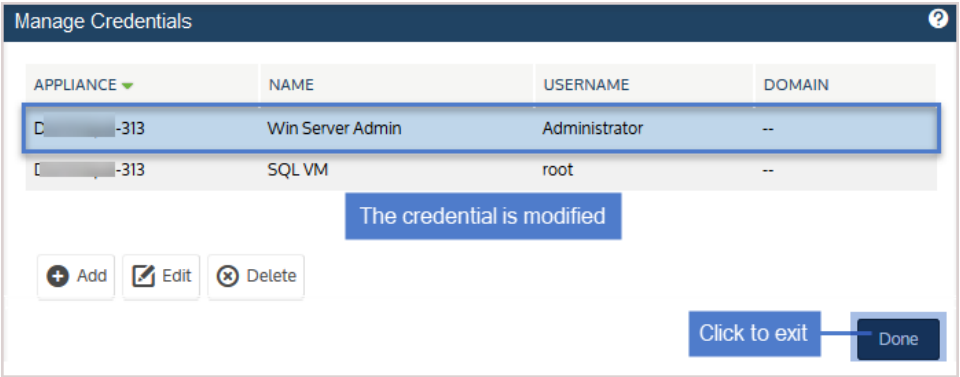

To remove a credential

You cannot remove a credential that is being used by an asset. Before removing a credential, remove it from any assets using Edit Asset.

1 Click **Configure > Protected Assets > Manage Credentials**.

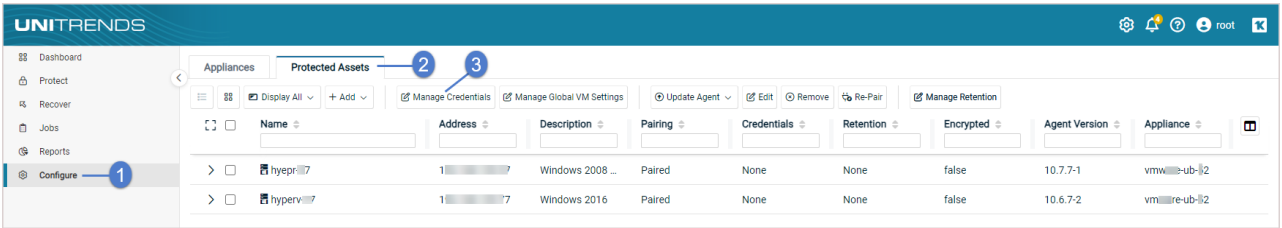

2 Select the credential, click **Delete**, then **Confirm**.

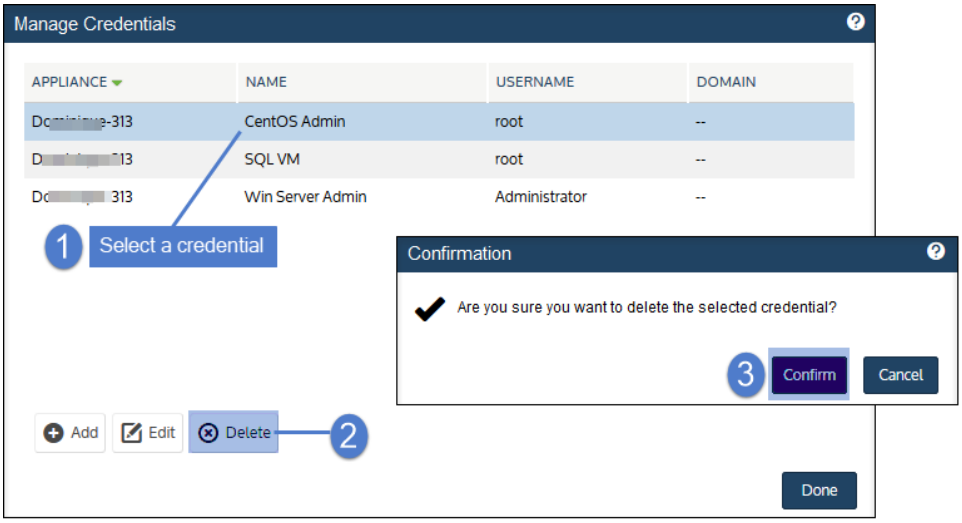

3 Click **Done** to exit.

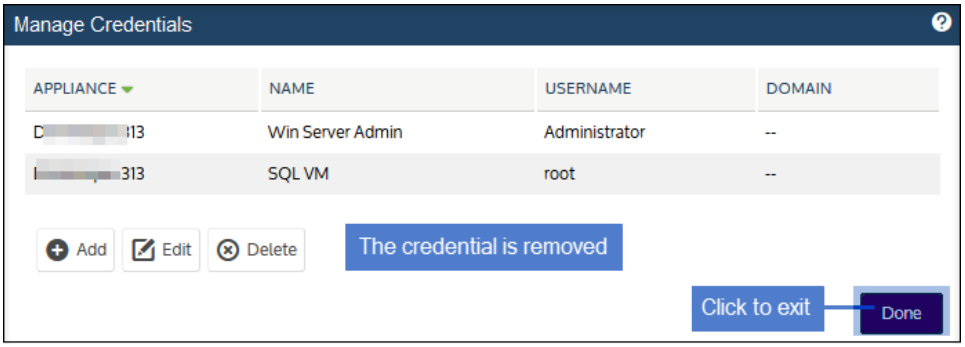

To apply a credential to an asset

1 Select **Configure > Protected Assets**.

2 Select the asset and click **Edit**.

> **Note:** To apply the credential to a VM or application, expand the host asset to view and select the VM, application instance, or database.

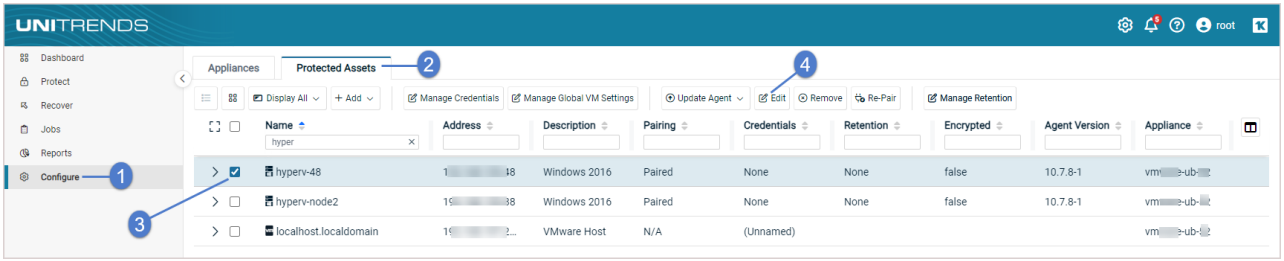

3 (If needed) Do these steps to create the credential (if using an existing credential, skip this step):

- Click **Manage Credentials > Add**.
- Enter details and click **Save**.

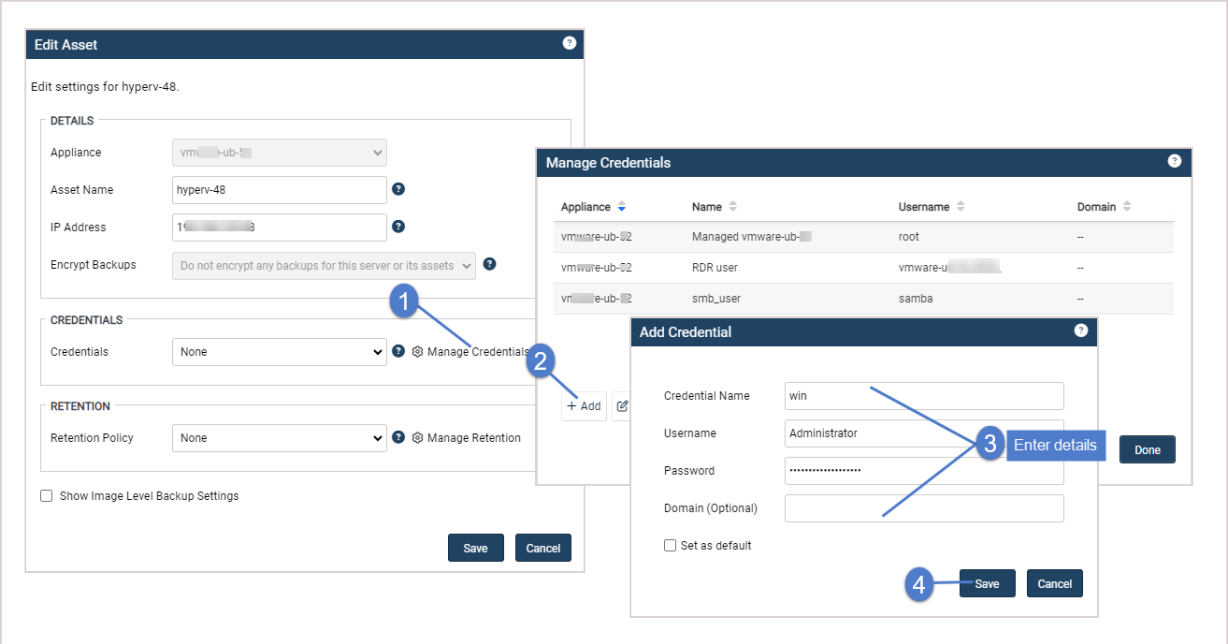

- The credential is added. Click **Done**.

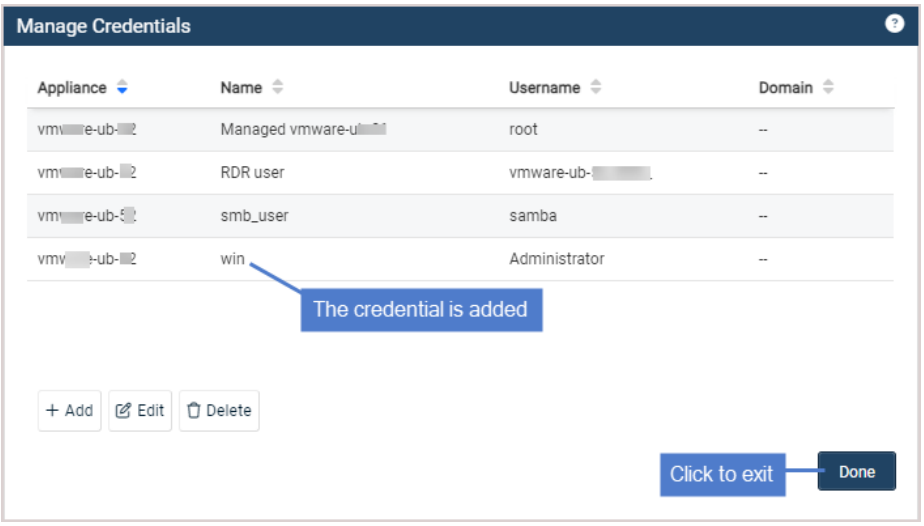

4 Select the credential from list and click **Save**.

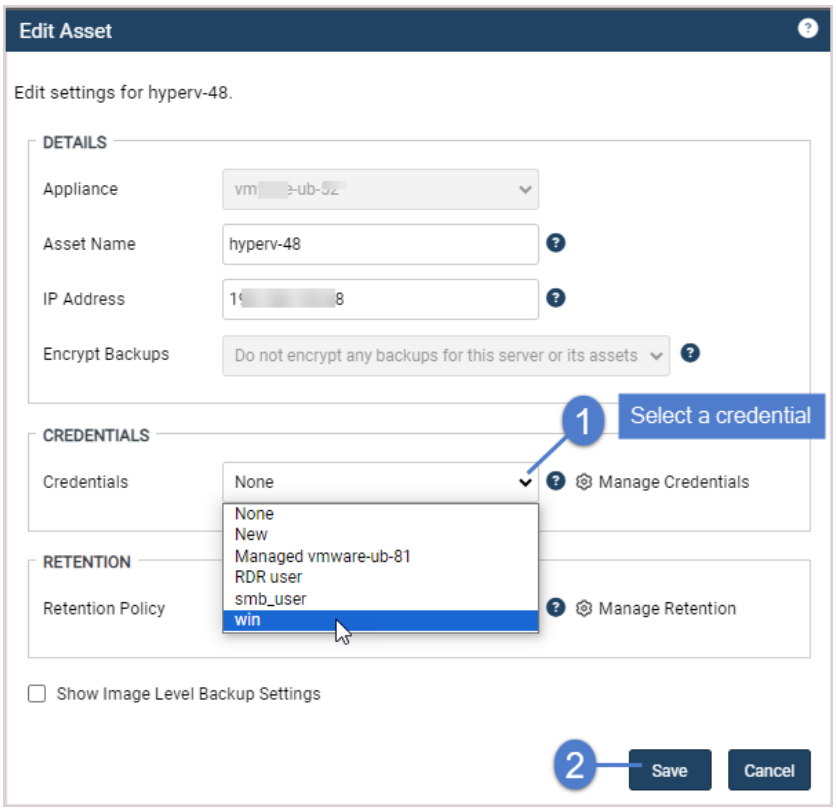

The credential is applied to the asset:

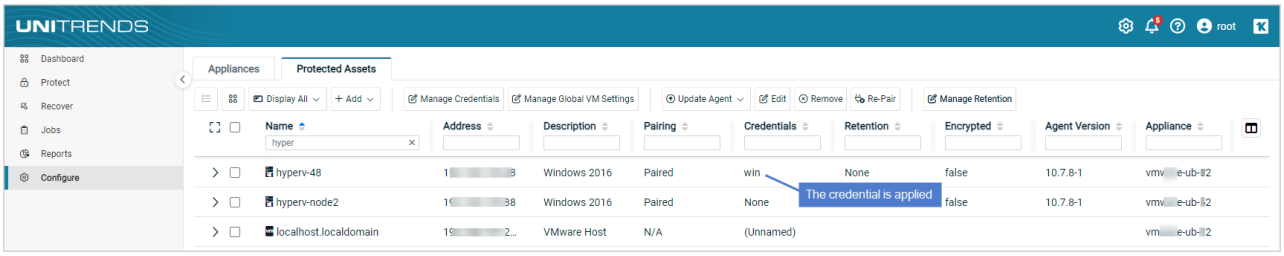

Managing retention with long-term data management

Retention settings assure that the necessary recovery points are available on your appliance. Appliances that were originally imaged with release 10.7.8 or higher are configured with a default backup retention policy of 30 days. This 30-day policy is applied to each protected asset. You can also create your own custom policies to quickly apply your own retention settings to assets.

Note: Long-term retention is not enabled on appliances that have been upgraded from a pre-10.3 release. If your appliance was originally imaged with a pre-10.3 Unitrends release, see "Switching to long-term retention" on page 335 to enable this feature.

Default retention policy

Appliances are configured with a default, pre-loaded backup retention policy of 30 days. The 30-day default retention policy is applied to each protected asset to ensure that 7 daily backups and 4 weekly backups are retained.

Notes:

- The 30-day default retention policy applies to appliances imaged with release 10.7.8 or higher. This default policy does not apply to appliances that were originally imaged with an earlier release. Release 10.7.8 was generally available on September 6th, 2023.
- For appliances imaged with release 10.7.11 or higher, the 30-day default retention policy enables the appliance to purge any backup that is no longer held by the retention policy. If needed, you can opt to retain an asset's last available recovery point by unchecking the **Delete Final Backup** box in the Edit Retention Policy dialog. Release 10.7.11 was generally available on December 13th, 2023.
- Upgrading an appliance does not modify its retention policies in any way.
- Holds applied to individual backups override the asset's retention policy. Once the backup is no longer on hold, it is eligible for deletion by the policy. (For more on holding backups, see "Placing backups on hold" on page 634.)

The default retention policy is shown below. To use different retention settings, you can do any of the following:

- Modify the settings of the default *Pre-loaded retention policy* as shown below.
- Create a new custom policy and designate it as the default policy, as described in "To add a long-term retention policy to an appliance" on page 333.

  Notes:

  - An appliance cannot have more than one default policy at a time. Edit the existing default policy to clear the **Set as default policy** box before designating a new default policy.
  - The newly designated default policy is applied to any new assets and copied assets that are added to the appliance. If needed, edit any existing assets to apply the newly designated default policy.
- Add multiple custom policies to apply different settings to groups of assets. Once a custom policy is applied to an asset, it overrides the default policy.

To view or edit the default retention policy, go to **Configure > Protected Assets**, click **Manage Retention**, select **Pre-loaded retention policy**, and click **Edit**, as shown here:

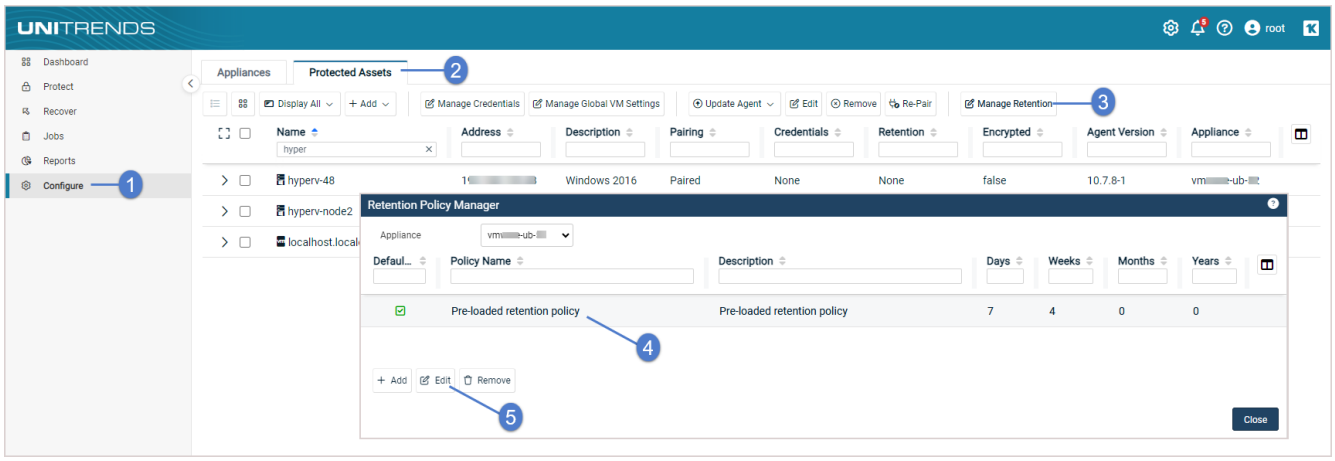

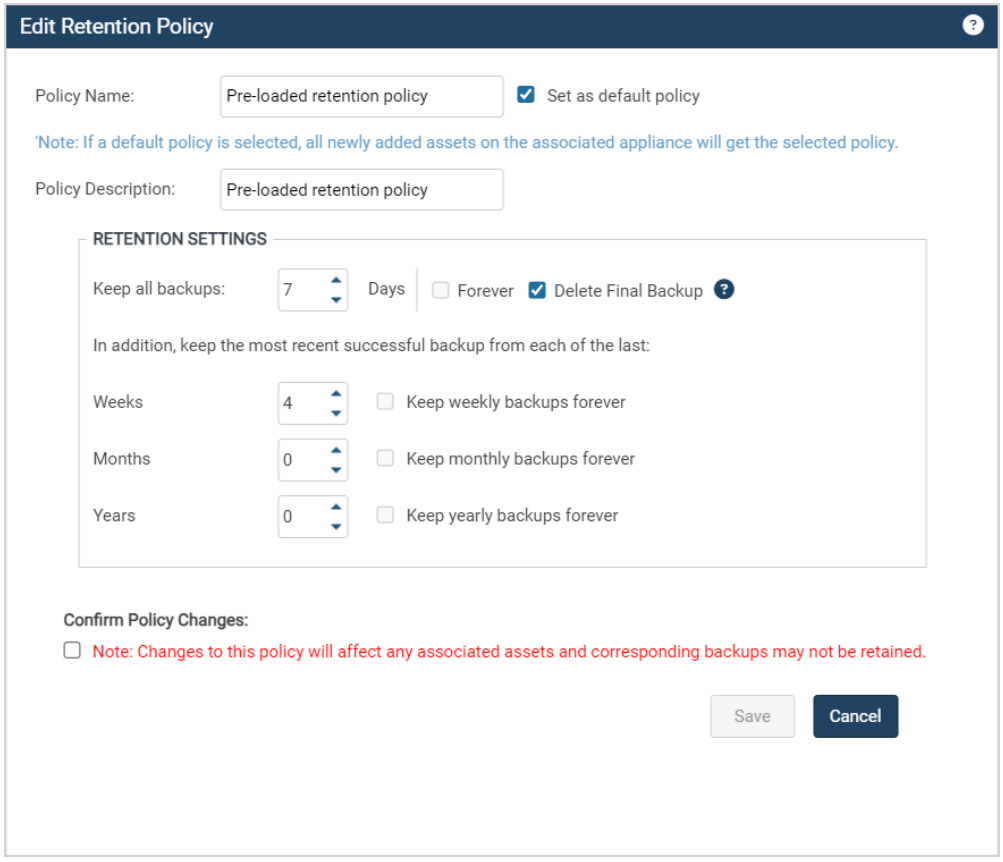

Custom policies

To use your own custom retention settings, you can quickly create policies that hold backups for a specified number of days. You can create multiple policies and customize them to achieve different RPOs and RTOs for your assets. Each policy you create can be applied to multiple assets. These policies automatically retain and purge backups as necessary to maintain a customized inventory of daily, weekly, monthly, and yearly retention points.

To start managing retention with custom policies:

Step 1: Review the "Long-term retention policy settings"

Step 2: Add a policy as described in "To add a long-term retention policy to an appliance" on page 333

Step 3: Apply the policy to assets as described in "To apply a long-term retention policy to a protected asset" on page 334

Long-term retention policy settings

Long-term retention policies are configured with the following settings:

Retention setting	Description
Policy Name	Enter a name for the policy.
Set as default policy	Check the **Set as default policy** box on the Add Retention Policy or Edit Retention Policy dialog to designate the policy as the default policy of the appliance. Any new assets and new copied assets that are added to the appliance receive the default policy. **Notes:** • An appliance cannot have more than one default policy at a time. If needed, edit the existing default policy to clear the **Set as default policy** box before designating a new default policy. • Applying a non-default policy to an asset overrides the default policy. • If you have designated a new default policy, you can apply it to existing assets as described in "To apply a long-term retention policy to a protected asset" on page 334. • For copied assets, you can use the **Set as global policy** option to override the default policy. Use this method to have different default policies for local assets versus copied assets.
Set as global policy	The Configure > Copied Assets tab displays only for appliances that are receiving backup copies from another Unitrends appliance. The tab lists all assets whose backup copies are stored on this appliance. If your appliance is receiving backup copies from another appliance, you can create or edit the policies used for copied assets by going to **Configure > Copied Assets** and clicking **Manage Retention**. When adding or editing a policy for copied assets, you can check the **Set as global policy** box to designate the policy as the global policy for copied assets. Any newly added backup copy source appliance receives this policy, which determines how long its hot copies are retained. As copied assets are added for the new source appliance, this global policy is automatically applied. **Notes:**

Retention setting	Description
	• A target appliance cannot have more than one global policy at a time. • Global policies are applicable to any new copied assets that are added to the appliance. (Global policies do not apply to local assets.) • The appliance's default policy applies to any copied assets unless you use the **Set as global policy** box to override the default policy for copied assets.
Policy description	Enter a brief description of the policy. This can be a note regarding the policy's intended use case or a brief summary of its retention settings.
Days	All backups from the last *N* number of days are retained. Days end at midnight, appliance time. Or check the **Forever** box to retain all backups. Check the **Delete Final Backup** box to automatically delete data when it reaches the end of its retention policy, even if it's the last available recovery point. Leave this box unchecked to retain the last available recovery point.
Weeks	The most recent successful backup from each of the last *N* number of weeks is retained. Weeks end on Sunday of the calendar week. Or check the **Forever** box to keep weekly backups forever.
Months	The most recent successful backup from each of the last *N* number of months is retained. Months end on the final day of the calendar month. Or check the **Forever** box to keep monthly backups forever.
Years	The most recent successful backup from each of the last *N* number of years is retained. Years end on the final day of the calendar year. Or check the **Forever** box to keep yearly backups forever.

Notes:

- If the most recent successful backup is an incremental or differential, the complete backup group is retained.
- Redundant backups are not retained if high-frequency retention intervals are configured to overlap their low-frequency counterparts. For example, a retention policy specifying 52 weeks and 12 months *does not* retain a total of 52 + 12 backups for the preceding 12-month period.
- An asset must have at least one successful backup for each retention point to be considered compliant.
- A yellow alert displays in the UI if any asset is, for any reason, not compliant with its associated long-term retention policy. This alert is automatically dismissed if compliance is achieved.

- Retention compliance is not achievable if the asset's scheduled backups are less frequent than the policy's most frequent retention interval. For example, a retention policy specifying 4 weeks and 0 days cannot maintain a complete inventory of compliant retention points for an asset that is backed-up on a bi-weekly basis.
- Imported backup copies are retained for a default period of 72 hours regardless of retention settings applied to the asset from which they originated.

To add a long-term retention policy to an appliance

1. Log in to the backup appliance and select **Configure > Protected Assets**.
2. Click **Manage Retention**.
3. Click **Add**.

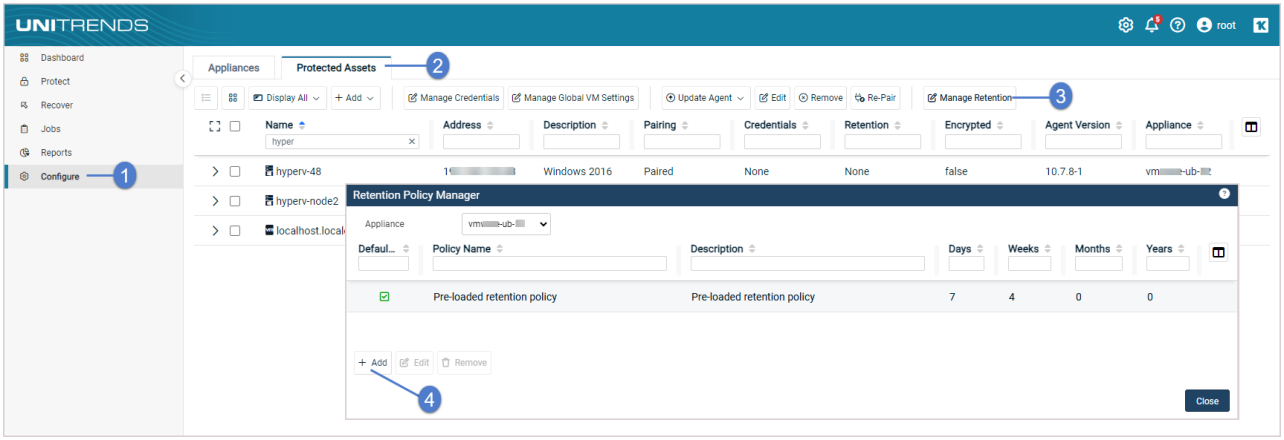

4. Enter the retention policy settings. For details, see "Long-term retention policy settings" on page 331.
5. Click **Add**. The long-term retention policy is created.

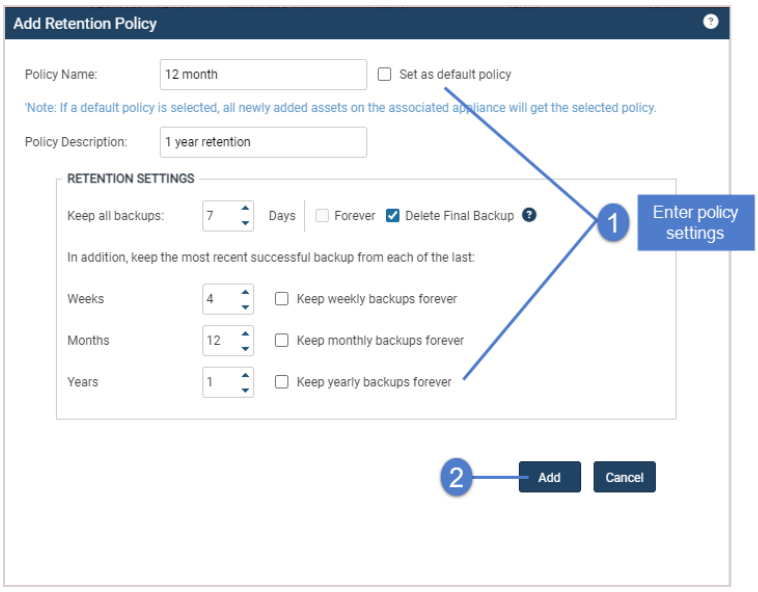

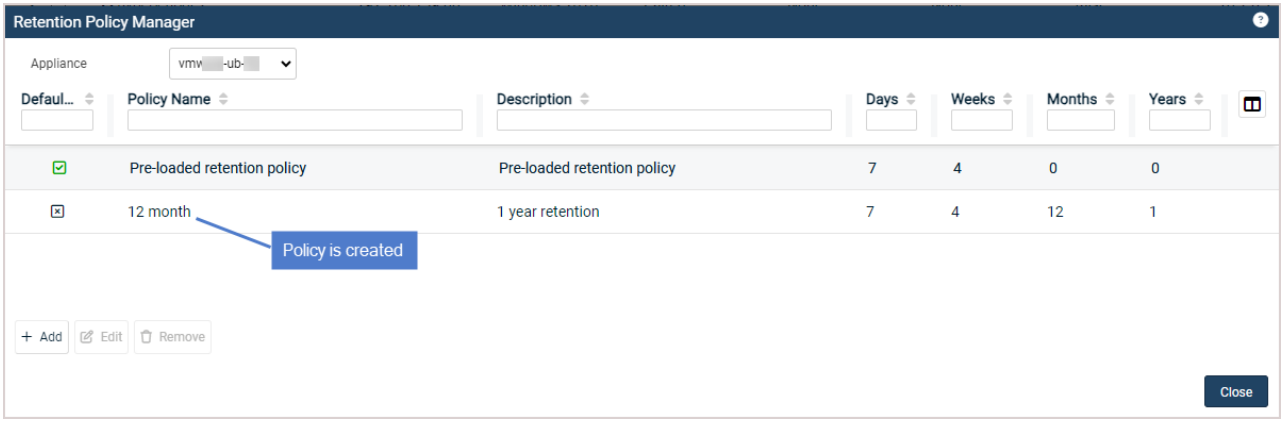

To apply a long-term retention policy to a protected asset

1. Log in to the backup appliance and select **Configure > Protected Assets**.
2. Select the desired asset.
3. Click **Edit**.
4. Select a policy from the **Retention Policy** dropdown.

   Note: If the selected policy will delete one or more of the asset's backups, a dialog listing these backups displays. Click **Close** to continue.

5. Click **Save**. The long-term retention policy is applied.

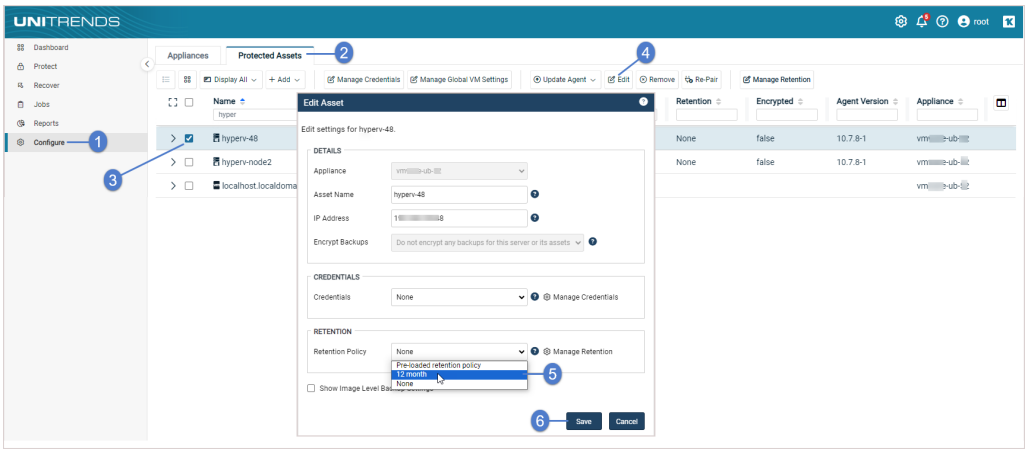

To edit a long-term retention policy

1. Log in to the backup appliance and select **Configure > Protected Assets**.
2. Click **Manage Retention**.
3. Select the policy you wish to edit.
4. Click **Edit**.

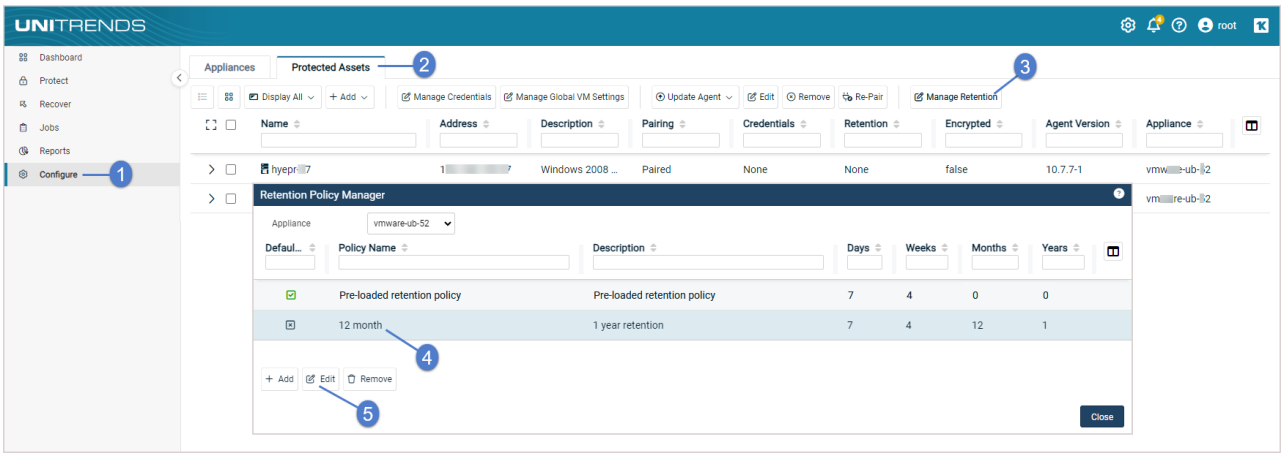

5 Modify the retention policy settings. For details, see "Long-term retention policy settings" on page 331.

6 Select the **Note** checkbox.

7 Click **Save**. Changes to the policy are committed.

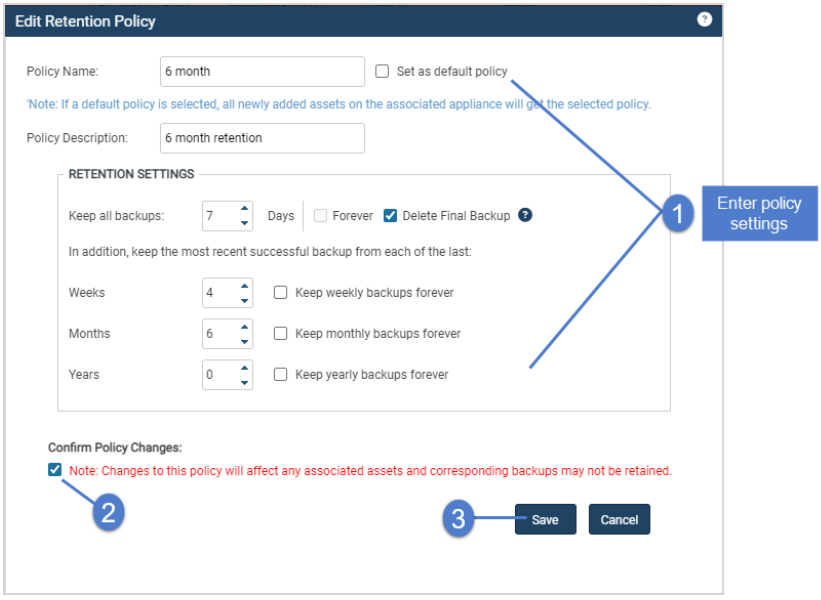

Switching to long-term retention

If your appliance was originally imaged with a pre-10.3 Unitrends release, the appliance utilizes legacy asset-level retention until the switch to long-term retention is manually performed. Long-term retention policies provide a more precise, granular, and space-efficient alternative to legacy asset-level retention settings. A long-term retention policy automatically retains and purges backups as necessary to maintain a customized inventory of weekly, monthly, and yearly retention points. A Unitrends appliance can maintain multiple long-term retention policies that can each be applied to multiple assets.

Note: Long-term retention is enabled by default on appliances that were originally imaged with release 10.3 or later. Do not run the "To switch to long-term retention" procedure if your appliance was originally imaged with release 10.3 or later.

To switch to long-term retention, review the "Considerations for switching to long-term retention", then run the "To switch to long-term retention" procedure.

Once you have switched to long-term retention, set up retention policies as described in "Managing retention with long-term data management" on page 328.

Considerations for switching to long-term retention

Migrating from legacy asset-level retention to long-term retention impacts a number of Unitrends features, which are detailed in the table below. Carefully review the considerations in the table before switching to long-term retention.

WARNING! Once the switch to long-term retention is performed, it cannot be reversed from the web UI.

Feature	Details
Legacy asset-level retention settings	The switch to long-term retention nominally voids all legacy asset-level retention settings; however, any backups held in accordance with these settings remain preserved in a legal hold state. When applying a long-term retention policy to an asset that previously used legacy asset-level retention settings, the *keep all backups* value of this new policy must be equal to or greater than the *keep backups for* value specified in the original retention settings.
SLA policies	Following the switch to long-term retention, all SLA policy retention settings are voided and the SLA policy feature loses its retention functionality. An asset can be assigned both a long-term retention policy and an SLA policy.
GFS	If GFS is enabled on your appliance, the switch to long-term retention directly translates your GFS retention settings into a default long-term retention policy titled *Migrated policy*. If you are using legacy asset-level retention policies in addition to GFS retention settings, ensure that these policies specify backup hold periods shorter than or equal to the *daily* GFS setting. Switching to long-term retention is not permitted unless this condition is satisfied. **Notes:** • On target appliances, the migrated policy is both global and default. • The *weeks* setting of the migrated policy is set to 4 by default.
Backup-level holds	Holds applied to individual backups in the Backup Catalog are not impacted by the switch to long-term retention.
Cold backup copy retention	Job-level cold backup copy retention settings are not impacted by the switch to long-term retention.

To switch to long-term retention

Appliances upgraded from a pre-10.3 release do not have access to long-term data management functionality until the switch to long-term retention is manually performed. Long-term retention is enabled by default on all appliances that were originally imaged with version 10.3 or later.

CAUTION! Before proceeding, ensure that you have reviewed the "Considerations for switching to long-term retention" on page 336 and understand how your assets may be impacted.

1. Log in to the backup appliance and select **Configure > Protected Assets**.
2. Select **Manage Retention**.
3. Select **Switch to Long Term Retention**.
4. Click **Save**.
5. If you are sure you want to switch to long-term retention, click **Confirm**.

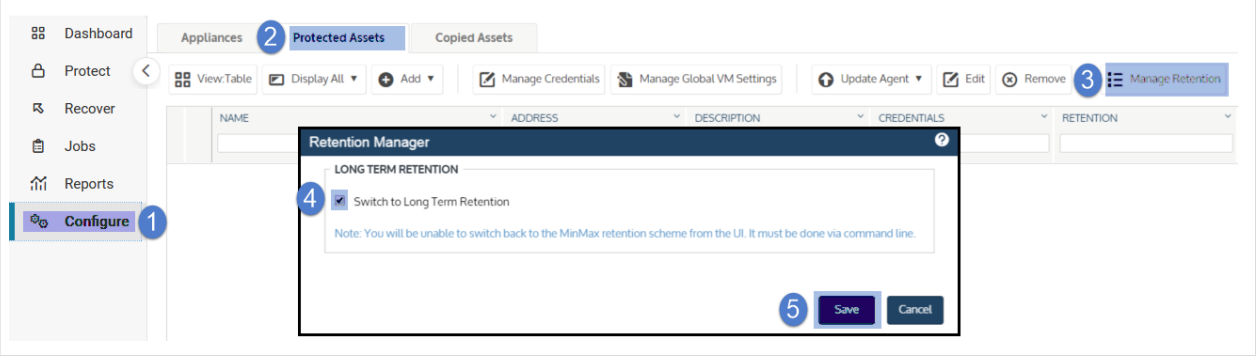

Managing retention with legacy asset-level retention settings

Retention settings are used to control how long backups are retained on the appliance. See the table below for a description of each setting. See "To apply retention settings to one asset" on page 338 to apply settings to individual assets.

Notes:

- You can apply retention settings to individual assets or to SLA policies. Because each asset can have only one retention policy, you cannot edit an asset's retention settings if the asset has been added to an SLA policy. For more on SLA policies, see "Creating SLA policies" on page 536.
- This retention scheme is not available on appliances that have switched to long-term retention.
- Sharepoint, Oracle, and Cisco UCS assets are not compatible with this retention scheme. To customize retention options for these asset types, see "Managing retention with long-term data management" on page 328

Retention setting	Description
Minimum Retention	Minimum retention settings.
Keep backups for *N* days	Number of days backups must be retained. Backups that are younger than *N* days are not purged for any reason, including at the expense of new, incoming backups. The age of a backup is determined by the last backup in the group, e.g., the last incremental before a new full.
Warn when less than *N* days of backups remain	Use this option to receive an email notification if this asset has less than *N* days of backups stored on the appliance.
Maximum Retention	Maximum retention setting.
Delete backups after *N* Days	Number of days after which the appliance will delete backups. Backups are eligible to be deleted once the full has exceeded this limit. At this point, the full and all associated incrementals and differentials in the group are deleted.

To apply retention settings to one asset

Use this procedure to apply retention settings to one asset.

Note: Because each asset can have only one retention policy, you cannot edit an asset's retention settings if the asset has been added to an SLA policy. For more on SLA policies, see "Creating SLA policies" on page 536

1. Select **Configure > Protected Assets**.
2. Select the desired asset.
3. Click **Edit**.
4. Click **Manage Retention**.
5. Define settings as desired and click **Save**.

Secure agent pairing for Windows and Linux agents

Beginning in Windows agent release 10.6.6 and Linux agent release 10.7.5, a secure pairing is automatically established between the appliance and the Windows or Linux agent on each of its protected assets. This pairing enables Transport Layer Security (TLS) to encrypt data and authenticate connections between appliances and agents. Communication between appliances and agents is only allowed if there is a matching (paired) certificate. (For pairing requirements, see "Windows agent requirements" on page 362 or "Requirements for secure pairing of Unitrends Linux agents" on page 391.)

The **Configure > Protected Assets** tab enables you to view the pairing status of each protected asset and to re-pair an asset to its appliance.

Re-pairing requires the asset to be in pairing mode. For Windows assets, you can use the Unitrends Agent Pairing Utility to enable pairing mode for a specified amount of time, known as the *pairing window*. For Linux assets, you will use instructions in the Unitrends Agent Pairing - Linux KB article to enable pairing mode.

See these topics for details on working with secure agent pairing:

- "Agent pairing statuses"
- "Updating the asset's pairing status on the Configure > Assets page"
- "To enable pairing mode on a Linux asset"
- "To enable pairing mode on a Windows asset"
- "To extend pairing mode on a Windows asset"
- "To re-pair an asset"
- "To disable pairing mode on a Windows asset"

Agent pairing statuses

On the **Configure > Protected Assets** tab, the Pairing column shows the asset's secure agent pairing status:

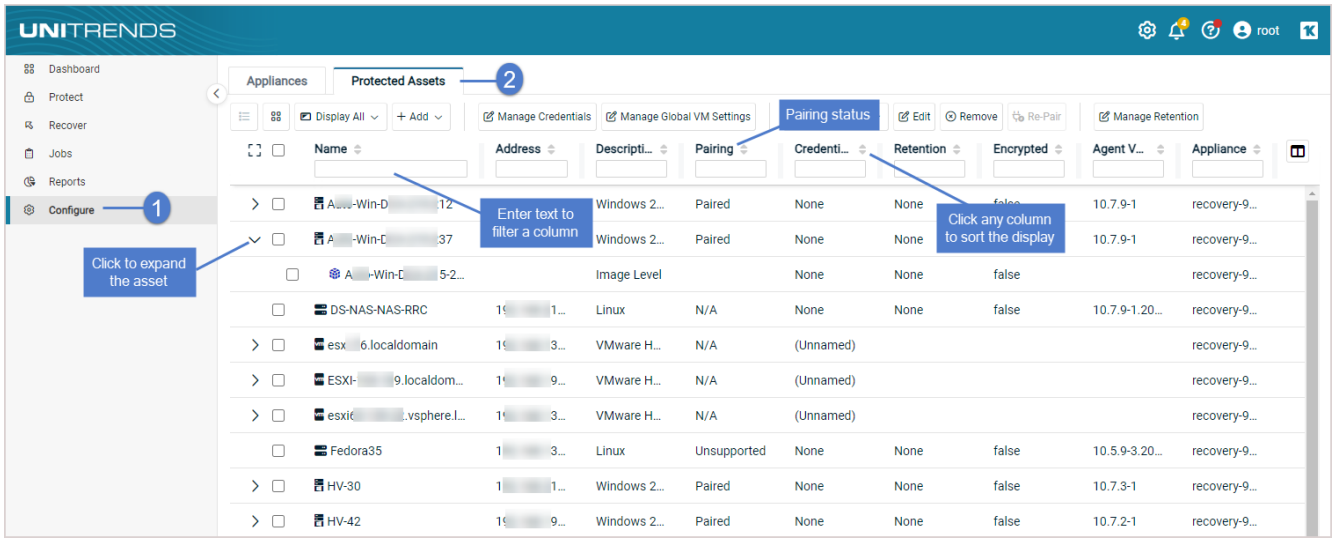

Pairing statuses include:

- Paired – The agent has been paired.
- Failed – The agent is not paired. Hover to see the error message. One of the following errors has occurred and you must re-pair the asset to its appliance:
  - The agent pairing time window has expired.
  - The agent can't save pairing keys.

- The agent pairing request failed.

  To re-pair the asset, place the asset in pairing mode as described in "To enable pairing mode on a Linux asset" or "To enable pairing mode on a Windows asset", then run the "To re-pair an asset" procedure.

- Unsupported – Pairing is disabled or the agent pairing version is not compatible.
- N/A – Not applicable for this asset.
- Remote – This asset resides on a managed appliance. Log in to its appliance directly to see the asset's pairing status.

Updating the asset's pairing status on the Configure > Assets page

The asset's pairing status is updated when a backup runs, during an inventory sync, or any time the asset is re-saved.

- To manually update the pairing status of all applicable assets, click ⚙ and select **Inventory Sync**.

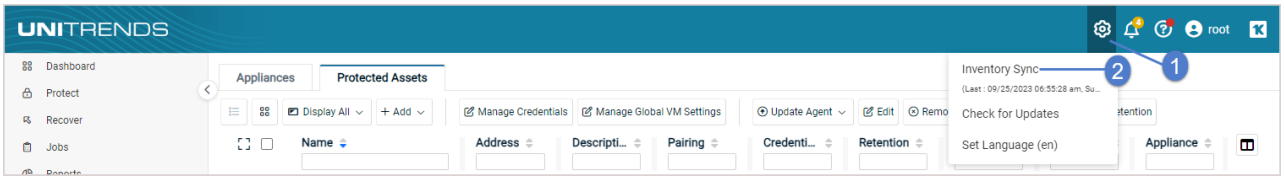

- To manually update the pairing status of one asset, select the asset, click **Edit**, then click **Save** in the Edit Asset dialog.

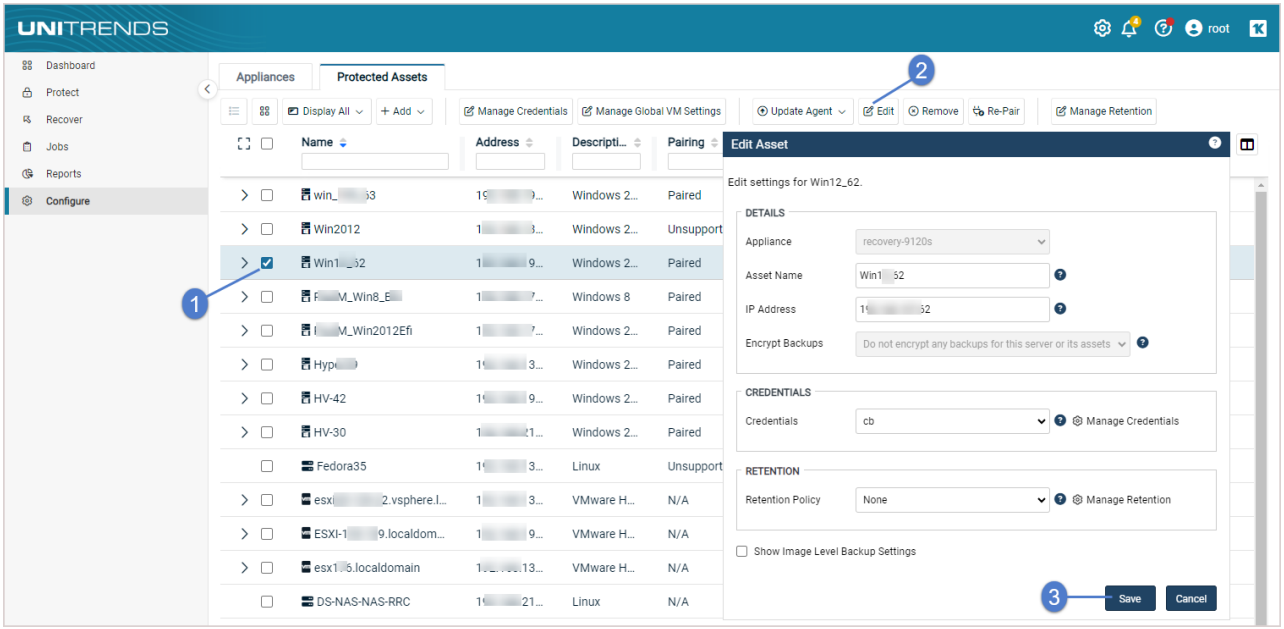

To enable pairing mode on a Linux asset

See this KB article to enable pairing mode on the Linux asset: Unitrends Agent Pairing - Linux.

To enable pairing mode on a Windows asset

In most cases, a secure agent pairing is required for communication between a Unitrends appliance and the Windows assets it protects. Upon installing the agent for the first time, the asset remains in pairing mode for 25 hours. During this 25-hour window, you can add the asset to an appliance, which establishes a secure pairing between its agent and the appliance.

If needed, you can use this procedure to manually enable pairing mode on an asset for a specified period of time. Pairing mode must be enabled to:

- Add the asset to a Unitrends appliance. (For example, use to add the asset to one or more appliances after the initial 25-hour pairing window has closed.)
- Re-establish pairing between the asset and the appliance. (For example, use if the asset's pairing status changed to *Failed* due to an error condition. After the error condition is resolved, use this procedure to enable pairing mode, then run the "To re-pair an asset" procedure.)
- Establish pairing between the asset and the appliance after upgrading the agent from an older version that did not support pairing. (Use this procedure to enable pairing mode, then run the "To re-pair an asset" procedure.)

This procedure is for Windows assets running agent version 10.7.9 or later. If your asset is running an older agent version, install the latest agent before running this procedure.

Note: If you are running a pre-10.7.9 agent version and cannot upgrade, use the procedure in this KB article to enable pairing on the Windows asset: Unitrends Agent Pairing - Windows.

1. Log in to the Windows asset as a user that has full access to all files and folders on the system (e.g., local administrator).
2. Launch the Unitrends Agent Pairing Utility from the Windows Start menu by selecting **Unitrends key manager GUI**:

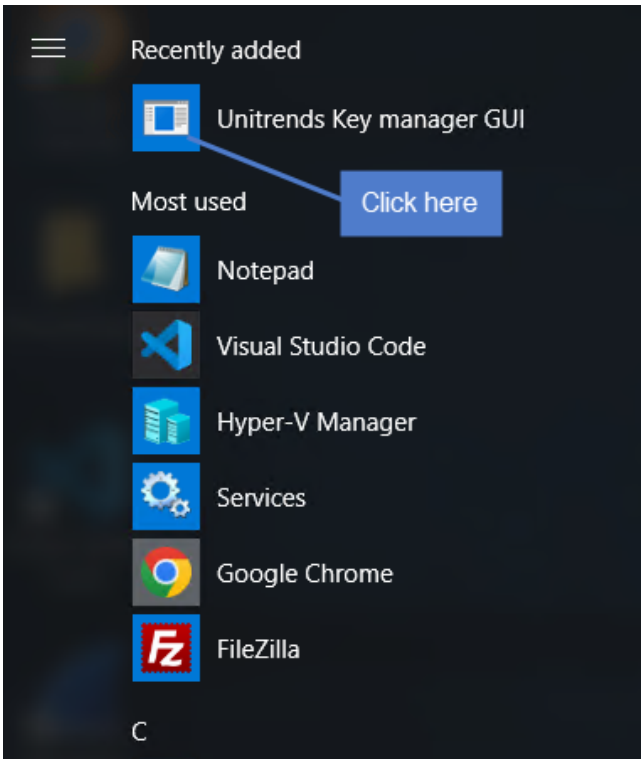

3 Check the following information:

- Pairing status – *Disabled* indicates the asset is not in pairing mode. *Enabled* indicates the asset is already in pairing mode.
- Time remaining – *00:00:00* indicates that the asset is not in paring mode. Any other number indicates the asset is in pairing mode and shows the number of hours, minutes, and seconds remaining in the pairing window. (The asset exits pairing mode when time remaining reaches *00:00:00*).

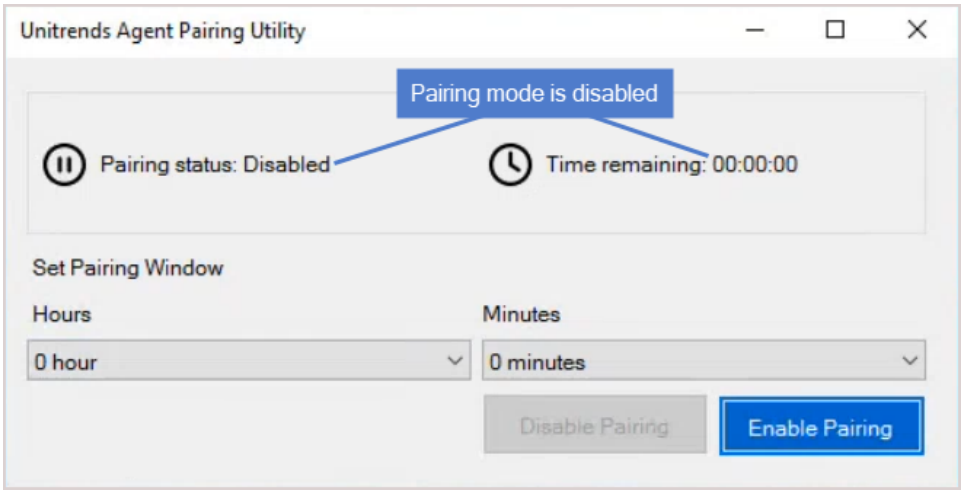

4 To enable pairing mode, select the number of hours and minutes that the asset will remain in pairing mode, then click **Enable Pairing**.

Note: If your asset is already in pairing mode, you can opt to extend the pairing window. Simply select the number of hours and minutes, then click **Extend Time** to add the hours and minutes to the time remaining.

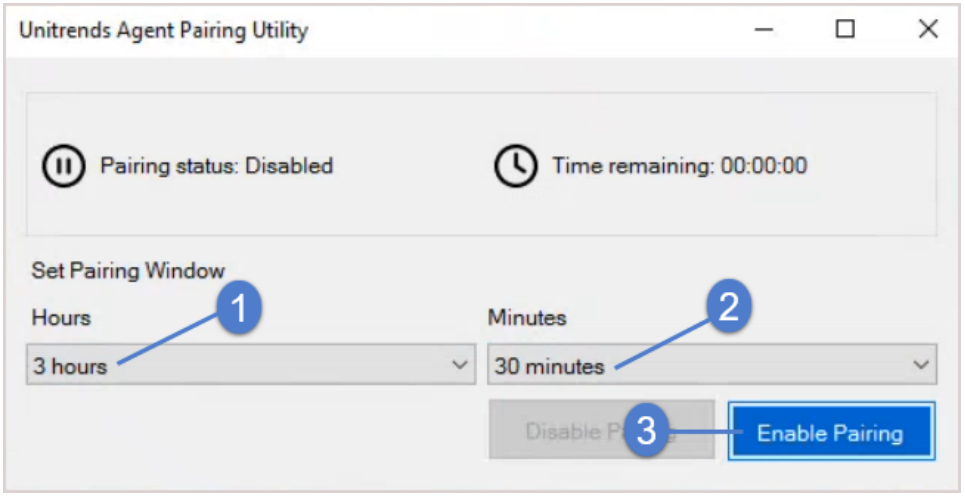

5 The asset enters pairing mode and the time remaining in the pairing window displays. You can now pair the asset to a Unitrends appliance as described in "To re-pair an asset" on page 345.

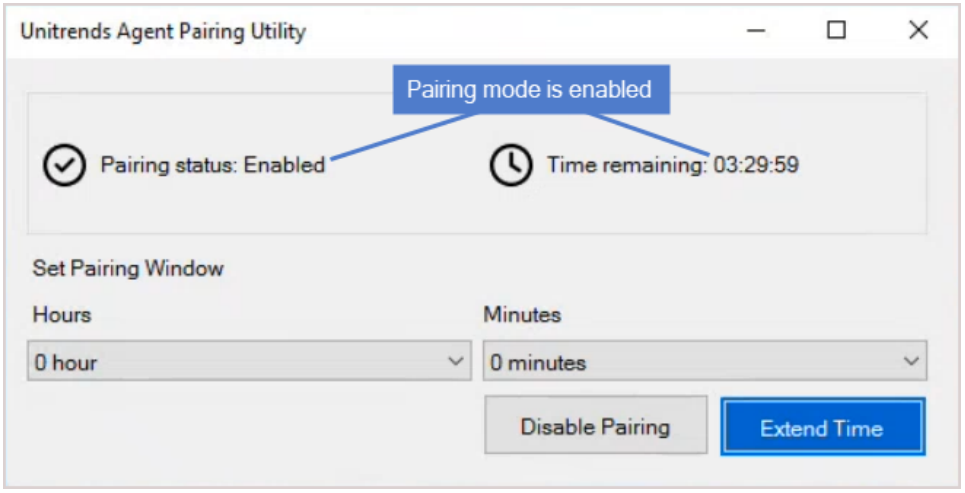

To extend pairing mode on a Windows asset

For an asset that is currently in pairing mode, use this procedure to extend the amount of time until the asset exits pairing mode.

Note: This procedure is for Windows assets running agent version 10.7.9 or later. If your asset is running an older agent version, install the latest agent before running this procedure.

1. Log in to the Windows asset as a user that has full access to all files and folders on the system (e.g., local administrator).
2. Launch the Unitrends Agent Pairing Utility from the Windows Start menu by selecting **Unitrends key manager GUI**:

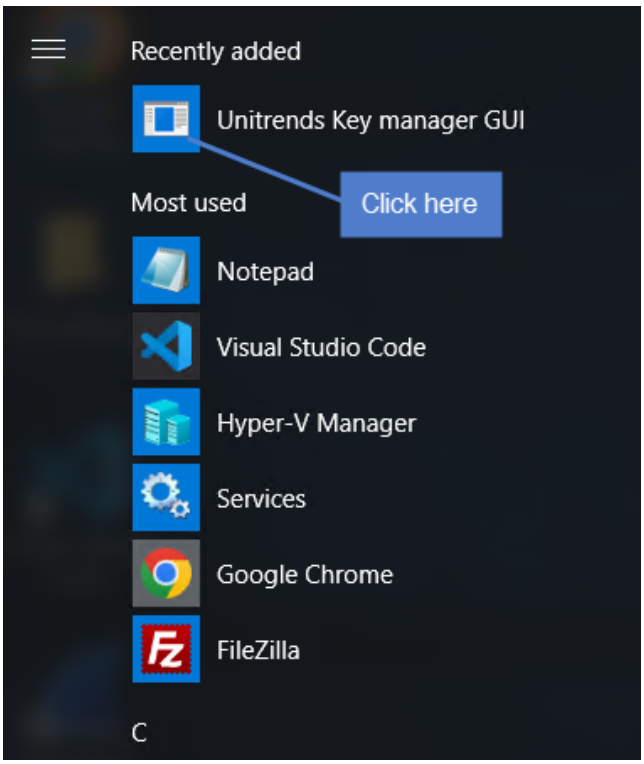

3. The utility indicates that the asset is currently in pairing mode (*Pairing status: Enabled*) and shows the time remaining until the asset exits pairing mode. In our example, *22:27:11* indicates that there are 22 hours, 27 minutes, and 11 seconds remaining.

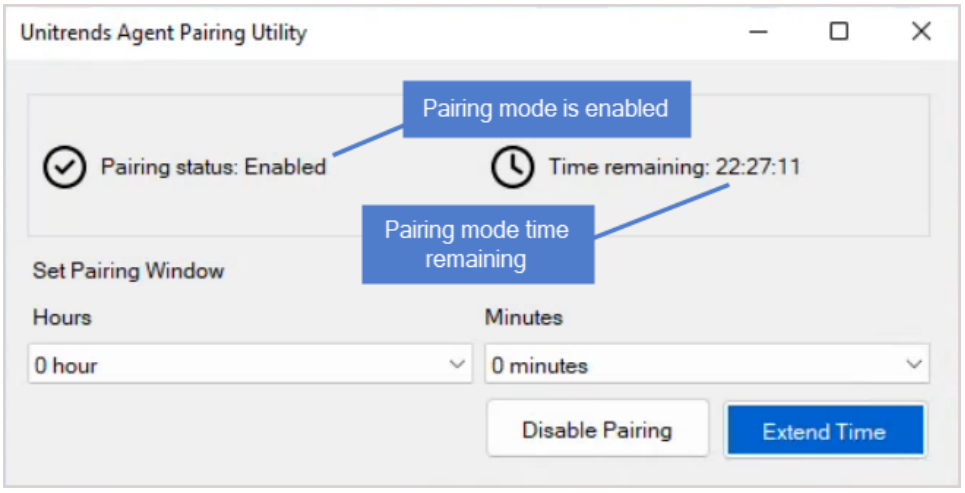

4. To extend the time remaining, select the number of hours and minutes to add, then click **Extend Time**:

Note: If the pairing window has expired, enable pairing mode instead of extending time: select the number of hours and minutes to define the pairing window, then click **Enable Pairing**.

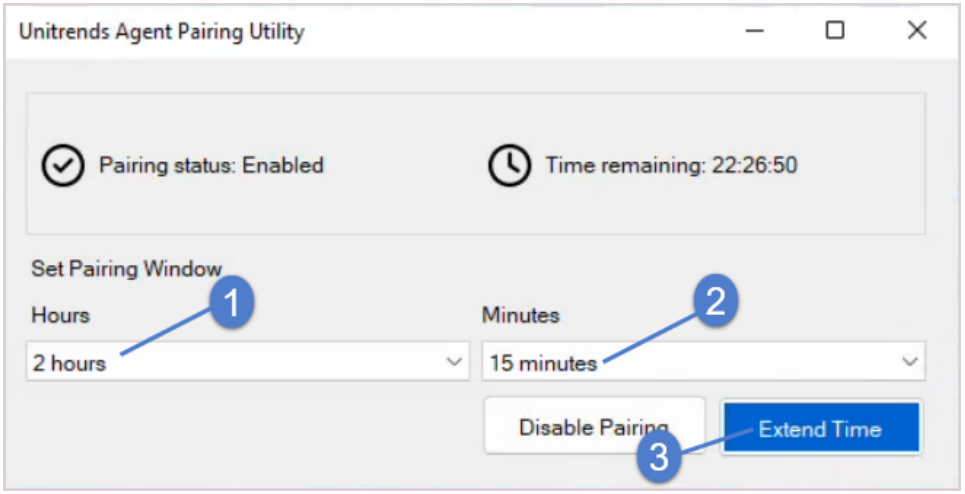

5 The pairing mode time remaining is extended:

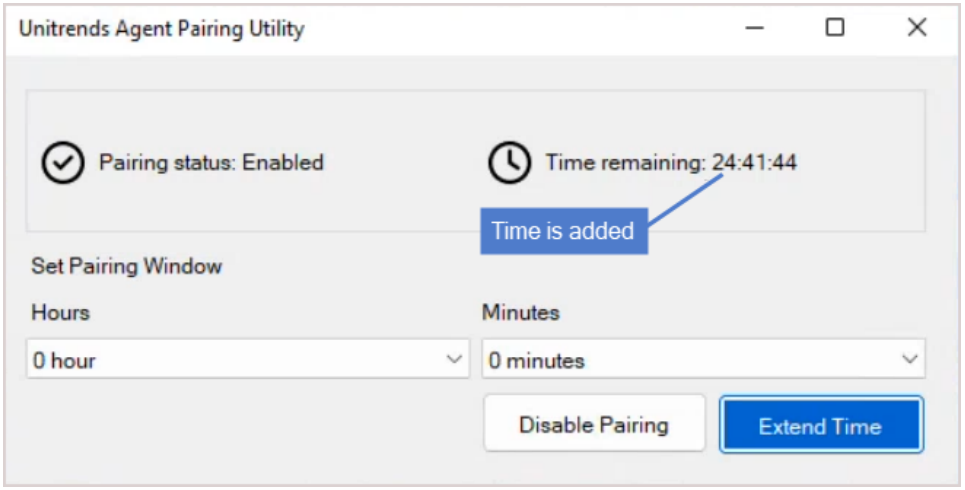

To re-pair an asset

Use this procedure to re-establish pairing between an asset and its backup appliance. You will typically run this procedure in the following cases:

- The asset's pairing status is *Failed* due to an error. Once the error is resolved, you must re-pair the asset to resume communication between the asset and the appliance.
- The asset's pairing status is *Unsupported* only because the asset was running an older agent version and you've updated the agent to a supported version (Windows agent 10.6.6+ or Linux agent 10.7.5+.)

IMPORTANT! Before running this procedure, ensure that the asset is in pairing mode (see "To enable pairing mode on a Linux asset" or "To enable pairing mode on a Windows asset" above).

1 On the **Configure > Protected Assets** tab, select one or more assets to re-pair.

2 Click **Re-Pair**.

3 Click **Confirm**.

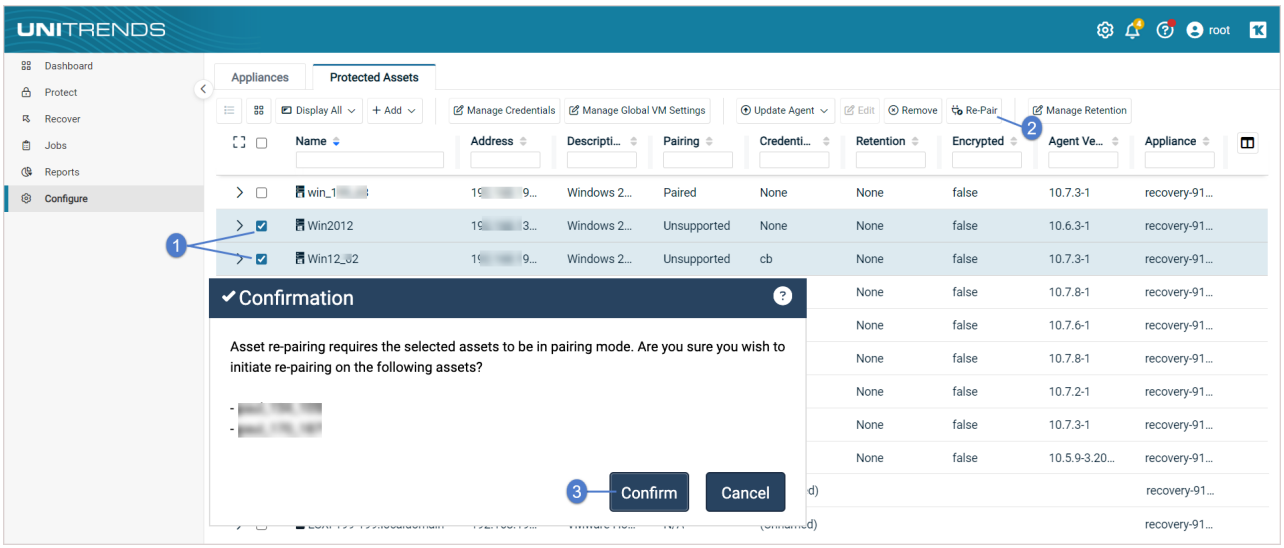

4 The pairing operation begins and you see the message *Initializing Pairing*. When finished:

- The asset's pairing status is updated to *Paired* if the operation was successful.
- The asset's pairing status is unchanged if the operation was unsuccessful. (Click  to view the *SSL pairing failed* alert.)

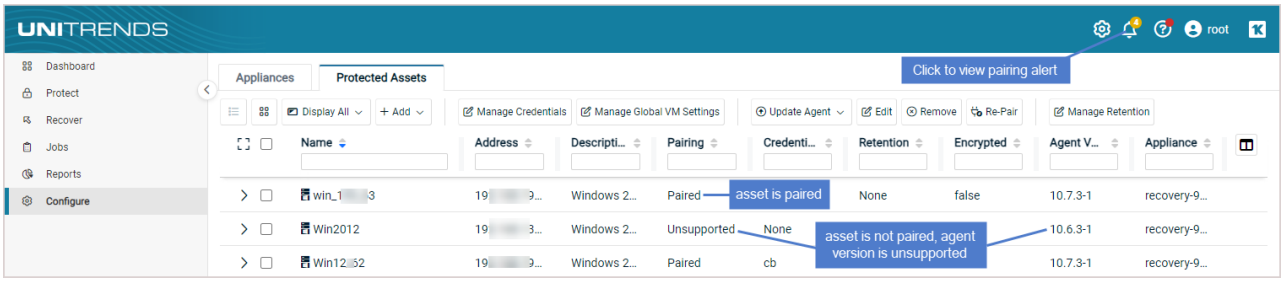

To disable pairing mode on a Windows asset

The Windows asset exits pairing mode automatically when its pairing window closes, but you can manually close the pairing window by running this procedure.

Notes:

- This procedure is for Windows assets running agent version 10.7.9 or later. If your asset is running an older agent version, install the latest agent before running this procedure.

- This procedure closes the asset's pairing window only. It does not impact the asset's pairing status on the **Configure > Protected Assets** tab in any way.

1. Log in to the Windows asset as a user that has full access to all files and folders on the system (e.g., local administrator).
2. Launch the Unitrends Agent Pairing Utility from the Windows Start menu by selecting **Unitrends key manager GUI**:

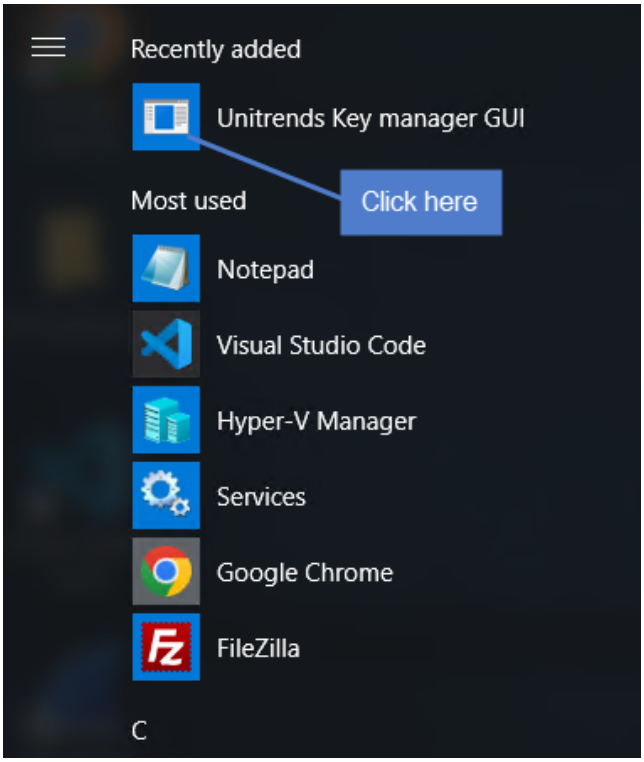

3. To exit pairing mode, click **Disable Pairing**:

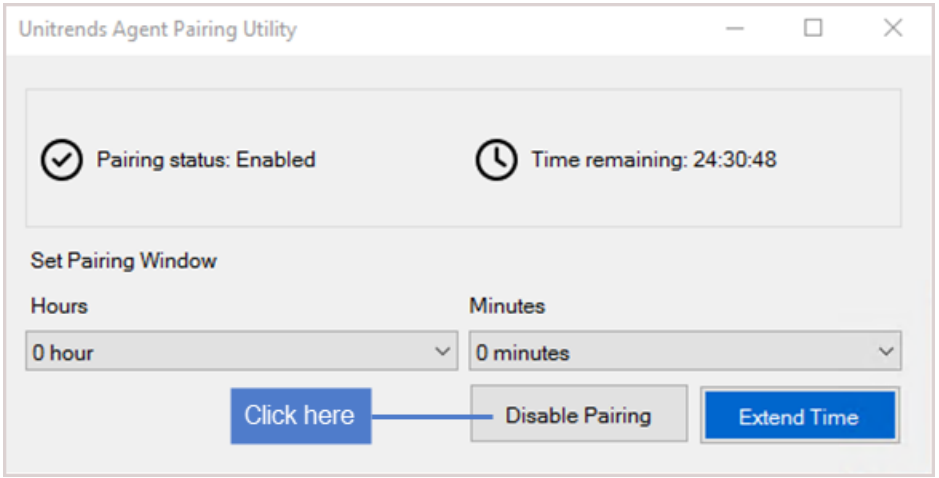

4 Pairing mode is disabled:

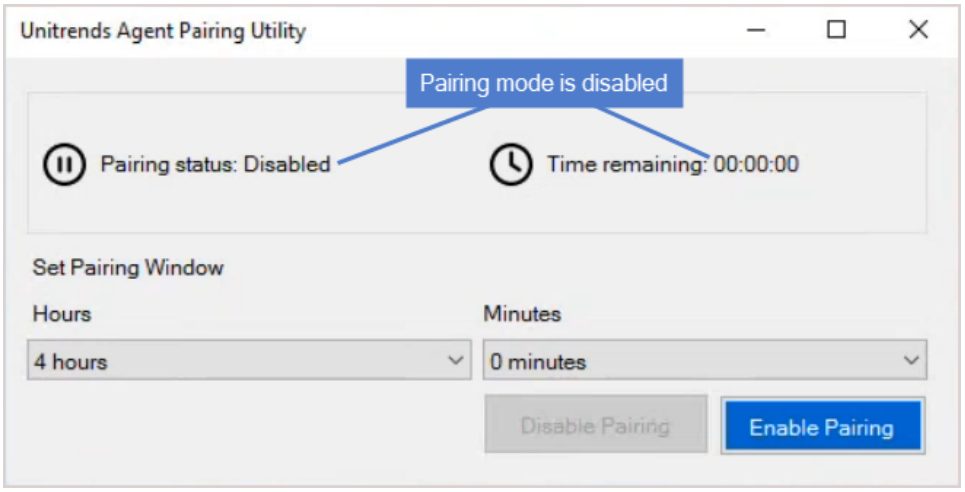

Grouping assets in custom folders

On the Protect page, you can create folders in the inventory tree to customize how protected assets are grouped and displayed in the UI. You can then assign users to groups so they can quickly locate the assets they need to work with. Once folders are created and the **Show Groups** option is selected on the Protect page, assets are presented in folder groups in these areas:

- On the Protect page
- In the Create Backup Jobs dialog
- In the Create Backup Copy Jobs dialog

You can opt to view folder groups or hide them by clicking the **Show Groups** or **Hide Groups** icon on the Protect page. The selected view determines how the inventory tree displays on the page and dialogs listed above. When users hide groups, they can then see all assets on the appliance, even ones in groups to which they have not been assigned.

Note: Folder groups do not display in the Backup Catalog or on reports.

The asset grouping feature is supported on appliances running version 9.0.0-13 or higher. Groups can be created or edited by Unitrends users that have administrator or superuser privileges only. The following assets can be grouped:

- Agent-based assets, such as Windows or Linux servers
- Hyper-V virtual hosts
- VMware virtual hosts
- AHV virtual hosts

See the following topics for details:

- "Unitrends users and asset groups" on page 349
- "Working with asset groups" on page 350

- "To create a top-level asset group" on page 351
- "To create an inner group" on page 353
- "To view or hide asset groups" on page 358
- "To edit an asset group" on page 356
- "To delete an asset group" on page 360

Unitrends users and asset groups

For usability, you can use asset groups to customize which assets a user can see in Show Groups view. A user's privilege level determines which groups the user can see and whether the user can add, edit, or delete groups:

- Users that have administrator or superuser privileges can see all assets and groups.
- Users that have administrator or superuser privileges cannot be removed from any group.
- While in Show Groups view, users that have manage or monitor privileges can see ungrouped assets and assets in their assigned groups only. These users cannot see assets in groups to which they have not been assigned. To see all assets, the user must switch to Hide Groups view.
- Only users with administrator or superuser privileges can add, edit, or delete asset groups.
- You can assign users when creating a new group or by editing a group. A Unitrends user account must be created before the user can be assigned to a group. For details on creating users, see "Users and roles" on page 119.

In this example the root superuser sees all asset groups:

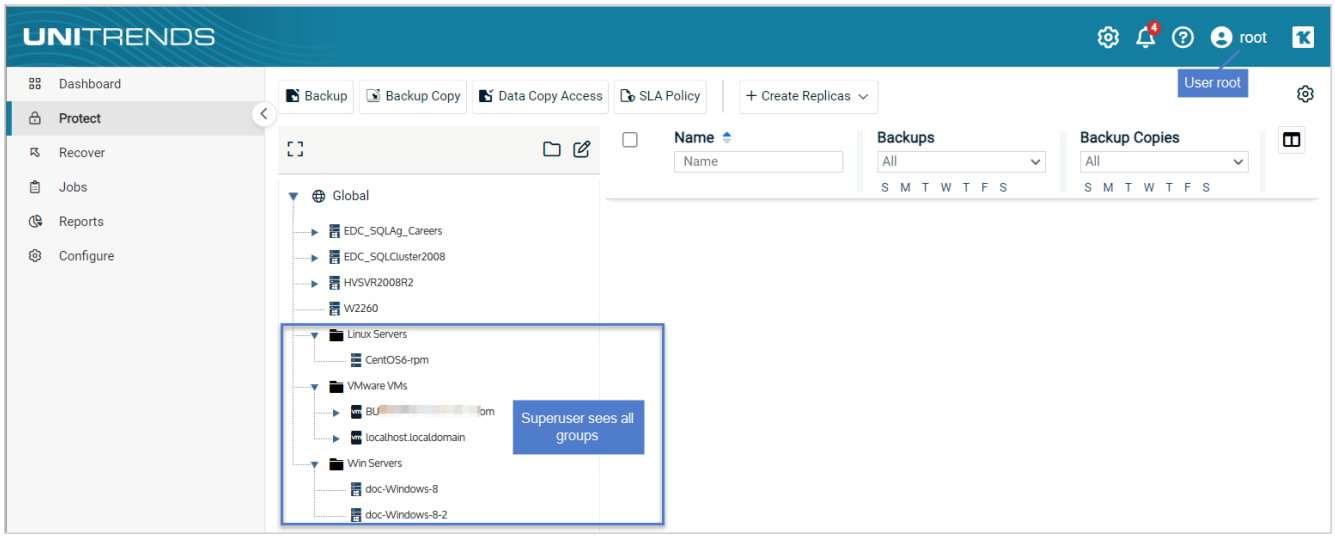

User *jmorgan* sees only their assigned group(s):

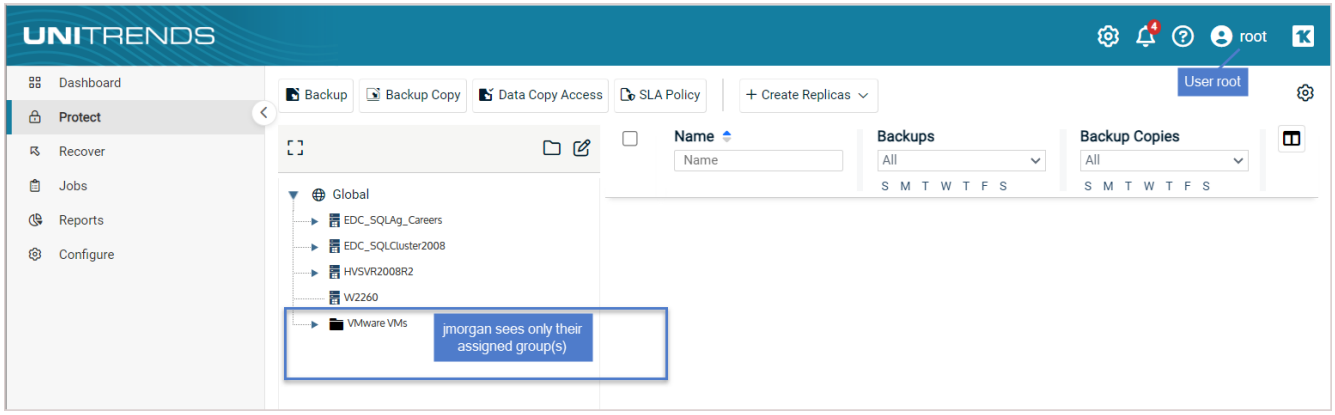

Working with asset groups

Consider the following when working with asset groups:

- You can create, edit, and delete asset groups without impacting the original inventory tree. Simply click **Hide Groups** on the Protect page to return to the original inventory tree view. Click **Show Groups** to return to group view.
- You can create groups at different levels in the inventory tree. For example, you can create a top-level group that contains assets and inner groups.
- An inner group can be assigned one or more assets from its parent group.
- Assets can be assigned to one group only. Assigning an asset to an inner group moves the asset from its parent group to the child inner group.
- When an asset is removed from a top-level group, or a top-level group is deleted, the assets are moved to their original place in the inventory tree.
- When an asset is removed from an inner group, or an inner group is deleted, the assets are moved to the parent group.

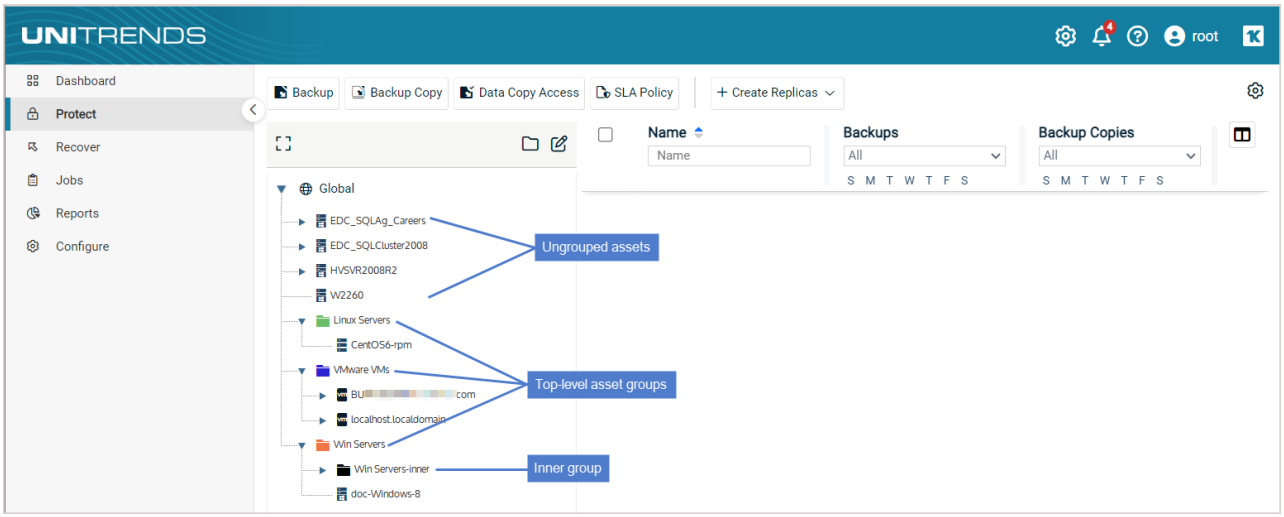

To create a top-level asset group

1 Log in to the UI as a user with administrator or superuser privileges.

2 Select **Protect** or **Protect > Protected Assets**.

3 Click the **Show Groups** icon located above the inventory tree.

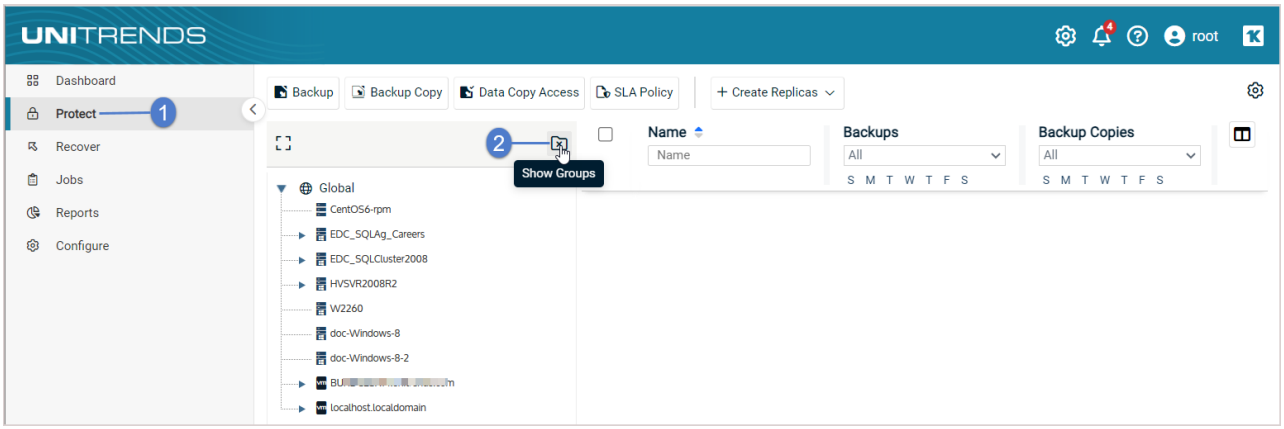

4 Click the **Manage Groups** pencil icon.

5 Click **Add**.

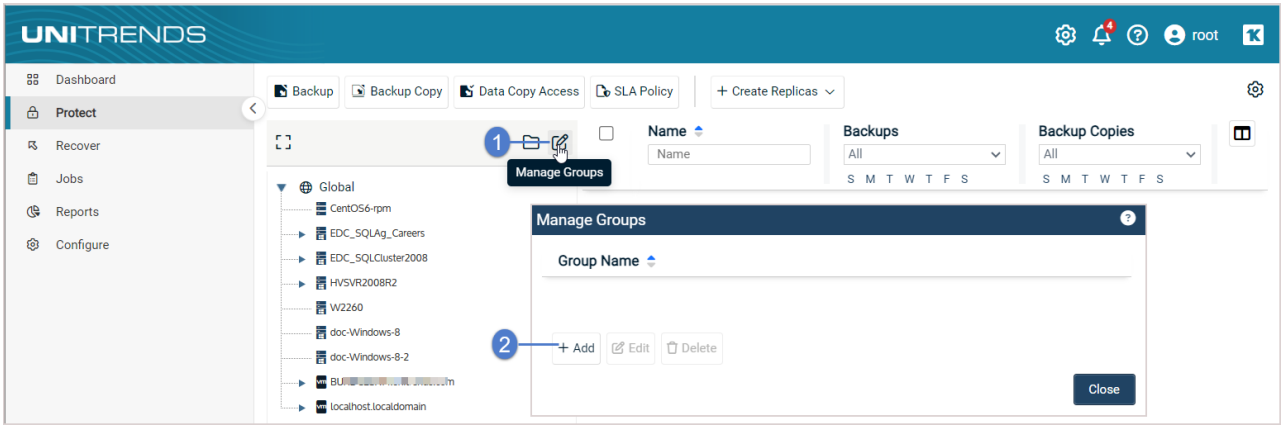

6 Enter a unique Group Name.

7 (Optional) Click the color drop-down to select a display color for this group's folder.

8 On the **Users** tab, click to select users to add to the group.

> **Note:** You cannot create users here. You can add existing Unitrends users to the group. For details on creating users, see "Users and roles" on page 119.

- Unitrends users display in the list.

- If you have created Unitrends Active Directory (AD) users, these do not display. To add an existing Unitrends AD user, click **Add AD User**, enter the AD username (without *@domain*), and click **Save**.
- Users with administrator or superuser privileges are automatically added to every group. You cannot remove these users from the group.

9 On the **Not in Group** tab, click to select assets to add to the group.

10 Click **Save**.

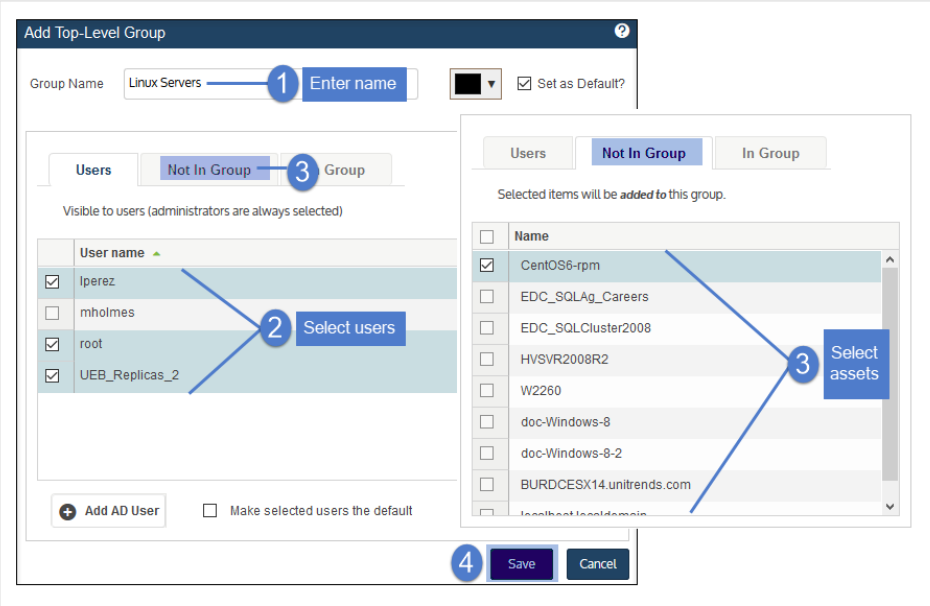

11 The group is added. Click **Close** to exit.

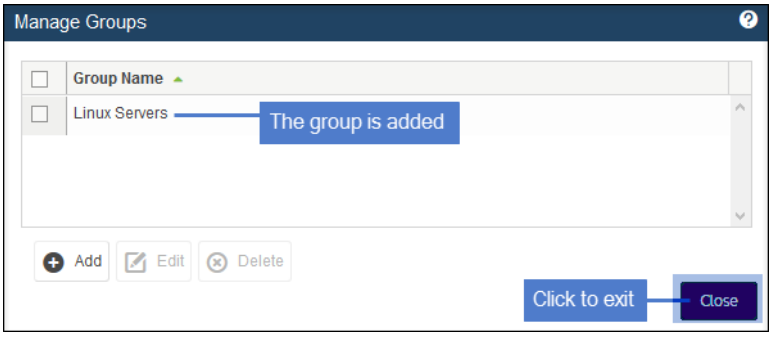

12 The new group displays in the inventory tree on the Protect page. Expand the folder to view the assigned assets.

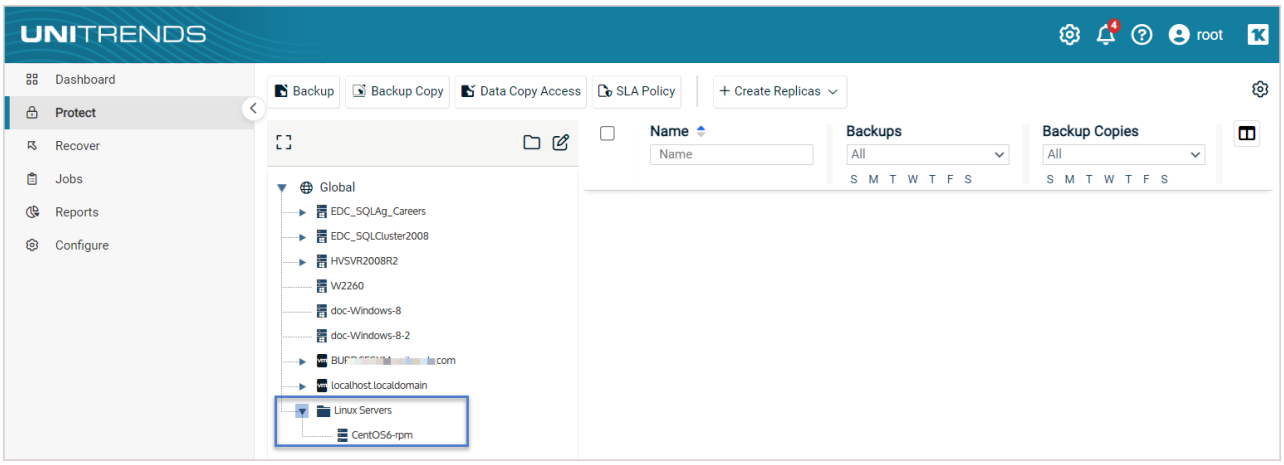

To create an inner group

1. Log in to the UI as a user with administrator or superuser privileges.
2. Select **Protect** or **Protect > Protected Assets**.
3. Click the **Show Groups** icon located above the inventory tree.

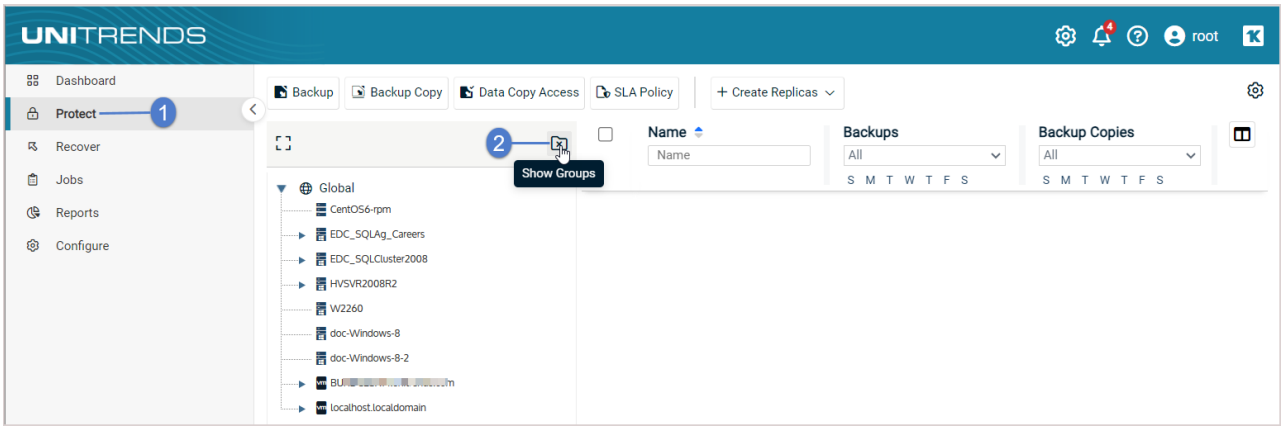

4. In the inventory tree, select the top-level group that will contain the new inner group.
5. Click the **Manage Groups** pencil icon.
6. Click **Add**.

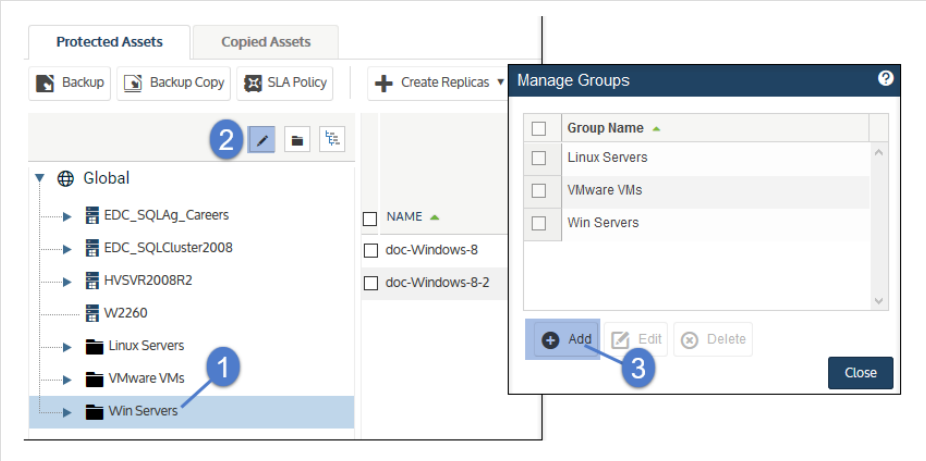

7 (Optional) Modify the Group Name. This name must be unique.

8 (Optional) Click the color drop-down to select a display color for this group's folder.

9 (Optional) On the **Users** tab, click to select users to add to the group.

> **Note:** You cannot create users here. You can add existing Unitrends users to the group. To create users, see "Users and roles" on page 119.

- Unitrends users display in the list.
- If you have created Unitrends Active Directory (AD) users, these do not display. To add an existing Unitrends AD user, click **Add AD User**, enter the AD username (without *@domain*), and click **Save**.
- Users with administrator or superuser privileges are automatically added to every group. You cannot remove these users from the group.

10 On the **Not in Group** tab, click to select assets to add to the group.

These assets are moved from the parent group to this inner group.

> **Note:** If you do not see any assets, verify that you selected the outer group in the inventory tree on the Protected Assets tab. You cannot create an inner group by selecting the outer group in the Manage Groups dialog

11 Click **Save**.

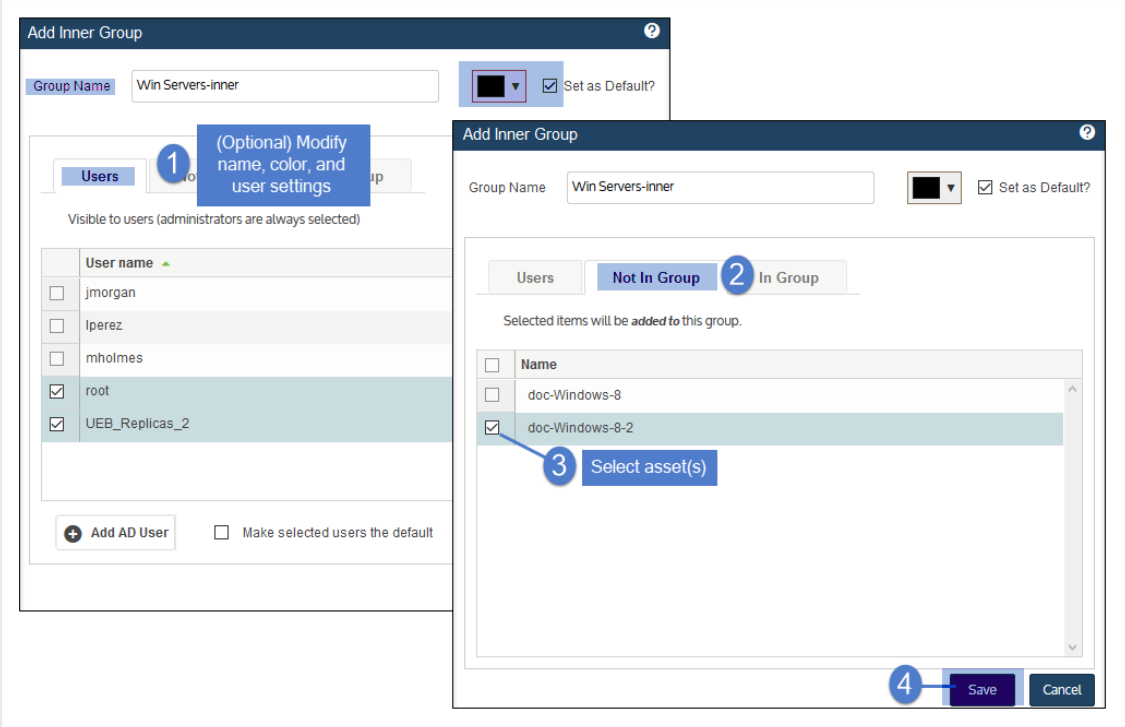

12 Click **Close**.

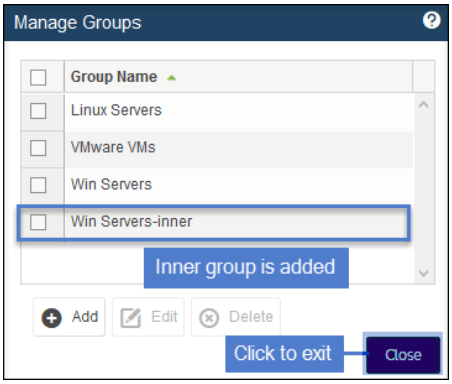

The new group displays in the inventory tree on the Protect page. Expand the folder to view the assigned assets.

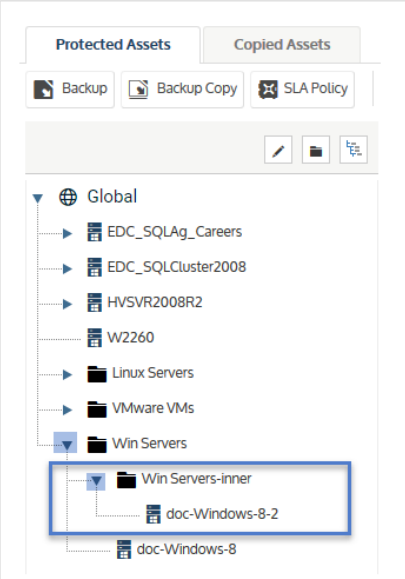

To edit an asset group

1. Log in to the UI as a user with administrator or superuser privileges.
2. Select **Protect** or **Protect > Protected Assets**.
3. Click the **Show Groups** icon located above the inventory tree.

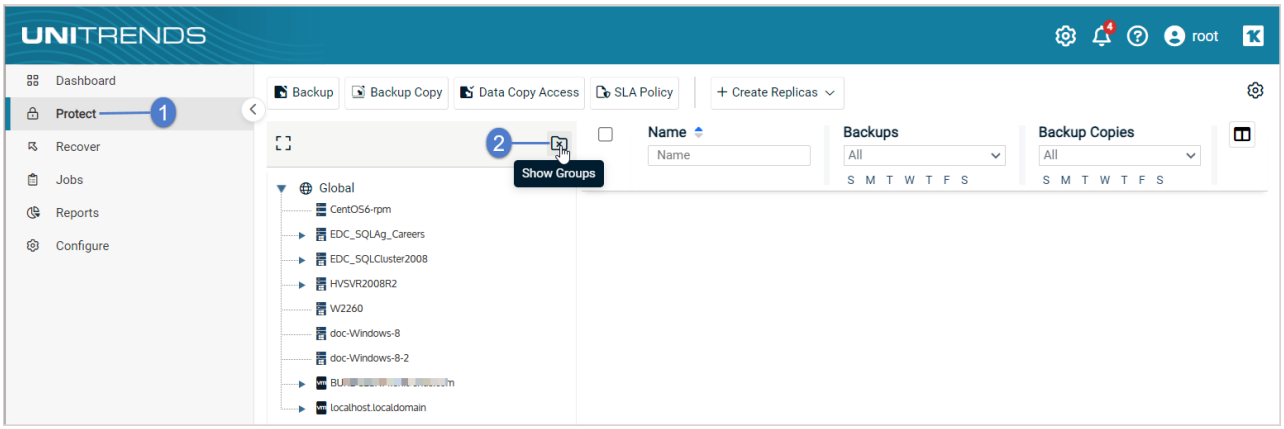

4. Click the **Manage Groups** pencil icon.
5. In the Group Name list, select the group to edit.
6. Click **Edit**.

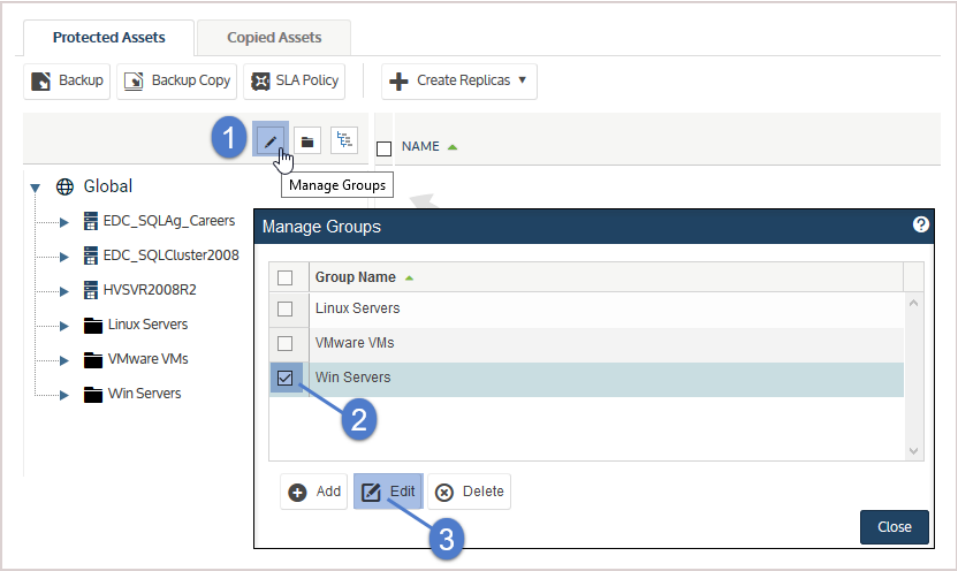

7 Modify options as needed:

- The Group Name must be unique.
- Click the color drop-down to modify the display color for this group's folder.
- On the **Users** tab, click to select users to add or remove.

  Users with administrator or superuser privileges are automatically added to every group. You cannot remove these users from the group.
- On the **Not in Group** tab, click to select assets to add to the group.
- On the **In Group** tab, click to select assets to remove from the group.

  Assets removed from a top-level group are moved to their original place in the inventory tree.

  Assets removed from an inner group are moved to the parent group.

8 Click **Save**.

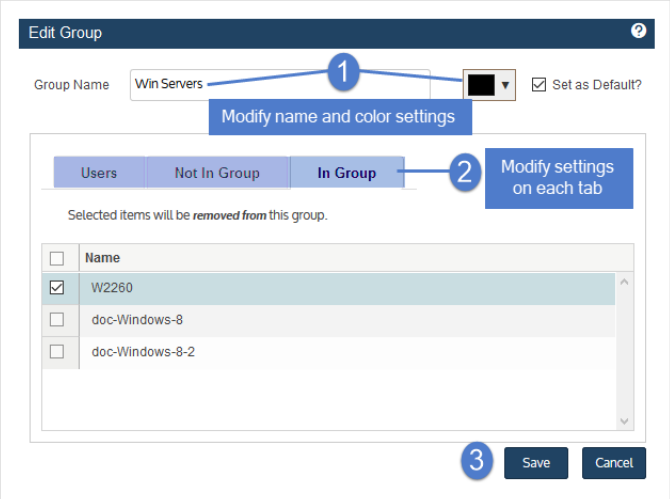

9 Click **Close** to exit.

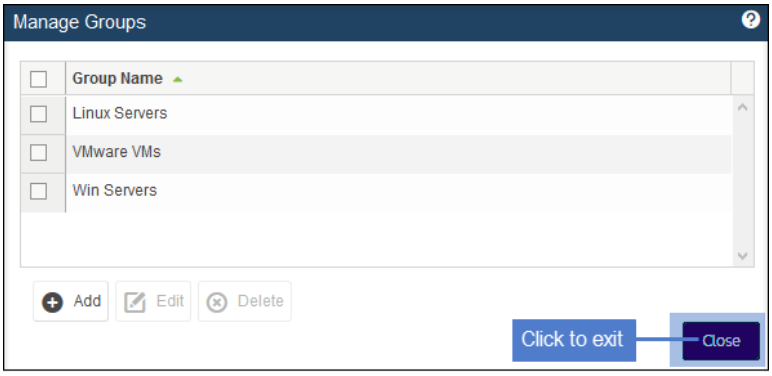

To view or hide asset groups

You can toggle the UI display to view or hide asset groups. In Show Groups view, any asset that is assigned to a group displays under its group folder. In Hide Groups view, any asset that is assigned to a group returns to its original location in the inventory tree (and no asset group folders display).

To view asset groups

Go to **Protect** or **Protect > Protected Assets** and click the **Show Groups** icon located above the inventory tree:

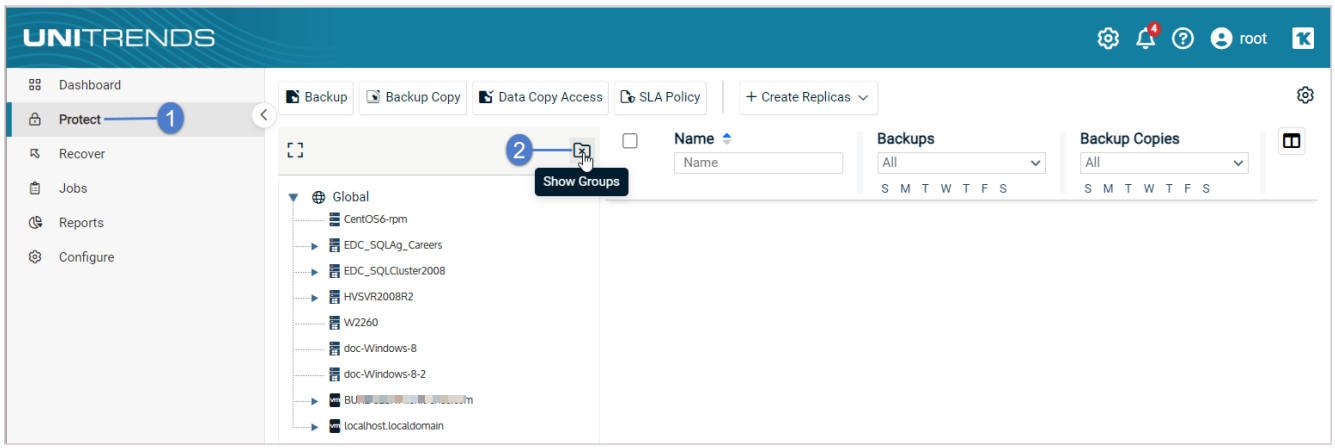

Grouped assets now display under their asset group folders:

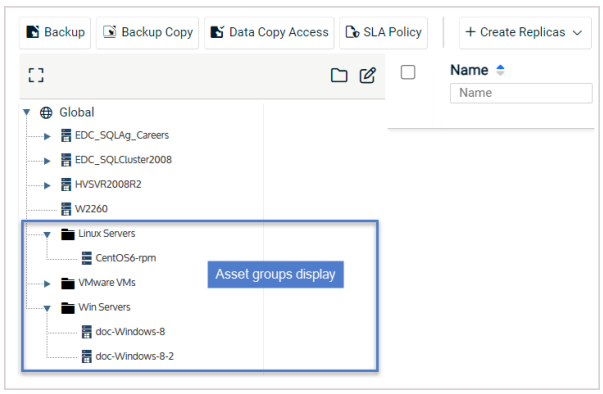

To hide asset groups

Go to **Protect** or **Protect > Protected Assets** and click the **Hide Groups** icon located above the inventory tree:

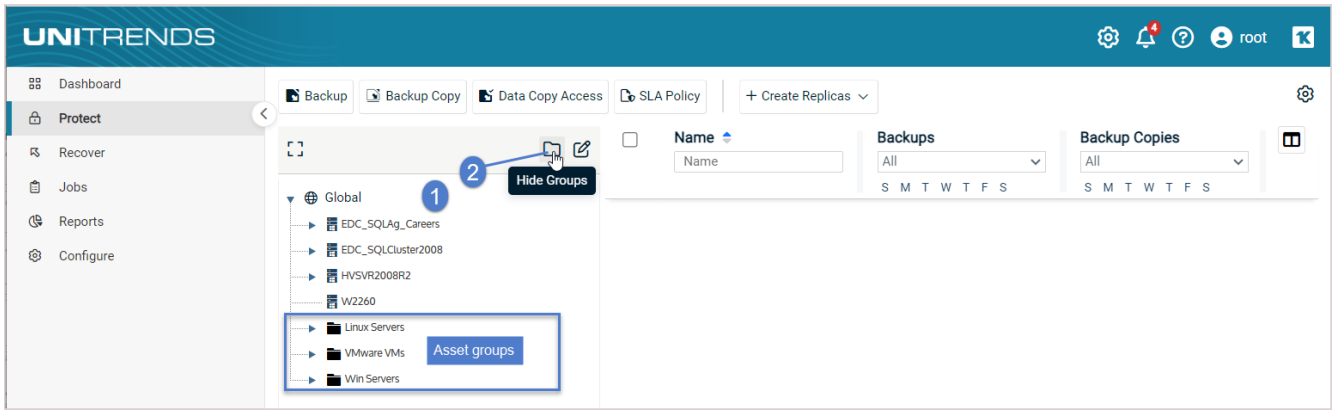

Asset groups no longer display. Assets from the groups now display individually in the main inventory tree.

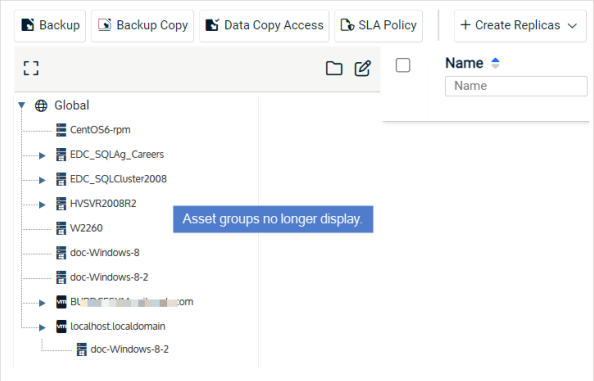

To delete an asset group

1. Log in to the UI as a user with administrator or superuser privileges.
2. On the **Protect** page, click the **Show Groups** icon located above the inventory tree.

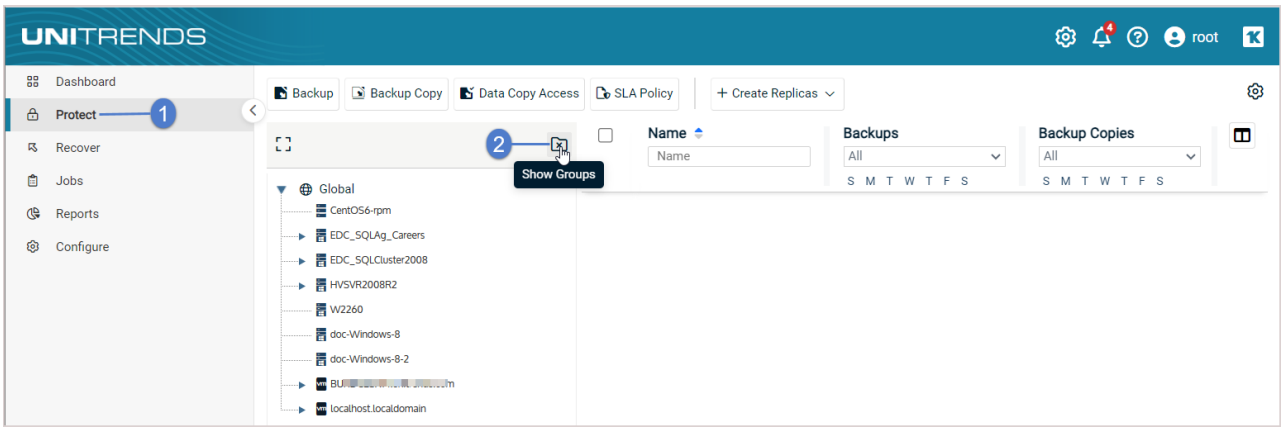

3. Click the **Manage Groups** pencil icon.
4. In the Group Name list, select the group to delete.
5. Click **Delete**.
6. Check the **I understand...** box to confirm, then click **Delete Group**.

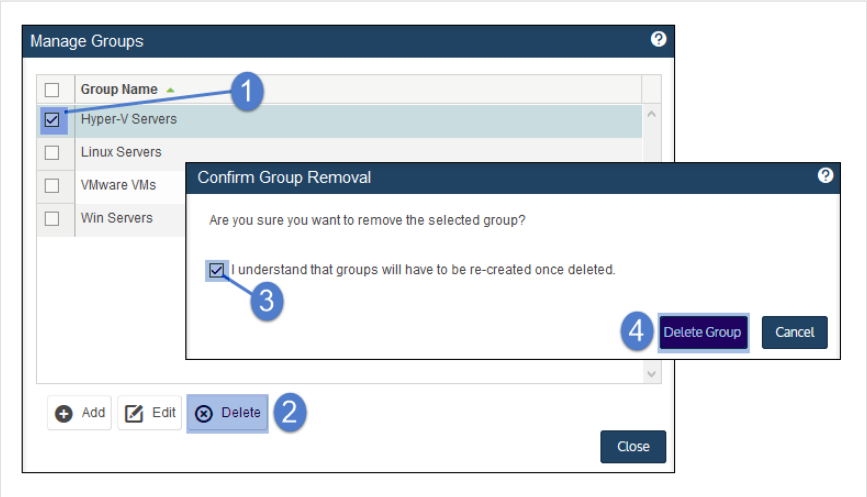

7 Assets in the group you deleted are moved to:

- Their original place in the inventory tree if you deleted a top-level group.
- To the parent group if you deleted an inner group.

Unitrends agents

Before you can protect a physical asset, you must install the Unitrends agent. (You can also opt to install the agent on virtual machines if you prefer to use file-level protection.) For most Windows assets, the appliance can push-install the agent when you add the asset. For other physical assets, you must install the Unitrends agent manually before you add the asset.

Note: A Unitrends agent is not used to protect iSeries assets, For details on iSeries, see "iSeries Backups Overview and Procedures" on page 767.

Agent installation procedures vary by operating system. See the following topics for details:

Operating system	Agent install procedure
Microsoft Windows	"Installing the Windows agent" on page 362
Linux	"Installing and updating the Linux agent" on page 387
CentOS	"Installing and updating the Linux agent" on page 387
Debian	"Installing and updating the Linux agent" on page 387
Red Hat	"Installing and updating the Linux agent" on page 387

Operating system	Agent install procedure
SUSE	"Installing and updating the Linux agent" on page 387
Ubuntu	"Installing and updating the Linux agent" on page 387
Solaris	"Installing and updating the Solaris agent" on page 401
Novell Netware	"Installing and updating the Novell Netware agent" on page 399
Mac	"Installing and updating the Mac agent" on page 398
AIX	"Installing and updating the AIX agent" on page 396
UnixWare	"Installing and updating the UnixWare agent" on page 402
HP-UX	"Installing and updating the HP-UX agent" on page 397

Installing the Windows agent

To protect a Windows asset, you must install a lightweight agent on the Windows machine. This Windows agent is required to run file-level, image-level, and application backups. Depending on the asset's operating system and configuration, the agent can be push-installed by the appliance or installed manually. Push install is recommended to reduce setup time. Once a Windows asset is set up for push install, agent updates can also be pushed, reducing maintenance time.

To install the Windows agent, start by reviewing the "Windows agent requirements", then install the agent as described in "Push-installing the Windows agent" on page 365 or "Manually installing the Windows agent" on page 366.

Note: If you have trouble installing the agent, look at the application messages in the Windows event viewer to address the error. For details, see this KB article: Troubleshooting Windows event IDs.

Windows agent requirements

The following requirements must be met before installing the Windows agent:

- The Unitrends appliance(s) protecting the Windows asset must be running an equal or higher version than the agent that will be installed. Beginning in release 10.8.1, the agent installer enforces version compatibility by raising an error if the appliance version is older than the agent.
- Windows administrative privileges for the user installing the agent.
- Approximately 1100 MB of free space on the Windows system drive, usually volume *C:*.

- Single Instance Storage (SIS) on Windows Storage Server 2008 is not supported and must be disabled for the agent to properly perform backups.
- The Windows Volume Shadow Copy Service (VSS) framework must be installed.
- To protect Exchange, SQL Server, or Hyper-V, the following VSS writers are required:
  - VSS Exchange Writer is required for the Exchange agent.
  - VSS SQL Writer is required for the SQL Server agent.
  - VSS Hyper-V Writer is required for the Hyper-V agent.
- Agent operations use ports 1743, 1745, and 888. Ensure that these ports are not in use on the Windows asset.

  Note: Unitrends does not officially support backup through firewalls. For details, see this KB article: Backup fails through Router, DMZ, or Firewall.

- Secure agent pairing requirements – Beginning in release 10.6.6, a secure pairing is automatically established between the appliance and the Windows agent on each of its protected assets. This pairing enables Transport Layer Security (TLS) to encrypt data and authenticate connections between appliances and agents. Communication between appliances and agents is only allowed if there is a matching (paired) certificate. These secure pairing requirements must be met to protect Windows assets with agent release 10.6.6 or higher:
  - The Unitrends appliance must be running release 10.6.6 or higher.
  - The Windows asset must be running agent release 10.6.6 or higher.
  - The Unitrends appliance version must be equal to or higher than the Windows agent version.

    IMPORTANT! Be sure to upgrade your Unitrends appliance before upgrading your Windows agents.
    - Jobs will fail if you attempt to protect a 10.6.6 or higher agent with an appliance that is running an older release.
    - You cannot add an asset that is running a 10.6.6 or higher agent to an appliance that is running an older release. If you attempt this, you receive an error similar to: *Failed to save client: Registration for client* assetName *failed. The Unitrends System could not connect to the Unitrends Agent on* assetName*. Please ensure that the Agent software is installed on* assetName*, the Agent service is running (if applicable), and no firewall settings are preventing access*.
  - If upgrading from a pre-10.6.9 agent release, backups may fail until the pairing completes successfully (this can take up to two hours).
  - To protect Hyper-V clusters, SQL clusters, or file server clusters with the secure agent pairing feature:
    - The cluster must be running agent version 10.6.9 or higher.
    - The Unitrends appliance and cluster must be running in the same time zone.

- The secure agent pairing feature is not used to protect Windows XP, 2003, and Vista. To protect a Windows XP, 2003, or Vista asset, install agent version 10.6.7. The 10.6.7 agent detects the asset's OS version and disables secure agent pairing.

See "Secure agent pairing for Windows and Linux agents" on page 338 for details on working with this feature.

In addition to the general "Windows agent requirements" on page 362, the following prerequisites must be met to push-install the Windows agent:

Item	Description
Windows versions supported	These Windows versions are supported: • Windows Server – 64-bit Windows 2008 R2 and later versions listed in the Compatibility and Interoperability Matrix are supported. (32-bit versions are not supported.) • Windows Workstation – 64-bit Windows 7 and later versions listed in the Compatibility and Interoperability Matrix are supported. (32-bit versions are not supported.) Windows agent push is NOT supported for Azure virtual machines. **Note:** A known Microsoft issue may prevent successful agent push install on Windows 7 and Windows 2008 R2 systems. To resolve this issue, see the following Microsoft knowledge base article: The "Untrusted publisher" dialog box appears when you install a driver in Windows 7 or Windows Server 2008 R2.
Credentials	Trust credentials must be defined for the Windows server asset on the backup appliance. See "Managing asset credentials" on page 322 for details.
Windows environment	The Windows machine must be configured as described in "Windows configuration requirements for push installation" on page 364.

Windows configuration requirements for push installation

The following Windows configuration settings are required for the agent push feature:

Note: For troubleshooting steps, see Troubleshooting Agent Push.

- *Workstation* and *Server* services must be running and set to automatic restart.
- The Windows asset must be able to access the appliance's Samba share:
  - SMB 2.0 – The SMB 2.0 security option is enabled by default on Unitrends appliances that were originally imaged or deployed with version 10.4.8 or higher. SMB 2.0 must be enabled on the Windows asset.

  - SMB 1.0 – The SMB 1.0 security option is enabled by default on appliances that were originally imaged or deployed with a pre-10.4.8 version. SMB 1.0 must be enabled on the Windows asset.

    Note: Upgrading from a pre-10.4.8 version does not change the SMB 1.0 setting. (To configure your appliance to use SMB 2.0, see How Unitrends supports SMBv2.)

- For Windows 7 and later, *Network discovery* and *Printer and File Sharing* must be enabled for the current network profile (in **Control Panel > Network and Sharing Center**).
- Trust credentials entered in the Add Asset dialog in the Unitrends UI must have administrative privileges. On systems with user account controls (UAC) enabled, at least one of the following must also apply:
  - The trust credentials entered are for a domain administrator account.
  - The trust credentials entered are for a system local administrator account. Being a different member of the Administrators group is insufficient, it must be the built-in account to bypass UAC. If the administrator account is disabled, enable it by executing the following in an elevated command prompt:
    `net user administrator /active:yes`
  - The Registry key *LocalAccountTokenFilterPolicy* exists and is set to **1** (to use a local administrator that is not the 'Administrator' account).
- Verify Remote IPC and Remote Admin shares are enabled. These shares should be enabled with File and Printer Sharing, but verifying is a good idea if you're still having trouble. To verify, issue the following command from an elevated command prompt and check the output for ADMIN$ and IPC$: **net share**
- Firewall rules must allow inbound and outbound traffic between both machines. Default Windows firewall rules limit many services to the subnet. If the backup appliance is outside the Windows asset's subnet, modify firewall *Printer and File Sharing* settings (TCP ports 139 and 445) to allow communication between the systems.

Push-installing the Windows agent

Add the asset as described in "To add an agent-based asset" on page 289. The following are installed automatically (assuming all "Windows agent requirements" on page 362 have been met):

- The Windows agent.
- For Hyper-V servers or Windows servers with the Hyper-V role enabled, the Hyper-V CBT driver is installed. This driver is used for faster Hyper-V incremental backups. (You do not need to reboot to enable this driver.)
- For Microsoft SQL and Exchange servers, SQL and Exchange components are installed. You can then run application backups to protect these databases. For more information, see "Exchange backup requirements and considerations" and "SQL backup requirements and considerations".
- For Windows assets that are eligible for image-level backups, the Windows Volume CBT driver is installed or updated if an update is available. This driver is used to enable incremental image-level backups. To enable this driver, you must reboot the Windows asset after installing the Volume CBT driver for the first time or after updating from a pre-10.3.3 agent version. (The last driver update was in agent version 10.3.3. If you are updating agent version 10.3.4 or later, a reboot is not required.)

Note: If you see this error, the Windows asset is paired to an appliance running an older version than the agent you're trying to install: *Appliance version validation failed. Please verify all paired appliance versions meet or exceed the agent*. Update the appliance to enable the push-install.

Manually installing the Windows agent

Use the procedures in this section to install the agent manually.

To run these procedures, you will download the applicable agent MSI file from the Unitrends Downloads page (https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads) to the Windows machine.

You can then install the agent by launching the installer or from the command line. In most cases, you will use the installer. You will need to use the command line procedure to install on a Windows 2008 server that was deployed with the server core option, or to install to multiple Windows machines by using Windows Group Policy. See the following topics for details:

- "Agent installer for Windows" on page 366
- "Command-line installer for Windows agents" on page 377
- "Agent deployment using Group Policy" on page 379

Agent installer for Windows

The agent installer loads the applicable components in *Unitrends_Agentx86.msi* or *Unitrends_Agentx64.msi* onto the Windows asset during installation.

The following are included with the Windows agent (*Unitrends_Agentx86.msi* or *Unitrends_Agentx64.msi*):

- For Hyper-V servers or Windows servers with the Hyper-V role enabled, the Hyper-V CBT driver is installed. This driver is used for faster Hyper-V incremental backups. (You do not need to reboot to enable this driver.)
- For Microsoft SQL and Exchange servers, SQL and Exchange components are installed. You can then run application backups to protect these databases. For more information, see "Exchange backup requirements and considerations" and "SQL backup requirements and considerations".
- (Optional) For Windows assets that are eligible for image-level backups, the installer provides the option to install the Windows Volume CBT driver. This driver is used to enable incremental image-level backups. To enable this driver, you must reboot the Windows asset after installing the Volume CBT driver for the first time or after updating from a pre-10.3.3 agent version. (The last driver update was in agent version 10.3.3. If you are updating agent version 10.3.4 or later, a reboot is not required.)

Consider the following before installing the Windows agent:

- For Microsoft Vista assets, administrator privileges are required to install the agent. You must log in as a user that has administrator privileges on the Vista server to install the agent. Members of the Administrator group that have not been assigned administrator privileges are not able to install the agent.
- For Windows Server assets, the agent performs backup and recovery of the system state, including support for ISS, COM+, Cluster Database, and Active Directory. The agent must be installed on the Windows server while logged in using the local system Administrator account. If the local system Administrator account cannot be used

for the installation, the Windows User Account Control facility must be disabled. Once the agent has been installed, User Account Control can be re-enabled.

To install the Windows agent

1. Log in to the Windows asset as a user that has full access to all files and folders on the system (e.g., local administrator).
2. Download the agent MSI file from https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads.
   - For the 64-bit agent, click the **MSI** link in the Windows row.
   - For the 32-bit agent, click the **Link** in the Legacy Agents row. On the Legacy Agents page, click the **Link** in the 32-bit Agents row. On the 32-bit Agents page, click the **MSI** link in the Microsoft Windows row.

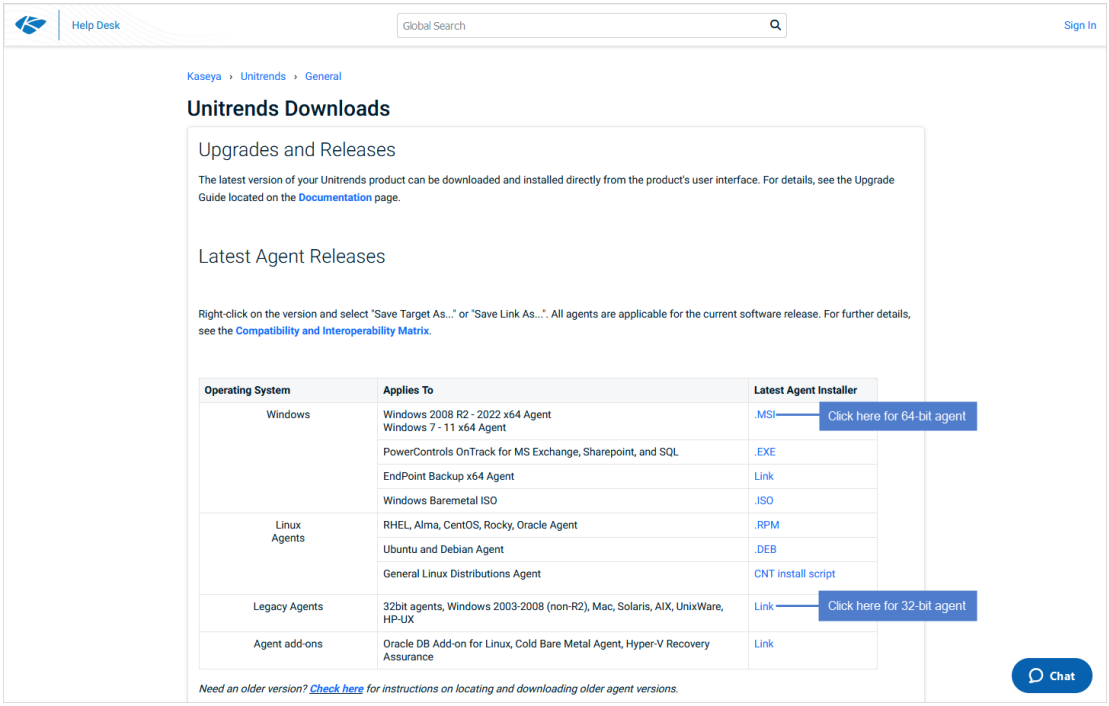

3. Double-click the MSI file to launch the installer.

   MSI file names:

   - *Unitrends_Agentx64.msi* - agent for 64-bit Windows assets
   - *Unitrends_Agentx86.msi* - agent for 32-bit Windows assets

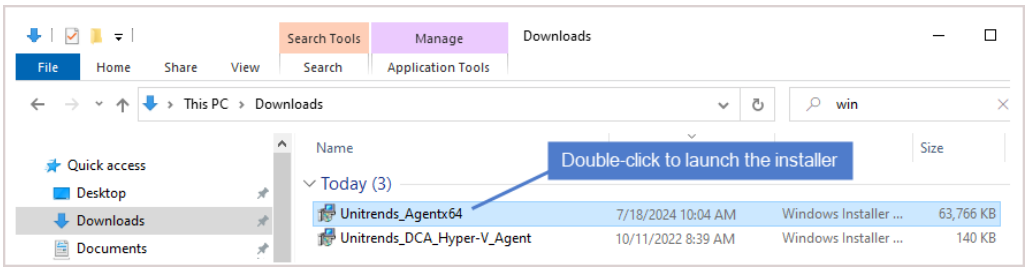

4 Click **Next** to proceed.

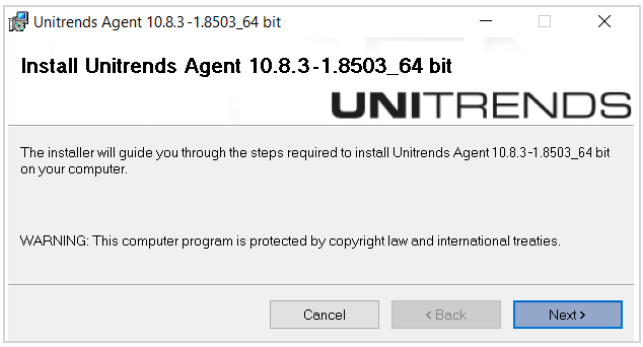

5 Click **Next** to begin the installation process. The installation can be interrupted at any time by clicking **Cancel**.

Note: If you receive the message *The currently installing agent version is newer than the appliance version*, click **Yes** to proceed with the installation or click **No** to exit. If you proceed with the installation, be sure to upgrade the older appliance as soon as possible. Running an appliance version older than the agent is not supported and can cause undesirable results.

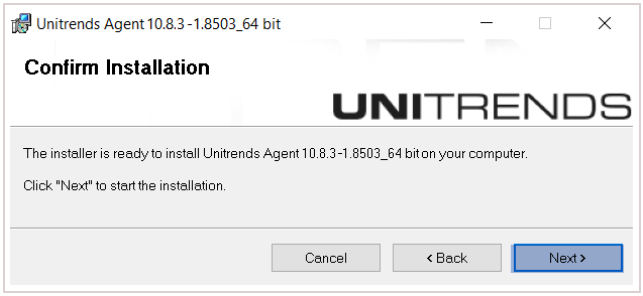

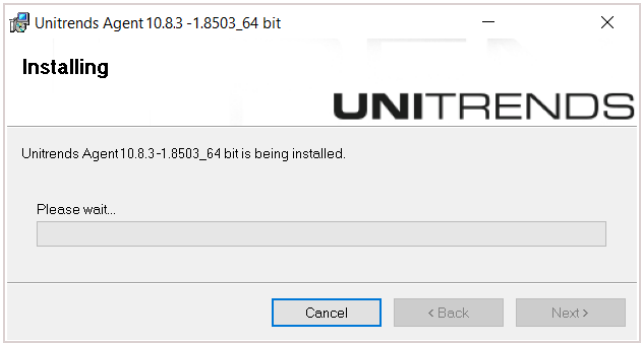

6 (Recommended) Click **Yes** to include the Windows Volume CBT driver. (This driver enables the option to run incremental image-level backups. You can run file-level backups and full image-level backups without installing this driver.)

Note: If the latest Windows Volume CBT driver is already installed, the Unitrends CBT Driver Install window does not display.

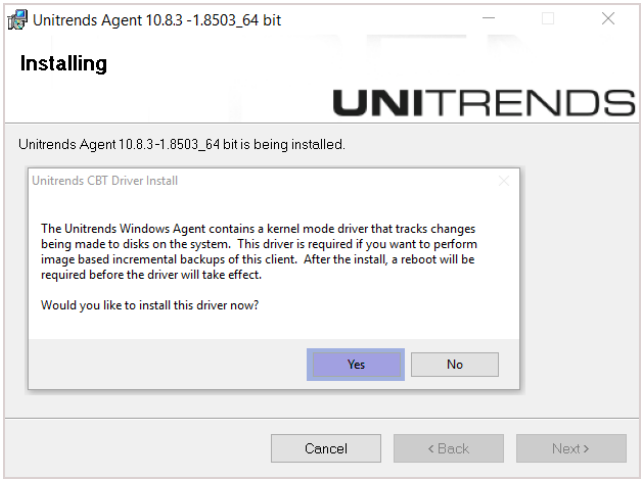

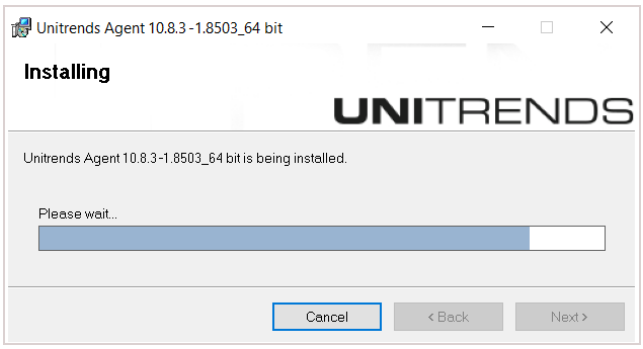

Note: The agent is installed to the *\PCBP* directory on the Windows system drive, usually volume *C:* (e.g., *C:\PCBP*).

7 After the agent is installed, do one of the following:

- If you did not opt to install the Volume CBT driver, installation is complete. Click **Done** to exit the installer.

  OR

- If you opted to install the Volume CBT driver, click **Next** and continue with the next step in this procedure.

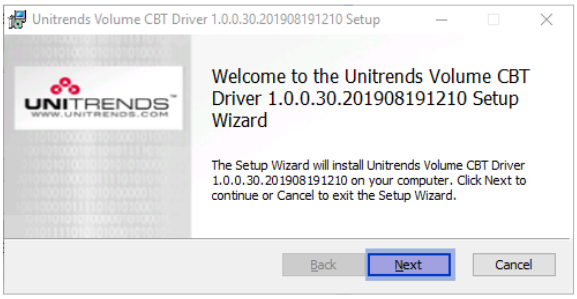

Notes:

- If this Volume CBT driver version has already been installed, the Unitrends Volume CBT Filter Driver Setup installer does not display. Windows agent installation is complete. If needed, reboot the Windows asset to enable the existing Volume CBT driver.
- If an older Volume CBT driver version has already been installed, you are given the option to install this driver version and advised as to whether the new driver is required for subsequent incrementals.
- If a newer Volume CBT driver version has already been installed, you are given the option to install this older driver version.

8 Click **Install** to begin the installation process (or click **Back** to review or modify data). The installation can be interrupted at any time by clicking **Cancel**.

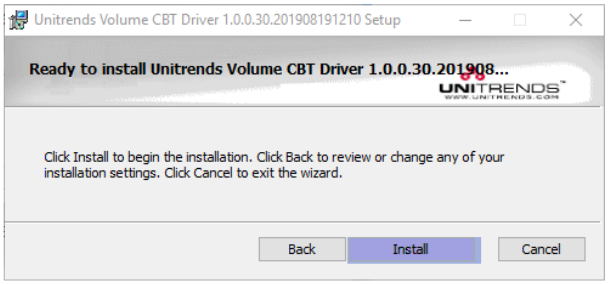

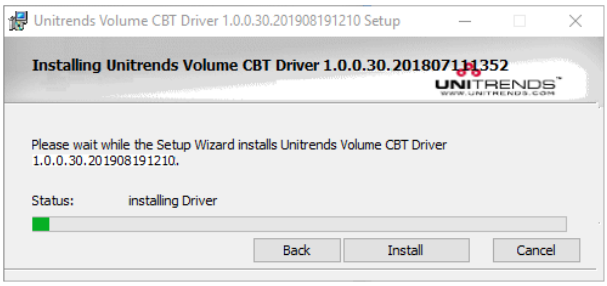

9 Click **Finish** to exit the installer.

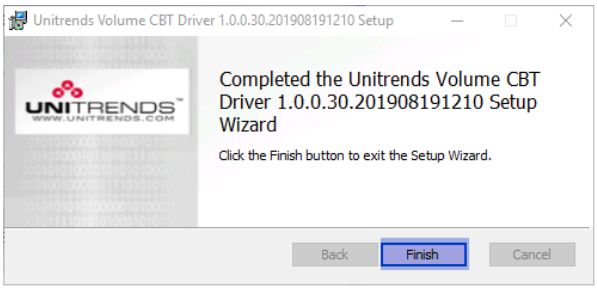

10 (If needed) If installing the driver for the first time or updating from a pre-10.3.3 agent release, you must reboot the Windows asset to enable the Volume CBT driver. Click **Yes** to reboot now or **No** to reboot at a later time.

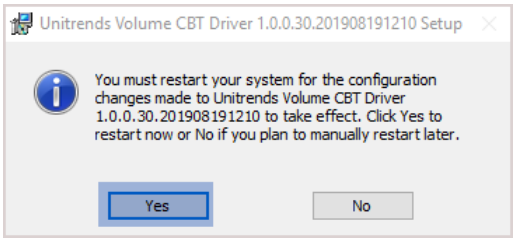

Notes:

- If the Volume CBT driver has not been installed or has not been enabled, image-level incrementals are not supported. Any scheduled incremental is automatically promoted to a full backup. If you attempt to run an on-demand incremental, you receive a message indicating that only full backups are supported.
- An image-level incremental is automatically promoted to a full backup in these other cases:
  - There is a problem detected with the Volume CBT driver.
  - A newer Volume CBT driver version is installed but has not been enabled. To run incremental backups, enable the new driver by rebooting the Windows asset.
  - The version of the Volume CBT driver that was included with the Windows agent is greater than the version that is enabled on the Windows asset. To run incremental backups, you must install the new driver and then reboot the Windows asset (see the next bullet for details).

- The installer for the associated Volume CBT driver is placed in *C:\PCBP\Installers*. You can install this driver at any time by running the installer, called *uvcbt.msi*. After installing the driver, you must enable it by rebooting the Windows asset.
- The Windows agent and the Windows Volume CBT driver are installed as separate, independent packages. Uninstalling the Windows agent does not uninstall the Windows Volume CBT driver. (To uninstall the Windows Volume CBT driver, use the Windows Control Panel Add/Remove Programs feature to remove *uvcbt.msi*.)

To install the Windows bare metal agent

The bare metal agent is needed only for image-based bare metal protection. For most assets, you can use the newer unified bare metal protection feature, which does not require the bare metal agent. (For details, see "Windows Bare Metal Protection and Recovery" on page 1207.)

Note: To install *Unitrends_BareMetal.msi* on assets that are running User Account Control (UAC), special installation is required. Use the "Installing the bare metal agent on a Windows asset running User Account Control" procedure instead. UAC is enabled by default on Windows Vista, Windows Server 2008, and Windows Server 2012.

1. Log in to the Windows asset as a user that has full access to all files and folders on the system (e.g., local administrator).
2. Download the agent MSI file from https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads:
   - Click the Agent Add-ons **Link**.

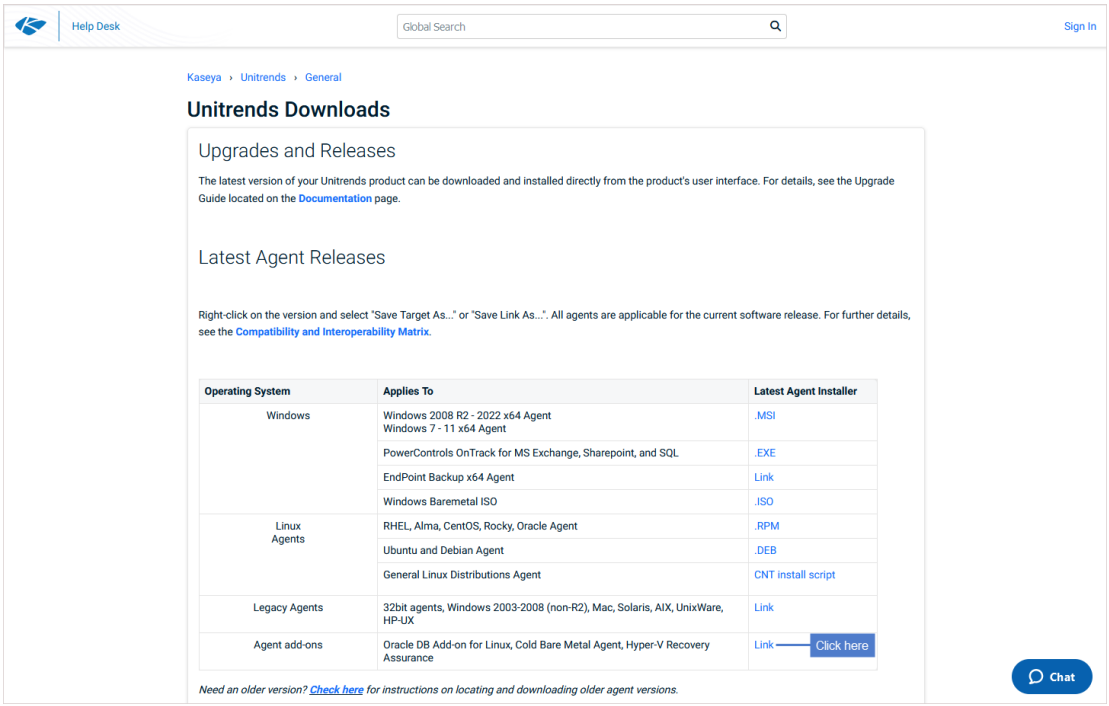

- On the Agent Add-ons page, click the Bare Metal Agent **MSI**.

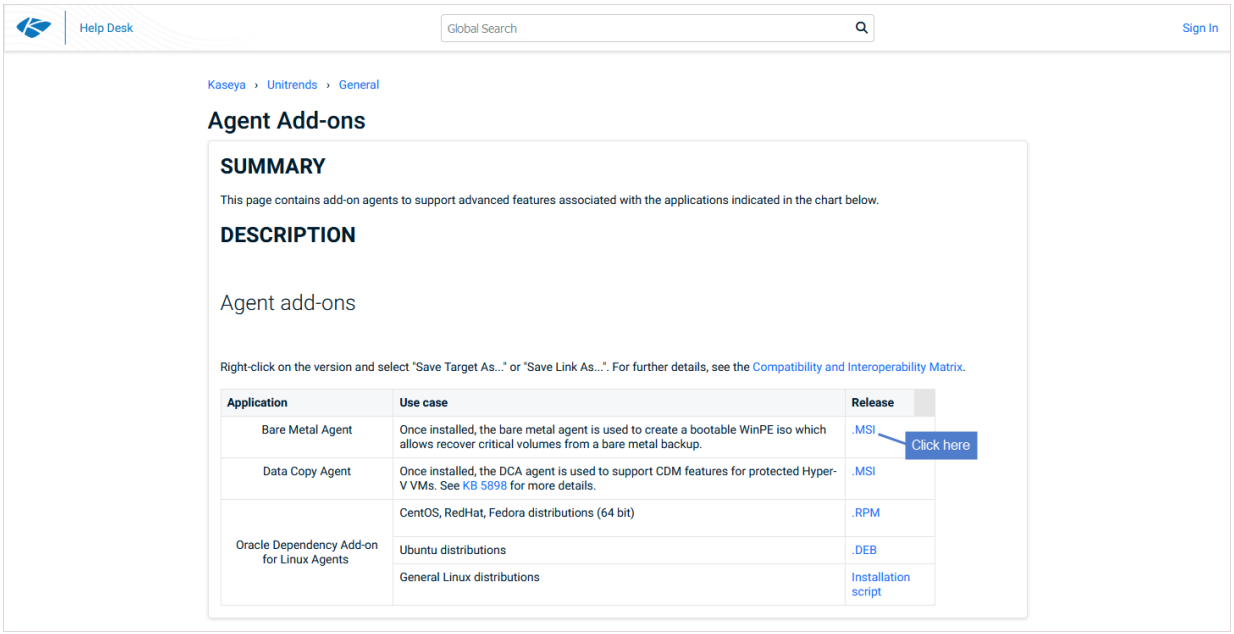

3 Double-click the *Unitrends_BareMetal.msi* file to launch the installer.

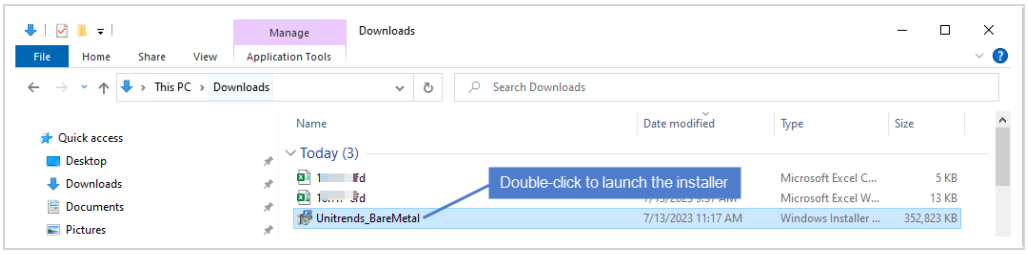

4 Click **Next** to proceed.

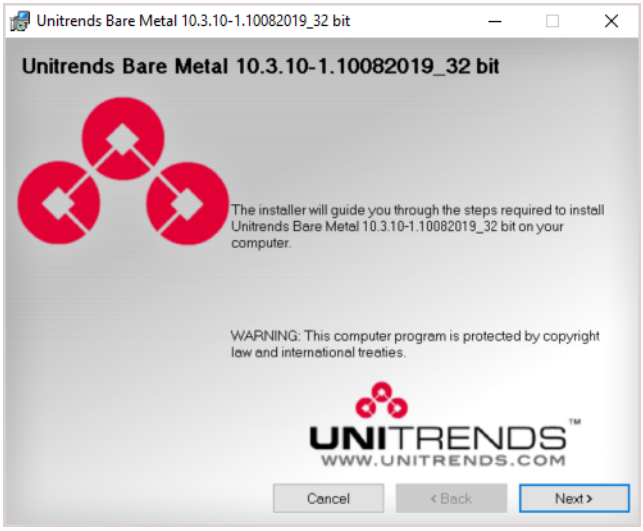

5 Click **Next** to install to the default location.

Installing to the default *C:\PCBP_BM* directory is strongly recommended. To install in another location (folder or volume), click **Browse** or manually enter the directory path.

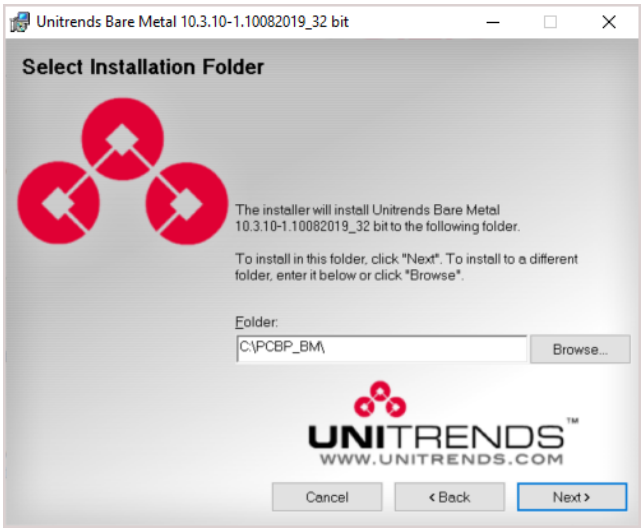

6 Click **Next** to begin the installation process. The installation can be interrupted at any time by clicking **Cancel**.

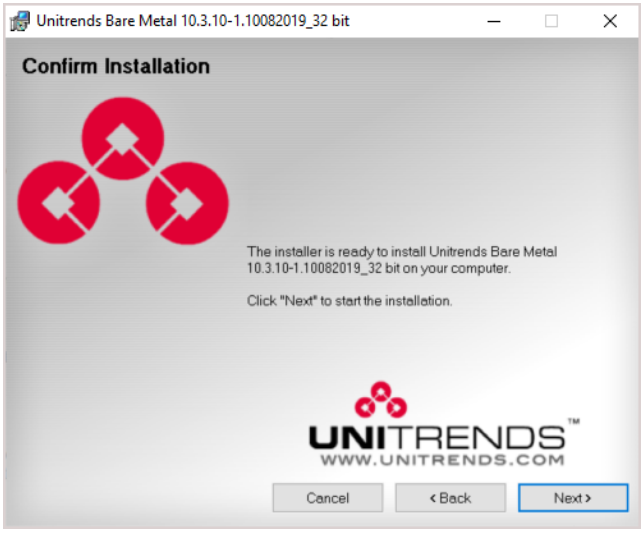

7 When installation is complete, click **Close** to exit the installer.

Installing the bare metal agent on a Windows asset running User Account Control

User Account Control (UAC) is enabled by default on Windows Vista, Windows Server 2008, and Windows Server 2012. To install *Unitrends_BareMetal.msi* on assets where UAC is enabled, you must invoke the installation with elevated privileges. See these topics for details:

- "Preparing to install on Vista"
- "Preparing to install on Windows Server 2012/2008"
- "To install the bare metal agent on a Windows asset running User Account Control"

Preparing to install on Vista

You must log in to the server as a user that has been assigned administrator privileges on the Vista server. Members of the Administrator group who have not been assigned administrator privileges on this Vista server cannot install the bare metal agent.

Preparing to install on Windows Server 2012/2008

Do one of the following to install on Windows Server 2012/2008:

- Log in using the local system Administrator account.
- Disable User Account Control (UAC) before you install the bare metal agent. (After you have installed the agent you can re-enable UAC.).

To install the bare metal agent on a Windows asset running User Account Control

1 Download *Unitrends_BareMetal.msi* and save it to the Windows machine. To download the file:

- Go to https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads.
- On the Unitrends Downloads page, click the Agent add-ons **Link**.

Operating System	Applies To	Latest Agent Installer
Windows	Windows 2008 R2 - 2022 x64 Agent Windows 7 - 11 x64 Agent	.MSI
	PowerControls OnTrack for MS Exchange, Sharepoint, and SQL	.EXE
	EndPoint Backup x64 Agent	Link
	Windows Baremetal ISO	.ISO
Linux Agents	RHEL, Alma, CentOS, Rocky, Oracle Agent	.RPM
	Ubuntu and Debian Agent	.DEB
	General Linux Distributions Agent	CNT install script
Legacy Agents	32bit agents, Windows 2003-2008 (non-R2), Mac, Solaris, AIX, UnixWare, HP-UX	Link
Agent add-ons	Oracle DB Add-on for Linux, Cold Bare Metal Agent, Hyper-V Recovery Assurance	Link — Click here

- On the Agent Add-ons page, click the Bare Metal Agent **MSI**.

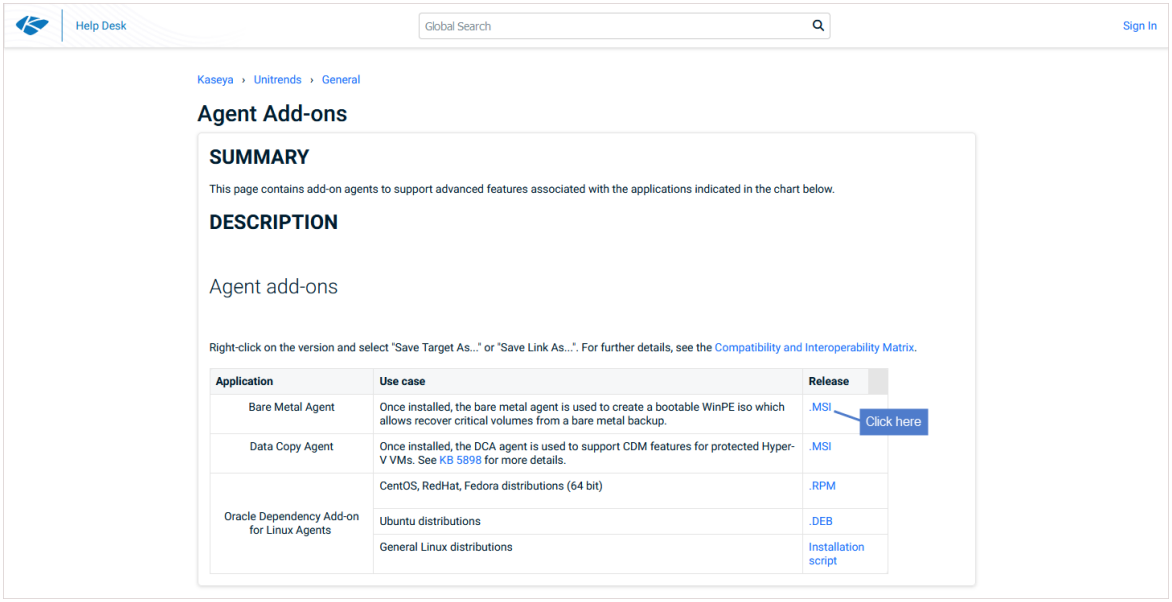

2 Log in to the Windows machine and select **Start > All Programs > Accessories**.

3 Right click **Command Prompt** and select **Run as administrator**.

4 Select **Yes** in the **UAC** window to continue.

5 Issue this command to install the agent, where *FullInstallPath* is the full path of the location where you saved *Unitrends_BareMetal.msi*:

```
Msiexec /package C:\FullInstallPath\Unitrends_BareMetal.msi
```

For directories with spaces in the name, add quotes to the command. For example, to download *Unitrends_BareMetal.msi* to *C:\Program Files*, use this command:

```
Msiexec /package "C:\Program Files\Unitrends_BareMetal.msi"
```

6 Exit the Command Prompt window and restart the server to apply the changes.

7 For Windows Server 2012/2008 only — After you have installed the agent, locate the Unitrends Agent entry on the Windows Start menu, right click on the **Unitrends Agent Menu**, and select **Run as administrator**. (This is required following the initial installation only.)

8 Additional configuration is required for bare metal protection of Vista and Windows Server 2012/2008 assets. Contact support for assistance (see "Support for Unitrends appliances" on page 25).

Command-line installer for Windows agents

With the command-line installer, you can install, remove, and repair the Windows agent. This method installs the following:

- The Windows agent.
- The Windows Volume CBT driver (used to enable incremental image-level backups). To enable the Volume CBT driver, you must reboot the Windows server after the agent has been installed.
- For Hyper-V servers or Windows servers with the Hyper-V role enabled, the Hyper-V CBT driver is installed. This driver is used for faster Hyper-V incremental backups. (You do not need to reboot to enable this driver.)
- For Microsoft SQL and Exchange servers, SQL and Exchange components are installed. You can then run application backups to protect these databases. For more information, see "Exchange backup requirements and considerations" and "SQL backup requirements and considerations".

The agent installer utilizes the **msiexec** command to manage the Windows agent from the command line. However, not all of the msiexec default parameters are supported with the installer. The following msiexec parameters are available:

/i - Installs and updates the software.

/f - Repairs the software.

/uninstall - Removes the software.

/quiet - Installs software in quiet mode with no user interaction.

/l* - Enables logging.

FORCE_BOOT - If set to True, restarts the Windows machine after installing the agent. If set to False, does not restart after installing the agent.

The following table describes the optional command-line parameters. These parameters are case sensitive and must be entered in upper case on the command line. The values specified for the parameter are not case sensitive. These options can also be used in conjunction with Microsoft's Group Policy methodology to deploy mass agent installations.

Parameter name (case sensitive)	Parameter value	Default value
CheckVersion	Yes | No	Yes

Parameter name (case sensitive)	Parameter value	Default value
		Yes enforces appliance/agent version compatibility. If the appliance version is older than the agent, an error is raised and the agent is not installed.
USNAPS	True | False	False
BARE_METAL	True | False	True
REMOTE_ADMIN	True | False	True
ODM	True | False	True
SQL_AGENT	True | False	True if SQL Server is installed on the asset, otherwise False.
EXCHANGE_ AGENT	True | False	True if Microsoft Exchange server is installed on the asset, otherwise False.
IP		127.0.0.1
FIREWALL	True | False	False

Following are Windows agent installer command-line examples:

Note: You must install the agent to the *\PCBP* directory under the Windows system drive, usually *C:*.

- Example 1 — Install *Unitrends_Agentx64.msi* with default values, where *C:* is the Windows system drive:

```
msiexec /i "C:\PCBP\Unitrends_Agentx64.msi"
```

- Example 2 — Install *Unitrends_Agentx64.msi* in quiet mode with override version checking (agent is installed without checking appliance version), where *C:* is the Windows system drive:

```
msiexec /quiet CHECKVERSION=No /i "C:\PCBP\Unitrends_Agentx64.msi"
```

- Example 3 — Install *Unitrends_Agentx64.msi* with default values and turn on logging, where *C:* is the Windows system drive and *C:\temp\Unitrends.log* is the path of the log file named *Unitrends.log*:

```
msiexec /quiet /l* "C:\temp\Unitrends.log" /i "C:\PCBP\Unitrends_Agentx64.msi"
```

- Example 4 — Uninstall *Unitrends_Agentx64.msi*, where *C:* is the Windows system drive:

```
msiexec /quiet /uninstall "C:\PCBP\Unitrends_Agentx64.msi"
```

Agent deployment using Group Policy

With the command-line installer and optional parameters (described in "Command-line installer for Windows agents" on page 377 above), you can use Microsoft Group Policy to deploy mass agent installations. See the Microsoft Group Policy documentation for details on downloading and using the Group Policy software.

To deploy the Windows agent by using Group Policy

1. An Active Directory domain is needed. Begin by creating a new Group Policy Object (see Microsoft documentation for details).
2. When the object has been created, select and edit it. If using the Group Policy Management Console, this action invokes the Group Policy Object Editor.
3. Determine whether the Group Policy Object will be a computer configuration or a user configuration. Depending on the configuration selected, expand the Software Settings folder and select the **Software Installation** option.
4. Right-click **Software Installation**, point to **New**, and then click **Package**.
5. In the Open dialog box, type the full Universal Naming Convention (UNC) path of the shared installer package. For example:

   ```
   \\fileserver\share\fileName.msi
   ```

   IMPORTANT! Do not use the **Browse** button to access the location. Make sure to type in the UNC path to the shared installer package.
6. Click **Open**.
7. Click **Assigned**, and then click **OK**. The package is listed in the right pane of the Group Policy window.
8. Close the **Group Policy** snap-in, click **OK**, and then quit the Active Directory Users and Computers snap-in.
9. Link the Group Policy Object to the domain by dragging the object to the domain name.
10. Double click the Group Policy Object name to add computers or users to the object.
    - If the computer configuration was selected, the agent will be installed on the specified computers when the computers are restarted.
    - If the user configuration was selected, the agent will be installed on any computer in the domain where the specified users log in to the domain.
    - Once installed, an entry for the application displays in the Add/Remove Programs interface of the Microsoft Windows operating system.
11. For assets running Windows Vista and earlier operating systems, reboot the asset. (A reboot is not required for newer operating systems.)

Updating and removing the Windows agent

Windows agent updates can be pushed to assets from the appliance or installed manually.

Updating the agent updates all applicable components in *Unitrends_Agentx86.msi*, *Unitrends_Agentx64.msi*, or *Unitrends_BareMetal.msi* during installation.

Notes:

- For servers that are using Windows deduplication, you must run a new full backup after upgrading the agent.
- In some other cases, a new full backup is needed after upgrading to the agent. In these cases, the appliance cannot run an incremental or differential until a full backup runs. If you attempt an on-demand incremental or differential, the appliance promotes the backup to a full. If the next job is a scheduled incremental or differential, the job fails. After this failure, the appliance promotes the next scheduled run to a full. Once this full succeeds, subsequent incrementals and differentials run as scheduled.
- The Windows agent includes a Windows Volume CBT driver to enable incremental image-level backups. If installing the agent manually, you can opt to include this driver. If push-installing agent updates, this driver is installed or updated as needed on assets that are eligible for image-level backups. To enable this driver, you must reboot the Windows asset after the Volume CBT driver has been installed or updated. (The last driver update was in agent version 10.3.3. If you are running a pre-10.3.3 agent version or if you have never installed the Volume CBT driver, a reboot is required.)
- The Windows bare metal agent is required only if you need to run bare metal backups that can be used to perform image-based bare metal recovery of your Windows asset. For most Windows assets, regular Windows backups are used to perform integrated bare metal recovery, which is the recommended approach. To see if bare metal backups are required for your Windows asset, see "Which bare metal method should I use?" on page 1207. The bare metal agent must be installed manually (push install is not supported).

See the following topics for details:

- "Push installing agent updates" on page 380
- "Manually updating and removing Windows agents" on page 382
- "Manually installing and uninstalling the Hyper-V CBT driver" on page 384

Push installing agent updates

Pushing updates to Windows assets greatly reduces administration time and ensures that the latest protection software is running on your assets.

Push install update notifications

Any time an agent update is available for Windows assets, a notification displays in the Alerts area of the Global menu. Note that alerts display only for assets that meet the push install requirements described in "Windows agent requirements".

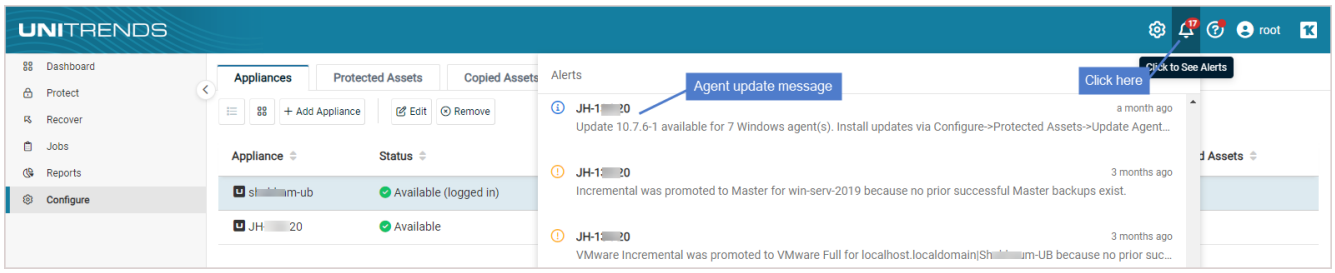

To push install agent updates

Use this procedure to update the Windows core agent.

1 Select **Configure > Protected Assets**. All protected assets display.

2 Check the Agent Version to see the agent that is currently installed.

3 Do one of the following:

- To install available updates on all eligible assets, select **Update Agent > Update All**.

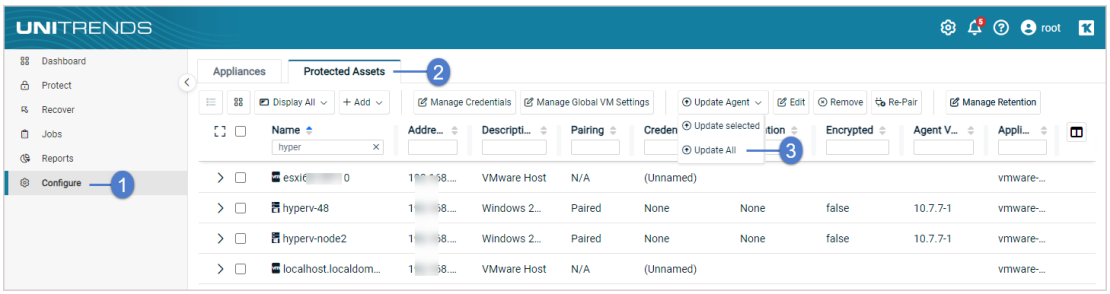

- To install available updates on a subset of eligible assets, check boxes to select assets to update, then select **Update Agent > Update selected**.

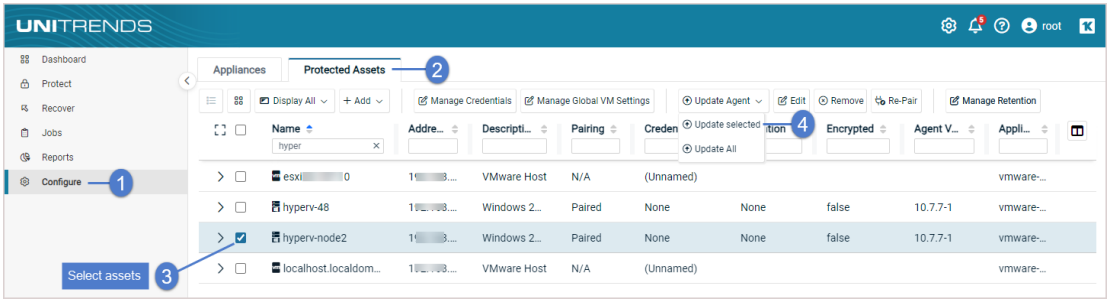

4 Updated Windows agents are installed on all selected Windows assets meeting these conditions:

- Trust credentials are valid.
- No backup or recovery job is currently in progress or scheduled to run soon for the asset.
- Push update requirements have been met (see "Windows agent requirements").
- Updates are available for the asset (asset is not running the latest agent release).

- The Unitrends appliance(s) protecting the Windows asset is running an equal or higher version than the agent that will be installed.

  Note: If you see this error, the Windows asset is paired to an appliance running an older version than the agent you're trying to install: *Appliance version validation failed. Please verify all paired appliance versions meet or exceed the agent*. Update the appliance to enable the push-install.

If applicable, the following are also updated (if new versions are present):

- For Hyper-V servers or Windows servers with the Hyper-V role enabled, the Hyper-V CBT driver is updated. This driver is used for faster Hyper-V incremental backups. (You do not need to reboot to enable this driver.)
- For Microsoft SQL and Exchange servers, SQL and Exchange components are updated. These are used to run application backups for these databases. For more information, see "Exchange backup requirements and considerations" and "SQL backup requirements and considerations".
- For Windows assets that are eligible for image-level backups, the Windows Volume CBT driver is installed or updated if an update is available. This driver is used to enable incremental image-level backups. To enable this driver, you must reboot the Windows server after the Volume CBT driver has been installed or updated. (The last driver update was in agent version 10.3.3. If you are running a pre-10.3.3 agent version, a reboot is required.)

5 If the Windows server uses Windows deduplication, run a new full backup.

Manually updating and removing Windows agents

Use these procedures to manually update, remove, or repair the Windows agent:

- "To manually update the Windows agent" on page 382
- "To uninstall the Windows agent" on page 382
- "To repair the Windows agent" on page 384

To manually update the Windows agent

Install the latest agent version as described in "Manually installing the Windows agent" on page 366.

Notes:

- If you are running Hyper-V incremental backups and are upgrading from a pre-10.1.0-3 agent version, you must manually uninstall the older Windows agent before installing the latest agent. In all other cases, it is not necessary to uninstall existing agent software.
- Uninstall the older agent by using the Windows Add/Remove Programs interface. For details, see "To uninstall the Windows agent" on page 382.

To uninstall the Windows agent

Note: If you installed the agent by using the command line method or by using Group Policy, remove or repair the agent by using the command line method. See "Command-line installer for Windows agents" on page 377 for details.

1 In the Windows Add/Remove Programs interface, select the **Unitrends Agent** in the list.

2 Click **Uninstall**, then **Yes** to confirm.

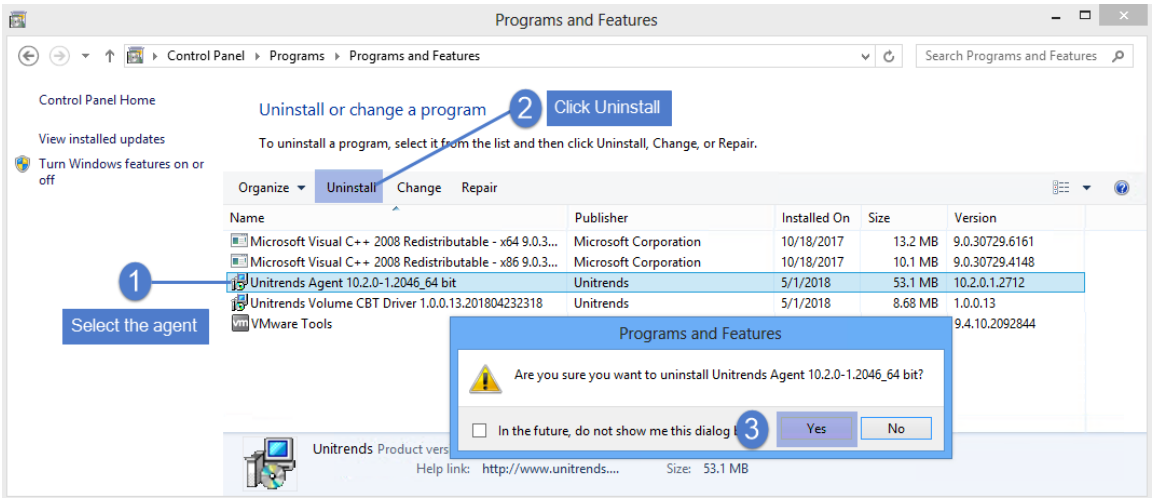

The agent is uninstalled:

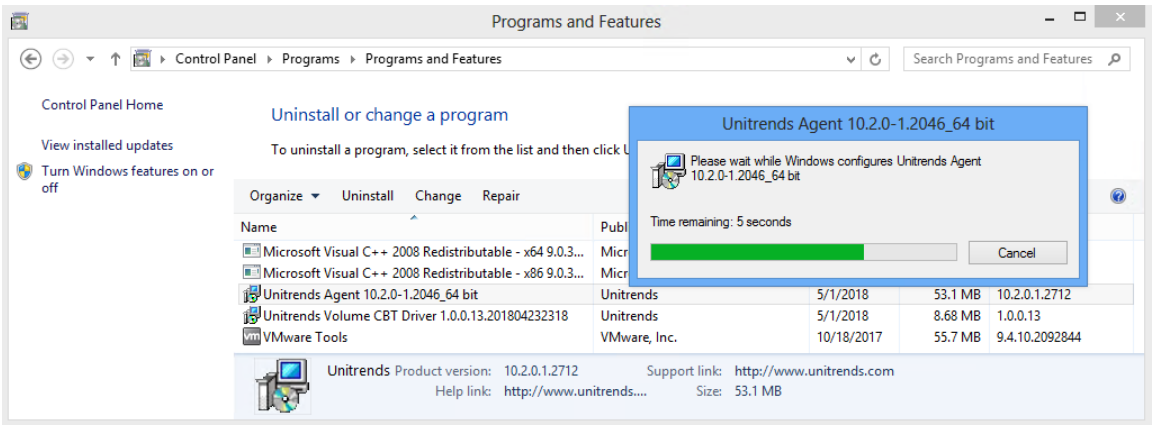

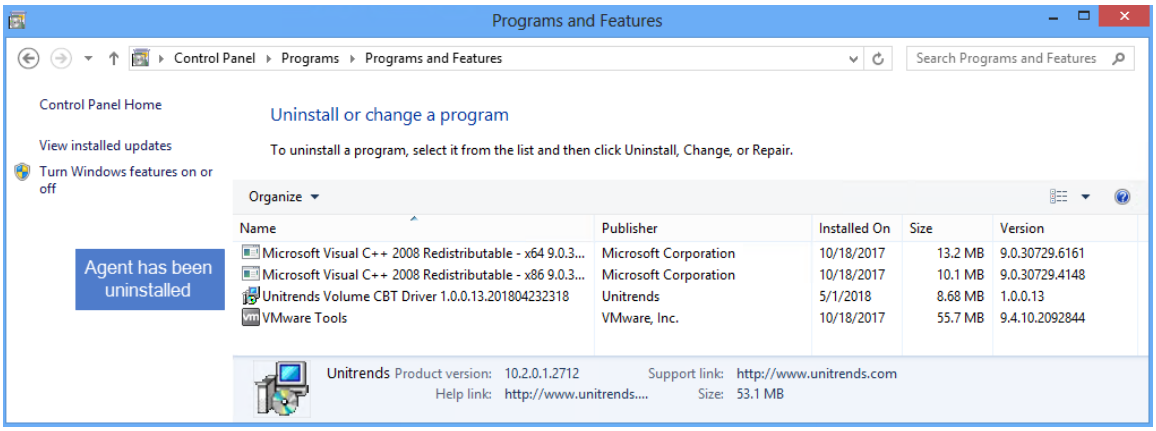

3 If needed, repeat this procedure to remove the Unitrends Volume CBT Driver.

To repair the Windows agent

Note: If you installed the agent by using the command line method or by usingGroup Policy, repair the agent by using the command line method. See "Command-line installer for Windows agents" on page 377 for details.

1 In the Windows Add/Remove Programs interface, select the **Unitrends Agent** in the list.

2 Click **Repair**.

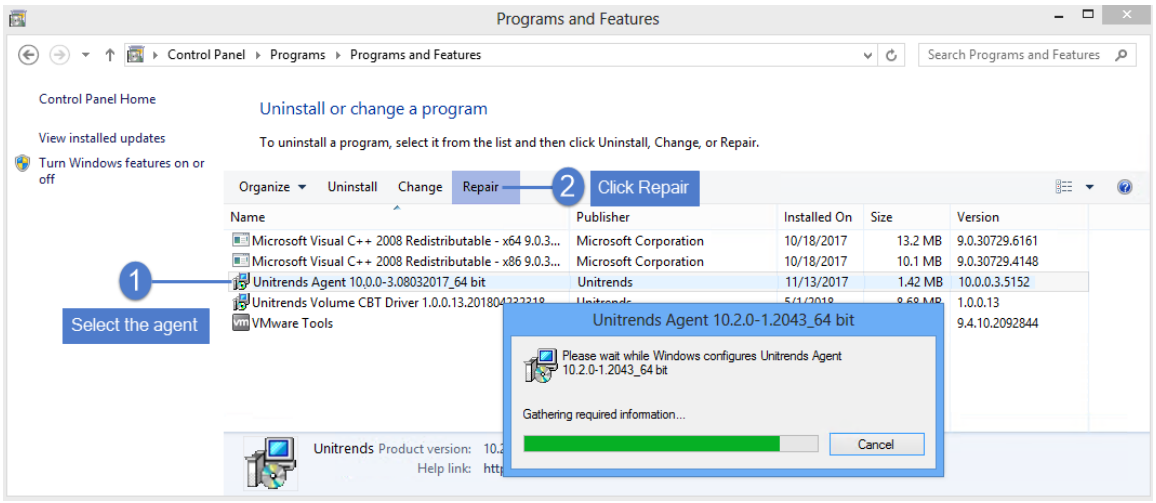

The agent is repaired and updated:

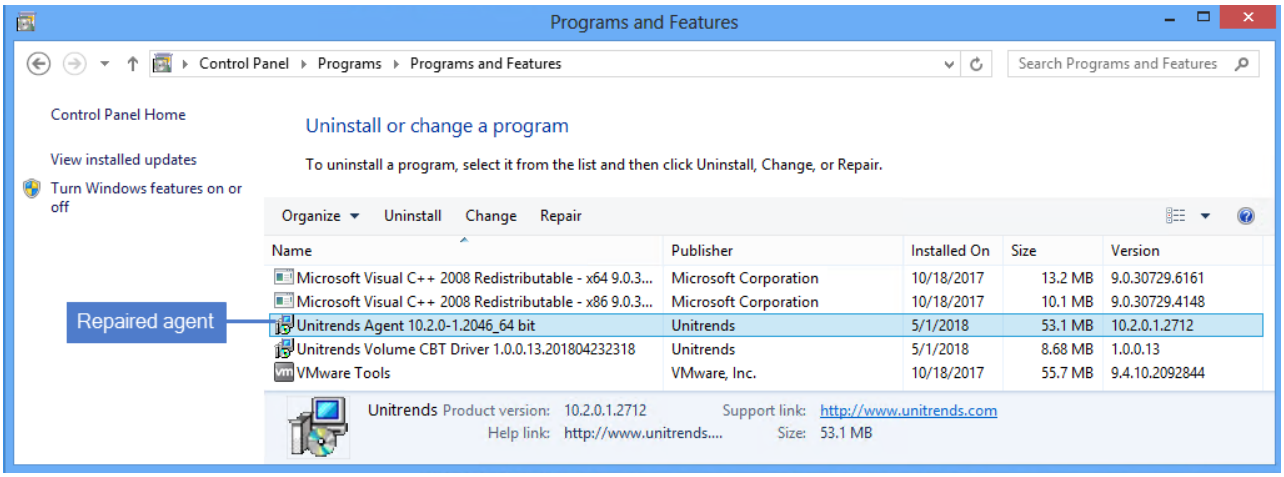

Manually installing and uninstalling the Hyper-V CBT driver

The Hyper-V CBT driver greatly increases performance for Hyper-V incremental backups. This driver is automatically installed with the Windows core agent on Hyper-V servers and Windows servers that have the Hyper-V role enabled. In most cases, you should not need to manually install or remove this driver. Use these procedures in the following cases:

- Customer Support has indicated that you need to reinstall or remove the driver.

- Your Hyper-V environment is using the Microsoft Hyper-V Replicas feature. This feature is not compatible with the Hyper-V CBT driver and you must uninstall the driver. Once the driver has been uninstalled, Hyper-V incrementals are supported but do not use the driver.

Use these procedures to install and uninstall the Hyper-V CBT driver:

- "Locate the script"
- "To install the Hyper-V CBT driver manually by using Windows File Explorer" on page 386
- "To install the Hyper-V CBT driver manually by Windows command line" on page 386
- "To uninstall the Hyper-V CBT driver manually by using Windows File Explorer" on page 386
- "To uninstall the Hyper-V CBT driver manually by Windows command line" on page 387

Locate the script

Use Windows File Explorer or the Windows Command Prompt to locate the driver scripts. The install and uninstall scripts reside under the *\PCBP\Win2012.dir\<DriverFolder>* directory on the volume where the Windows core agent was installed. The DriverFolder name varies by version. In the following example, the Windows agent was installed on *C:* and the DriverFolder is *04-20-2017,9.2.0*:

File Explorer example:

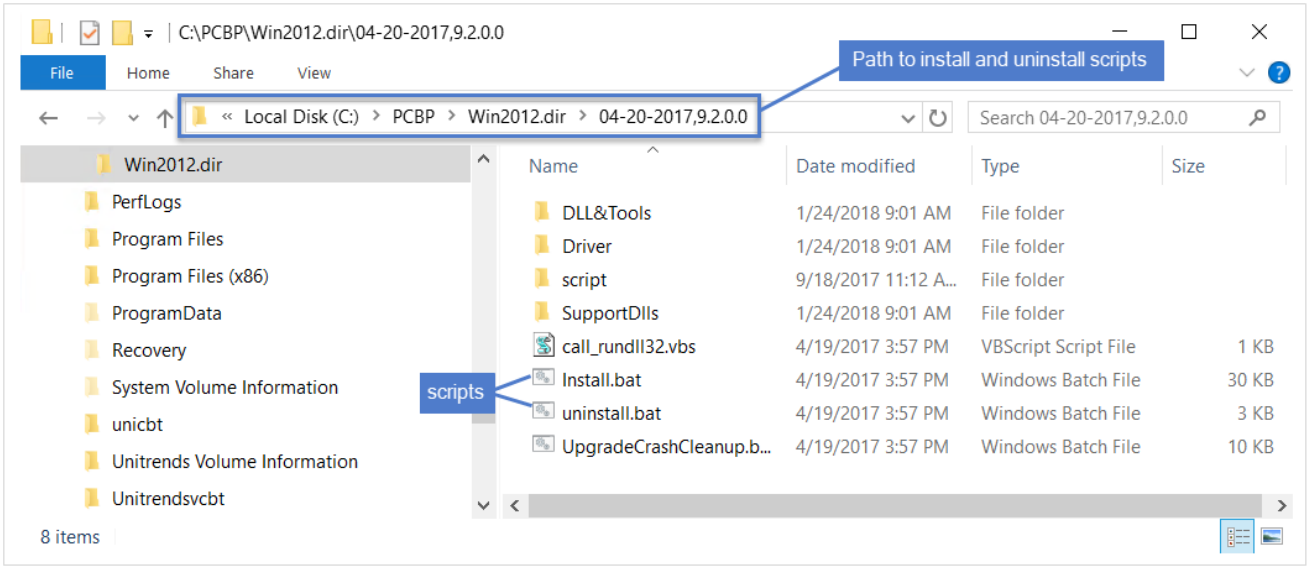

Command Prompt example:

1. Launch the Windows Command Prompt as a user with administrative privileges.
2. Change to the *\PCBP\Win2012.dir* directory under the volume where the Windows agent was installed (*C:* in this example):

```
# cd C:\PCBP\Win2012.dir
```

3 Issue this command to list all subdirectories:

```
# dir
```

4 View the command output and note the name of the driver directory. The directory name varies by driver version and is in the format *<Date>,<Version>*. For example, *04-20-2017,9.2.0.0*.

To install the Hyper-V CBT driver manually by using Windows File Explorer

1 Browse to the driver directory under *\PCBP\Win2012.dir* on the volume where the Windows agent was installed.

2 Double-click **Install.bat** to run the installer.

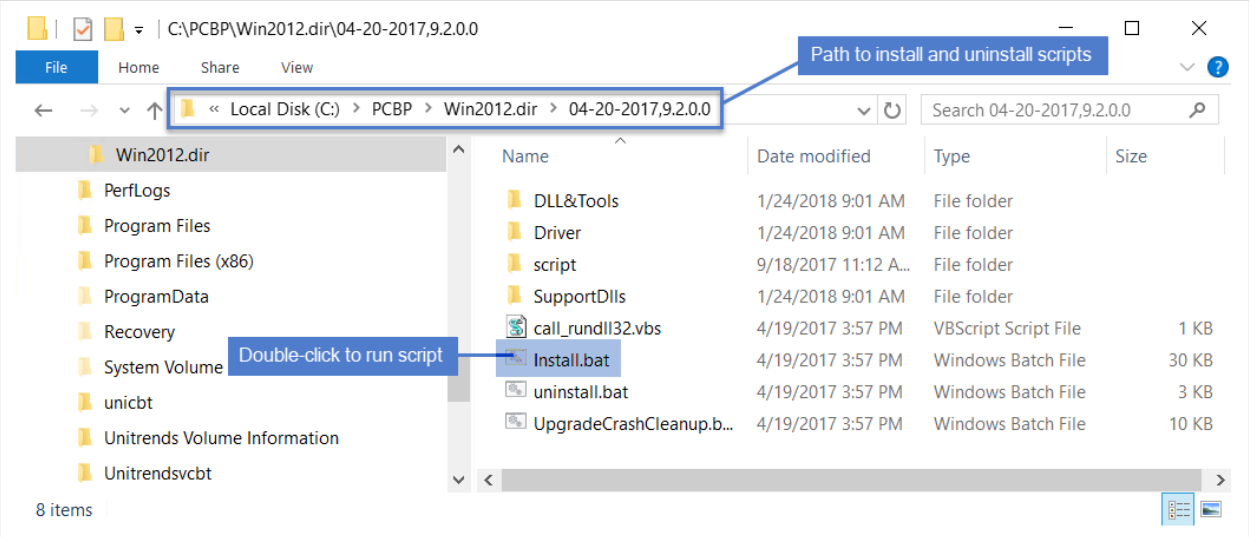

To install the Hyper-V CBT driver manually by Windows command line

1 Launch the Windows Command Prompt as a user with administrative privileges.

2 Change to the driver directory under *\PCBP\Win2012.dir* on the volume where the Windows agent was installed. The directory name varies by driver version. In this example, the directory is *04-20-2017,9.2.0.0* and the Windows agent was installed on volume *C:*.

```
# cd C:\PCBP\Win2012.dir\04-20-2017,9.2.0.0
```

3 Issue this command to run the install script:

```
# Install.bat
```

To uninstall the Hyper-V CBT driver manually by using Windows File Explorer

Note: The Hyper-V CBT driver is installed each time the Windows agent is installed or updated. Repeat this procedure as needed after updating or reinstalling the Windows agent.

1 Browse to the driver directory under *\PCBP\Win2012.dir* on the volume where the Windows agent was installed.

2 Double-click **uninstall.bat** to run the installer.

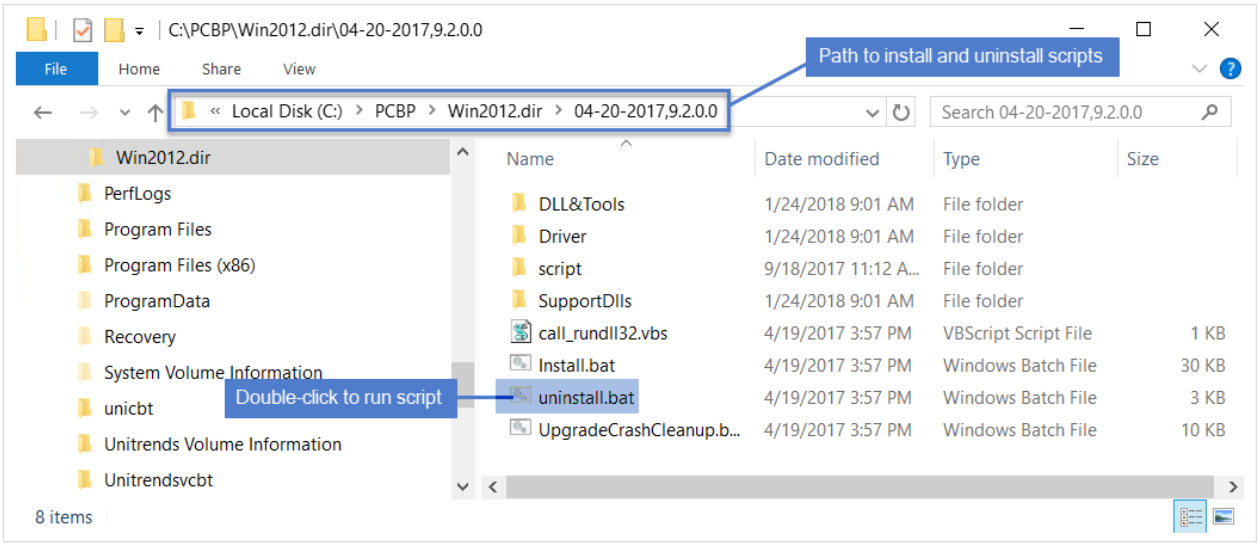

To uninstall the Hyper-V CBT driver manually by Windows command line

Note: The Hyper-V CBT driver is installed each time the Windows agent is installed or updated. Repeat this procedure as needed after updating or reinstalling the Windows agent.

1 Launch the Windows Command Prompt as a user with administrative privileges.

2 Change to the driver directory under *\PCBP\Win2012.dir* on the volume where the Windows agent was installed. The directory name varies by driver version. In this example, the directory is *04-20-2017,9.2.0.0* and the Windows agent was installed on volume *C:*.

```
# cd C:\PCBP\Win2012.dir\04-20-2017,9.2.0.0
```

3 Issue this command to run the uninstall script:

```
# unnstall.bat
```

Installing and updating the Linux agent

Unitrends protects most Linux distributions, including CentOS, Debian, Red Hat, SUSE, and Ubuntu. Before adding a Linux asset to the Unitrends appliance, you must install an agent.

Once the agent is installed, you can update to a newer agent version using these same installation procedures. It is not necessary to remove the old agent. If an agent is already installed, you have the option to save any custom agent settings during installation.

Use the procedures in this section to install or update the Linux agent:

- "Preparing to install the Linux agent" on page 388
- "Installing the Linux agent" on page 392

Preparing to install the Linux agent

Unitrends provides several Linux agent installers. Unitrends recommends using the RPM-based or dpkg-based installers when possible, so that needed dependencies are automatically installed with the agent. If these installers are not supported for your Linux distribution, use the GZEXE installers. With the GZEXE installers, you might need to install dependencies before installing the agent.

See "Linux distributions and agent installers" below to determine which installer to use for your Linux asset. You can download the agent installers from the Latest Agent Releases on the Unitrends Downloads page. You might not see an agent for the particular Linux distribution that you are using, but if it is a supported distribution listed in the Unitrends Compatibility and Interoperability Matrix, the standard Linux agent will work with your machine. For Oracle Linux assets, use the CentOS or Red Hat agent.

Linux distributions and agent installers

Linux distributions	Agent installers
• CentOS 64-bit • Oracle Linux 64-bit • Red Hat 64-bit	RPM-based installers. Automatically installs dependencies. For details, see "To install the Linux agent on CentOS, Oracle Linux, and Red Hat".
• All 64-bit distributions listed in the Unitrends Compatibility and Interoperability Matrix	GZEXE installers. For details, see "To install the Linux agent using GZEXE".
• Ubuntu 64-bit	dpkg-based installers. Automatically installs dependencies. For details, see "Installing the Linux agent on Ubuntu".
• All 32-bit distributions listed in the Unitrends Compatibility and Interoperability Matrix	Download the applicable agent from the 32-bit Agents page.

Where to access your agent on the Unitrends Downloads page:

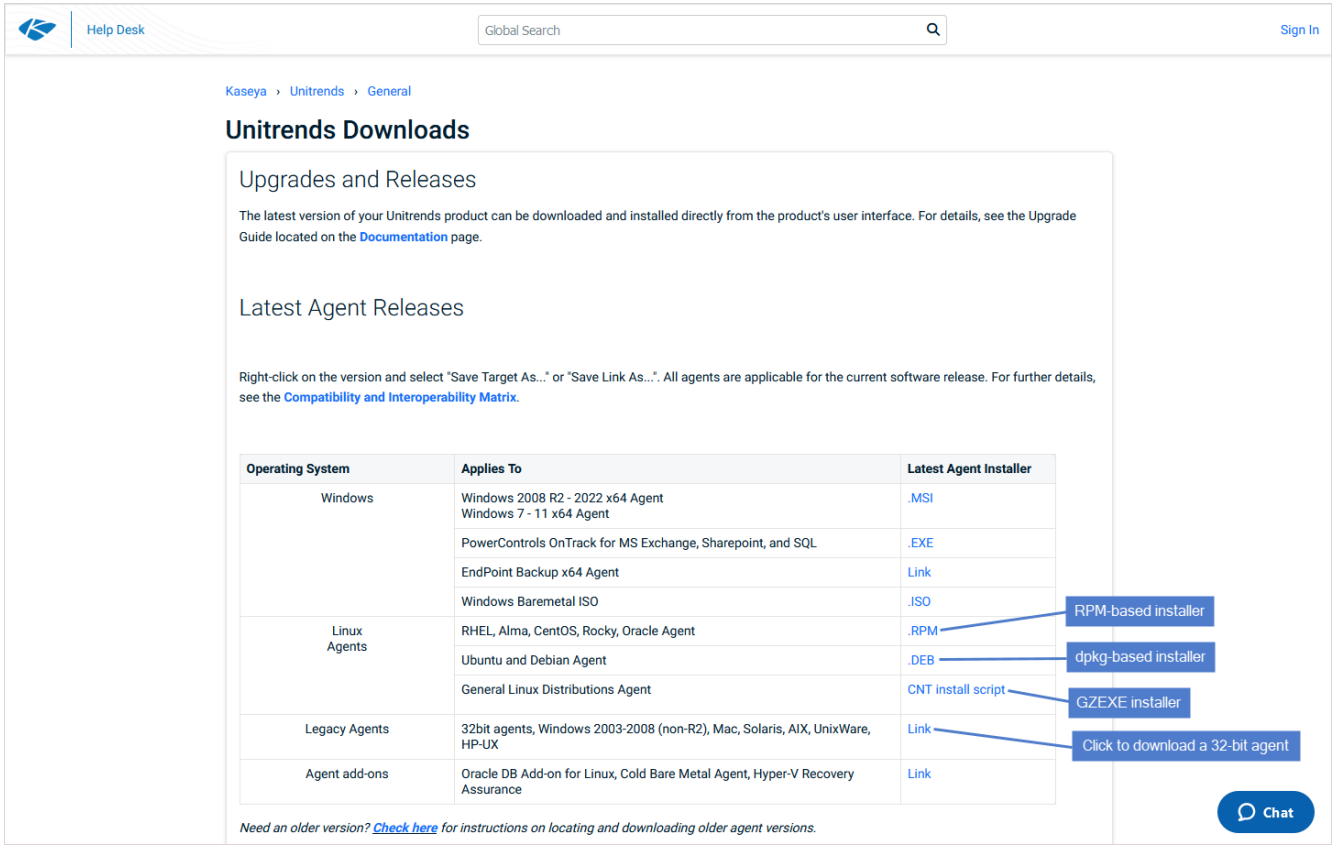

About Linux agent dependencies

When using GZEXE installers, you might need to install additional libraries. If this is the case, the installer stops the installation and lists the required dependencies. The dependencies it lists are the resources needed and not the name of the package you must install. The table below identifies the packages containing the commonly needed dependencies.

Dependencies by operating system

The following dependencies are required to protect Linux environments. Red Hat dependencies replace XINETD, which was a dependency for earlier versions.

Operating System	Dependencies
Red Hat 6 i386	• ed Packages are located on the installation media.
Red Hat 6 x86_64	• ed

Operating System	Dependencies
	• glibc.i686 • nss-softokn-freebl.i686 The following packages might need to be updated to match the version of a new dependency. • glibc.x86_64 (must match glibc.i686) • glibc-common.x86_64 (must match glibc.i686) • nss-softokn-freebl.x86_64 (must match nss-softokn-freebl.i686) Packages are located on the installation media.
Red Hat 8 x86_64 and Oracle Linux 8	• libwrap.so.0()(64bit) • libc.so.6 • libc.so.6(GLIBC_2.0) • libc.so.6(GLIBC_2.1) • libc.so.6(GLIBC_2.2) • libmenu.so.5()(64bit) • libncurses.so.5()(64bit) • glibc-common-2.17-260.el7.x86_64.rpm Packages are located on the installation media.
SUSE 15 SP1	• libcrypto.6 Packages are located on the installation media.
Oracle on Linux (for application backups only)	Samba is only used to protect Oracle with application backups. The Samba packages listed below only need to be installed if you wish to protect Oracle data with Unitrends application backups. For assistance with these packages, you can download the Oracle Dependency plug-in from the Agent Add-ons page. You must install the Linux agent before you can install this plug-in. Depending on your Linux distribution, use the Oracle Dependency for CentOS or Red Hat. Dependencies for Oracle application backups: • samba-client (for Oracle Linux 5 and CentOS 5) • cifs-utils (for most other Linux distributions)

Requirements for secure pairing of Unitrends Linux agents

Beginning in Linux agent release 10.7.5, a secure pairing is automatically established between the appliance and the Linux agent on each of its protected assets. This pairing enables Transport Layer Security (TLS) to encrypt data and authenticate connections between appliances and agents. Communication between appliances and agents is only allowed if there is a matching (paired) certificate.

This feature blocks any communication with Unitrends agent software that doesn't originate from a paired appliance (think of a Bluetooth headset, if it's not paired or in pairing mode, no one else can communicate with it). This eliminates the threat of a rogue appliance running backups or code against an agent.

To use the secure pairing feature, these requirements must be met:

- The Unitrends appliance must be running release 10.7.5 or higher.
- The Linux asset must be running agent release 10.7.5 or higher.
- The Unitrends appliance version must be equal to or higher than the Linux agent version.

  IMPORTANT! Be sure to upgrade your Unitrends appliance before upgrading your Linux agents.

  - Jobs will fail if you attempt to protect a 10.7.5 or higher agent with an appliance that is running an older release.
  - You cannot add an asset that is running a 10.7.5 or higher agent to an appliance that is running an older release. If you attempt this, you receive an error similar to: *Failed to save client: Registration for client* assetName *failed. The Unitrends System could not connect to the Unitrends Agent on* assetName*. Please ensure that the Agent software is installed on* assetName*, the Agent service is running (if applicable), and no firewall settings are preventing access*.
- The Linux agent listens for pairing requests on port 888. Ensure that port 888 is accessible on the Linux asset.
- The Linux asset must be running one of these versions:

  IMPORTANT! Do NOT install the 10.7.5 agent on other Linux versions that are not listed below. Instead, locate your Linux version on the Unitrends Downloads page and install the latest supported agent.

  - Alma Linux 9, 64-bit
  - CentOS 7, 64-bit
  - CentOS 9, 64-bit
  - Debian 10, 64-bit
  - OpenSUSE 42, 64-bit
  - Oracle Linux 8.1, 64-bit
  - RHEL 7, 64-bit
  - RHEL 8.4, 64-bit
  - RHEL 9, 64-bit

- Rocky Linux 9, 64-bit
- SLES 11 SP3, 64-bit
- SUSE 15, 64-bit
- Ubuntu 22.04, 64-bit

See "Secure agent pairing for Windows and Linux agents" on page 338 for details on working with this feature.

Installing the Linux agent

Installation procedures for the Linux agent vary by Linux distribution. See the following topics for instructions:

- "To install the Linux agent using GZEXE" on page 392
- "To install the Linux agent on CentOS, Oracle Linux, and Red Hat" on page 393
- "Installing the Linux agent on Ubuntu" on page 393

To install the Linux agent using GZEXE

This section explains how to install the agent using GZEXE installers, which are available for all supported Linux distributions. If the agent requires dependencies, the installer stops the installation and lists the required dependencies.

1. Save the applicable agent installer on the Linux machine that you want to add to the Unitrends appliance. You can download the installer from the Latest Agent Releases on the Unitrends Downloads page.
2. Open a terminal, and log in as root user.
3. Change directories to the location where you have saved the agent installer, and run the command **`ls -l`** to view the installer file and determine whether you have execute permission. If necessary, add execute permission using the command:

   ```
   # chmod +x <file_name>
   ```

4. Perform one of the following depending on whether you are using a 32-bit or 64-bit installer:
   - For a 32-bit installer, run the command:

     ```
     # ./lnx32_cnt
     ```

   - For a 64-bit installer, run the command:

     ```
     # ./lnx64_cnt
     ```

5. If necessary, install any required dependencies. The installer notifies you of any dependencies the agent needs. The dependencies listed are the resources needed and not the name of the package that you must install. For more about locating and installing dependencies, see "About Linux agent dependencies" on page 389.

   Run the applicable command from step 4 above after installing the dependencies.

6 (Optional) To protect Oracle databases, install the Oracle Dependency from the Agent Add-ons page.

7 Enter the hostname for the backup appliance that will protect the asset.

8 If you are using a firewall, configure it to allow the Unitrends appliance to communicate with the Linux machine. For details, see "Configuring a Linux firewall to communicate with the Unitrends appliance" on page 395.

9 Add the Linux asset to your Unitrends backup appliance to begin protecting it. For instructions, see "To add an agent-based asset" on page 289.

To install the Linux agent on CentOS, Oracle Linux, and Red Hat

For CentOS, Oracle Linux, and Red Hat assets, you can use RPM-based installers that often automatically install the necessary dependencies if connected to a remote repository.

1 Save the applicable agent installer on the Linux machine that you want to add to the Unitrends appliance. You can download the installer from the Latest Agent Releases on the Unitrends Downloads page.

> **Note:** For Oracle Linux assets, download the CentOS or Red Hat agent installer.

2 Open a terminal, and log in as root user.

3 Change directories to the location where you have saved the agent installer.

4 Perform one of the following depending on whether you are using a 32-bit or 64-bit installer:

- For a 32-bit asset, run the command:

```
# yum localinstall --nogpgcheck unitrends-linux-agent-<release>.<build_date>.i386.rpm
```

- For a 64-bit asset, run the command:

```
# yum localinstall --nogpgcheck unitrends-linux-agent-<release>.<build_date>.x86_64.rpm
```

5 Install any required dependencies. The installer notifies you of any dependencies the agent needs. The dependencies listed are the resources needed and not the name of the package that you must install. For more about locating and installing dependencies, see "About Linux agent dependencies" on page 389.

6 (Optional) To protect Oracle databases, iinstall the Oracle Dependency from the Agent Add-ons page.

7 If you are using a firewall, configure it to allow the Unitrends appliance to communicate with the Linux machine. For details, see "Configuring a Linux firewall to communicate with the Unitrends appliance" on page 395.

8 Add the Linux asset to your Unitrends backup appliance to begin protecting it. For instructions, see "To add an agent-based asset" on page 289.

Installing the Linux agent on Ubuntu

For Ubuntu assets, you can use dpkg-based installers that often automatically install all necessary dependencies if connected to a remote repository. You can choose to install the agent using core utilities or the GDebi tool. If you install using core utilities, you must run two commands if the necessary dependencies have not been installed on your Ubuntu machine. If you use the GDebi tool, one command installs the agent and all necessary dependencies.

For instructions, see the following topics:

To install the Linux agent on Ubuntu using core utilities

For Ubuntu assets, you can use dpkg-based installers that install all necessary dependencies.

Note: This procedure might require you to run two commands. The first command installs the agent if the necessary dependencies are already installed on the asset. If the agent requires dependencies, the second command in this procedure installs them and then installs the agent. If you have installed the GDebi tool on the asset, you can use it to install the agent using only one command. For details, see "To install the Linux agent on Ubuntu using GDebi" on page 394.

1. Save the applicable agent installer on the Linux machine that you want to add to the Unitrends appliance. You can download the installer from the Latest Agent Releases on the Unitrends Downloads page.
2. Open a terminal and change directories to the location where you saved the agent installer.
3. Perform one of the following:
   - For the 32-bit installer, run the command:

     ```
     # sudo dpkg -i unitrends-linux-agent-<release>-<build_date>.i386.deb
     ```

   - For the 64-bit installer, run the command:

     ```
     # sudo dpkg -i unitrends-linux-agent-<release>-<build_date>.amd64.deb
     ```

4. If the installer stopped because the agent requires dependencies, run the following command to install them:

   ```
   # sudo apt-get install -f
   ```

5. If you are using a firewall, configure it to allow the Unitrends appliance to communicate with the Linux machine. For details, see "Configuring a Linux firewall to communicate with the Unitrends appliance" on page 395.
6. Add the Linux asset to your Unitrends backup appliance to begin protecting it. For instructions, see "To add an agent-based asset" on page 289.

To install the Linux agent on Ubuntu using GDebi

To install the agent with this procedure, you must have installed the GDebi package on your Ubuntu assets. Installation of the agent using GDebi requires only one command. To install the agent using core utilities, see "To install the Linux agent on Ubuntu using core utilities" on page 394.

1. Save the applicable agent installer on the Linux machine that you want to add to the Unitrends appliance. You can download the installer from the Latest Agent Releases on the Unitrends Downloads page.
2. Open a terminal and change directories to the location where you saved the agent installer.
3. Perform one of the following depending on whether you are using a 32-bit or 64-bit installer:

- To install the 32-bit agent, run the following command:

```
# sudo gdebi unitrends-linux-agent-<release>-<build_date>.i386.deb
```

- To install the 64-bit agent, run the following command:

```
# sudo gdebi unitrends-linux-agent-<release>-<build_date>.amd64.deb
```

4 If you are using a firewall, configure it to allow the Unitrends appliance to communicate with the Linux machine. For details, see "Configuring a Linux firewall to communicate with the Unitrends appliance" on page 395.

5 Add the Linux asset to your Unitrends backup appliance to begin protecting it. For instructions, see "To add an agent-based asset" on page 289.

Configuring a Linux firewall to communicate with the Unitrends appliance

If you are protecting a Linux machine with a firewall, you must configure the firewall to allow communication with the Unitrends appliance before you can add the Linux machine as an asset.

To configure the Linux firewall

1 Modify the Linux machine's firewall settings to allow port 1743 and ports 1745 through 1845.

> **Note:** If you have altered the general configuration variable dataport_count from its default value of 100, this new value is added to 1745 to produce your required port range (e.g., If dataport_count is set to 75, your firewall must allow port 1743 and ports 1745 through 1820).

2 Open a terminal or text editor with root access and log in as user root.

3 Run the following command:

```
# /usr/bp/bin/bputil -p "Configuration Options" data 1745 /usr/bp/bpinit/master.ini
```

4 (If needed) If the Linux asset will be running agent version 10.7.5 or higher, you must also modify the Linux machine's firewall settings to allow port 888. (Beginning in Linux agent release 10.7.5, a secure pairing between the agent and appliance is required. This pairing is established over port 888.)

Removing the Linux agent

Use one of the following commands to uninstall the Unitrends agent.

- For an agent installed with the GZEXE installer, issue the **uninstall** command from the directory where the agent is installed. For example, to remove the agent from the default install location, enter:

```
# /usr/bp/uninstall
```

- For an agent installed with the RPM-based installer, issue this command:

```
# yum remove unitrends-linux-agent
```

- For an agent installed with the dpkg-based installer, issue this command:

```
# sudo apt-get remove unitrends-linux-agent
```

Installing and updating the AIX agent

Before adding an AIX asset to the Unitrends appliance, you must install an agent.

Once the agent is installed, you can update to a newer agent version using these same installation procedures. It is not necessary to remove the old agent. If an agent is already installed, you have the option to save any custom agent settings during installation.

Use the procedures in this section to install or update the AIX agent.

Preparing to install the AIX agent

The AIX agent enables you to back up, verify, and restore AIX server data. The AIX agent cannot be used to protect encrypted file systems.

Before installing an AIX agent:

- Make sure your AIX system is running a supported version listed in the Unitrends Compatibility and Interoperability Matrix.
- Add the name and IP address of the Unitrends appliance to the hosts file on the AIX server.

To install the AIX agent

1. Log in to the AIX machine as user root.
2. Download the AIX agent from the Unitrends Downloads page (https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads) into the */tmp* folder on the AIX machine.
3. Execute the following commands to begin the installation:

```
# chmod 755 /tmp/aix5_cnt
# /tmp/aix5_cnt
```

4. Follow the steps on the screen. You will be asked to specify the directory location where the agent will be installed. If this is a reinstall, you will be asked if you wish to overwrite certain files. Type the interrupt character or press return to continue. The files will be moved to their permanent location.
5. When the configuration process finishes, you will be prompted to reboot the AIX machine to complete the installation.

Once you have installed the agent, you are ready to add your AIX asset to the Unitrends appliance. For details, see "To add an agent-based asset" on page 289.

Removing the AIX agent

To uninstall the agent, log in to the AIX server and run the **uninstall** command from the directory where the agent is installed. For example, to remove the agent from the default install location, enter this command:

```
# /usr/bp/uninstall
```

Installing and updating the HP-UX agent

Before adding an HP-UX asset to the Unitrends appliance, you must install an agent.

Once the agent is installed, you can update to a newer agent version using these same installation procedures. It is not necessary to remove the old agent. If an agent is already installed, you have the option to save any custom agent settings during installation.

Use the procedures in this section to install or update the HP-UX agent.

Preparing to install the HP-UX agent

The HP-UX agent enables you to back up, verify, and restore HP-UX server data.

Before installing an HP-UX agent:

- Make sure your HP-UX system is running a supported version listed in the Unitrends Compatibility and Interoperability Matrix.
- Add the Unitrends appliance name to the local host table or set up the TCP/IP system to use DNS with the Unitrends appliance.

To install the HP-UX agent

1. Log in to the HP-UX machine as user root.
2. Download the HP-UX agent from the Unitrends Downloads page (https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads) to the HP-UX machine.
3. Change to the working directory where you have saved the agent, and run the command **`ls -l`** to view the installer file and determine whether you have execute permission. If necessary, add execute permission using the command:

   ```
   # chmod +x <file_name>
   ```

4. Begin the installation be executing the file:

   ```
   # ./<file_name>
   ```

5. Press **Enter** to accept the default directory (*/usr/bp*).
6. Enter the hostname of the Unitrends appliance.
7. If using a firewall, enter **`y`** when asked if the client and the server (backup appliance) are separated by a firewall. This forces data communication to use port 1745.
8. When prompted, press **Enter**. Your agent installation is complete.

Once you have installed the agent, you are ready to add your HP-UX asset to the Unitrends appliance. For details, see "To add an agent-based asset" on page 289.

Removing the HP-UX agent

To uninstall the agent, log in to the HP-UX server and run the **uninstall** command from the directory where the agent is installed. For example, to remove the agent from the default install location, enter this command:

```
# /usr/bp/uninstall
```

Installing and updating the Mac agent

Before adding a Mac asset to the Unitrends appliance, you must install an agent. The Mac agent enables you to back up, verify, and recover Mac server data.

Once the agent is installed, you can update to a newer agent version using these same installation procedures. It is not necessary to remove the old agent. If an agent is already installed, you have the option to save any custom agent settings during installation.

Use the procedures in this section to install or update the Mac agent.

Preparing to install the Mac agent

Before installing a Mac agent:

- Make sure your Mac system is running a supported version listed in the Unitrends Compatibility and Interoperability Matrix.
- Add the Unitrends appliance name to the local host table or set up the TCP/IP system to use DNS with the Unitrends appliance.

To install or update the Mac agent

1. Log in to the Mac machine as user root.
2. Download the applicable Mac agent from the Unitrends Downloads page (https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads) to the Mac machine.
3. Change to the working directory where you have saved the agent, and run the command **ls -l** to view the installer file and determine whether you have execute permission. If necessary, add execute permission using the command:

   ```
   # chmod +x <file_name>
   ```

4. Begin the installation be executing the file:

   ```
   # sudo ./<file_name>
   ```

5. When a list of distribution files is presented, press **Enter** to continue.
6. Specify the directory location where the agent will be installed or press **Enter** to accept the default directory (*/usr/local/bp* for the 9.0.0 agent or */usr/bp* for 8.0.0 and earlier agents).
7. Enter the hostname of the Unitrends appliance.

8 If using a firewall, enter **y** when asked if the client and the server (backup appliance) are separated by a firewall. This forces data communication to use port 1745.

9 When prompted, press **Enter**. Your agent installation is complete.

Once you have installed the agent, you are ready to add your Mac asset to the Unitrends appliance. For details, see "To add an agent-based asset" on page 289.

Removing the Mac agent

To uninstall the agent, log in to the Mac server and run the **uninstall** command from the directory where the agent is installed. For example:

- To remove agent version 9.0.0 or later from the default install location, enter this command:

```
# /usr/local/bp/uninstall
```

- To remove an earlier agent version from the default install location, enter this command:

```
# /usr/bp/uninstall
```

Installing and updating the Novell Netware agent

Before adding a Novell Netware asset to the Unitrends appliance, you must install an agent.

Once the agent is installed, you can update to a newer agent version using these same installation procedures. It is not necessary to remove the old agent. If an agent is already installed, you have the option to save any custom agent settings during installation.

Use the procedures in this section to install or update the Novell Netware agent.

Preparing to install the Novell Netware agent

The Novell Netware agent enables you to back up, verify, and restore Novell Netware server data.

Before installing a Novell Netware agent:

- Make sure your Novell Netware machine is running a supported version listed in the Unitrends Compatibility and Interoperability Matrix.
- Add the Unitrends appliance name to the local host table or set up the TCP/IP system to use DNS with the Unitrends appliance.
- Verify that the Novell Storage Manager Service (NMS) package has been installed on the Novell Netware machine. This package is installed by default with Novell 6.5. SP3 and above. If running a prior version of Novell, you must install this package separately.

Novell Netware agent restrictions and limitations

Protecting Novell NetWare with Unitrends has the following restrictions:

- When restoring NetWare client backups from a backup copy, the backup copy must be restored in its entirety.

- File-level backups using TSA must be restored to a NetWare client.

Protecting Novell NetWare version 5.1 with Unitrends does not support the following functionality:

- TSA GroupWise backups
- Bare Metal Optimizer
- eDirectory backups
- Servers with legacy file system (LFS) volumes only supports DOS 8.3 filenames.

Protecting Novell NetWare version 6.0 with Unitrends does not support eDirectory backups.

To install the Novell NetWare agent

The Unitrends agent must first be installed to a Windows Server and then pushed to the Novell client. The following procedures guide you through the process.

1. Mount the Novell Server on the Windows Novell client .
2. Copy the Novell agent *bp_nov.exe* file to the Windows Novell client. (Download the agent from the Unitrends Downloads page at https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads.)
3. On the Windows Novell client, run *bp_nov.exe*.
4. When asked for the destination drive and directory on the NetWare server, enter:

   ```
   # <mapped_drive_letter>:\TMP\BP
   ```

5. Select **Full Installation**.
6. The Unitrends Novell NetWare agent is copied to your Novell server.
7. After the copy is complete, go to the Novell server and run the following command at a console prompt:

   ```
   # SYS:\TMP\BP\bpinstall.ncf
   ```

8. Go to the **Backup Professional Installation** screen.
9. The installation screen asks where you want to install. Select **Yes** to select the default *sys:\bp*.
10. If the version of Novell supports eDirectory backups, you are asked: *Do you want to install the eDirectory Backup Before Command? (Press Y or N).*

    This requires the dsbk utility to be installed. If it is not present, select **N** or download it before continuing.
11. The installation provides an option to configure the GroupWise database paths. If the database backups will be managed outside of the system agent, this configuration may be skipped.
12. Select **Enter** to accept the default ports and autoexec.ncf settings.
13. To load the protection software on the Novell sever, run the following command:

    ```
    # LOAD SYS:\BP\bps.nlm
    ```

During the LOAD process, SMS-TSAs based backups and restores are enabled. It is recommended to use the credentials for the full context administrative user account. At this time the Admin password must be provided to enable SMS-TSA based backup. You are given an option to store the password in an encrypted state on the Novell server. This allows the bps service to auto-log in when there is a reboot or if the service is ever manually loaded.

14 If the password changes, you will be prompted for the password the next time the bps service loads.

If you choose not to store the password, you will be prompted to enter it whenever the bps service loads.

The Unitrends Novell NetWare agent has been installed on your Novell server. You are ready to add your Novell Netware asset to the Unitrends appliance. For details, see "To add an agent-based asset" on page 289.

Removing the Novell Netware agent

To uninstall the agent from a Novell Netware machine

1 Stop any running backups on the Novell Netware machine.

2 Run the **uninstall** command from the directory where the agent is installed. For example, enter this command to remove the agent from the default install location:

```
# sys:\bp\bpinstl uninstall
```

Installing and updating the Solaris agent

Before adding a Solaris asset to the Unitrends appliance, you must install an agent.

Once the agent is installed, you can update to a newer agent version using these same installation procedures. It is not necessary to remove the old agent. If an agent is already installed, you have the option to save any custom agent settings during installation.

Use the procedures in this section to install or update the Solaris agent.

Preparing to install the Solaris agent

The Solaris agent enables you to back up, verify, and restore Solaris data.

Before installing a Solaris agent:

- Make sure your Solaris system is running a supported version listed in the Unitrends Compatibility and Interoperability Matrix.
- Add the Unitrends appliance name to the local host table or set up the TCP/IP system to use DNS with the Unitrends appliance.

To install the Solaris agent

1 Log in to the Solaris machine as user root.

2 Place the agent file, *solaris8_cnt*, in the */tmp* directory on the Solaris system. (Download the agent from the Unitrends Downloads page at https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads.)

3 Grant execute permission to the file by running the following command:

```
# chmod 711 /tmp/solaris8_cnt
```

4 Begin the installation by executing the file:

```
# ./tmp/solaris8_cnt
```

5 Enter **y** to continue the installation and press **Enter** to continue.

6 Press **enter** to accept the default installation directory (*\usr\bp*) or enter the full path where you prefer the software be installed. Respond with a **y** when asked if the directory can be created.

7 If this is a reinstall, you will be asked if you wish to overwrite certain files. Type the interrupt character or press **Enter** to continue. Once the files have been moved to their permanent location, you will be given a chance to review the release notes.

The agent is installed and you are ready to add your Solaris asset to the Unitrends appliance. For details, see "To add an agent-based asset" on page 289.

Removing the Solaris agent

To uninstall the agent, log in to the Solaris server and run the **uninstall** command from the directory where the agent is installed. For example, enter this command to remove the agent from the default install location:

```
# /usr/bp/uninstall
```

Installing and updating the UnixWare agent

Before adding a UnixWare asset to the Unitrends appliance, you must install an agent.

Once the agent is installed, you can update to a newer agent version using these same installation procedures. It is not necessary to remove the old agent. If an agent is already installed, you have the option to save any custom agent settings during installation.

Use these procedures to install or update the UnixWare agent:

- "Preparing to install the UnixWare agent" on page 402
- "To install the UnixWare agent" on page 403

Preparing to install the UnixWare agent

The Unitrends agent for UnixWare enabls you to backup, verify, and restore UnixWare data.

Before installing a UnixWare agent:

- Make sure your UnixWare system is running a supported version listed in the Unitrends Compatibility and Interoperability Matrix.
- Add the Unitrends appliance name to the local host table or set up the TCP/IP system to use DNS with the Unitrends appliance.

To install the UnixWare agent

1 From a terminal window on the UnixWare system, download the agent file, *svr4_cnt*, from the Unitrends Downloads page (https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads).

2 In binary mode, copy the agent

```
# /bp/<release_number>/svr4_cnt
```

to the */tmp* directory.

3 Install the agent as shown below. Follow the prompts and accept the default values.

```
# cd /tmp; chmod 755 svr4_cnt
# ./svr4_cnt
To CONTINUE with installation type y
Please press ENTER to continue
Please press ENTER to continue
[Default: /usr/bp ] Enter directory:
Please press ENTER to continue
(99) Complete Installation
[Default: none ] Enter email address for this computer's backup summariesEnter:
[Default: ]Enter the hostname of the Backup Professional Server:
[Default: no ]Is this client and server separated by a firewall? (y/n):
Please press ENTER to continue
This completes the UnixWare Installation.
```

The agent is installed and you are ready to add your UnixWare asset to the Unitrends appliance. For details, see "To add an agent-based asset" on page 289.

Removing the UnixWare agent

To uninstall the agent, log in to the UnixWare server and run the **uninstall** command from the directory where the agent is installed. For example, enter this command to remove the agent from the default install location:

```
# /usr/bp/uninstall
```

Copied Assets

The Copied Assets tab displays only for appliances that are receiving backup copies from another Unitrends appliance. The tab lists all assets whose backup copies are stored on this appliance. From this tab you can view, edit, and remove copied assets using the buttons across the top of the tab. See these topics for details:

- "Viewing copied assets " on page 404
- "Managing retention of copied assets with long-term data management" on page 404
- "Switching to long-term retention" on page 410

- "Managing retention of copied assets with legacy asset-level retention" on page 412
- "Removing copied assets " on page 414

Viewing copied assets

To view all copied assets

1 On the **Configure > Appliances** page, click the **Copied Assets** tab.

2 Use these options to customize your view:

- Select an item from the **Display** list to filter the display by asset type.
- Click on any column to sort alphabetically (a to z). To reverse the order (z to a), click again.

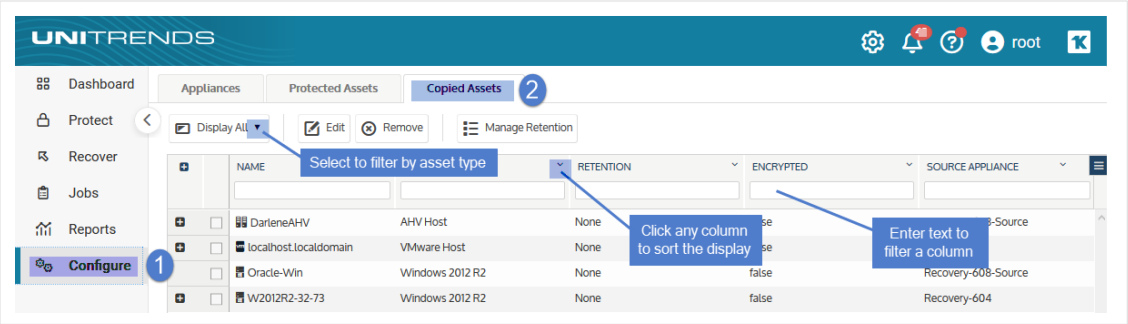

Managing retention of copied assets with long-term data management

Retention settings assure that the necessary recovery points are available on your appliance. Appliances are configured with a default backup retention policy of 30 days. This 30-day policy is applied to each protected asset. To use your own custom retention setting, the long-term data management feature enables you to quickly create policies that hold backups for a specified number of days. You can create multiple policies and customize them to achieve different RPOs and RTOs for your assets. Each policy you create can be applied to multiple assets. These policies automatically retain and purge backups as necessary to maintain a customized inventory of weekly, monthly, and yearly retention points.

Notes:

- The 30-day default retention policy applies to appliances imaged with release 10.7.8 or higher. This default policy does not apply to appliances that were originally imaged with an earlier release. Upgrading an appliance that was imaged with a pre-10.7.8 release does not modify its retention policies in any way.
- The 30-day default retention policy ensures that 7 daily backups and 4 weekly backups are retained for each protected asset.

To start managing retention with long-term data management:

Note: Long-term retention is not enabled on appliances that have been upgraded from a pre-10.3 release. If your appliance was originally imaged with a pre-10.3 Unitrends release, see "Switching to long-term retention" on page 410 to enable this feature.

Step 1: Review the "Long-term retention policy settings "

Step 2: Add a policy as described in "To add a long-term retention policy to a backup copy target appliance" on page 407

Step 3: Apply the policy to assets as described in "To apply a long-term retention policy to a copied asset" on page 408

Long-term retention policy settings

Long-term retention policies are configured with the following settings:

Retention setting	Description
Policy Name	Enter a name for the policy.
Set as default policy	Check the **Set as default policy** box on the Add Retention Policy or Edit Retention Policy dialog to designate the policy as the default policy of the appliance. Any new assets and new copied assets that are added to the appliance receive the default policy. **Notes:** • An appliance cannot have more than one default policy at a time. If needed, edit the existing default policy to clear the **Set as default policy** box before designating a new default policy. • Applying a non-default policy to an asset overrides the default policy. • If you have designated a new default policy, you can apply it to existing assets as described in "To apply a long-term retention policy to a protected asset" on page 334. • For copied assets, you can use the **Set as global policy** option to override the default policy. Use this method to have different default policies for local assets versus copied assets.
Set as global policy	The Configure > Copied Assets tab displays only for appliances that are receiving backup copies from another Unitrends appliance. The tab lists all assets whose backup copies are stored on this appliance. If your appliance is receiving backup copies from another appliance, you can create or edit the policies used for copied assets by going to **Configure > Copied Assets** and clicking **Manage Retention**. When adding or editing a policy for copied assets, you can check the **Set as global**

Retention setting	Description
	policy box to designate the policy as the global policy for copied assets. Any newly added backup copy source appliance receives this policy, which determines how long its hot copies are retained. As copied assets are added for the new source appliance, this global policy is automatically applied. **Notes:** • A target appliance cannot have more than one global policy at a time. • Global policies are applicable to any new copied assets that are added to the appliance. (Global policies do not apply to local assets.) • The appliance's default policy applies to any copied assets unless you use the **Set as global policy** box to override the default policy for copied assets.
Policy description	Enter a brief description of the policy. This can be a note regarding the policy's intended use case or a brief summary of its retention settings.
Days	All backups from the last *N* number of days are retained. Days end at midnight, appliance time. Or check the **Forever** box to retain all backups. Check the **Delete Final Backup** box to automatically delete data when it reaches the end of its retention policy, even if it's the last available recovery point. Leave this box unchecked to retain the last available recovery point.
Weeks	The most recent successful backup from each of the last *N* number of weeks is retained. Weeks end on Sunday of the calendar week. Or check the **Forever** box to keep weekly backups forever.
Months	The most recent successful backup from each of the last *N* number of months is retained. Months end on the final day of the calendar month. Or check the **Forever** box to keep monthly backups forever.
Years	The most recent successful backup from each of the last *N* number of years is retained. Years end on the final day of the calendar year. Or check the **Forever** box to keep yearly backups forever.

Notes:

- If the most recent successful backup is an incremental or differential, the complete backup group is retained.
- Redundant backups are not retained if high-frequency retention intervals are configured to overlap their low-frequency counterparts. For example, a retention policy specifying 52 weeks and 12 months *does not* retain a total of 52 + 12 backups for the preceding 12-month period.

- An asset must have at least one successful backup for each retention point to be considered compliant.
- A yellow alert displays in the UI if any asset is, for any reason, not compliant with its associated long-term retention policy. This alert is automatically dismissed if compliance is achieved.
- Retention compliance is not achievable if the asset's scheduled backups are less frequent than the policy's most frequent retention interval. For example, a retention policy specifying 4 weeks and 0 days cannot maintain a complete inventory of compliant retention points for an asset that is backed-up on a bi-weekly basis.
- Imported backup copies are retained for a default period of 72 hours regardless of retention settings applied to the asset from which they originated.

To add a long-term retention policy to a backup copy target appliance

1 Log in to the backup copy target appliance and select **Configure > Copied Assets**.

2 Click **Manage Retention**.

3 Click **Add**.

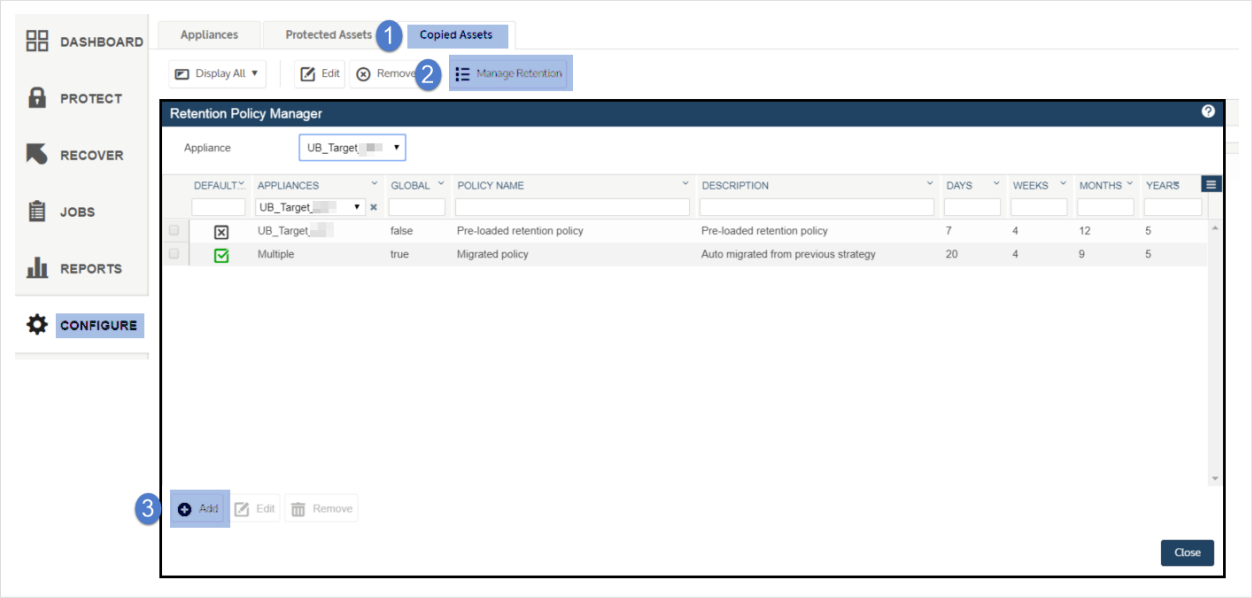

4 Enter the retention policy settings. For details, see "Long-term retention policy settings " on page 405

5 Click **Assign to Source**.

6 Select the local appliance in addition to any source appliances you wish to assign the policy to. An assigned policy can be applied to the source's copied assets. If desired, you can designate the policy as the source's default policy by double clicking the adjacent value under the Default column and selecting **Yes**.

Note: Assigned policies are only applicable to assets that reside on the target appliance. This includes the source backup appliance's copied assets and the target appliance's local protected assets, but not the source appliance's local protected assets.

7 Click **Okay**.

8 Click **Add**. The long-term retention policy is created.

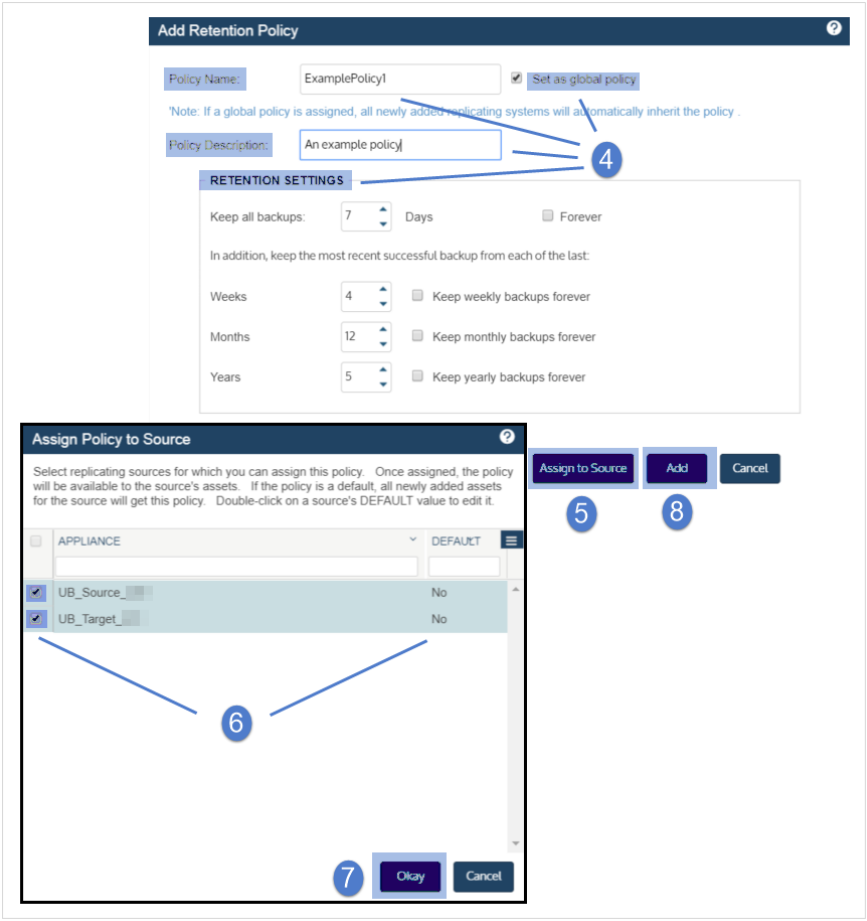

To apply a long-term retention policy to a copied asset

1 Log in to the backup copy target appliance and select **Configure > Copied Assets**.

2 Select the desired asset.

3 Click **Edit**.

4 Select a policy from the **Retention Policy** dropdown.

> **Note:** If the selected policy will delete one or more of the asset's backups, a dialog listing these backups displays. Click **Close** to continue.

5 Click **Save**. The long-term retention policy is applied.

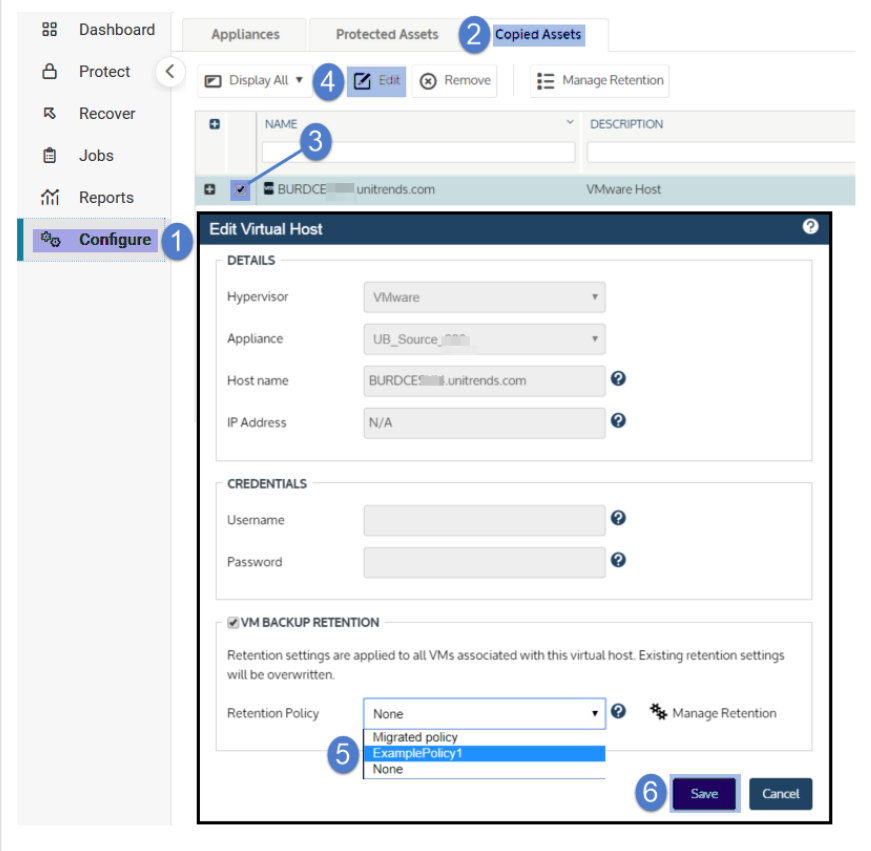

To edit a long-term retention policy

1. Log in to the backup appliance and select **Configure > Protected Assets**.
2. Click **Manage Retention**.
3. Select the policy you wish to edit.
4. Click **Edit**.

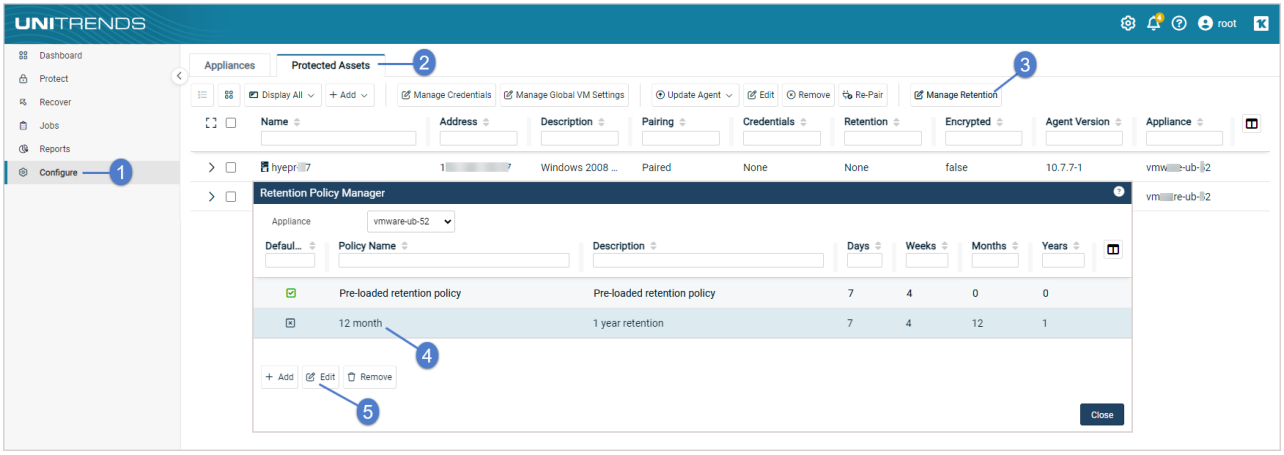

5 Modify the retention policy settings. For details, see "Long-term retention policy settings " on page 405

6 Select the **Note** checkbox.

7 Click **Save**. Changes to the policy are committed.

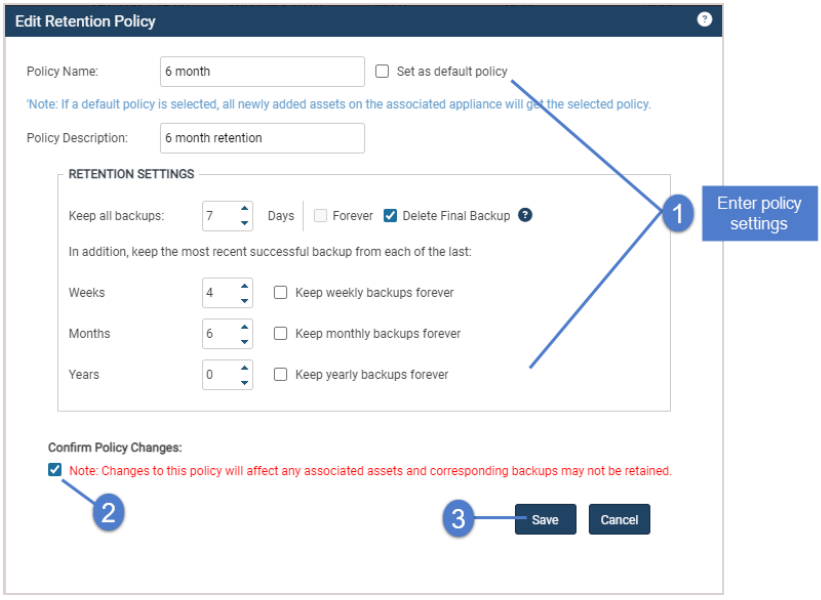

Switching to long-term retention

If your appliance was originally imaged with a pre-10.3 Unitrends release, the appliance utilizes legacy asset-level retention until the switch to long-term retention is manually performed. Long-term retention policies provide a more precise, granular, and space-efficient alternative to legacy asset-level retention settings. A long-term retention policy automatically retains and purges backups as necessary to maintain a customized inventory of weekly, monthly, and yearly retention points. A Unitrends appliance can maintain multiple long-term retention policies that can each be applied to multiple assets.

Note: Long-term retention is enabled by default on appliances that were originally imaged with release 10.3 or later. Do not run the "To switch to long-term retention" procedure if your appliance was originally imaged with release 10.3 or later.

To switch to long-term retention, review the "Considerations for switching to long-term retention", then run the "To switch to long-term retention" procedure.

Once you have switched to long-term retention, set up retention policies as described in "Managing retention of copied assets with long-term data management" on page 404.

Considerations for switching to long-term retention

Migrating from legacy asset-level retention to long-term retention impacts a number of Unitrends features, which are detailed in the table below. Carefully review the considerations in the table before switching to long-term retention.

WARNING! Once the switch to long-term retention is performed, it cannot be reversed from the web UI.

Feature	Details
Legacy asset-level retention settings	The switch to long-term retention nominally voids all legacy asset-level retention settings; however, any backups held in accordance with these settings remain preserved in a legal hold state. When applying a long-term retention policy to an asset that previously used legacy asset-level retention settings, the *keep all backups* value of this new policy must be equal to or greater than the *keep backups for* value specified in the original retention settings.
SLA policies	Following the switch to long-term retention, all SLA policy retention settings are voided and the SLA policy feature loses its retention functionality. An asset can be assigned both a long-term retention policy and an SLA policy.
GFS	If GFS is enabled on your appliance, the switch to long-term retention directly translates your GFS retention settings into a default long-term retention policy titled *Migrated policy*. If you are using legacy asset-level retention policies in addition to GFS retention settings, ensure that these policies specify backup hold periods shorter than or equal to the *daily* GFS setting. Switching to long-term retention is not permitted unless this condition is satisfied. **Notes:** • On target appliances, the migrated policy is both global and default. • The *weeks* setting of the migrated policy is set to 4 by default.
Backup-level holds	Holds applied to individual backups in the Backup Catalog are not impacted by the switch to long-term retention.
Cold backup copy retention	Job-level cold backup copy retention settings are not impacted by the switch to long-term retention.

To switch to long-term retention

Appliances upgraded from a pre-10.3 release do not have access to long-term data management functionality until the switch to long-term retention is manually performed. Long-term retention is enabled by default on all appliances that were originally imaged with version 10.3 or later.

CAUTION! Before proceeding, ensure that you have reviewed the "Considerations for switching to long-term retention" on page 411 and understand how your assets may be impacted.

1. Log in to the backup appliance and select **Configure >Protected Assets**.
2. Select **Manage Retention**.
3. Select **Switch to Long Term Retention**.
4. Click **Save**.
5. If you are sure you want to switch to long-term retention, click **Confirm**.

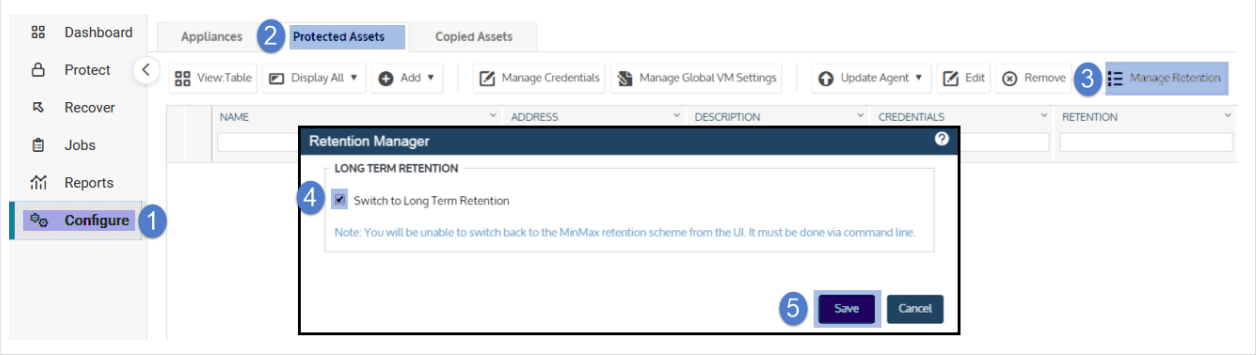

Managing retention of copied assets with legacy asset-level retention

To edit retention of a copied asset with legacy asset-level retention

Notes:

- This retention scheme is not available on appliances that have switched to long-term retention.
- Sharepoint, Oracle, and Cisco UCS assets are not compatible with this retention scheme. To customize retention options for these asset types, see "Managing retention of copied assets with long-term data management" on page 404.

1. Log in to the backup copy target appliance.
2. On the **Configure > Appliances** page, click the **Copied Assets** tab, select the copied asset.
3. Click **Edit > Manage Retention**.

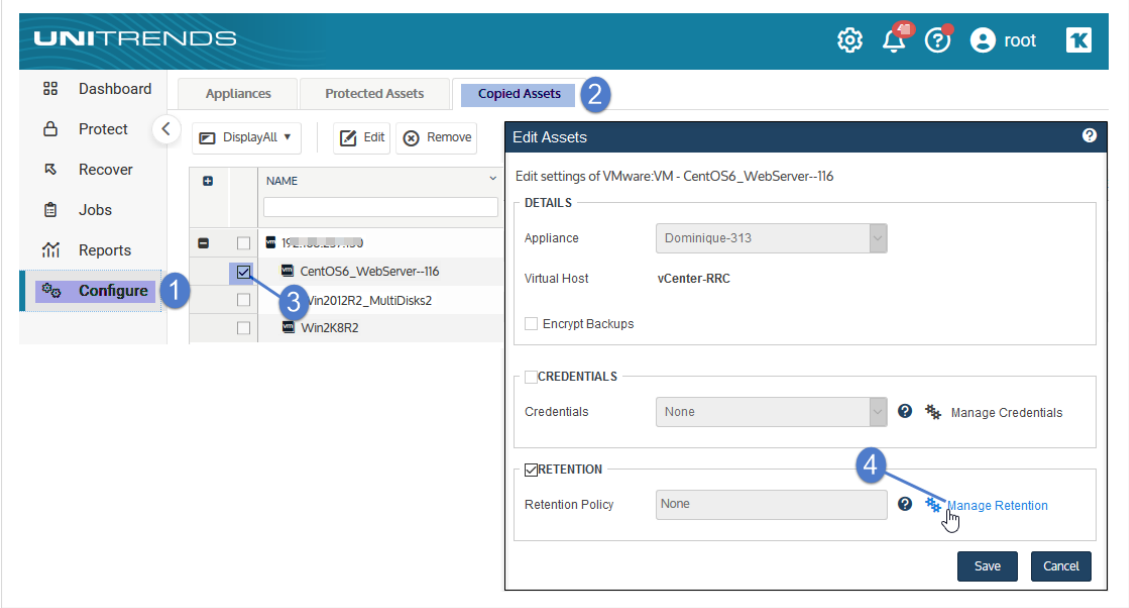

4 Modify retention settings and click **Save**. (See the table below for descriptions of each setting.)

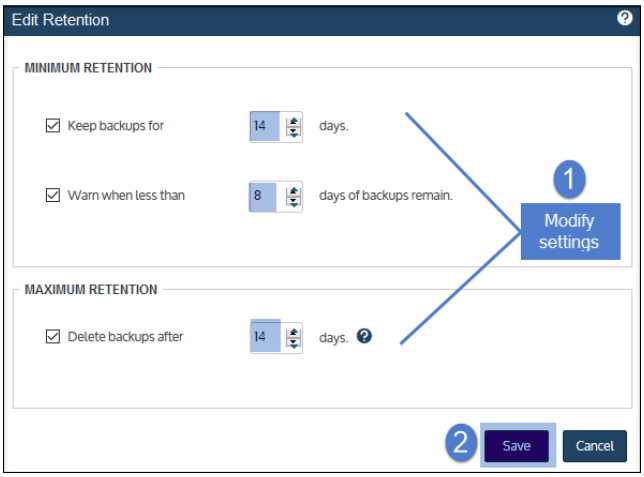

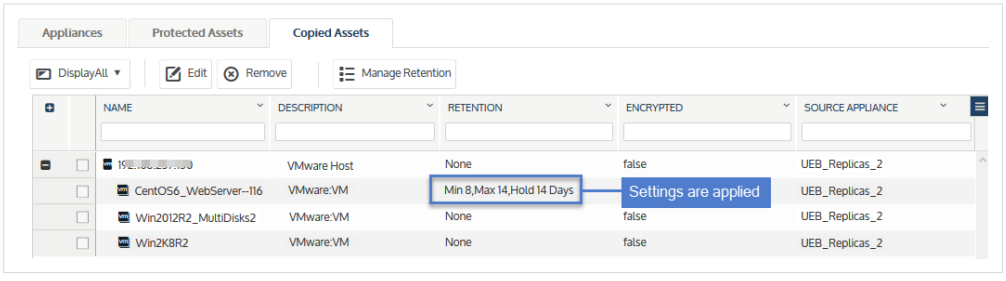

Retention settings

Retention setting	Description
Minimum Retention	Minimum retention settings.
Keep backups for *N* days	Number of days backups must be retained. Backups that are younger than *N* days are not purged for any reason, including at the expense of new, incoming backups. The age of a backup is determined by the last backup in the group, e.g., the last incremental before a new full.
Warn when less than *N* days of backups remain	Use this option to receive an email notification if this asset has less than *N* days of backups stored on the appliance.
Maximum Retention	Maximum retention setting.
Delete backups after *N* Days	Number of days after which the appliance will delete backups. Backups are eligible to be deleted once the full has exceeded this limit. At this point, the full and all associated incrementals and differentials in the group are deleted.

Removing copied assets

To remove a copied asset

1. Log in to the backup target appliance.
2. On the **Configure > Appliances** page, click the **Copied Assets** tab.
3. Select the copied asset.
4. Click **Remove**.
5. Click **Confirm**.

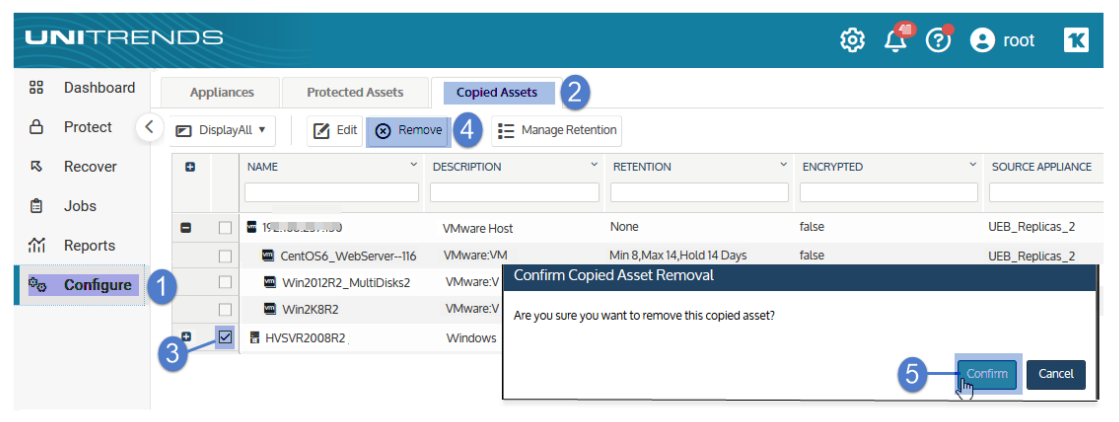

ConnectWise PSA integration

ConnectWise PSA integration is no longer supported on your Unitrends appliance. Instead, use the UniView Portal to integrate and manage ConnectWise PSA. UniView Portal is included with your Unitrends appliance license. To begin using UniView Portal, contact your Account Representative. See Integrating ConnectWise Manage in the UniView Portal Guide for integration details.

IMPORTANT! Upon upgrading ConnectWise PSA to release 2020.4 or higher, any existing ConnectWise PSA integration on the Unitrends appliance is disabled. Be sure to set up your ConnectWise PSA integration in the UniView Portal before upgrading.

This page is intentionally left blank.

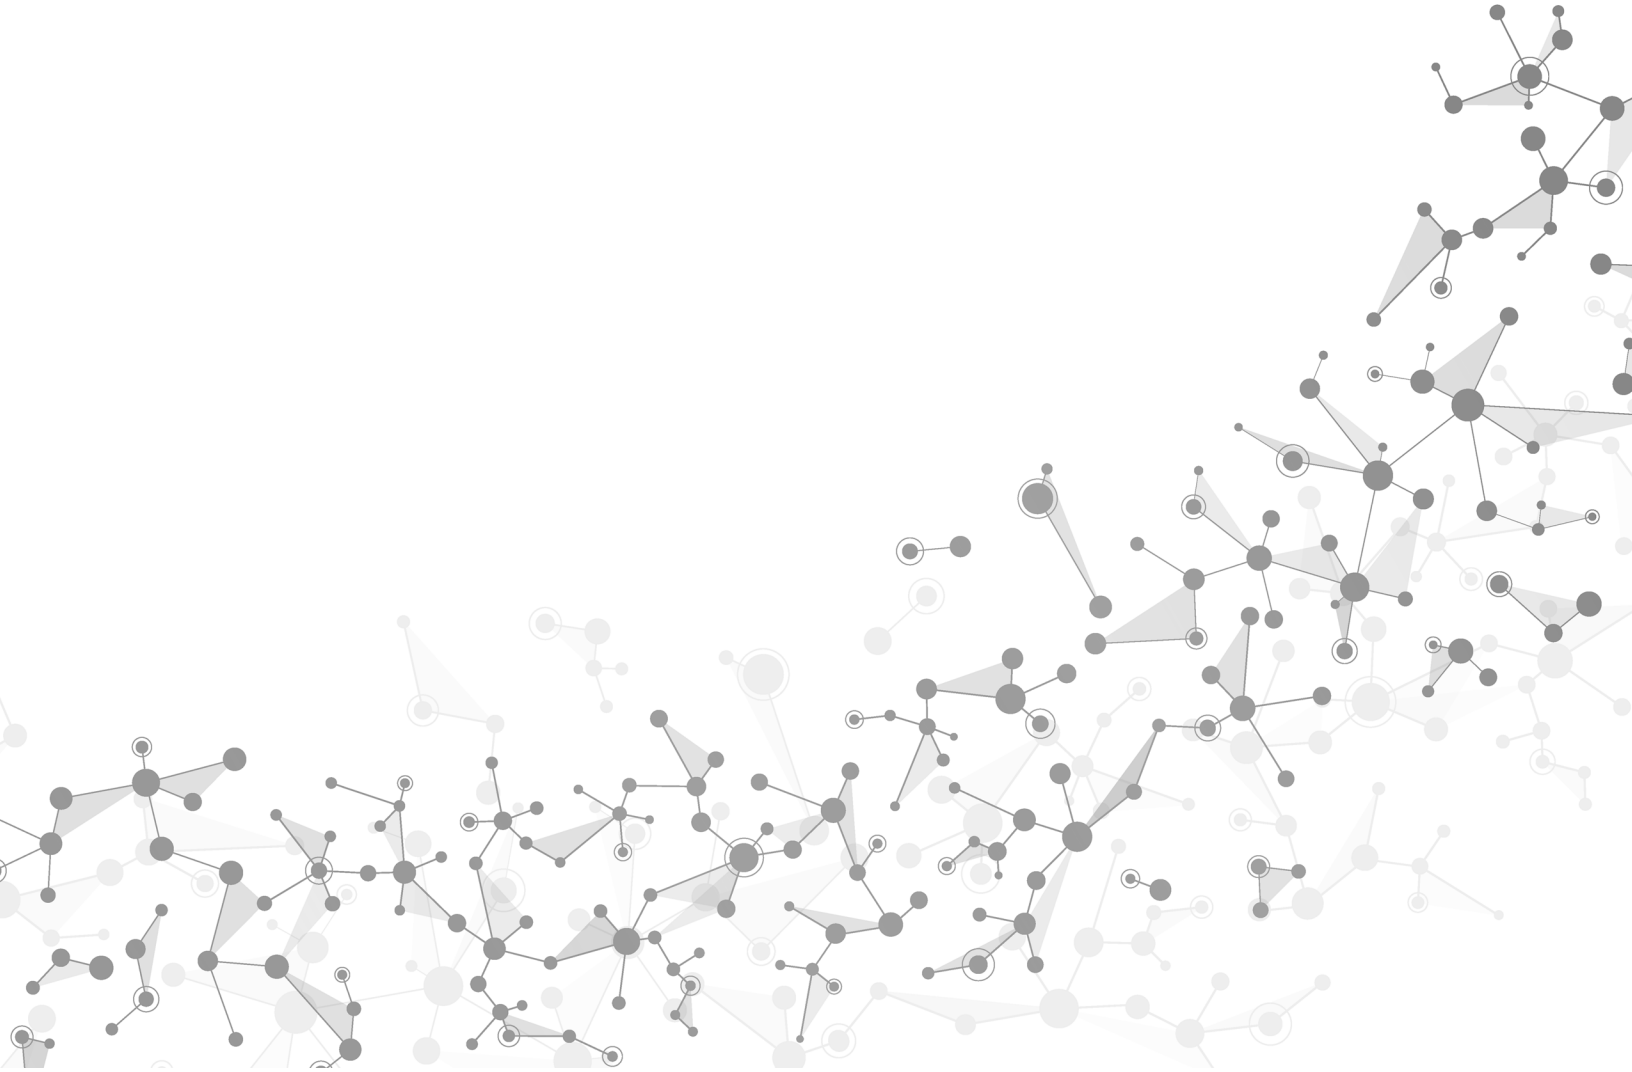

Chapter 4: Remote Appliance Management

Unitrends enables you to manage multiple Unitrends Backup, Recovery MAX, Recovery Series, and ION/ION+ appliances from a single interface. For most tasks, this eliminates the need to log in to each individual Unitrends appliance in your environment. Remote appliance management also enables you to consolidate daily email reporting from your appliances into a single Management Status report. For instructions on configuring email reports, see "Email reporting" on page 117.

For further information on remote appliance management, see the following topics:

- "Remote appliance management limitations"
- "Remote appliance management procedures"

Remote appliance management limitations

This feature can be used for appliances running on CentOS 7 or higher only.

Note: Support for managing CentOS 6 appliances using the Remote Appliance Management feature has been discontinued. As an alternative, please utilize UniView.

To check your appliance's OS version, click on **? > About**, as shown here:

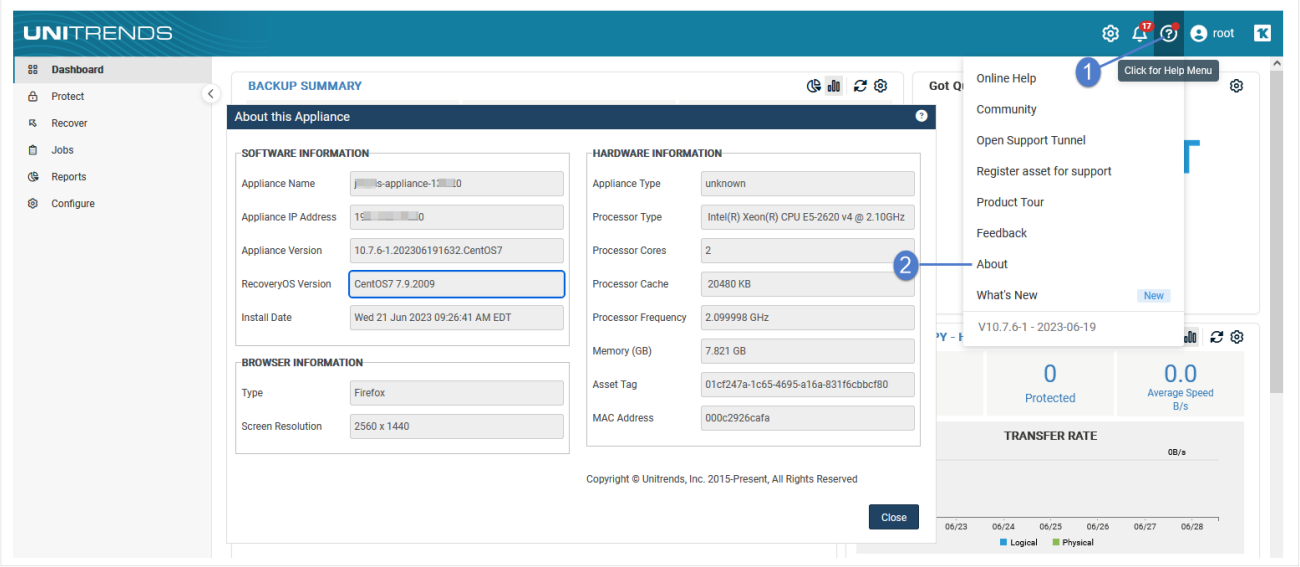

The following can only be performed by logging directly into the managed appliance:

- Configuring hot backup copy operations.
- Creating SLA policies.
- Changing users.

- Creating replicas and modifying recovery points.
- Changing hostname and FQDN.
- Changing the OS password.

Remote appliance management procedures

Use these procedures to manage existing appliances, add appliances you want to manage from another appliance's UI, and remove appliances you no longer want to manage. For details on configuring individual appliance settings, see "Appliance settings" on page 105.

To view appliances

Click **Configure** to view appliances on the Appliances tab. You can view appliances as a table or list. The list view is better for small deployments, while the table view is better for larger deployments.

- To view appliances in a list, click **View: List**.

  Each row in the list shows the appliance's name, status, IP address, Unitrends software version, storage, and number of registered assets.

- To view appliances in a table, click **View: Table**.

  Table view displays appliances in tiles on the left. Each tile includes the appliance's name and status. Additionally, *Logged In* displays in the tile of the appliance you are currently logged in to. Select a tile to view details about the appliance, such as name, IP address, description, and Unitrends software version.

The following appliances display on this tab:

- The appliance you are currently logged in to. Its status is *Available (logged in)*.
- Any additional appliances that this appliance is managing. Managed appliances display with the status *Available*.
- If the appliance you are logged in to has been configured as a backup copy target, its source Unitrends appliance displays. The source appliance can be in the following statuses:
  - *Pending* means the backup copy request is pending.
  - *Not Available* means the appliance is configured as a backup copy source only and cannot be managed from this UI. (To enable management, simply click **Edit**, check **Enable Management of this appliance**, supply **User Name** and **Password** credentials, and click **Save**.)

– *Available* means you can manage the source appliance from this UI.

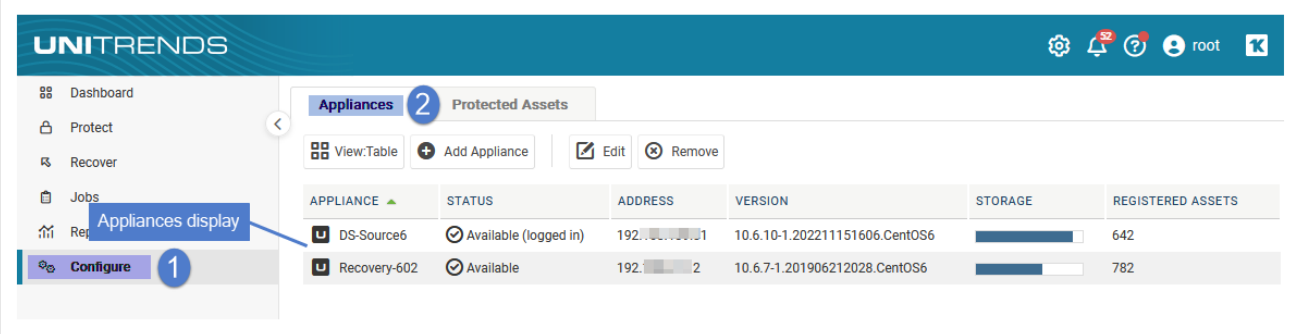

To add an appliance

1 On the **Configure > Appliances** page, click **Add Appliance**.

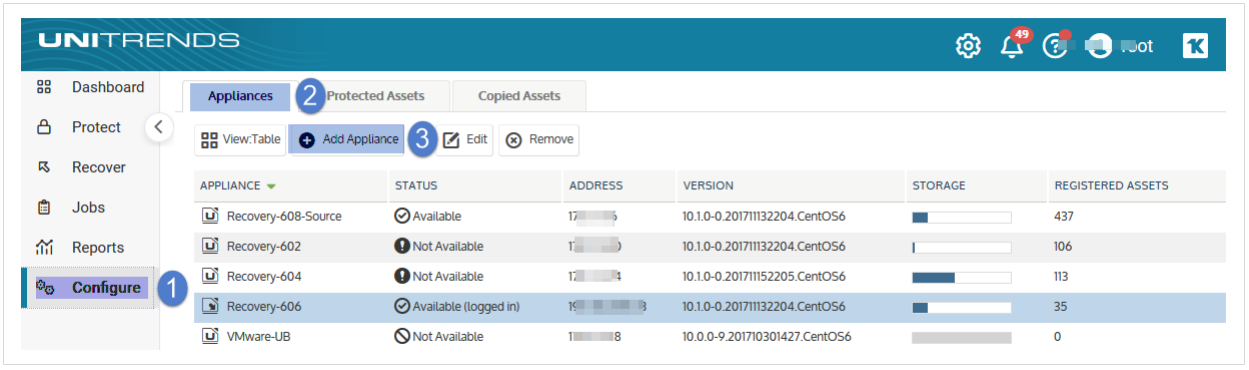

2 Enter the **Name** you want to use to identify the appliance.

3 Enter the **IP Address** of the appliance.

4 Enter the **User Name** and **Password** you used to configure the appliance.

5 Click **Save**.

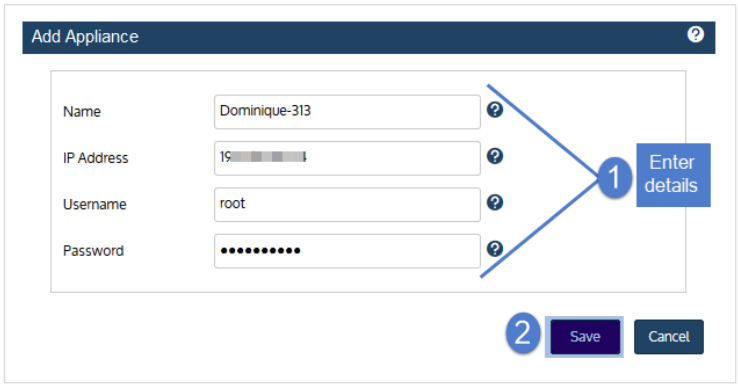

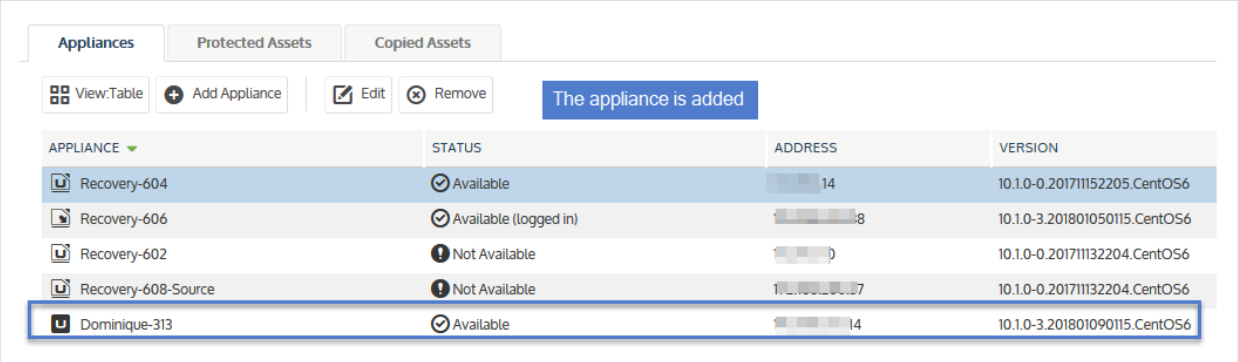

To manage remote appliances

Use the "To add an appliance" procedure above to add a remote appliance. Once you add a remote appliance, it displays on the **Configure > Appliances** tab in *Available* status. You can then manage operations on the remote appliance as your user role permits, with some exceptions (for example, you are unable to manage users and SLA policies on remote appliances). To manage an appliance from this UI, the appliance must have an *Available* status.

Appliance status information includes:

Status	Description
Available (logged in)	You are logged in to this appliance and can manage its operations.
Available	This is a remote appliance that can be managed from this UI.
Not Available	The appliance is configured as a backup copy source only and cannot be managed from this UI. To enable management, simply click **Edit**, check **Enable Management of this appliance**, supply **User Name** and **Password** credentials, and click **Save**.
Pending	This is a remote appliance that is requesting permission to send backup copies to the appliance you are logged in to. Click **!** to accept or deny the request.

To remove a managed appliance

1. Log in to the managing appliance.
2. Go to **Configure > Appliances** and select the appliance you want to remove.
3. Click **Remove**.
4. Click **Confirm**.

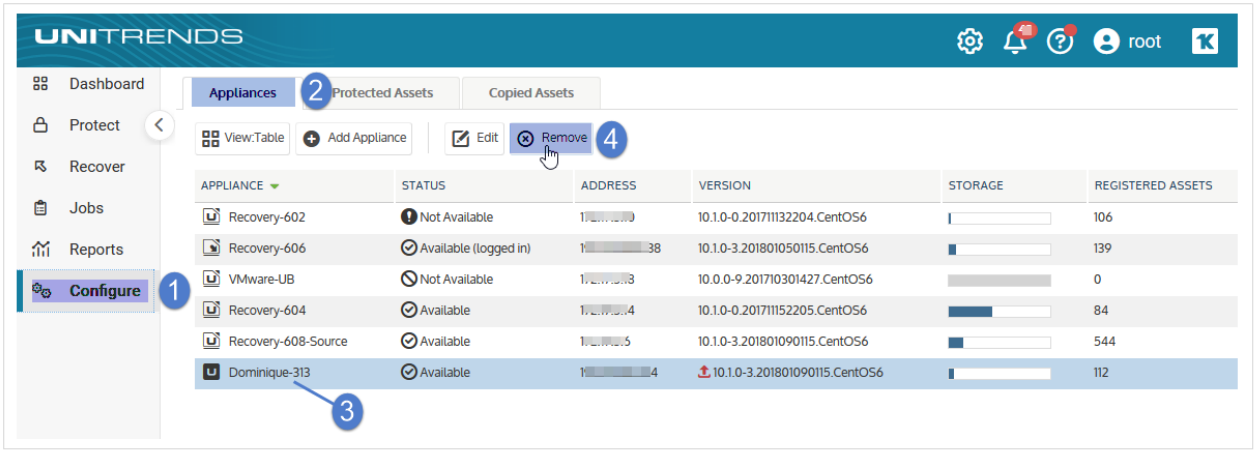

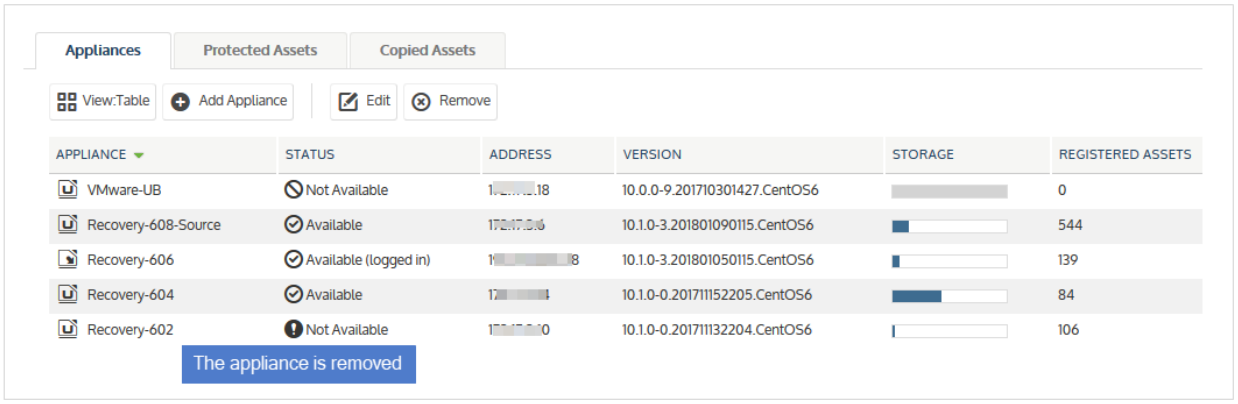

To install appliance updates

Make sure that there are no jobs running prior to upgrading your appliance. Once the upgrade begins, any running jobs terminate.

1 Log in to the appliance UI as a user with administrative credentials.

2 Click the gear icon and select **Check for Updates**.

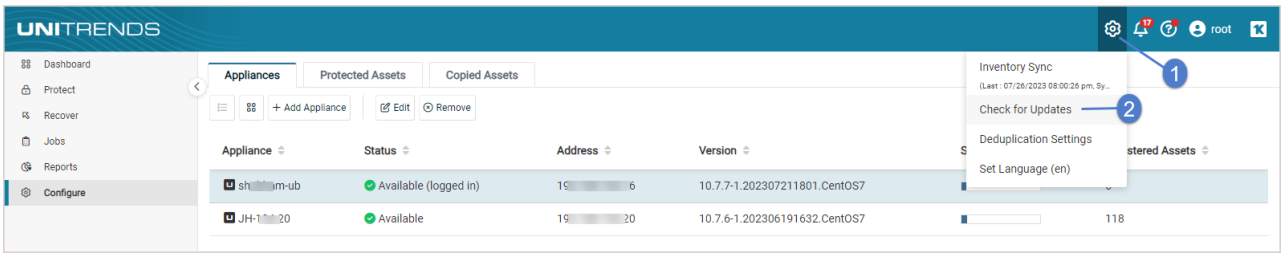

3 A list of available updates displays. Click **Apply Update** to begin the installation.

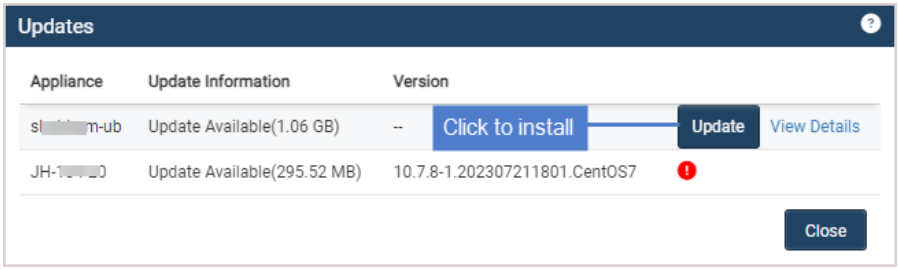

4 For some appliances, software maintenance is required with the update. If so, you see this message and the update will take some extra time:

Note: If you do not see this message, maintenance has already been performed on the appliance.

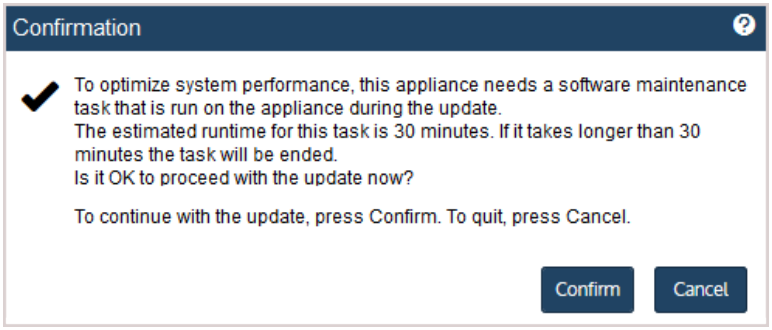

Do one of the following:

- Click **Confirm** to continue with the update.
- Click **Cancel** to quit. (You can then install the update at another time.)

5 During the upgrade, you see status messages as packages are installed. If you have trouble with the installation, see "Troubleshooting the upgrade" for tips. After the installation completes:

- Clear your browser cache, then close the browser.
- Open the browser and log back in to continue working with your appliance.

Note: If you receive a message indicating that you need to reboot the appliance to take advantage of the new kernel installed during the upgrade, you can either reboot now or reboot at a later time. If you reboot now, do the browser steps above after the appliance boots. (If you do not receive this message, the kernel was not updated and a reboot is not required.)

Troubleshooting the upgrade

In rare instances, your first attempt to update the Unitrends appliance might not be successful. See the following table for a description of upgrade issues and steps you can take to resolve them:

Issue	Next steps
The update times out because some of the packages did not install.	If the installation stops and you receive a message stating a package did not install successfully, in most instances you can resolve the issue by clicking the refresh arrows and attempting the update again. If necessary you can repeat this multiple times until the update completes. See Timeout error when upgrading a Unitrends appliance for more information.
The appliance is unable to download the update packages.	There are two possible solutions if your appliance is unable to download packages: • The appliance cannot reach the FTP or HTTP site - If you receive a message stating that the appliance is unable to download packages, this is the most likely cause. The FTP or HTTP site might be blocked by a firewall or some other restriction might be preventing you from reaching the site. To resolve this issue, you can download the update packages from the site you are not currently using (such as downloading from the HTTP site if you are currently using the FTP site, or vice-versa). For procedures, see Upgrade attempt fails because appliance cannot download update packages. • Your appliance is connected to a local network only - If your appliance is not connected to the Internet, you can update the software using an ISO image. For procedures, see How to upgrade the appliance via Unitrends' media.
An error message displays stating that the managing system must be updated.	To update the appliance, you must first update any other appliances that are managing it. Verify that any backup copy target appliance and any other managing appliances are running the latest release. Upgrade these appliances as needed. You can then upgrade any appliances that they are managing.
No data displays in the UI after installing appliance updates.	To resolve this issue: 1 Clear your browser cache, then close the browser. 2 Open the browser and log back in to continue working with your appliance.

To shutdown or restart an appliance

1 On the **Configure > Appliances** page, select the appliance and click **Edit.**

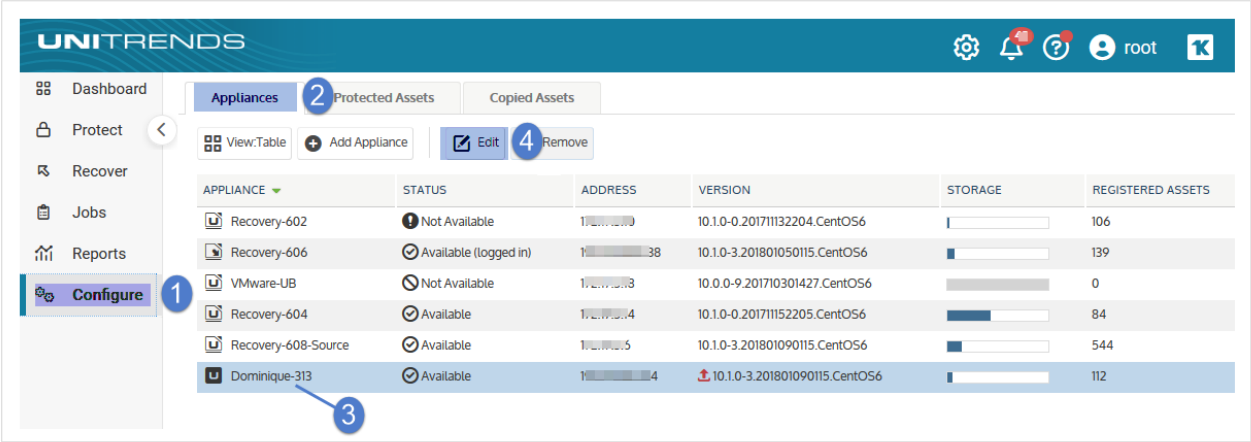

2 On the General tab, click **Shutdown/Restart**.

3 Enter the appliance root user password.

4 Click **Restart** to power down and restart the appliance OR click **Shutdown** to power down the appliance only.

5 Click **Restart** or **Shutdown** again to confirm.

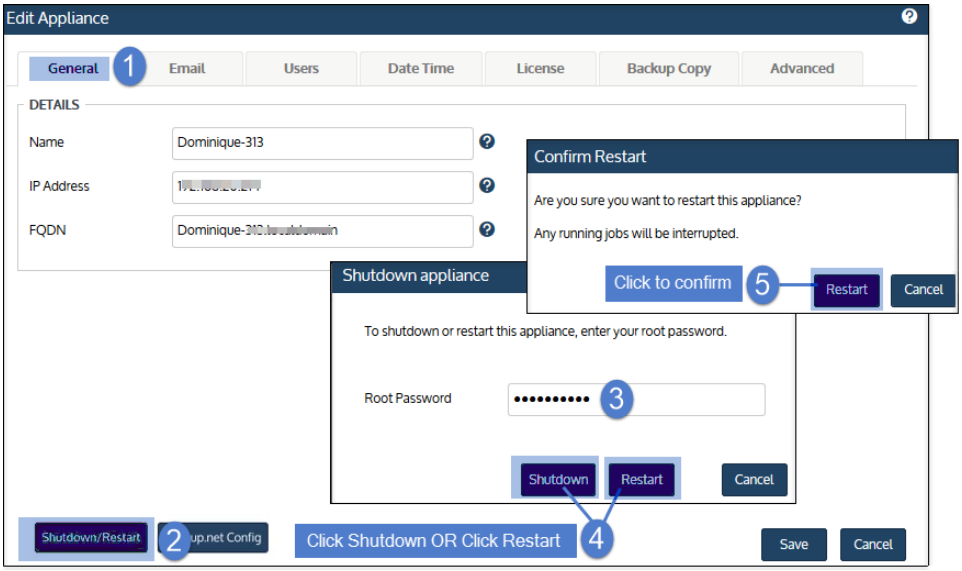

Chapter 5: Backup Administration and Procedures

These procedures provide instructions for creating and managing backup jobs and backup copy jobs. The procedures are run from the Jobs page of the Unitrends User Interface.

Note: Your Unitrends user account determines which procedures you can run. Procedures that you cannot run do not display or are disabled in the UI.

See the following topics for details:

- "Preparing for backups" and "About creating backup and backup copy jobs" on page 426 to determine how you will implement your protection strategy.
- "Creating backup jobs" on page 433 to create backup jobs manually.
- "Creating backup copy jobs" on page 491 to create backup copy jobs manually.
- "Creating SLA policies" on page 536 to create SLA policies. Creating an SLA policy automatically creates associated backup and backup copy jobs.
- "Managing scheduled jobs" on page 563 to manage existing job schedules.
- "Managing SLA policies" on page 589 to manage existing SLA policies.
- "Managing active jobs" on page 607 to manage jobs that are currently running.
- "Viewing recent jobs" on page 615 to view jobs that ran over the last 7 days.
- "Viewing system jobs" on page 620 to view system-level jobs that ran over the last 7 days.
- "Deleting backups and backup copies" on page 627 to delete backups and backup copies.

For information on generating reports on backups, backup copies, and SLA policies, see "Reports" on page 1307.

Preparing for backups

Before you run backups, it is recommended that you do the following to map out the best approach for your environment:

Step 1: Review your company's business continuity plan to determine what you need to protect and the maximum data loss (RPO) and downtime (RTO) thresholds for compliance. Note the RPO, RTO, and retention requirements for each asset you will protect. You will need this information to create your job schedules.

Step 2: Determine how you will protect your assets.

- See the "Protection Overview" on page 91 for information on the types of assets and data Unitrends protects, backup types and modes, best practices, and other key concepts.

> **Note:** For Windows, VMware, or Hyper-V assets that have aggressive RPOs, consider adding replicas or instant recovery to your protection strategy. For details, see "Windows file-level replicas" on page 993, "VM replicas" on page 876, and "Virtual machine instant recovery" on page 904.

- Review the applicable protection requirements and considerations. These vary by backup type. See these chapters for details:
  - "Host-level Backups Overview" on page 653 – *Host-level backups* protect VMware, Hyper-V, AHV, and XenServer virtual machines by leveraging hypervisor snapshots.
  - "File-level Backups Overview" on page 703 – *File-level backups* protect an asset's file system and operating system. You must install a Unitrends agent on the asset to run file-level backups.
  - "Windows Image-level Backups Overview" on page 709 – *Image-level backups* protect a Windows asset at the disk and volume level. You must install the Unitrends Windows agent on the asset to run an image-level backup.

    > **Note:** You can opt to protect a Windows asset with file-level backups, image-level backups, or both backup types. The Windows agent supports both backup methods.

  - "NAS Backups Overview" on page 724 – *NAS backups* protect data stored on a NAS device.
  - "Application Backups Overview" on page 733 – *Application backups* protect applications, such as Exchange, SQL, and Oracle. You must install a Unitrends agent on the host asset to run application backups.

Step 3: Add the assets to the backup appliance, as described in "Managing protected assets" on page 286.

Step 4: Add any backup copy targets to the appliance, as described in "Backup copy targets" on page 214. (Required only if you will by copying backups to a hot or cold target.)

Step 5: Review "About creating backup and backup copy jobs" to determine how you will create jobs.

About creating backup and backup copy jobs

Backup and backup copy jobs are initiated either by job schedules or on-demand by Unitrends users. Typically, schedules are used to implement a comprehensive protection strategy that aligns with the business continuity plan, and on-demand jobs are run in addition, as needed.

- To set up job schedules, you can:
  - Create each schedule manually by using the Create Backup Job and Create Backup Copy Job dialogs.
  - Use the SLA Policy Automation feature to define policies that the appliance uses to create the required backup and backup copy schedules.
  - Use a combination of both methods.
- To run an on-demand job, you can:
  - Create a new job and run it immediately, by using the Create Backup Job or Create Backup Copy Job dialogs.

- Run a scheduled job on-demand from the Job Manager tab.

Which method should I use to create backup and backup copy schedules?

Unitrends offers a wide range of backup types to protect over 100 versions of servers, storage, operating systems, hypervisors, and applications. The first step in creating a backup job or SLA policy is selecting the *type* of asset you want to protect (for example, *file-level assets* or *VMware assets*). The asset type you select determines which backup method the appliance uses and the type of backup that is created. A given backup job or SLA policy can protect one or more assets of the type you select. (When picking assets to include, only the ones that match this asset type display in the UI.)

To decide which scheduling method to use, start by identifying the assets you will protect. The manual job creation method can be used for all asset types (physical machines, virtual machines, applications, etc.). SLA Policy Automation is supported for these asset types only: file-level Windows, file-level Linux, image-level Windows, and VMware, Hyper-V, and AHV virtual machines. See the table below for a complete list of asset types, their associated backup types, and whether SLA policies are supported.

Note: You can opt to run file-level backups of a virtual machine by installing a Unitrends agent on the VM and adding it to the appliance as an agent-based asset. The appliance then treats the VM as a physical asset.

For asset types supported by both scheduling methods:

- SLA Policy Automation provides the quickest, simplest method for creating schedules that align with your business continuity plan's RPO and retention requirements.
- Creating schedules manually gives you granular control and supports additional features. For example, you can choose which backup modes to use and when to run full backups.

For a detailed comparison of the scheduling methods, see "Methods for scheduling jobs".

Backup type and SLA policy support by asset type

Asset type	Backup type	SLA Policy Automation supported?
File level	file-level	Yes for Windows and Linux No for other file-level assets
Image level	image-level	Yes for Windows only
VMware	host-level	Yes
Hyper-V	host-level	Yes
AHV	host-level	Yes

Backup type and SLA policy support by asset type		
Asset type	**Backup type**	**SLA Policy Automation supported?**
XenServer	host-level	No
NAS	NAS	No
NDMP	NAS	No
Exchange, Oracle, SQL, SharePoint, or Cisco UCS	application	No

Methods for scheduling jobs

You can create schedules manually, use the SLA Policy Automation feature, or use a combination of both methods. Each policy or schedule can be used for a single asset type. See the following table for details about each method:

Schedule creation method	Description
SLA Policy Automation	The backup appliance automatically controls the initiation and flow of backups and backup copies through a single, simple policy.
How it works	You create a policy with the following information: • Assets to protect • RPO (backup frequency) • Backup retention • Whether to run hot and/or cold backup copies The backup appliance automatically creates the backup and backup copy schedules needed for the settings you specified in the policy. See "SLA backup schedule" on page 430, "SLA cold backup copy schedule" on page 430, and "SLA hot backup copy schedule" on page 431 below for descriptions of the additional settings used to create each schedule type.
Benefits	SLA policies provide these benefits: • Quickest way to set up schedules and backup retention. Use one dialog to create a single policy (instead of manually creating backup and backup copy schedules, and then setting retention by creating a long-term data management

Schedule creation method	Description
	policy or by editing each asset individually). • Simplest way to align with RPO goals. Enter the desired RPO instead of manually calculating when backups of each asset should run and creating schedules based on your calculations. • Simplest way to copy backups to your hot and/or cold backup copy target. Check one box to copy to your hot target and another box to copy to your cold target.
Requirements and considerations	The following requirements and considerations apply to SLA Policy Automation: • The backup appliance must be running version 10.0 or higher (10.3 or higher for Windows image-level assets). • The following asset types are supported: file-level Windows, file-level Linux, image-level Windows, and VMware, Hyper-V, and AHV virtual machines. You must manually create job schedules for other asset types. • You must log in to the backup appliance directly to create SLA policies. You cannot create a policy for a managed appliance by logging in to its manager appliance. • An asset can be assigned to only one SLA policy. • A file-level asset can be assigned to one SLA policy and/or to one or more manually created backup schedules. • An image-level Windows asset can be assigned to one SLA policy and/or to one or more manually created backup schedules. • A virtual machine asset can be assigned either to one SLA policy or to one manually created backup schedule (to ensure that the VM exists in only one backup schedule). • SLA policy schedules do not support the auto-include assets option. You must manually create a backup schedule to use this option. • The hot backup copy option is supported only if a Unitrends appliance or the Unitrends Cloud has been added as a backup copy target. (For details on adding a hot target, see

Schedule creation method	Description
	"Backup copy targets" on page 214.) • The cold backup copy option is supported only if the cold target has been added to the appliance and this target is one of the following types: third-party cloud, NAS, or iSCSI. (For details on adding a cold target, see "Backup copy targets" on page 214.) • If multiple cold targets exist, the policy copies backups to the one that was added first. To copy to a different cold target, you must manually create a backup copy job instead. • Do not directly edit job schedules that were created by an SLA policy. Instead, modify the SLA policy itself. The appliance automatically modifies the policy's schedules based on the changes that you make. On the Job Manager tab, SLA policy schedule names display with the prefix *_SLA*, so you can easily distinguish them from manually created schedules.
SLA backup schedule	A policy's backup schedule uses the *incremental forever* backup mode (for details, see "Incremental forever backup groups" on page 98).
SLA cold backup copy schedule	A policy's cold backup copy schedule uses the following settings. For additional backup copy considerations, see "Preparing to create a backup copy job" on page 491. • Backup copy target – Uses the first third-party cloud, NAS, or iSCSI target that was added to the backup appliance. • Backups copied – Copies the last backups. Includes all backup modes (fulls, incrementals, etc.). • Frequency – Runs daily at 3:00 AM. • Retention – Option to set minimum retention (the length of time a copy must be retained before it can be deleted). • Encryption – Option to encrypt cold copies. (Encryption must first be configured on the appliance. For details, see "Encryption" on page 155.) • Email report – Job report is emailed when the job completes.

Schedule creation method	Description
	• Delete older copies – Older copies are deleted from the target as follows: – Copies are deleted only if there is not enough space for the copy job. – Copies that are held by a retention policy cannot be deleted. – Eligible copies are deleted until there is enough space for the job. – If deleting eligible copies cannot free adequate space for the entire job, the job fails and nothing is written to the target.
SLA hot backup copy schedule	Hot backup copies are sent to the target each time an eligible backup is created. For additional backup copy considerations, see "Preparing to create a backup copy job" on page 491.
Manually create each schedule	The backup appliance runs backup and backup copy jobs as indicated in user-created schedules.
How it works	You create the backup and backup copy schedules manually. To set backup retention, you create a long-term data management policy or edit each asset (setting retention is optional). The appliance runs jobs at the times and frequency indicated in the schedules.
Benefits	Manually creating schedules gives you more granular control and additional options. Manually-created backup schedules provide these benefits: • Supported for all asset types. • Can set retention by asset, by creating long-term data management policies or by using the Edit Asset dialog. This enables you to apply different retention settings to each asset. (An SLA policy's retention setting applies to all assets in the policy. If an asset is assigned to a policy, you cannot modify its retention setting by using the Edit Asset dialog.) • Can schedule any backup modes that are supported for the asset type (e.g., fulls, differentials, incrementals).

Schedule creation method	Description
	• Can control when full backups run. • Can choose whether to receive email reports when the job completes. • Can use the auto-include option to automatically add newly detected assets to the schedule. • If you have created separate backup storage areas, you can specify which backup target to use. • For VMware, supports scheduling with regular expression filters. (For details, see "To create a VMware backup schedule by using regular expression filters" on page 458.) Manually-created cold backup copy schedules provides these benefits: • Supported for all types of cold targets (eSATA, USB, tape, third-party cloud, attached disk, NAS, and SAN). • Can copy only fulls (or copy all backup modes). • Can specify a date range of backups to copy (or copy last backups). • Can specify the days and times when the jobs will run. • Can choose to fail the job without deleting copies if there is no more space available (or choose to delete older copies if there is no more space available). • These features are also supported and work just as they do with SLA policies: – Can encrypt backup copies. (Encryption must first be configured on the appliance. For details, see "Encryption" on page 155.) – Can set minimum retention (the length of time a copy must be retained before it can be deleted).
Requirements and considerations	Requirements and considerations vary by backup type and backup copy target. See these topics for details: • "Backup copy targets" on page 214

Schedule creation method	Description
	• "Host-level Backups Overview" on page 653 • "File-level Backups Overview" on page 703 • "Windows Image-level Backups Overview" on page 709 • "NAS Backups Overview" on page 724 • "Application Backups Overview" on page 733

Creating backup jobs

Use these procedures to manually create backup jobs.

These procedures assume that you have added to the appliance each asset you wish to protect. For details on adding assets, see these topics: "Managing protected assets" on page 286 and "Protecting SQL clusters and availability groups" on page 748 (provides instructions for adding assets to protect clustered SQL environments).

Unitrends also recommends that you review the following information to develop the best protection strategy for your environment: "Preparing for backups" on page 425 and "About creating backup and backup copy jobs" on page 426.

See the following topics to create backup jobs:

- "Selecting assets to back up" on page 433
- "Backup job procedures" on page 437

Notes:

- For Windows, Linux, VMware, Hyper-V, and AHV assets, you can create SLA policies instead of creating individual backups jobs. For details on how SLA policies work, see "Methods for scheduling jobs" on page 428.
- iSeries backup jobs are not created in the UI. The procedures in this topic do not apply to iSeries. To create an iSeries backup job, see "iSeries Backups Overview and Procedures" on page 767.

Selecting assets to back up

Any physical machine, virtual machine, or application you add to the appliance is an asset you can select to include in the jobs you create. When creating a backup job, you first select the type of asset you want to protect:

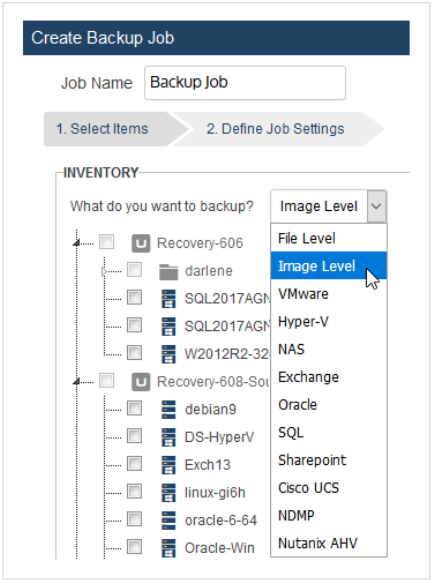

The inventory tree is filtered to display assets of the type you selected.

You can then select assets to include in the job. To select an asset, click on its check box. In some cases, you must expand nodes in the tree to view the asset you want to select. Assets display in the Inventory tree as follows:

Backup type	Description
File Level	Each file-level asset displays as a top-level node under its backup appliance. File-level assets include all physical machines (other than iSeries servers) and any virtual machine that is protected by installing a Unitrends agent and running file-level backups. 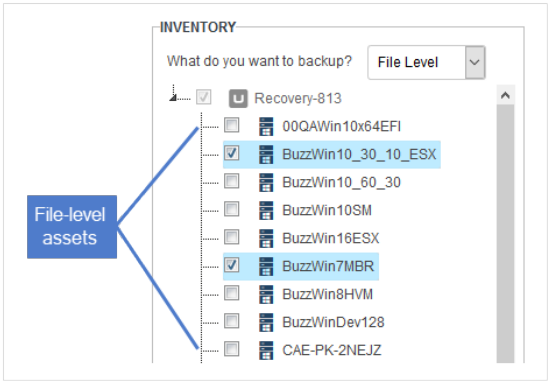
Image Level (Windows only)	Each Windows file-level asset displays as a top-level node under its backup appliance. Expand the Windows asset to select its image-level sub-node.

Backup type	Description
	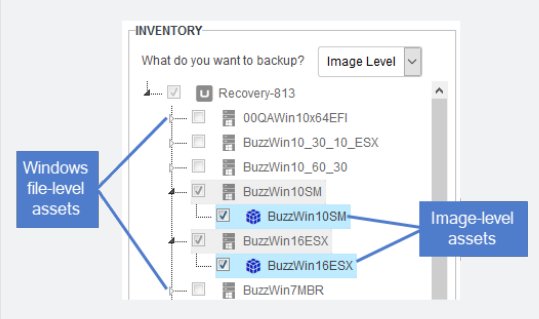
Host Level	Select one of the following in the **What do you want to backup** list: VMware, Hyper-V, XenServer, or Nutanix AHV. Each virtual host displays as a top-level node under its backup appliance. Each VM displays as a sub-node under its virtual host or under its Hyper-V instance. Select individual VMs or select a host to protect all of its VMs. VMware example: 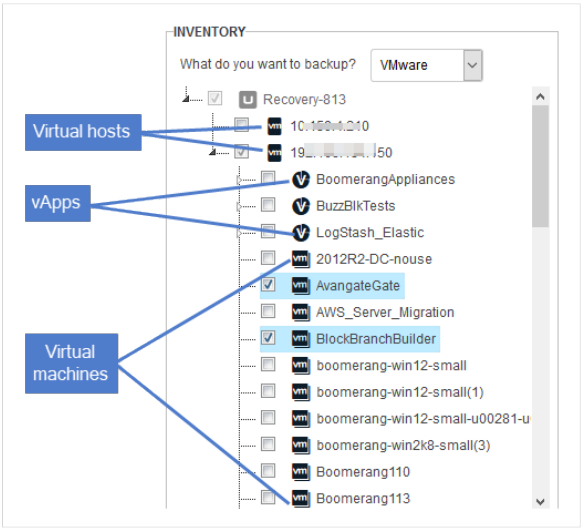 Hyper-V example:

Backup type	Description
	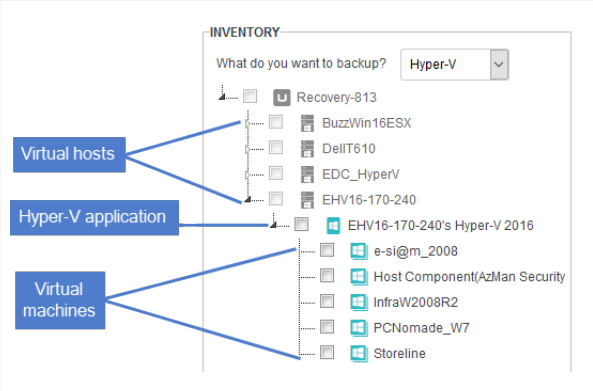
Application backups	Select one of the following in the **What do you want to backup** list: Exchange, Oracle, SQL, SharePoint, Cisco UCS, or NDMP. Each application displays as a sub-node under its host server asset. Instances and individual databases display under their application node. 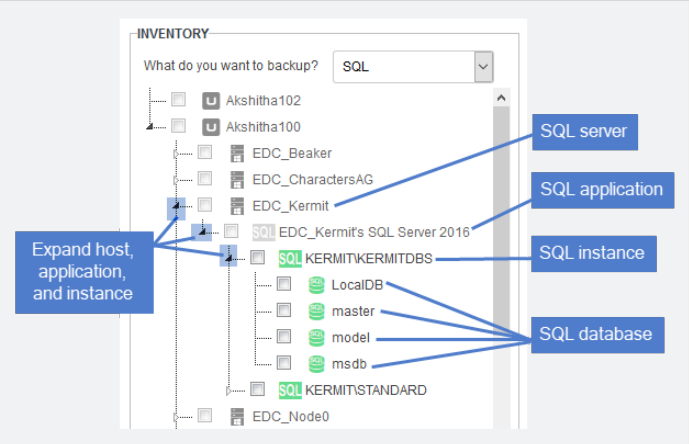
NAS backups	Each NAS asset displays as a top-level node under its backup appliance. 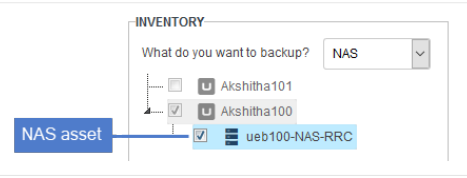

Backup job procedures

Use these procedures to create backup jobs:

To create a file-level backup job

Note: A file-level asset can be assigned to one or more manually created backup schedules and/or to one SLA policy. (SLA policies are supported for these file-level assets only: Windows and Linux.)

1 Select **Jobs > Active Jobs > Create Job > Backup**.

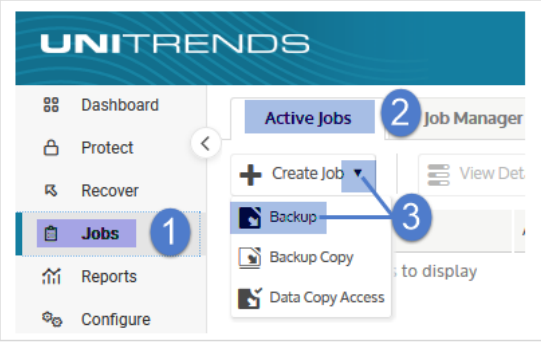

2 Select **File Level** in the **What do you want to backup?** list.

3 In the Inventory tree, check boxes to select the asset(s) you want to protect. Selected assets display in the Job Inventory Settings area.

To locate an asset by name, use the **Search** field below.

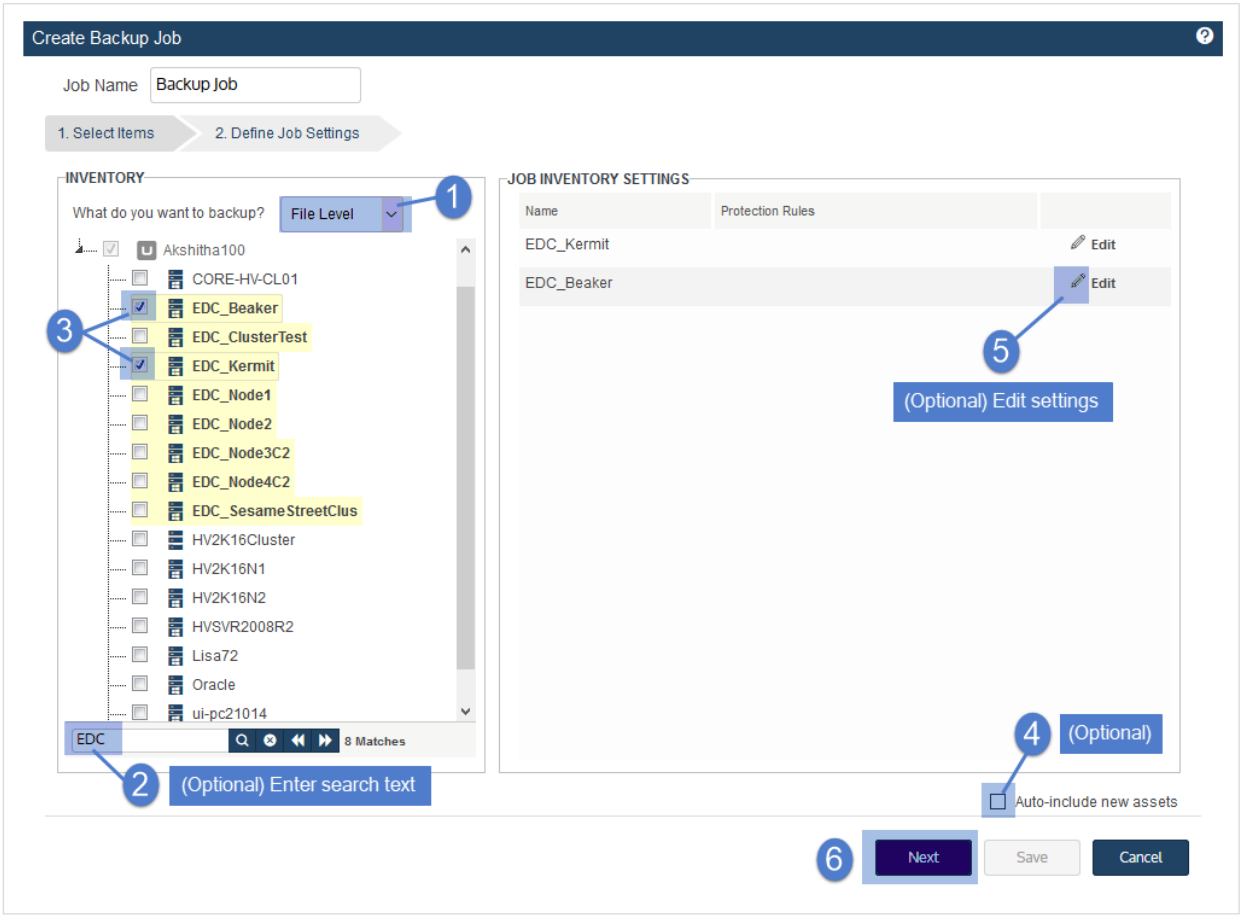

4 (Optional) Check the **Auto-include new assets** box to automatically add newly discovered file-level assets to the schedule.

5 (Optional) Edit Job Inventory Settings to apply options, such as data to include or exclude and commands to run pre- and/or post-backup: locate the asset in the list and click **Edit**, modify settings, then click **Save** to retain any changes.

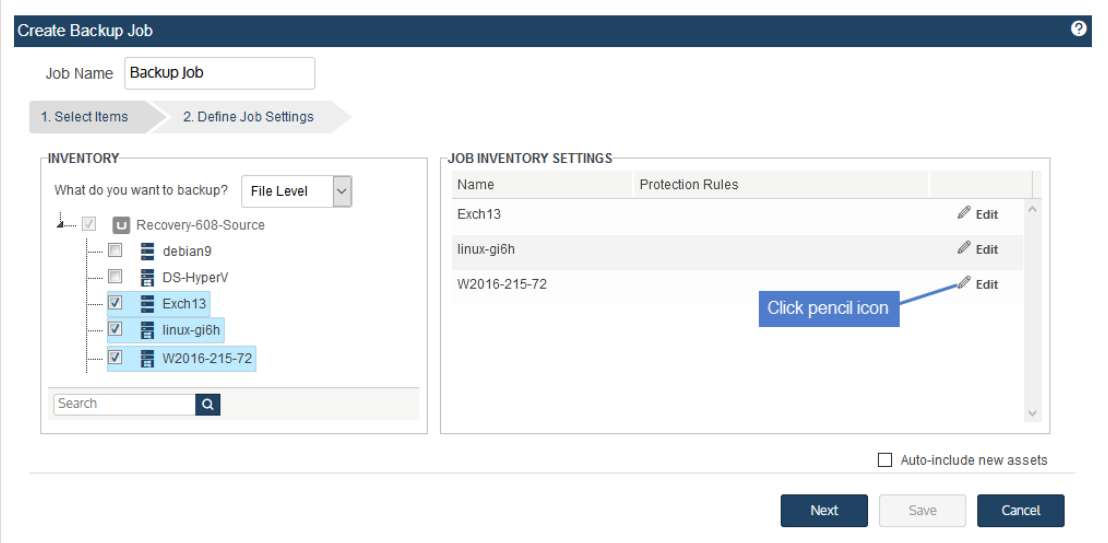

See the following for details and considerations:

File-level setting	Description
General considerations for including or excluding data from an asset's backups	Review the following before specifying data to include or exclude: • When creating a schedule, any inclusions or exclusions you add are applied to jobs run by that schedule only. Inclusions and exclusions are not applied automatically in other cases. • Running an on-demand backup of the asset does not automatically apply any inclusions or exclusions specified in the asset's schedule. To run an on-demand backup, do one of the following: – Create a one-time job that has the same inclusions and exclusions as in the schedule. – Manually run the schedule (select the schedule under **Jobs > Job Manager** and click **Run**). – Run a one-time Selective backup (so that a new full is not created). • If you specify both files to include and files to exclude, the inclusion is applied first. Any exclusions are then applied to the subset of included files.
Inclusion tab	Click to specify files, folders, or volumes to include in backups of this asset. • Data that does not meet the criteria you specify here is NOT included in the backup.

File-level setting	Description
	• Type in the full path (e.g., **C:/Documents**) or **Browse** the asset to specify data to include. (Wildcards are not supported.) • If you are running a full backup and include files or folders in the system drive (typically C:), do *not* check the System State box on the Advanced tab. Full backups fail if system state is excluded. • Run a new full backup upon creating or modifying included files. Example: 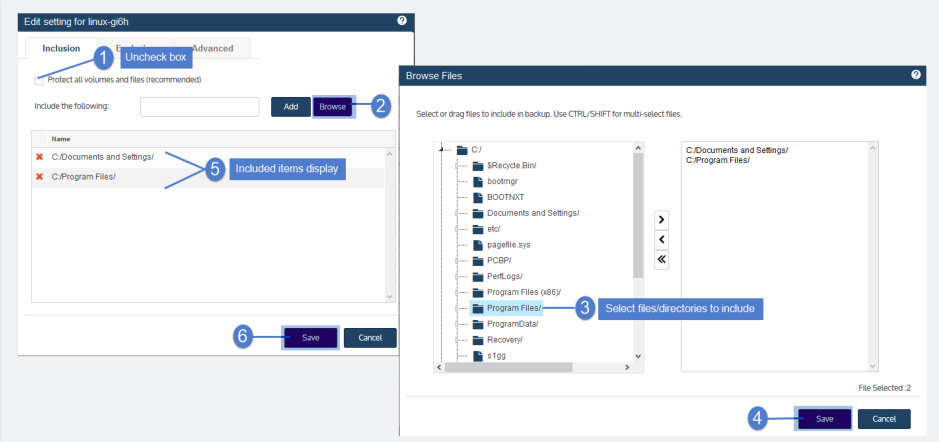
Exclusion tab	Click to specify files, folders, or volumes to exclude from backups of this asset. • Data that does not meet the criteria you specify here IS included in the backup. • To specify files to exclude, do any of the following: – Type in the full path (e.g., **C:/Documents**). – **Browse** the asset. – Enter a selection pattern. Wildcards are supported for Windows assets. Wildcards are not supported for these asset types: Linux, Unix, and NAS. See these rows below for usage examples: "Wildcard * usage", "Wildcard ? usage", and "Multiple wildcards". • Run a new full backup upon creating or modifying excluded files. Example:

File-level setting	Description
	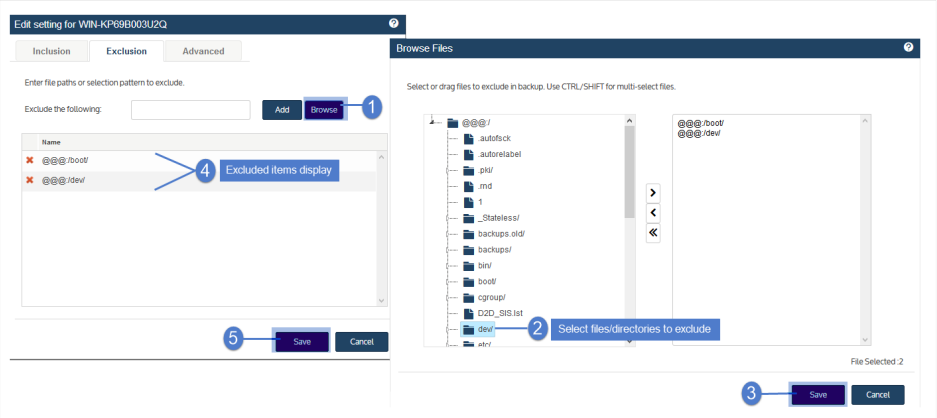
Wildcard * usage	An example of how to exclude all files with zero or more characters that match exclusion pattern: ***.txt** An example of how to exclude directories with zero or more characters and their contents within a specified path that match the exclusion pattern: **C:/windows/sys*** Limitations: • ***folder_abc** cannot be used to exclude all folders that match *folder_abc* on the protected asset. The full path must be provided. • If an entire directory is excluded, the directory name will still appear in the backup; however, its contents will be empty. • Multiple wildcard matches like the following are not supported: **C:**\abc.txt**
Wildcard ? usage	An example of how to exclude all files within specified path that matches a single character within exclusion pattern: **C:/PCBP/Lists.dir/pro_client?.spr** An example of how to exclude all directories and their contents within specified path that matches a single character within exclusion pattern: **C:/Programfiles/Case?/** Limitation: If an entire directory is excluded, the directory name itself will still appear in the backup; however its contents will be empty.
Multiple wildcards	An example that uses multiple “?” wildcards and only one * wildcard: **C:/?Log?/*.logs**

File-level setting	Description
	Limitation: If an entire directory is excluded, the directory name itself will still appear in the backup; however its contents will be empty.
Advanced tab	Use this tab to specify advanced options. See these rows below for details: "Advanced Exclusions", "Command to run Pre-Backup", and "Command to run Post-Backup". Example: 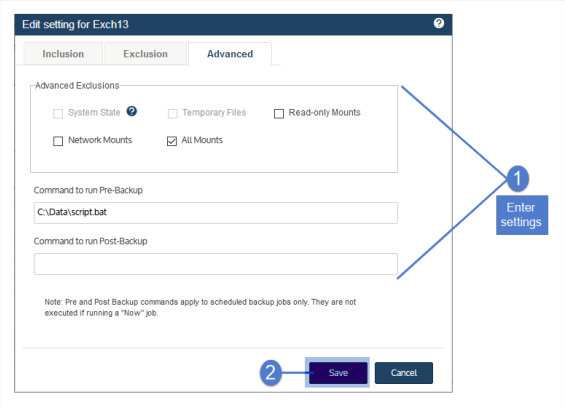
Advanced Exclusions	Check one or more boxes to exclude any of the following: system state, temporary files, read-only mounts, network mounts, or all mounts. Consider the following before applying advanced exclusions: • To perform bare metal recovery or use Windows replicas, the following must be included in the backup: system state and all boot and critical system (OS) disks/volumes. If you need these features for the asset, do not specify data to include or exclude unless you are sure these disks/volumes will be included. • If you are running a full backup and have selected files or folders in the system drive (typically C:) on the Inclusion tab, do *not* check the System State box on the Advanced tab. Full backups fail if system state is excluded. • Creating aliases for an asset - Adhere to the following when creating aliases for an asset: – You must include the system state on the asset whose backups contain the boot and critical OS volumes. – You must exclude the system state on the other aliased assets. This approach ensures you can perform bare metal recovery of the asset.

File-level setting	Description
	– Only one asset can include the system state. Disaster recovery of the asset fails if the system state is not included with the boot and OS volume or if the system state is included on aliased assets that do not include the boot and OS volume. **IMPORTANT!** For Windows assets, the backup must contain the system state, boot disk and any other system critical volumes to use the integrated bare metal recovery and Windows replica features. Be sure one of the aliased assets contains all of these disks to use these features.
Command to run Pre-Backup	To run a command or script on the asset before a scheduled backup starts, enter the full path to the command or script in the **Command to run Pre-Backup** field. For example, *C:\Data\script.bat* or */usr/jsmith/script.sh*. **Note:** Pre-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).
Command to run Post-Backup	To run a command or script on the asset after a scheduled backup completes, enter the full path to the command or script in the **Command to run Post-Backup** field. For example, *C:\Data\script.bat* or */usr/jsmith/script.sh*. **Note:** Post-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).

6 Click **Next**.

7 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

8 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

9 Click **Save**.

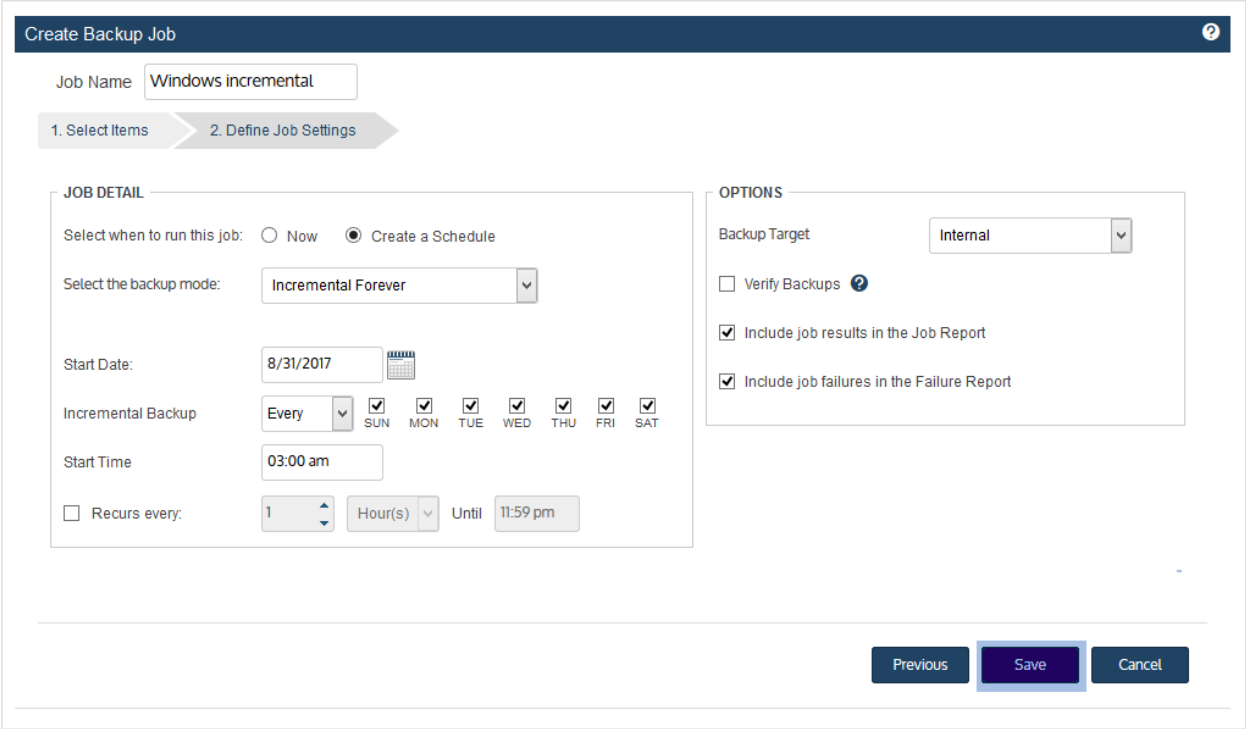

10 Click **OK** to close the Success message.

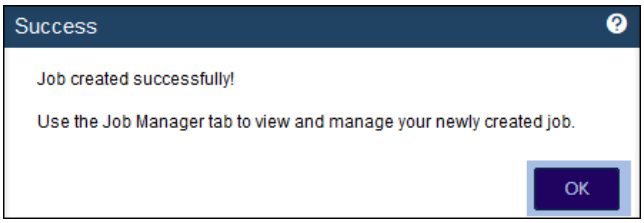

- If you created a schedule, the job runs at the date and times specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

Using the Custom backup mode in the Create Backup Job dialog

In most cases, the standard backup modes can be used to create backup schedules. If you need more granularity, you can opt to use the Custom backup mode.

After adding assets to the backup job, the Define Job Settings step displays. To use the Custom backup mode:

1. Select **Custom** in the **Select the backup mode** list.
2. Click the calendar icon.

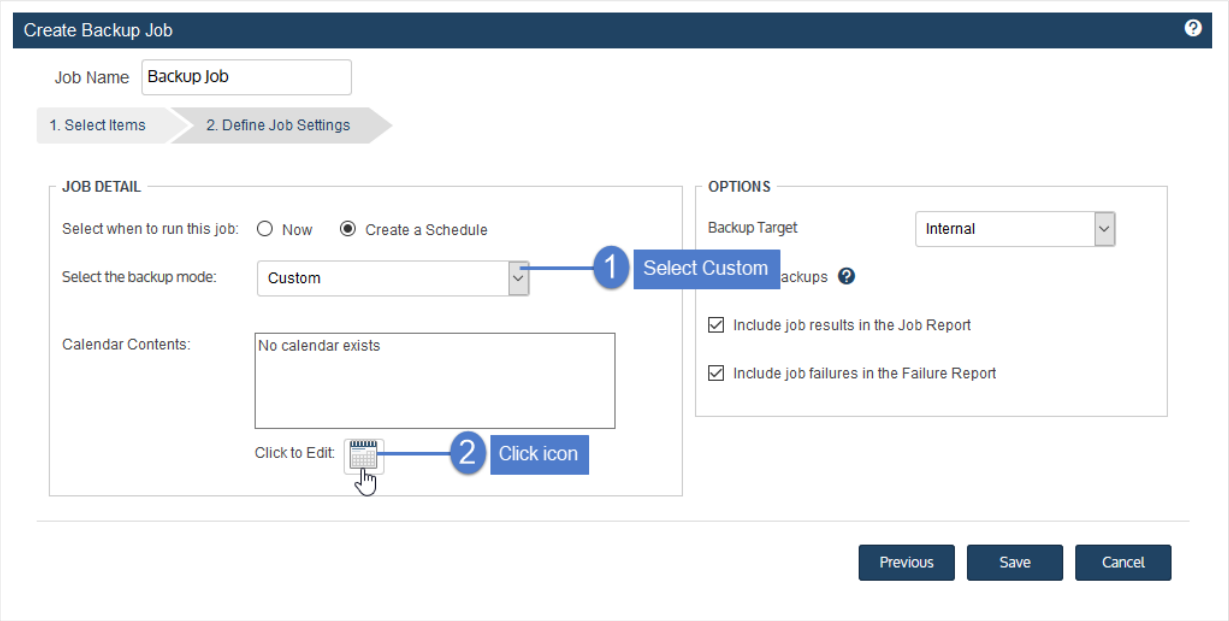

3. In the Calendar dialog, select a backup mode in the Backups area and drag it to a day on the calendar. (You cannot drag to a day in the past.)
4. In the Add Backup dialog, modify settings and click **Save**.

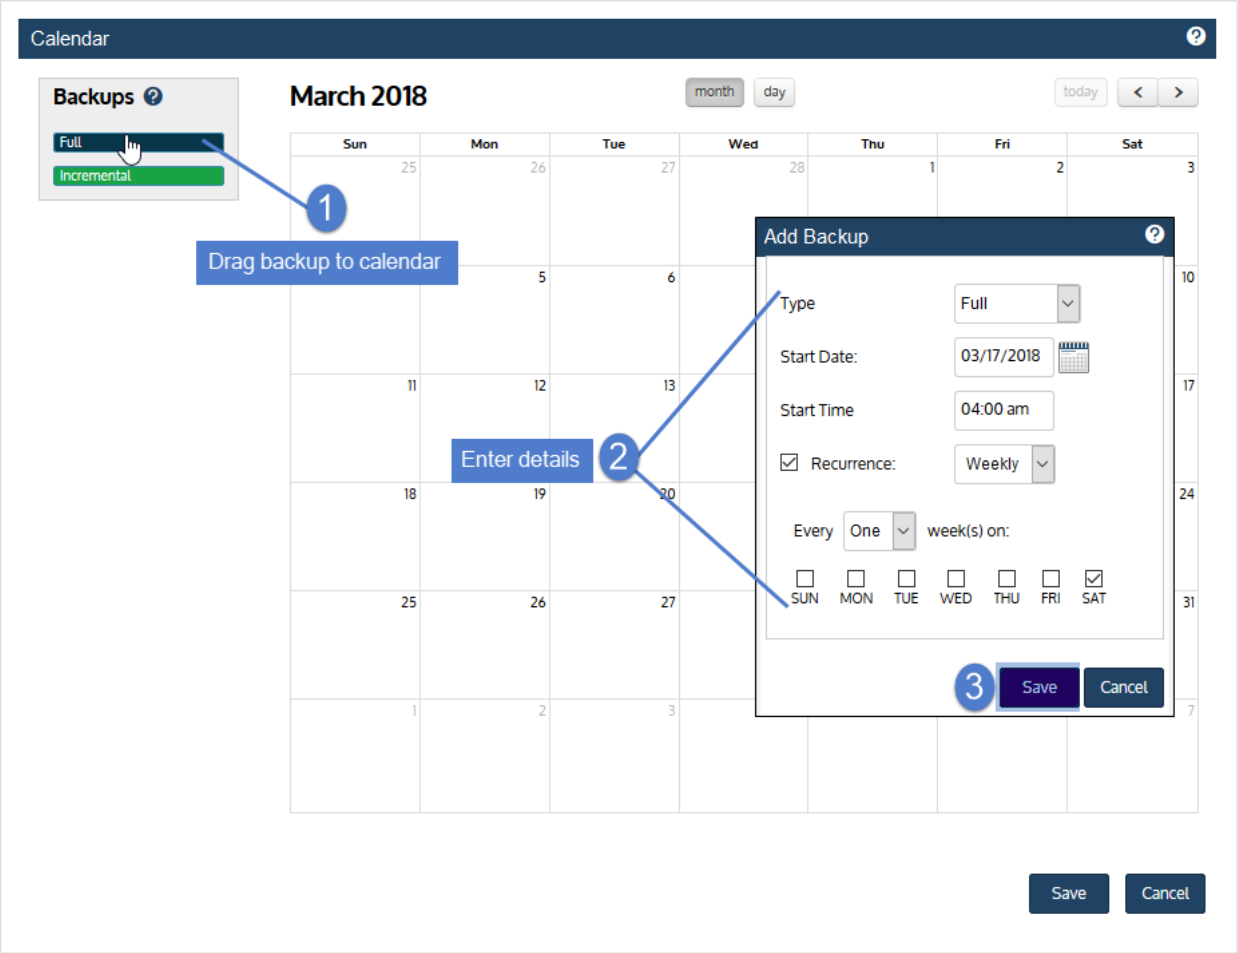

5 Repeat these steps to add other modes to the calendar.

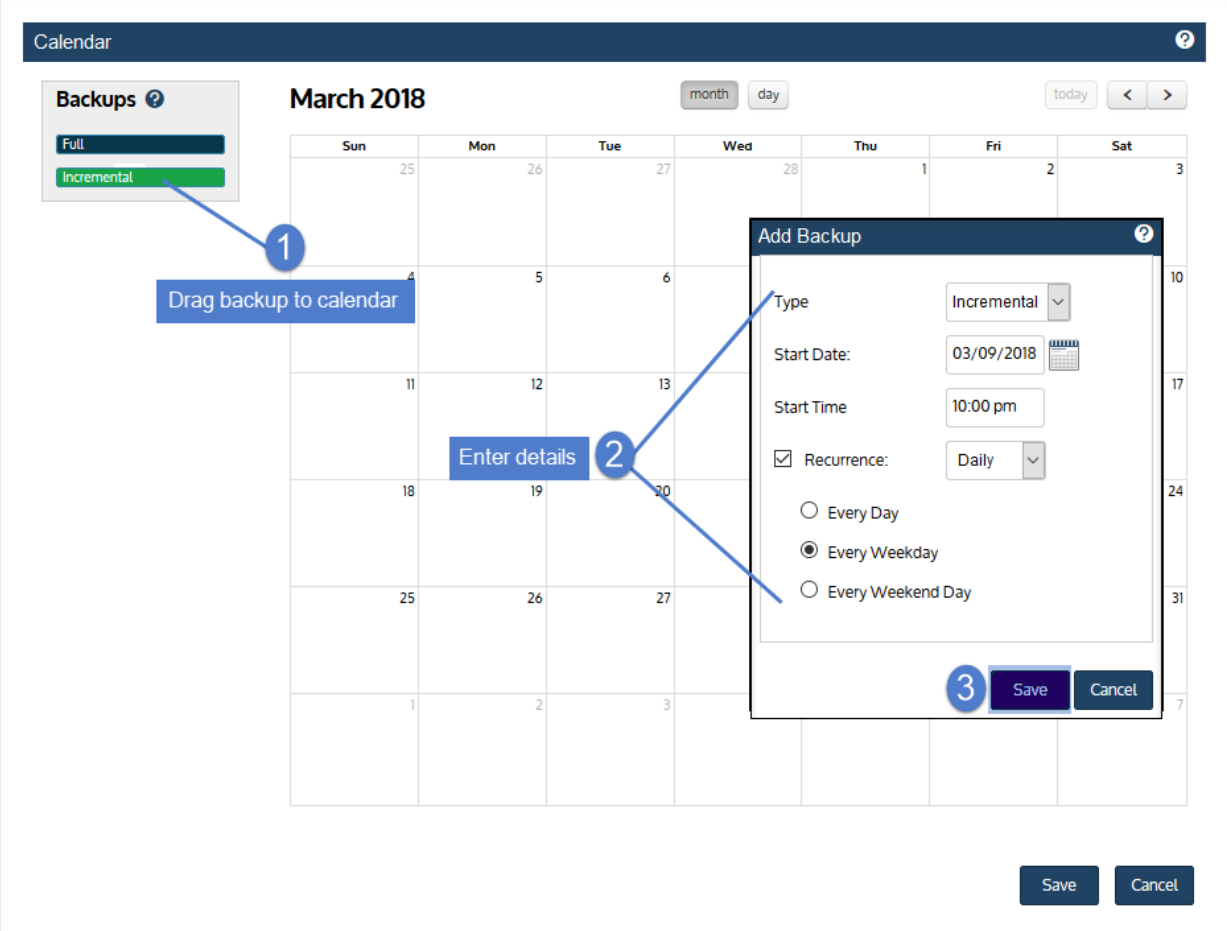

6 Click **Save** to save the settings and close the Calendar dialog.

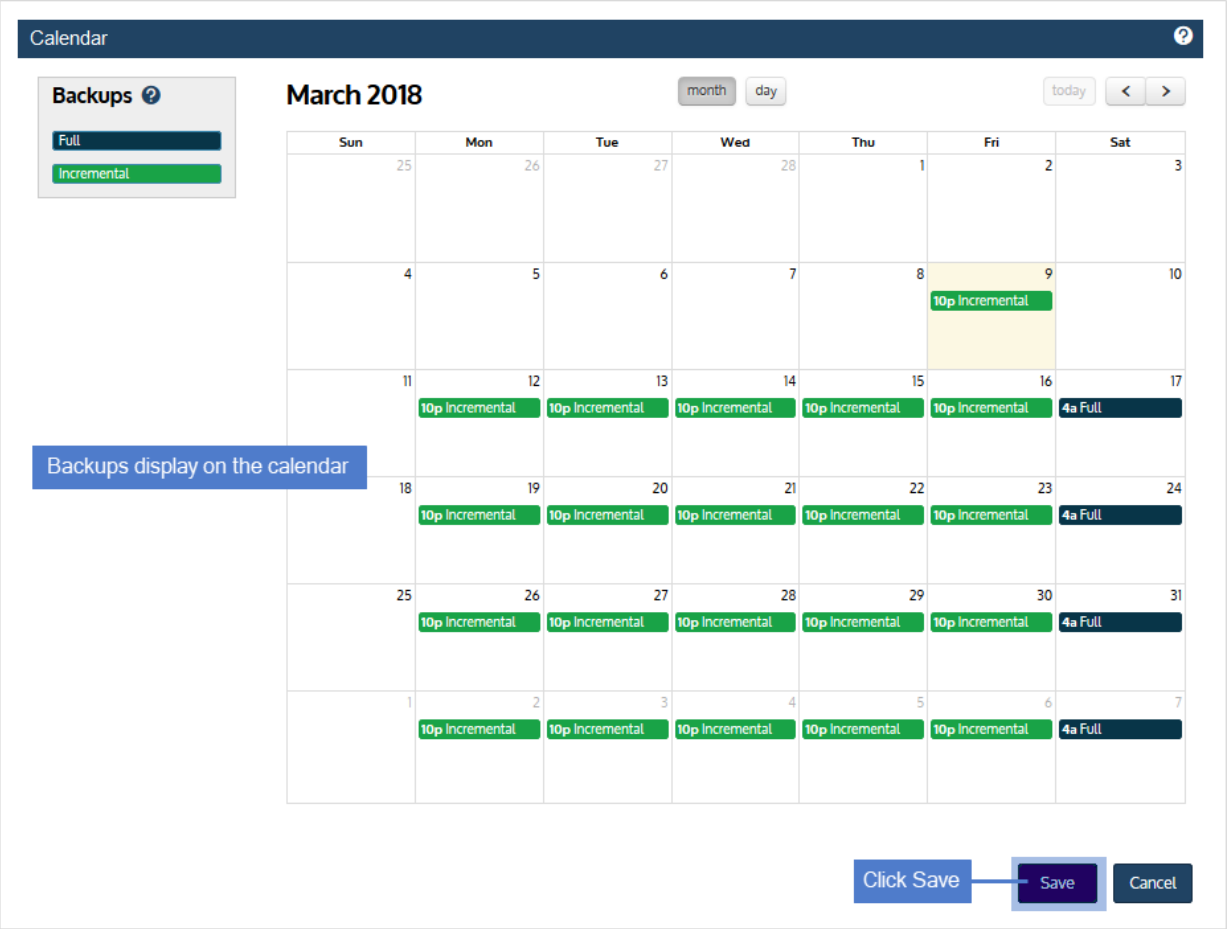

7 Click **Save** to save the schedule.

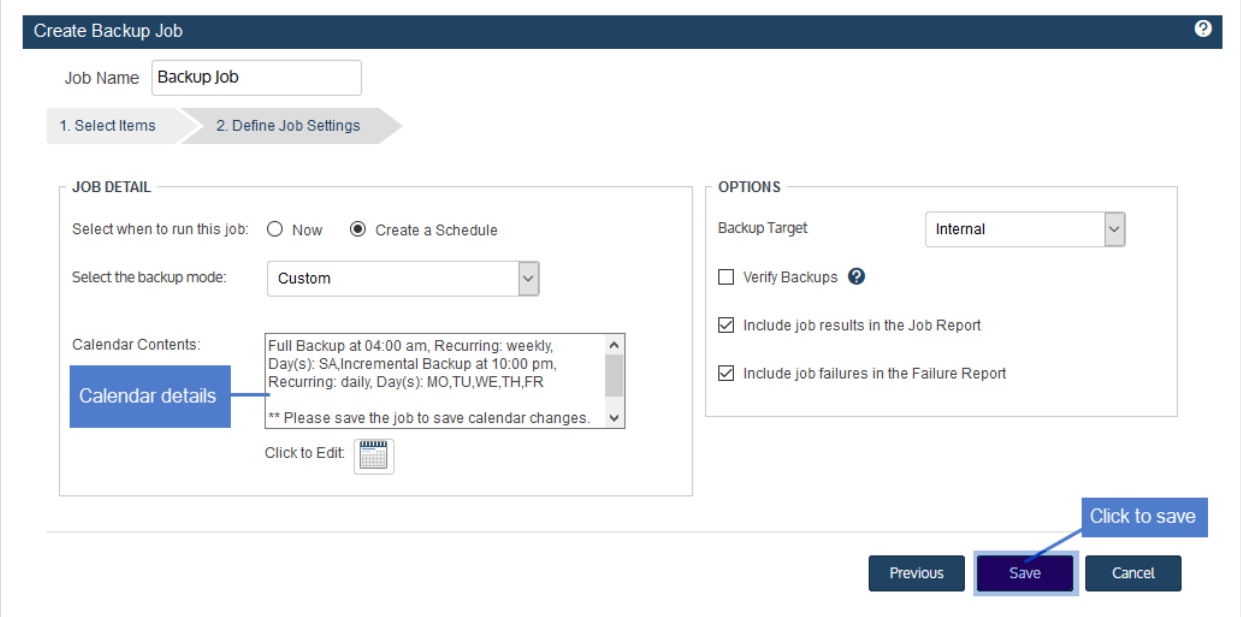

To create an image-level backup job

Notes:

- An image-level asset can be assigned to one or more manually created backup schedules and/or to one SLA policy.
- You can opt to use the application aware feature when running image-level backups. To use this feature, , check these boxes on the Edit Assets page before running the backup: **Show Image Level Backup Settings** and **Allow application aware**. For details, see "To edit an agent-based asset" on page 293. For more on protecting hosted applications with image-level backups, see "Windows Image-level Backups Overview" on page 709.
- Beginning in release 10.4.8, you have the option to index Windows image-level backups so you can quickly search for and recover individual files. To index an asset's image-level backups, check these boxes on the Edit Assets page before running the backup: **Show Image Level Backup Settings** and **Index Image-Level Backups**. For details, see "To edit an agent-based asset" on page 293.

1 Select **Jobs > Active Jobs > Create Job > Backup**.

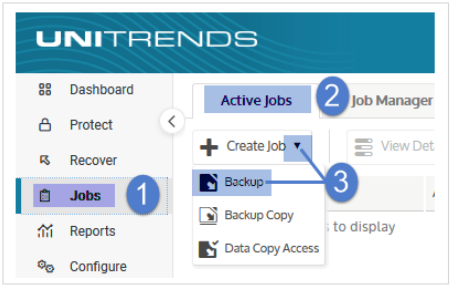

2 Select **Image Level** in the **What do you want to backup?** list.

3 In the Inventory tree, expand assets and check boxes to select the image-level assets you want to protect. Selected assets display in the Job Inventory Settings area.

- Windows assets running agent version 10.3 or higher can be selected in the list. Other assets are disabled.
- To locate an asset by name, use the **Search** field below.

Note: If image-level protection is not supported for an asset in the job, the appliance runs a file-level backup instead.

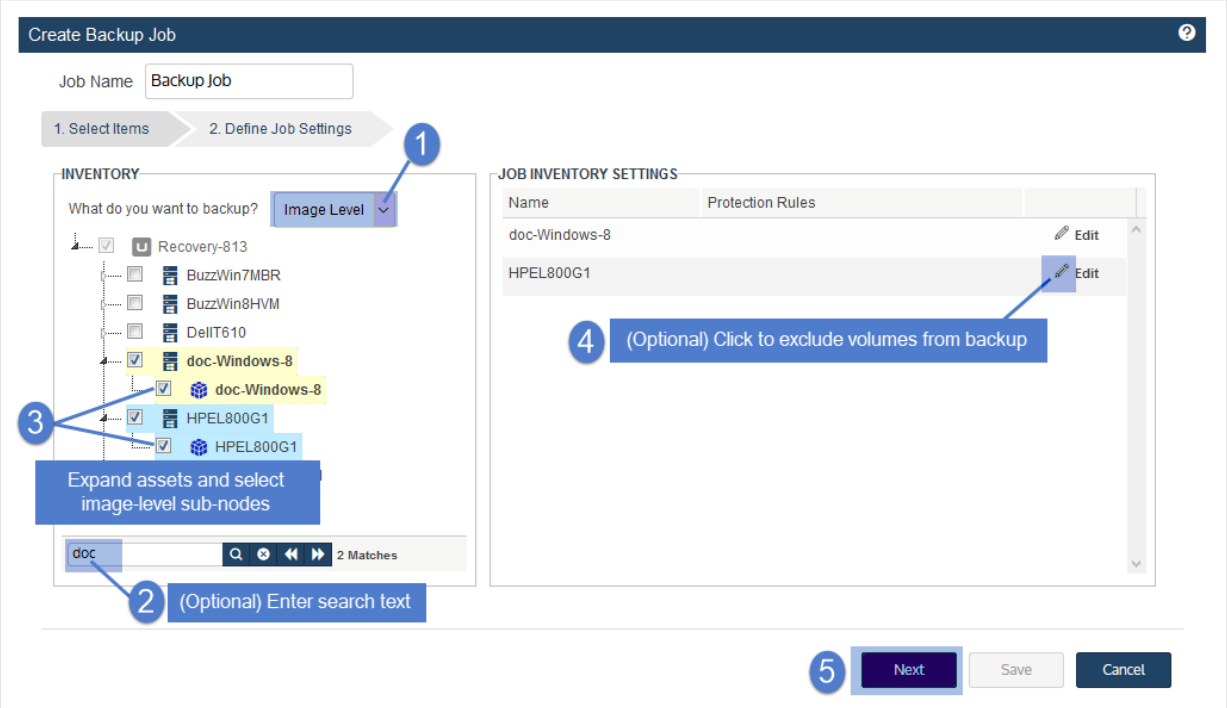

4 (Optional) Edit Job Inventory Settings to apply options, such as volumes to include or exclude and commands to run pre- and/or post-backup: locate the asset in the list and click **Edit**, modify settings, then click **Save** to retain any changes.

Notes:

- Critical system volumes are required for the image-level replicas feature and to recover the entire asset . Use care when omitting volumes from backup.
- When you recover the entire asset, any existing data on the target is overwritten or deleted. Volumes on the target disk that were excluded from backup may also be overwritten. For details, see "Windows unified bare metal recovery" on page 1209.
- To recover a SQL server, the master, model, and msdb system databases must be present in the image-level backup of the Windows asset. (These are included by default. If you want the recovered asset to

include a hosted SQL application, use care not to exclude these system databases from the image-level backup.)

- Image-level protection is not supported for read-only disks. You must exclude all volumes on read-only disks from the backup job or run file-level backups. Image-level backups fail if read-only volumes have not been excluded.
- Removable media is automatically excluded from image-level backups. (You do not need to exclude volumes on a read-only disk that resides on removable media.)

See the following for details:

Image-level setting	Description
General considerations	Review the following before specifying volumes to include or exclude: • When creating a schedule, any inclusions or exclusions you add are applied to jobs run by that schedule only. Inclusions and exclusions are not applied automatically in other cases. • Running an on-demand backup of the asset does not automatically apply any inclusions or exclusions specified in the asset's schedule. To run an on-demand backup, do one of the following: – Create a one-time job that has the same inclusions and exclusions as in the schedule. – Manually run the schedule (select the schedule under **Jobs > Job Manager** and click **Run**).
Inclusion tab	Click to specify volumes to include in backups of this asset. • Any volumes that you do not select here are NOT included in the backup. • Type in the volume (e.g., **C:/**) or **Browse** the asset to select volumes to include. • Run a new full backup upon creating or modifying included volumes.
Inclusion example:	

Image-level setting	Description
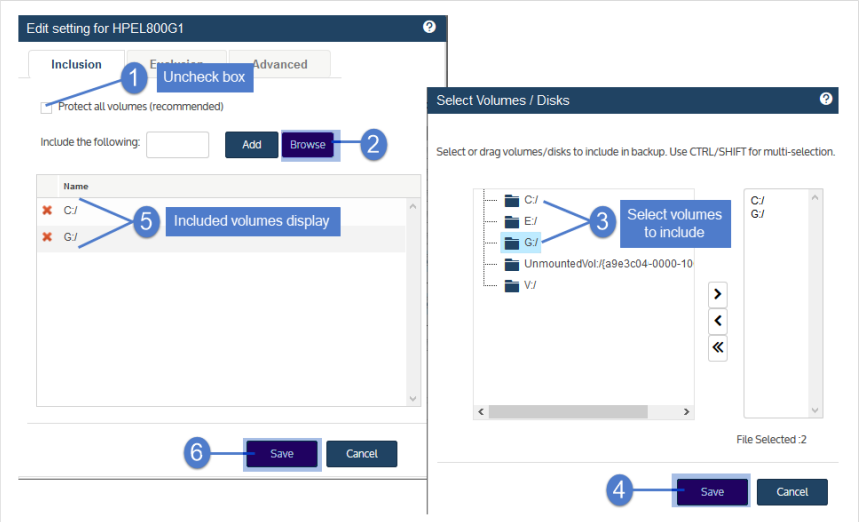	
Exclusion tab	Click to specify volumes to exclude from backups of this asset. • Any volumes that you do not select here are included in the backup. • Type in the volume (e.g., **C:/**) or **Browse** the asset to select volumes to exclude. • Run a new full backup upon creating or modifying excluded volumes.
Exclusion example:	

Image-level setting	Description
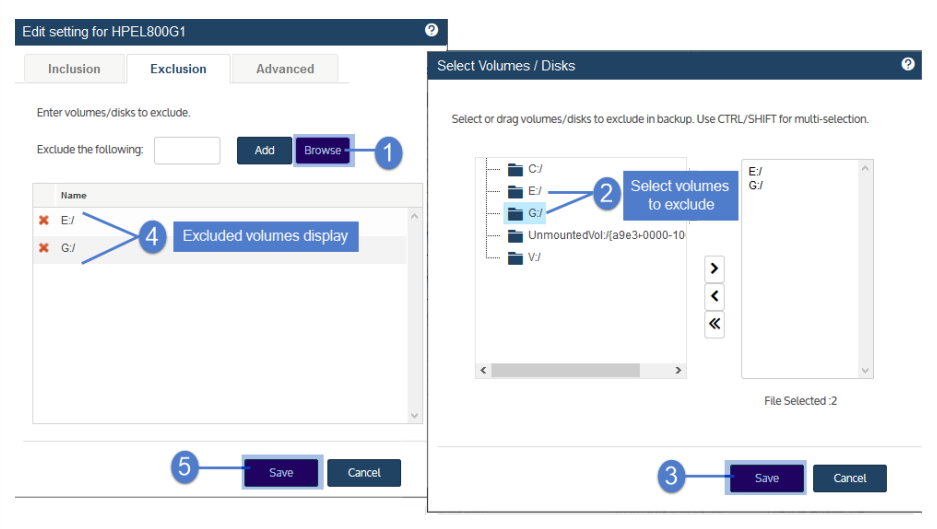	
Advanced tab	Use this tab to specify commands to run before or after the backup job. (Advanced Exclusion options are not supported for image-level backups.) See these rows below for details: "Command to run Pre-Backup" and "Command to run Post-Backup". Example: 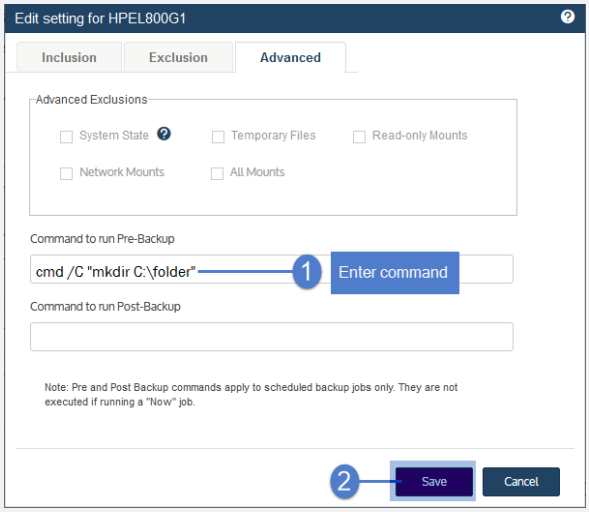
Command to run Pre-Backup	For image-level backup schedules, you cannot enter command syntax as you would inside the Windows Command Prompt. Instead, you must use this format:

Image-level setting	Description
	cmd /<argument> "<commandSyntax>" For example: *cmd /C "mkdir C:\folder"* To run a command or script on the asset before a scheduled backup starts, enter the command in the **Command to run Pre-Backup** field. **Note:** Pre-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).
Command to run Post-Backup	For image-level backup schedules, you cannot enter command syntax as you would inside the Windows Command Prompt. Instead, you must use this format: *cmd /<argument> "<commandSyntax>"* For example: *cmd /C "mkdir C:\folder"* To run a command or script on the asset after a scheduled backup completes, enter the command in the **Command to run Post-Backup** field. **Note:** Post-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).

5 Click **Next**.

6 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

7 Set remaining Job Details and Options, then click **Save**:

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

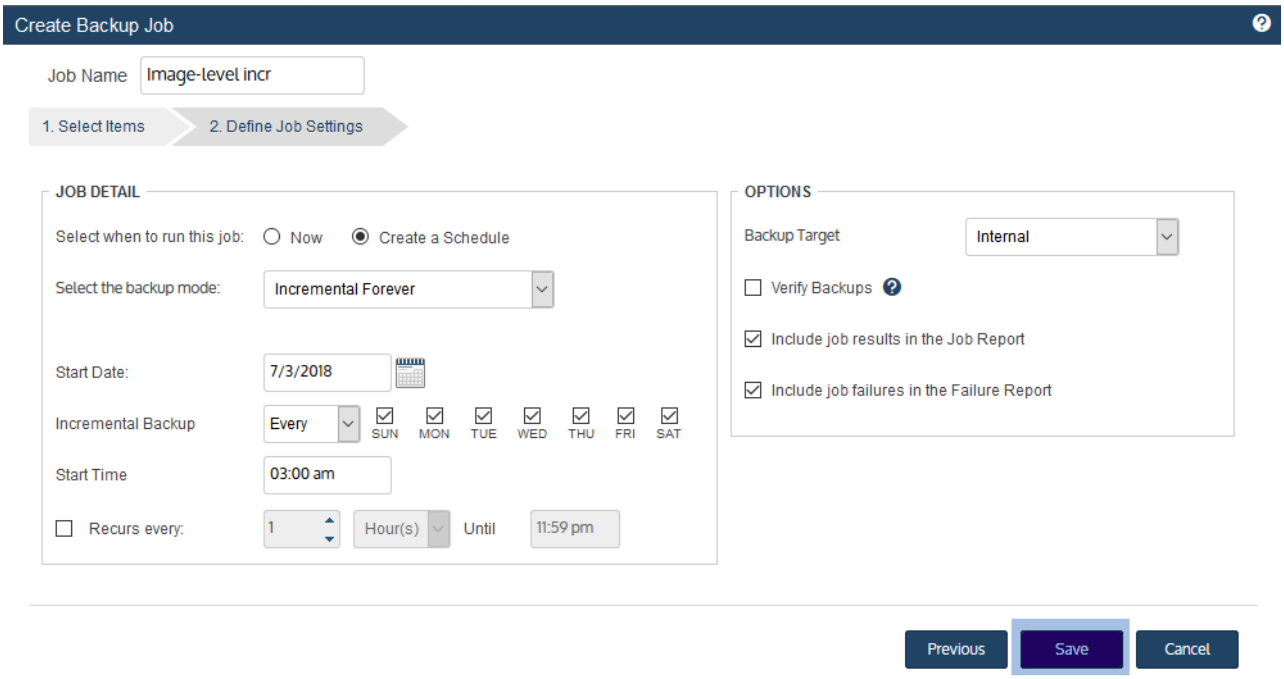

8 Click **OK** to close the Success message.

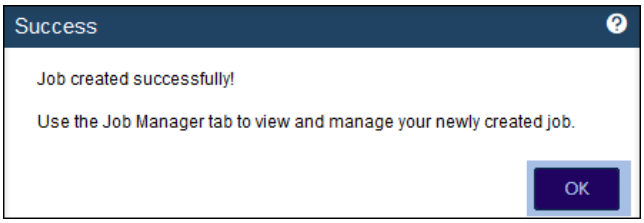

- If you created a schedule, the job runs at the date and times specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

To create a VMware backup job

Notes:

- A VMware asset can be assigned either to one manually created backup schedule or to one SLA policy (to ensure that the VM exists in only one backup schedule).
- To access newly added virtual machines, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

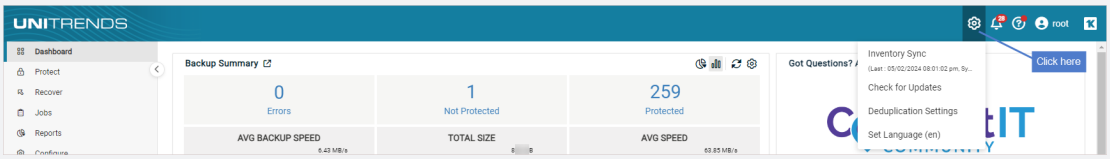

1 Select **Jobs** > **Active Jobs** > **Create Job** > **Backup**.

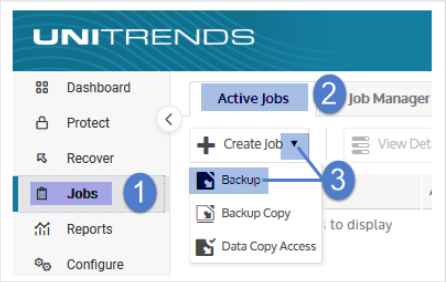

2 Select **VMware** in the **What do you want to backup?** list.

3 In the Inventory tree, expand the virtual host and check boxes to select virtual machines to protect. Selected VMs display in the Job Inventory Settings area.

- To locate an asset by name, use the **Search** field below.
- To view individual VMs, expand the virtual host and any vApps and resource pools.
- To quickly select multiple VMs, click a virtual host, vApp, or resource pool checkbox.
- To select one VM, click its checkbox.

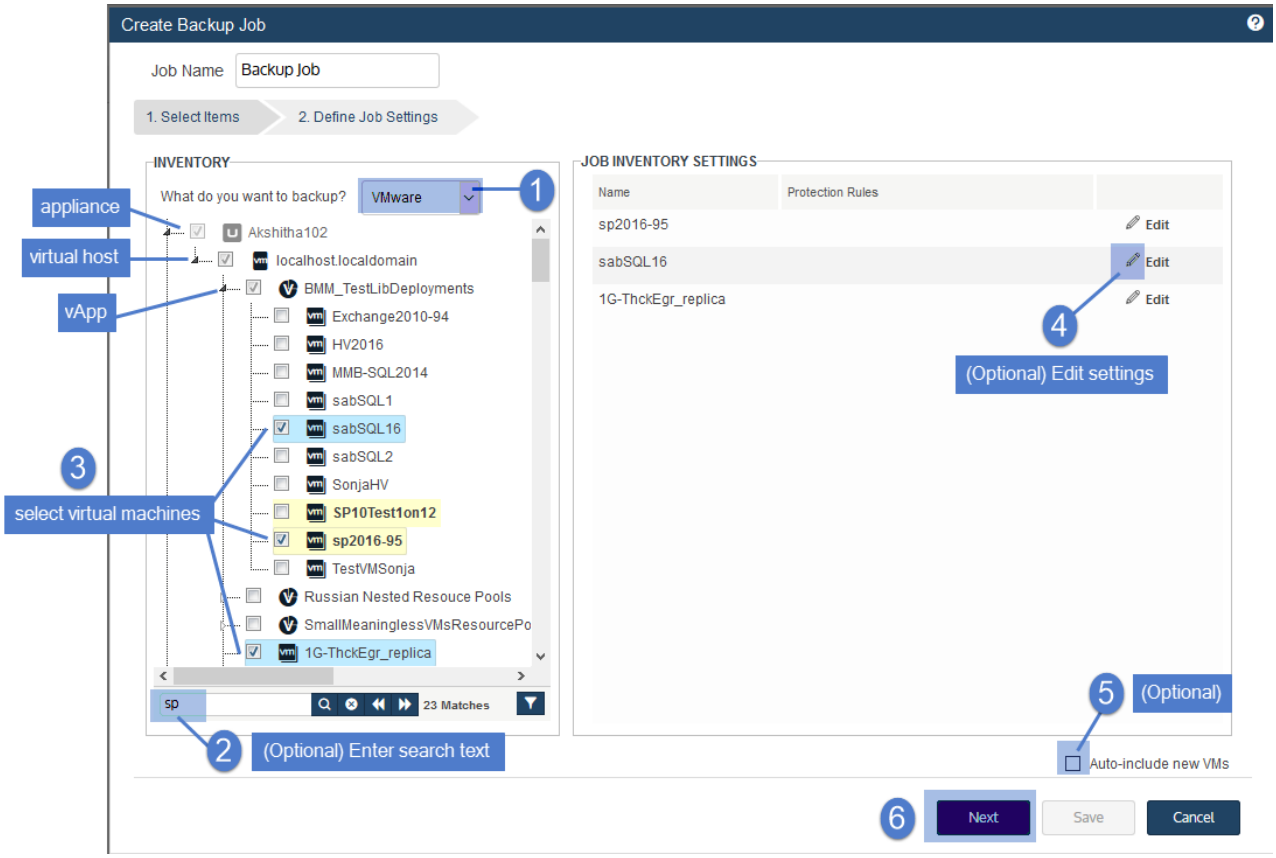

4 (Optional) Check the **Auto-include new VMs** box to automatically add newly discovered VMs to the schedule.

5 (Optional) Edit Job Inventory Settings to exclude VM disks from backup:

- Locate the VM in the Job Inventory Settings list.
- Click **Edit** to specify disks to exclude.
- Click **Save** to retain any changes.

> **Note:** Critical system volumes are required to recover the entire virtual machine. Use care when omitting disks from backup.

6 Click **Next**.

7 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

8 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

9 Click **Save**.

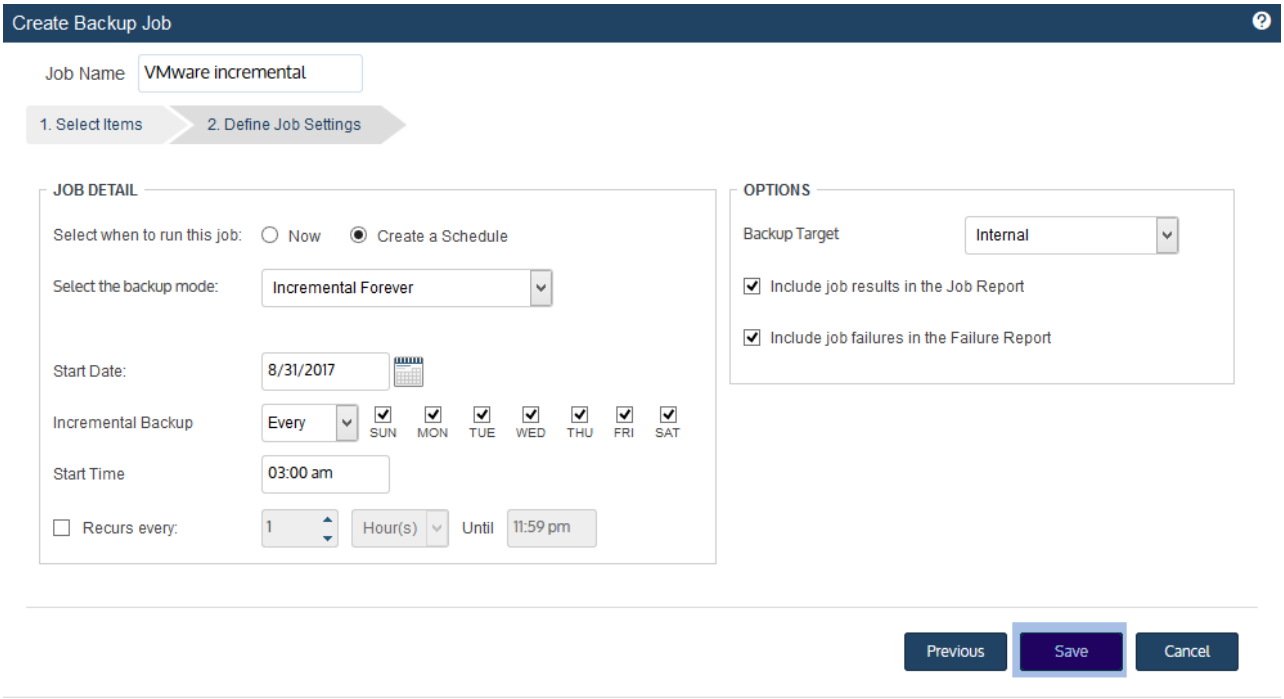

10 Click **OK** to close the Success message.

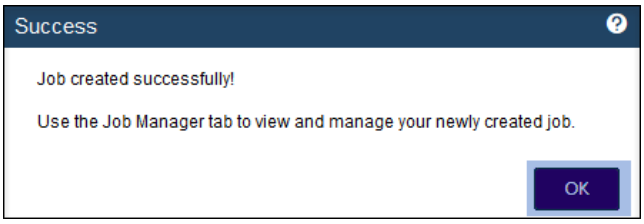

- If you created a schedule, the job runs at the date and times you specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

To create a VMware backup schedule by using regular expression filters

If you have a large virtual environment, creating filters for your backup schedules greatly reduces the overhead of adding VMs to your schedules and modifying schedules as your VM inventory changes. Filtered schedules automatically adjust to protect virtual machines that are created or deleted in your VMware environment. Once a VM is deleted from the hypervisor, it is automatically removed from the schedule. Any new VM that meets the filter criteria is automatically added to the schedule.

Note: Beginning in release 10.1, filtered schedules are updated each hour. This enhancement yields better performance, especially in environments with large numbers of VMs. If you are running backups at a greater frequency, any backup failures caused by the schedule not yet detecting a deleted VM can be ignored. Any run that does not back up a newly added VM can be ignored. The condition will resolve itself within an hour.

A filter consists of the following elements:

- A *name* that defines the VMware container type that will be searched. For example, ESX Servers, vApps, or VM DisplayName.
- A *filter string* that is the text that the filter searches for.
- An *action* that is applied to the container list and filter string to create the list of VMs to include in the schedule. For example, Equal, Contains, or Starts With.

Consider the following when working with filters:

- Filters are supported only for VMware backup schedules. Filters cannot be used for one-time backups.
- Filter combinations must be unique to a single schedule.
- Filters are logical “and” statements; “or” statements are not supported.

Notes:

- A VMware asset can be assigned either to one manually created backup schedule or to one SLA policy (to ensure that the VM exists in only one backup schedule).
- To access newly added virtual machines, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

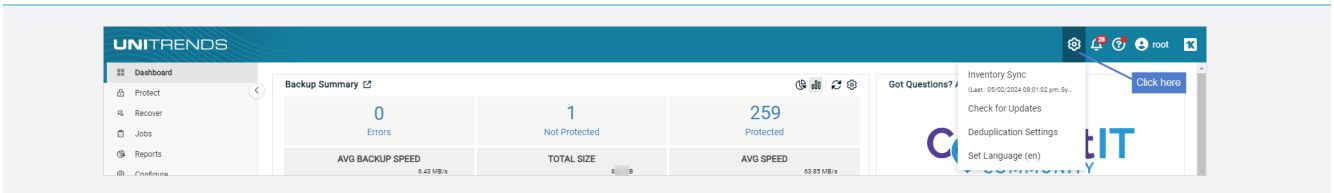

Use this procedure to create a schedule by using a regular expression filter:

1 Select **Jobs > Active Jobs > Create Job > Backup**.

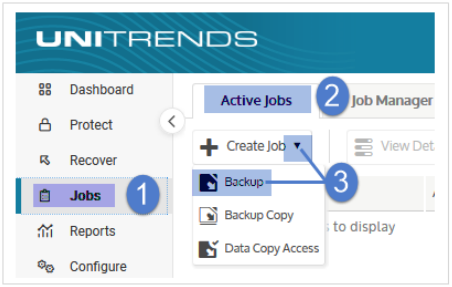

2 Select **VMware** in the **What do you want to backup?** list.

3 Click to select a virtual host in the Inventory tree.

4 Click the filter icon below the Inventory tree.

5 Add the filter:

- In the **Enter name** list, select a VMware container type.
- In the **Enter action** list, select an action.
- In the **Filter** field, enter a text string.
- Click the checkmark to apply.

Note: Filter actions for resource pools, vApps, and folders apply to all containers that meet the filter criteria and to any sub-containers. For example, the filter *Resource Pools Equal Nested Resource Pools* would apply to all VMs in these resource pools: Nested Resource Pools, Number1, Another Resource Pool, Number 2 and Number 3.

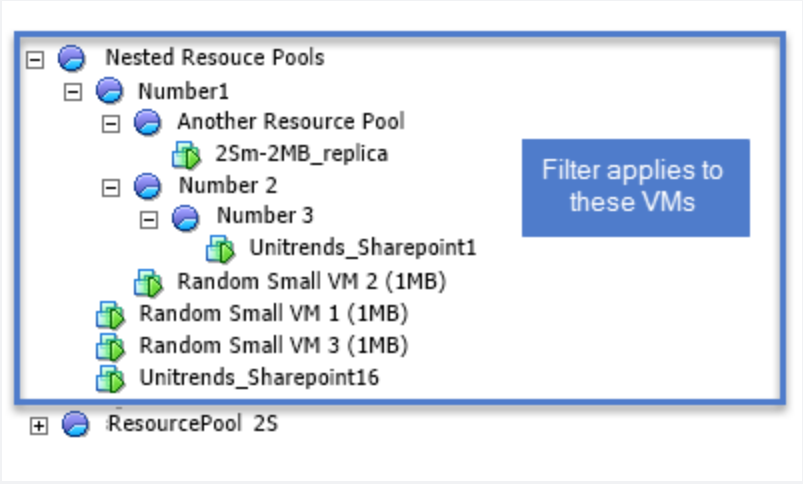

VMs meeting the filter criteria display in the Job Inventory Settings list.

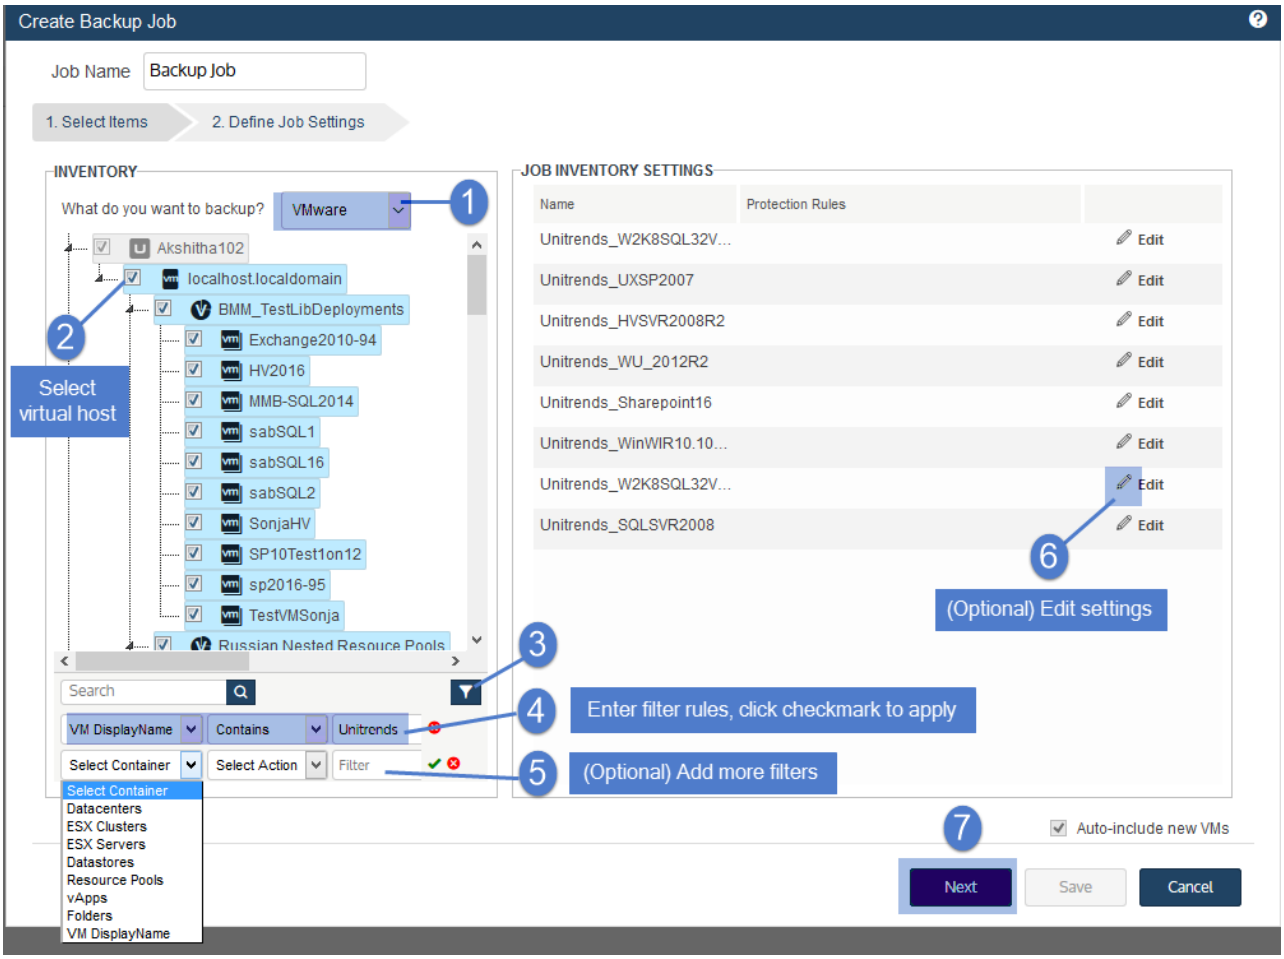

6 (Optional) Add more filters as needed to narrow the VM list.

7 (Optional) Edit Job Inventory Settings to exclude VM disks from backup:

- Locate the VM in the Job Inventory Settings list.
- Click **Edit** to specify disks to exclude.
- Click **Save** to retain any changes.

> **Note:** Critical system volumes are required to recover the entire virtual machine. Use care when omitting disks from backup.

8 Click **Next**.

9 Enter a unique Job Name and select **Create a Schedule**.

10 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup

mode in the Create Backup Job dialog" on page 445.

11 Click **Save**.

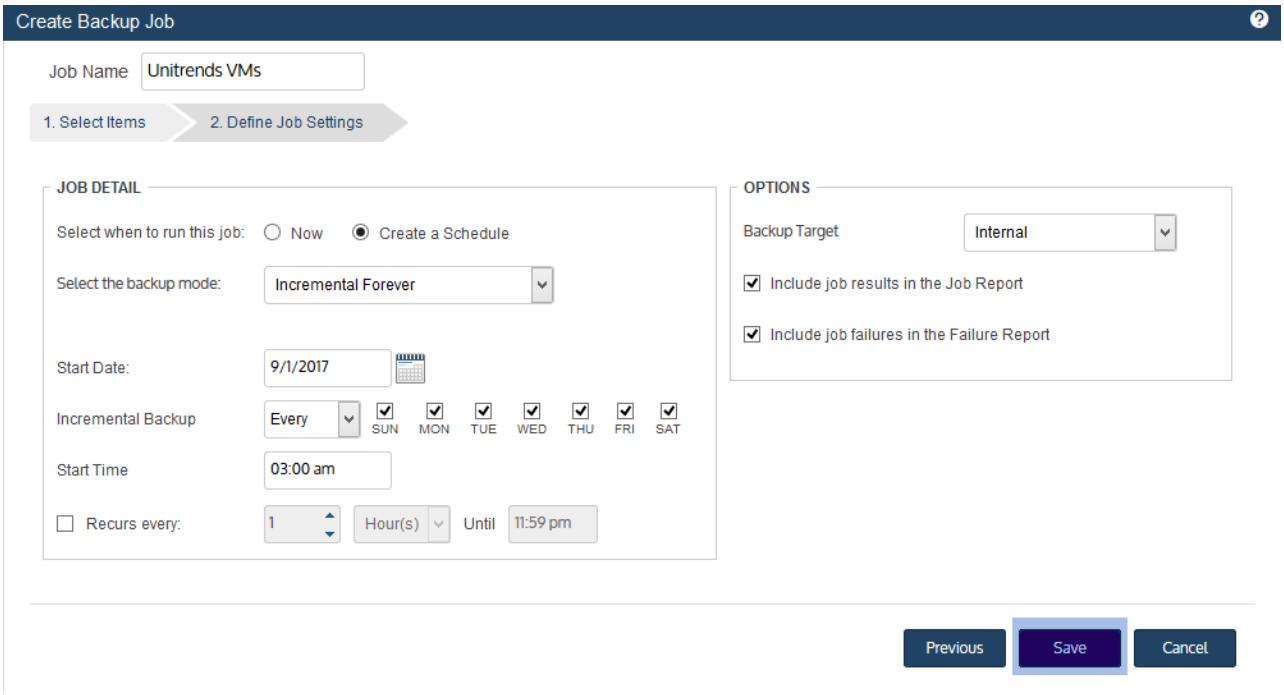

To create a Hyper-V backup job

Notes:

- A Hyper-V asset can be assigned either to one manually created backup schedule or to one SLA policy (to ensure that the VM exists in only one backup schedule).
- To access newly added virtual machines, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

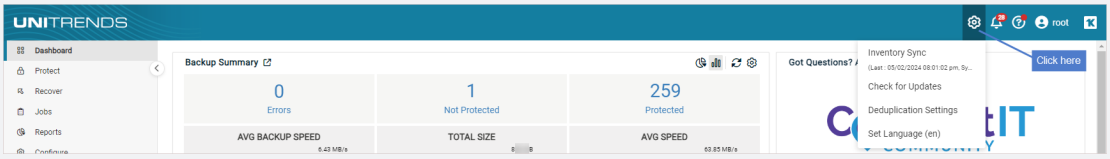

1 Select **Jobs > Active Jobs > Create Job > Backup**.

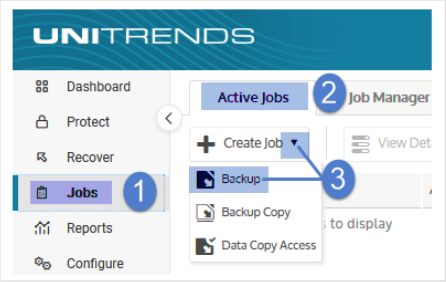

2 Select **Hyper-V** in the **What do you want to backup?** list.

3 In the Inventory tree, expand the Hyper-V server and application, then check boxes to select virtual machines to protect. Selected VMs display in the Job Inventory Settings area.

- To locate an asset by name, use the **Search** field below.
- To view individual VMs, expand the Hyper-V server and application.
- To quickly select all hosted VMs, click an application checkbox.
- To select one VM, click its checkbox.

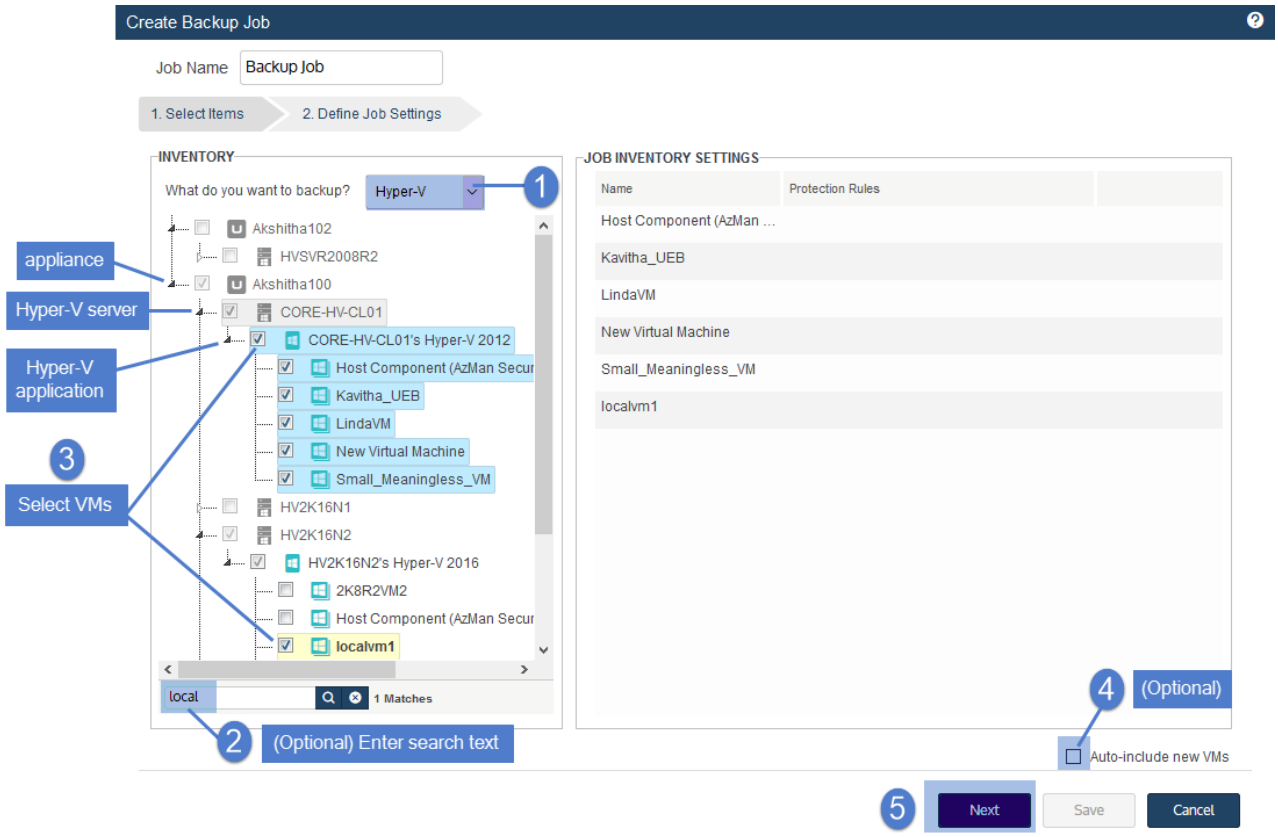

4 (Optional) Check the **Auto-include new VMs** box to automatically add newly discovered VMs to the schedule.

5 Click **Next**.

6 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

7 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

8 Click **Save**.

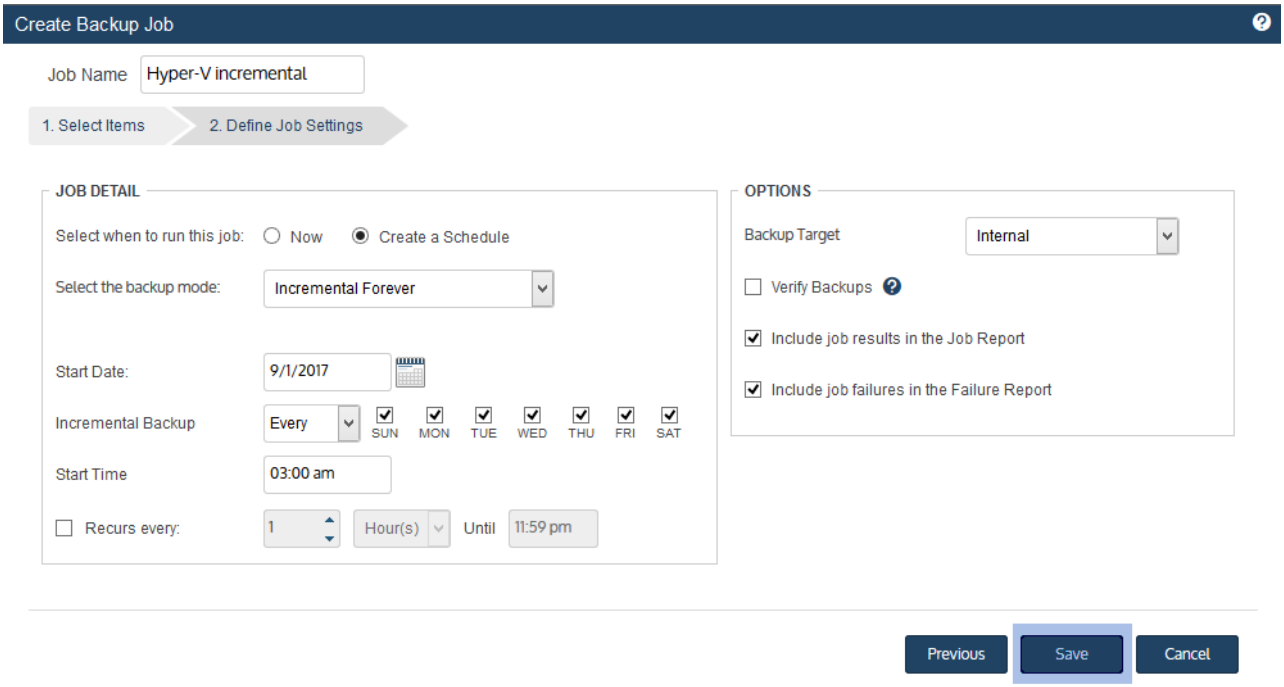

9 Click **OK** to close the Success message.

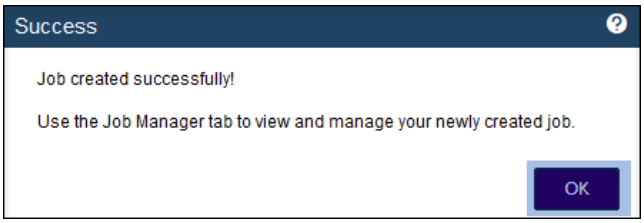

- If you created a schedule, the job runs at the date and times you specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

To create a Nutanix AHV backup job

Notes:

- An AHV asset can be assigned either to one manually created backup schedule or to one SLA policy (to ensure that the VM exists in only one backup schedule).
- To access newly added virtual machines, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

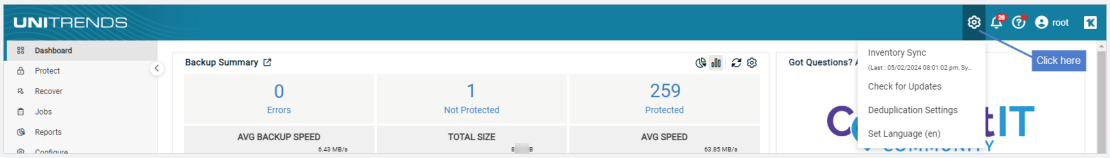

1 Select **Jobs > Active Jobs > Create Job > Backup**.

2 Select **Nutanix AHV** in the **What do you want to backup?** list.

3 In the Inventory tree, expand the AHV host cluster, then check boxes to select virtual machines to protect. Selected VMs display in the Job Inventory Settings area.

- To locate an asset by name, use the **Search** field below.
- To view individual VMs, expand the AHV host.
- To quickly select all hosted VMs, click the host checkbox.
- To select one VM, click its checkbox.

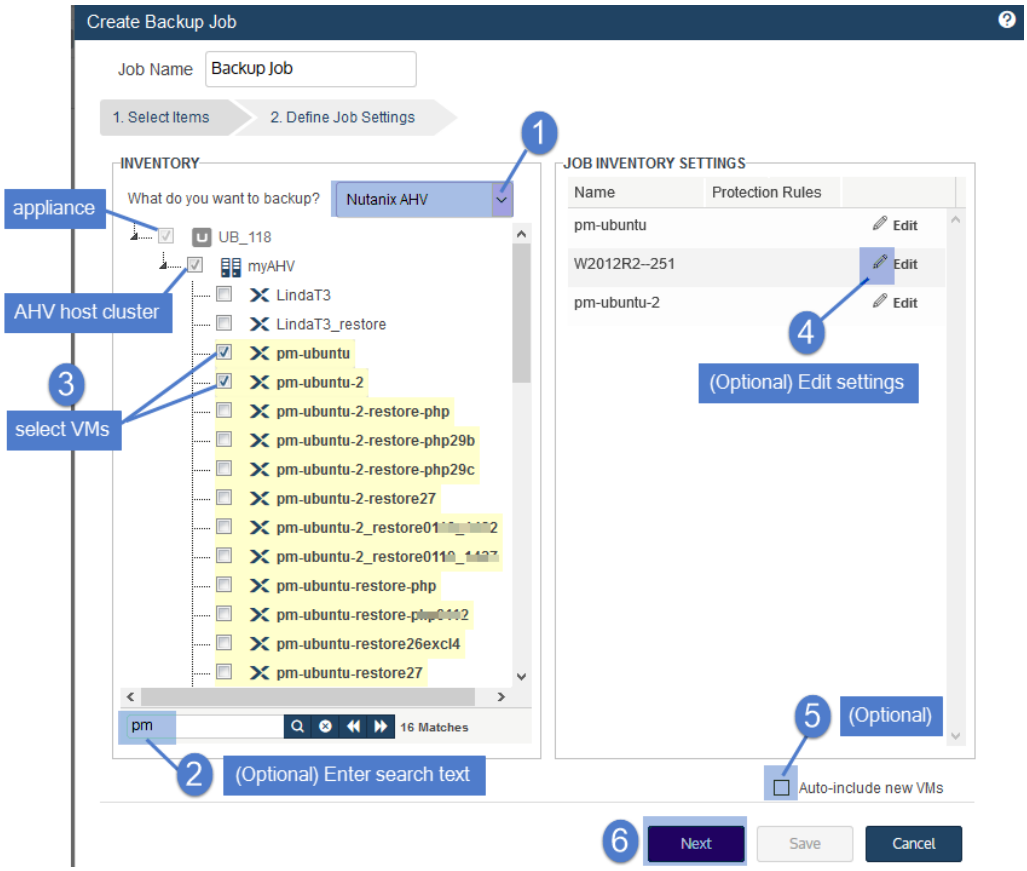

4 (Optional) Edit Job Inventory Settings to exclude VM disks from backup:

- Locate the VM in the Job Inventory Settings list.
- Click **Edit** to specify disks to exclude.
- Click **Save** to retain any changes.

> **Note:** Critical system volumes are required to recover the entire virtual machine. Use care when omitting disks from backup.

5 (Optional) Check the **Auto-include new VMs** box to automatically add newly discovered VMs to the schedule.

6 Click **Next**.

7 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

8 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

9 Click **Save**.

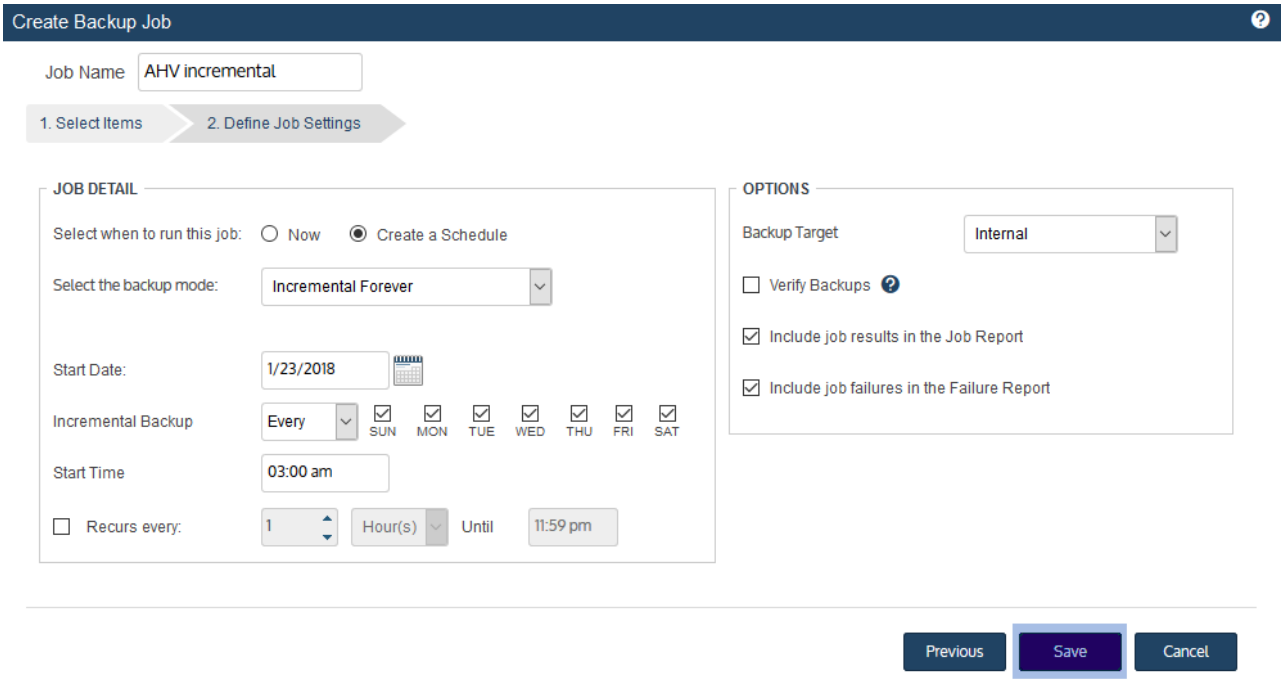

10 Click **OK** to close the Success message.

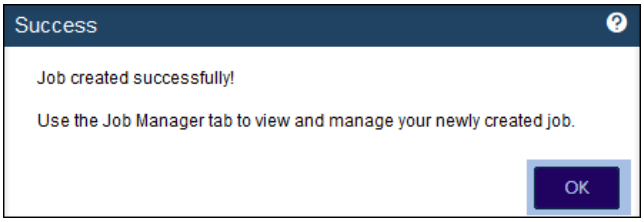

- If you created a schedule, the job runs at the date and times you specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

To create a XenServer backup job

Notes:

- A XenServer asset can be assigned to one backup schedule only.
- To access newly added virtual machines, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

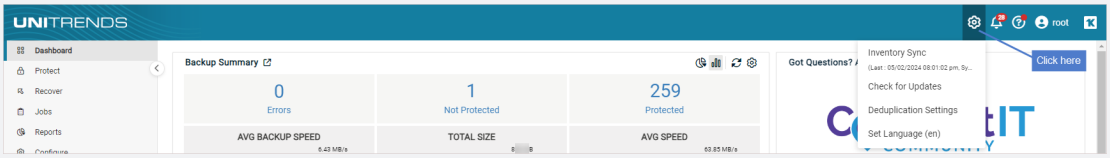

1 Select **Jobs** > **Active Jobs** > **Create Job** > **Backup**.

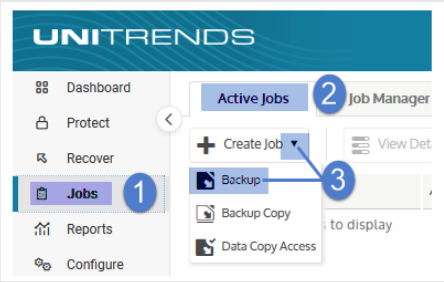

2 Select **XenServer** in the **What do you want to backup?** list.

3 In the Inventory tree, expand the virtual host and check boxes to select virtual machines to protect. Selected VMs display in the Job Inventory Settings area.

To locate the asset by name, use the **Search** field below.

4 (Optional) Edit Job Inventory Settings to exclude VM disks from backup:

- Locate the VM in the Job Inventory Settings list.
- Click **Edit** to specify disks to exclude.
- Click **Save** to retain any changes.

Note: Critical system volumes are required to recover the entire virtual machine. Use care when omitting disks from backup.

5 (Optional) Check the **Auto-include new VMs** box to automatically add newly discovered VMs to the schedule.

6 Click **Next**.

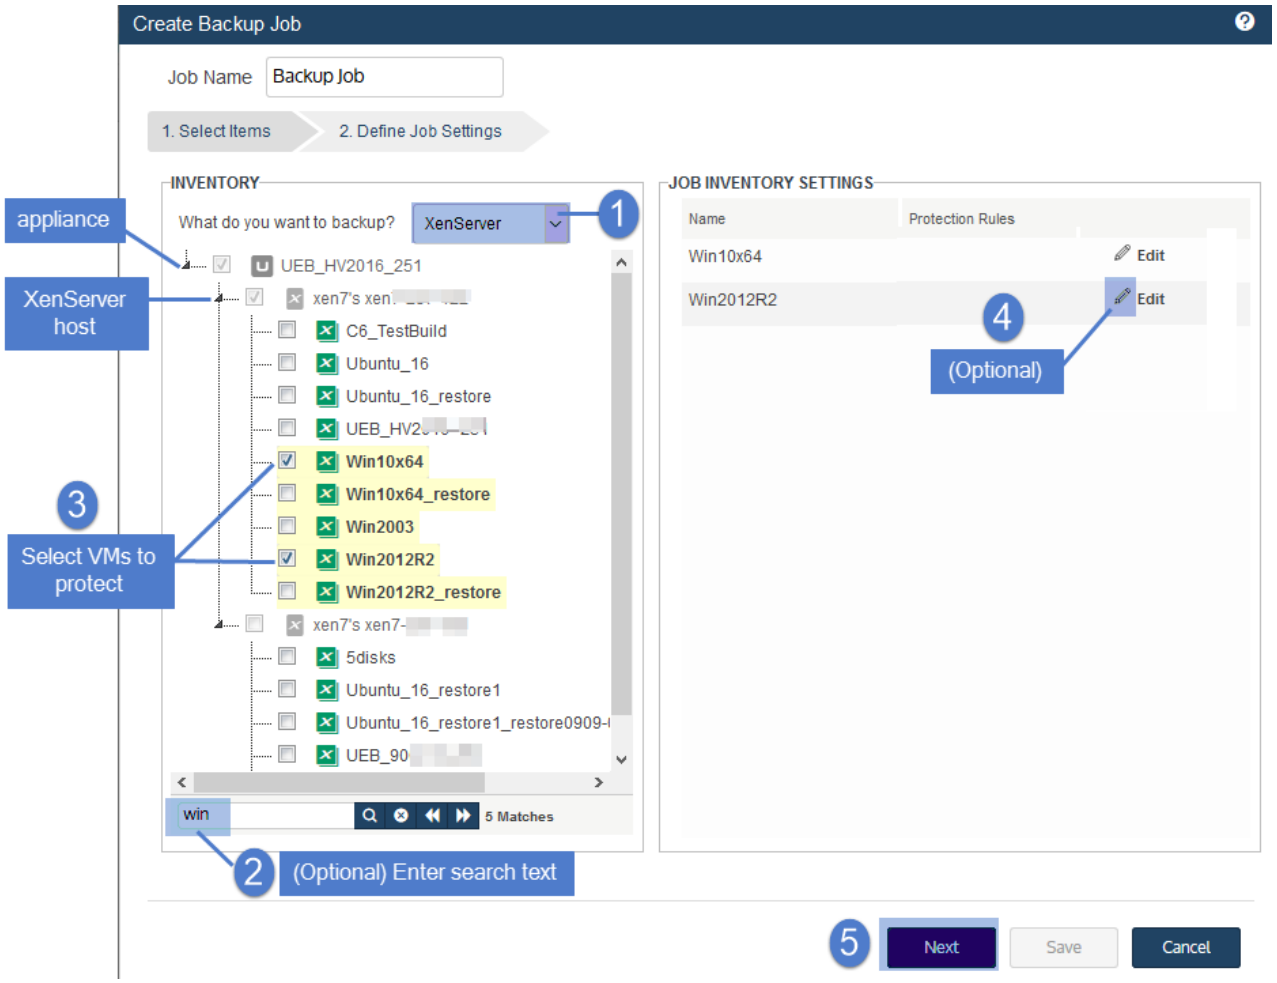

7 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

8 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

9 Click **Save**.

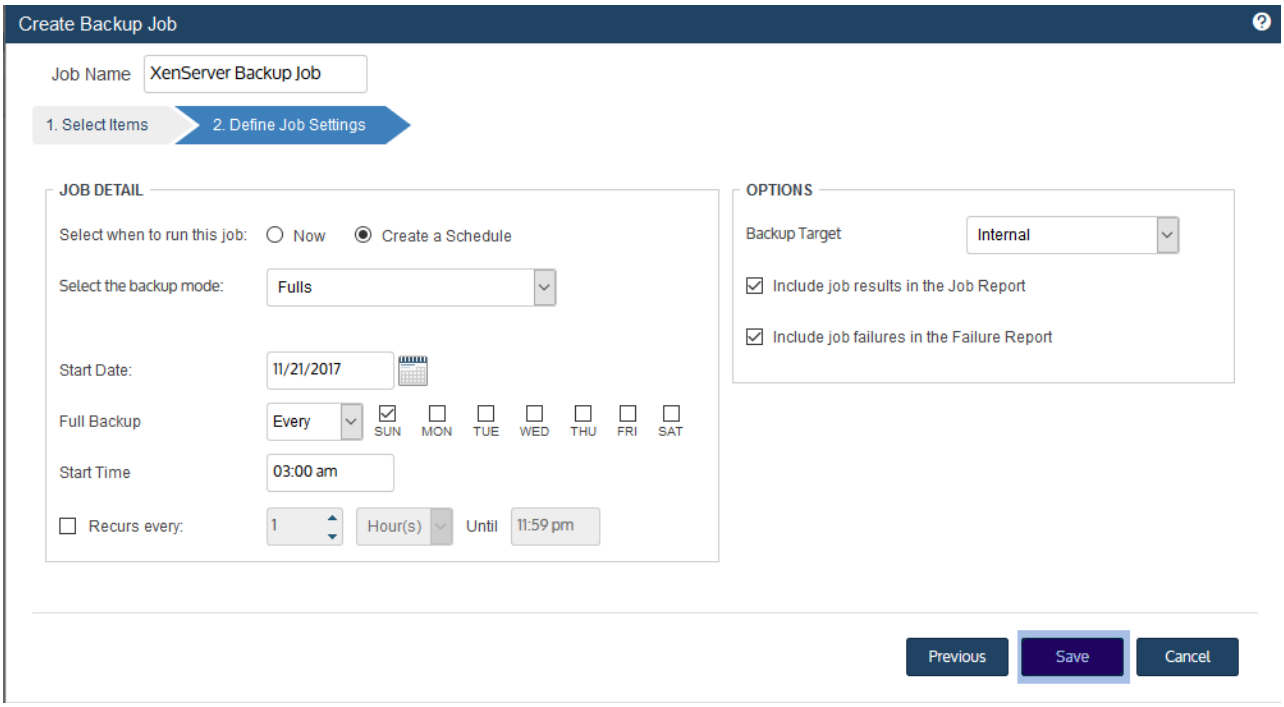

10 Click **OK** to close the Success message.

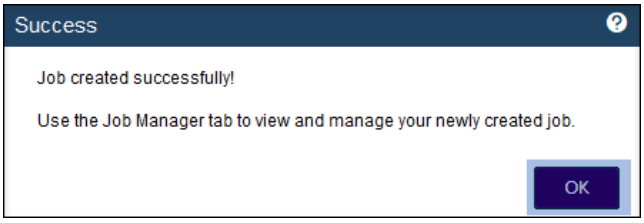

- If you created a schedule, the job runs at the date and times you specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

To create a NAS CIFS or NFS backup job

1 Select **Jobs > Active Jobs > Create Job > Backup**.

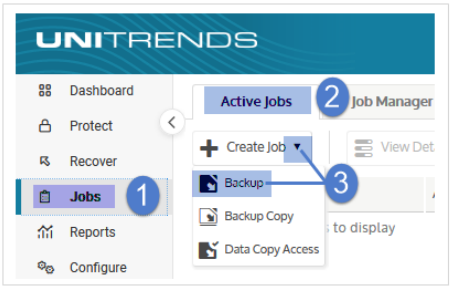

2 Select **NAS** in the **What do you want to backup?** list.

3 In the Inventory tree, check boxes to select the NAS assets to protect. Selected assets display in the Job Inventory Settings area.

To locate the asset by name, use the **Search** field below.

4 (Optional) Edit Job Inventory Settings to specify directories to include or exclude: locate the asset in the list and click **Edit**, modify settings, then click **Save** to retain any changes.

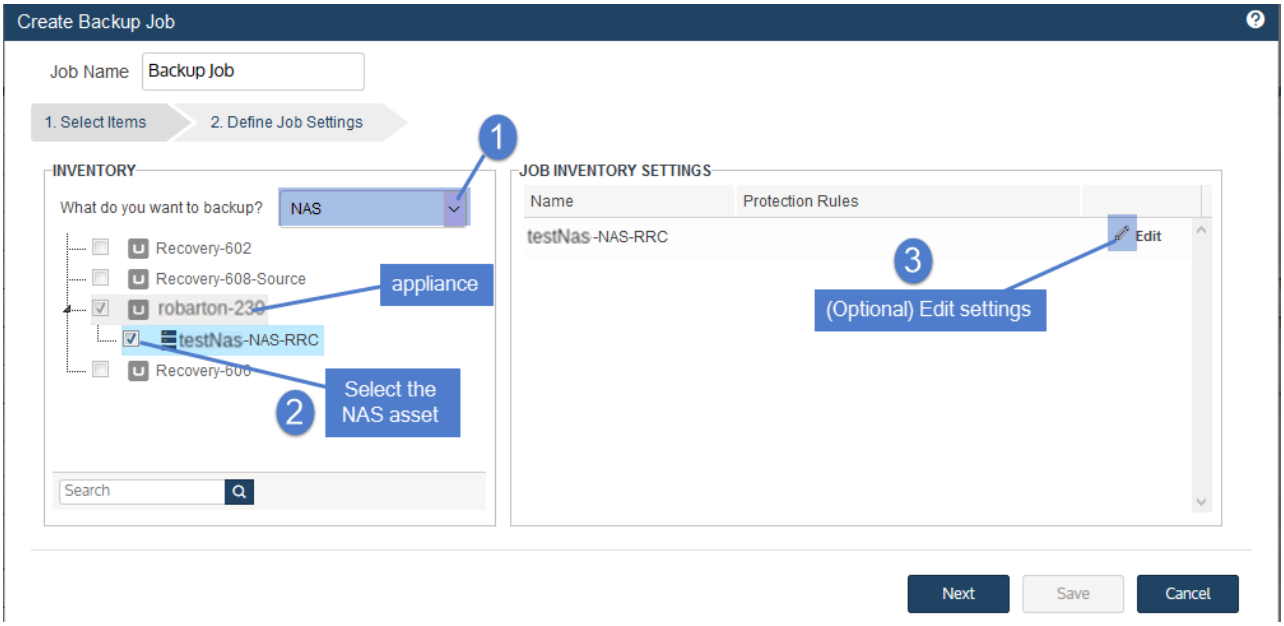

See the following for details and considerations:

File-level setting	Description
General considerations for including or excluding data from a NAS asset's backups	Review the following before specifying data to include or exclude: • When creating a schedule, any inclusions or exclusions you add are applied to jobs run by that schedule only. Inclusions and exclusions are not applied automatically in other cases. • Running an on-demand backup of the asset does not automatically apply any inclusions or exclusions specified in the asset's schedule. To run an on-demand backup, do one of the following: – Create a one-time job that has the same inclusions and exclusions as in the schedule. – Manually run the schedule (select the schedule under **Jobs > Job Manager** and click **Run**).

File-level setting	Description
	– Run a one-time Selective backup (so that a new full is not created). • If you specify both files to include and files to exclude, the inclusion is applied first. Any exclusions are then applied to the subset of included files.
Inclusion tab	Click to specify files to include in backups of this asset. • Data that does not meet the criteria you specify here is NOT included in the backup. • Type in the full path (e.g., **@@@:/mnt/nas/directory**) or **Browse** the asset to specify folders to include. (Wildcards are not supported.) • Run a new full backup upon creating or modifying included files. Example: 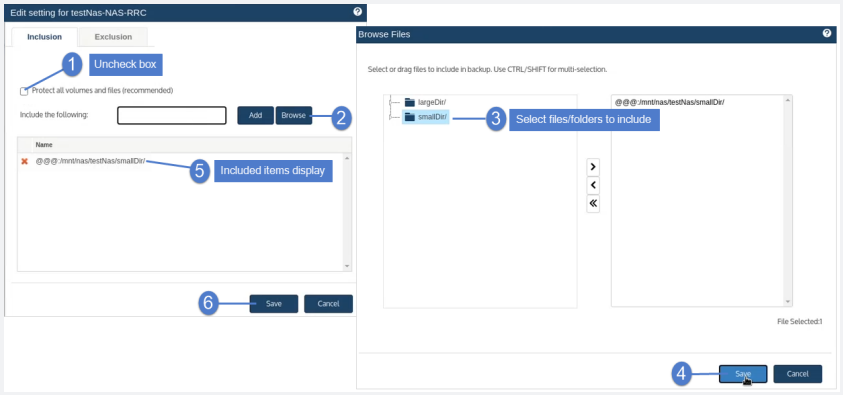
Exclusion tab	Click to specify files or folders to exclude from backups of this asset. • Data that does not meet the criteria you specify here IS included in the backup. • To specify files to exclude, do any of the following: – Type in the full path (e.g., **@@@:/mnt/nas/directory**). – **Browse** the asset. – Enter a selection pattern. Wildcards are supported. See these rows below for usage examples: "Wildcard * usage",

File-level setting	Description
	"Wildcard ? usage", and "Multiple wildcards". • Run a new full backup upon creating or modifying excluded files. Example: 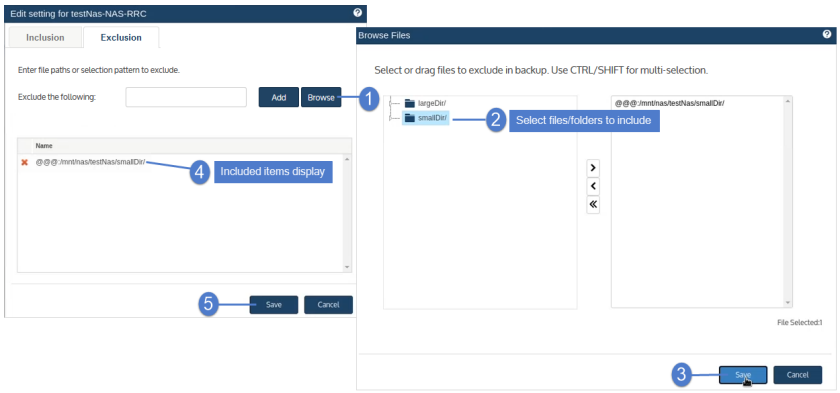
Wildcard * usage	An example of how to exclude all files with zero or more characters that match exclusion pattern: ***.txt** An example of how to exclude directories with zero or more characters and their contents within a specified path that match the exclusion pattern: **@@@:/mnt/nas/account*** Limitations: • ***folder_abc** cannot be used to exclude all folders that match *folder_abc* on the protected asset. The full path must be provided. • If an entire directory is excluded, the directory name will still appear in the backup; however, its contents will be empty. • Multiple wildcard matches like the following are not supported: **@@@:/mnt/nas/*/*/abc.txt**
Wildcard ? usage	An example of how to exclude all files within specified path that matches a single character within exclusion pattern: **@@@:/mnt/nas/Lists.dir/pro_client?.spr** An example of how to exclude all directories and their contents within specified path that matches a single character within exclusion pattern: **@@@:/mnt/nas/Lists.dir/Case?/** Limitation: If an entire directory is excluded, the directory name itself

File-level setting	Description
	will still appear in the backup; however its contents will be empty.
Multiple wildcards	An example that uses multiple "?" wildcards and only one * wildcard: **@@@:/mnt/nas/?Log?/*.logs** Limitation: If an entire directory is excluded, the directory name itself will still appear in the backup; however its contents will be empty.

5 Click **Next**.

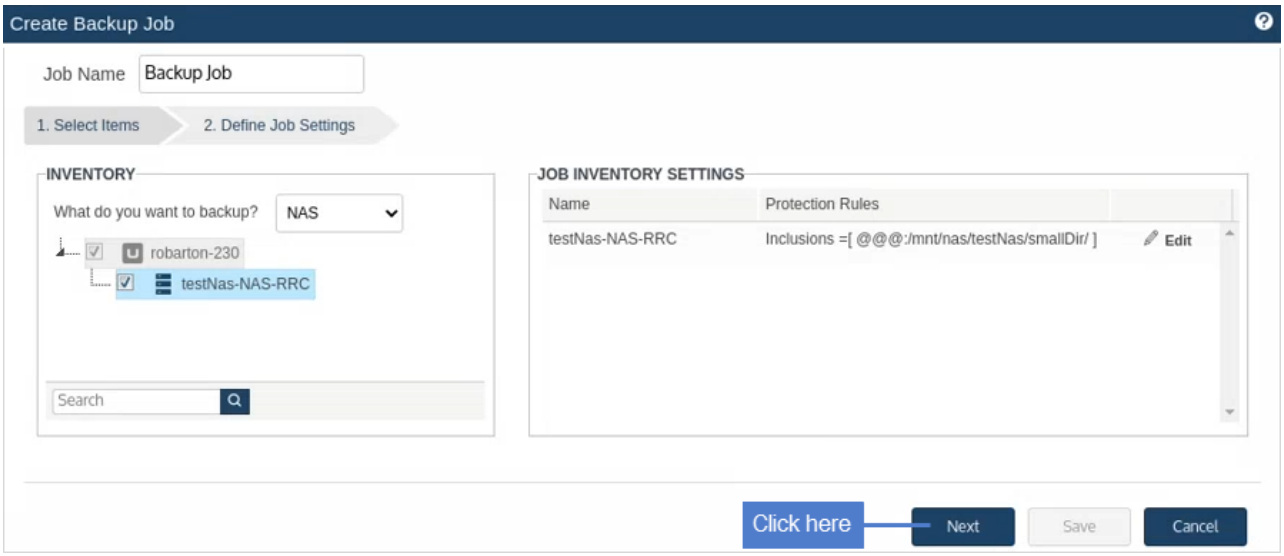

6 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

7 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

8 Click **Save**.

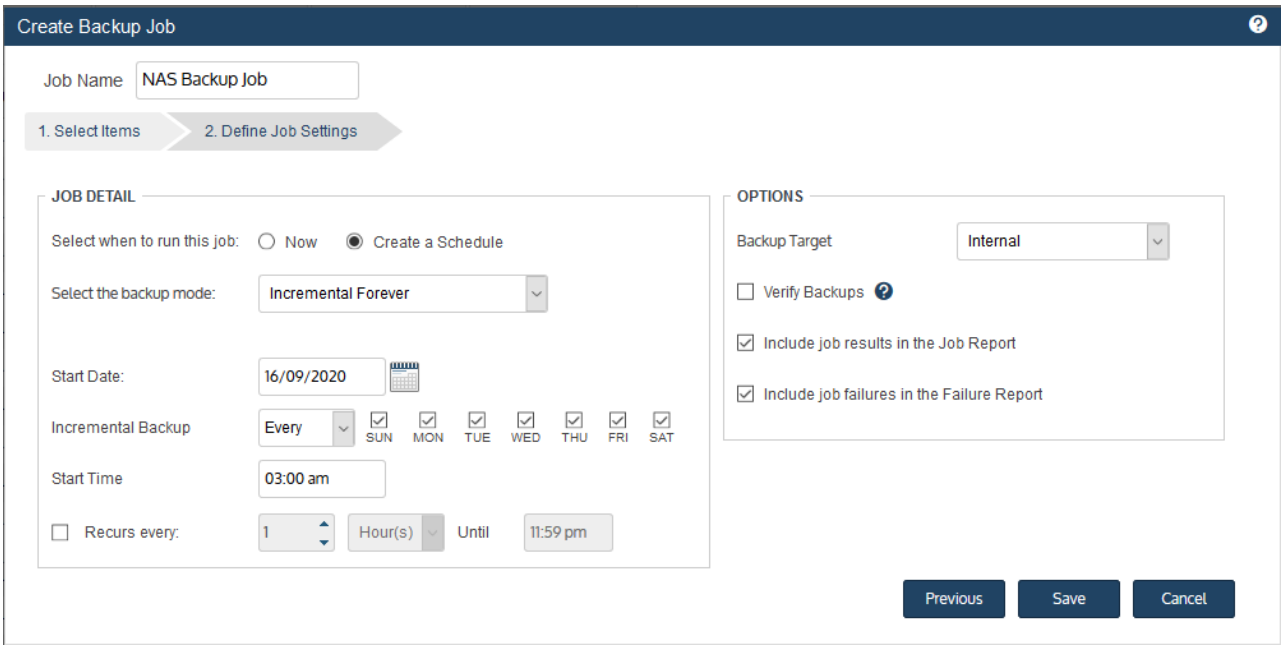

To create a NAS NDMP backup job

1 Select **Jobs > Active Jobs > Create Job > Backup**.

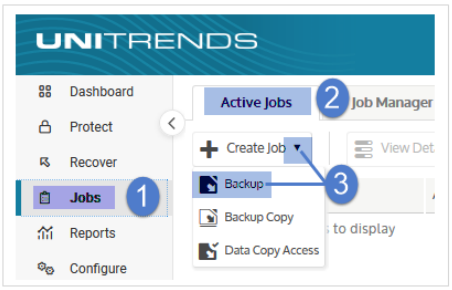

2 Select **NDMP** in the **What do you want to backup?** list.

3 In the Inventory tree, expand the desired NAS asset and check boxes to select volumes to protect. Selected volumes display in the Job Inventory Settings area.

To locate the asset by name, use the **Search** field below.

4 (Optional) Check the **Auto-include new volumes** box to add any newly discovered volumes to the schedule.

5 Click **Next**.

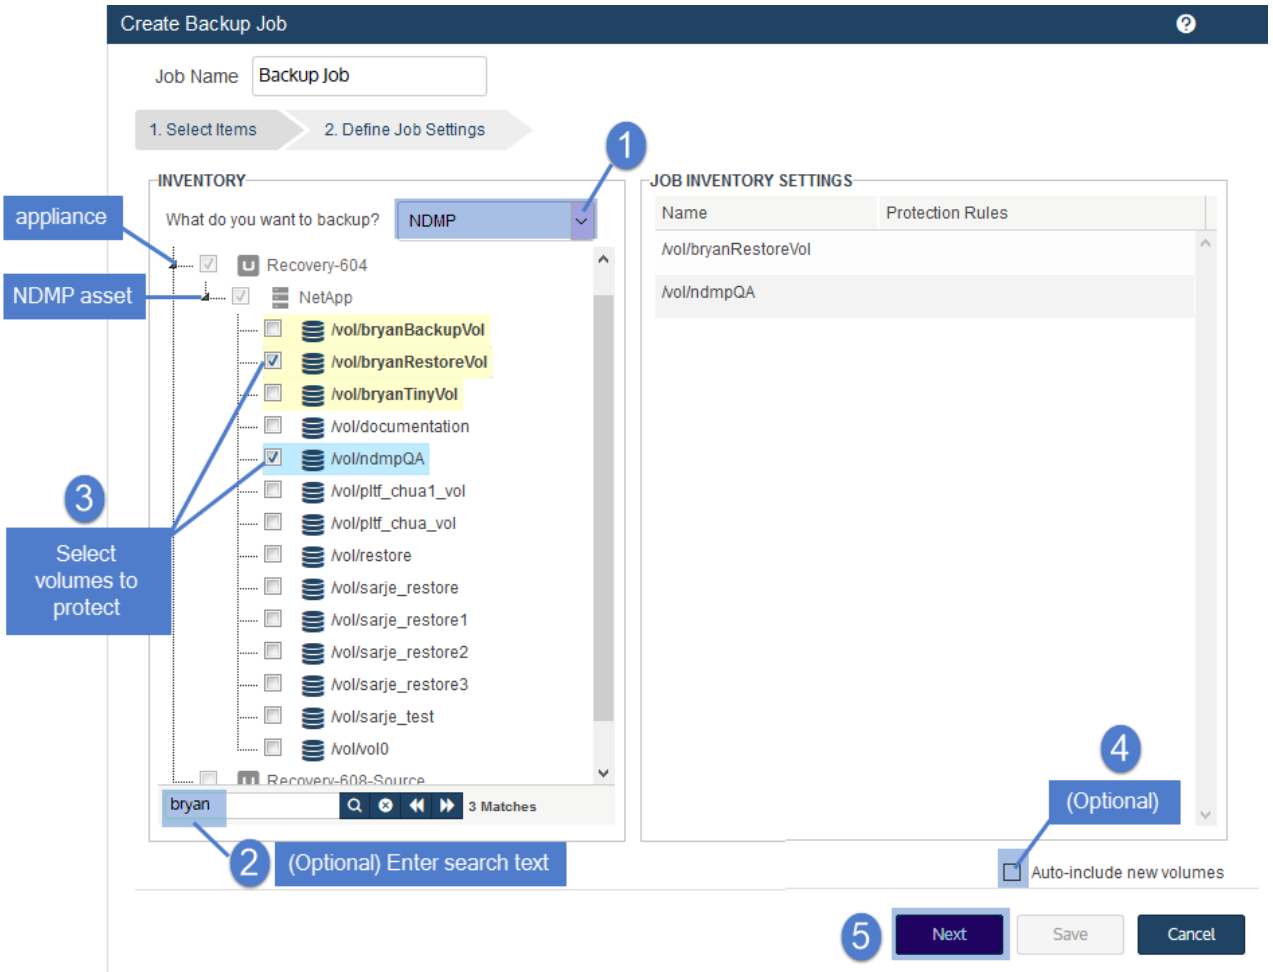

6 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you are creating a schedule, enter a unique job name.

7 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

8 Click **Save**.

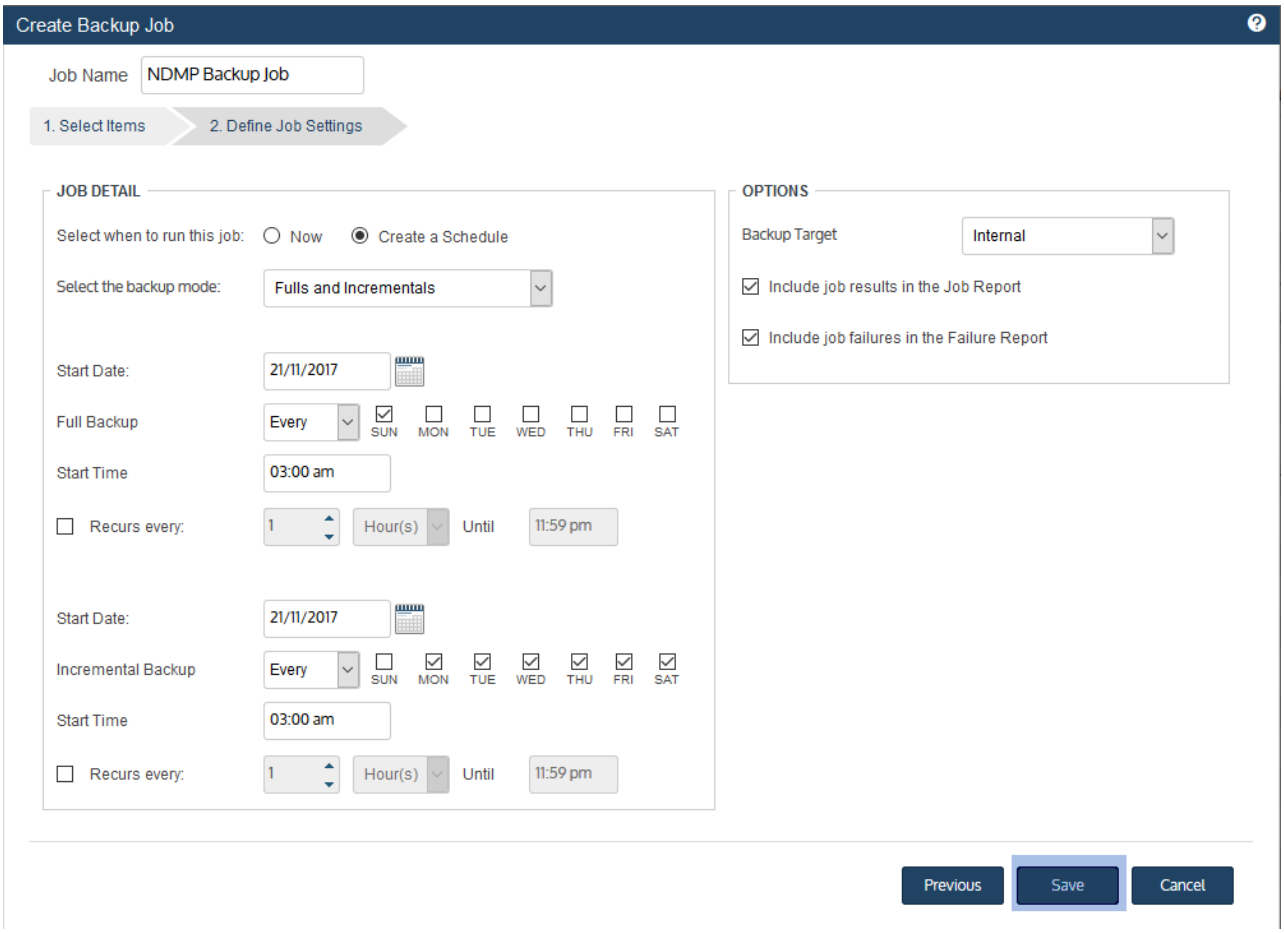

9 Click **OK** to close the Success message.

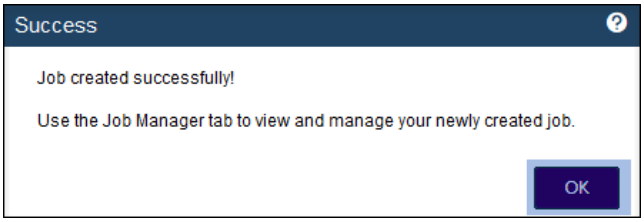

- Each selected volume is backed up in a separate job.
- If you created a schedule, jobs are queued at the times you specified and run as NDMP connections become available.
- If you chose Now, jobs are queued immediately and run as NDMP connections become available. Click **Active Jobs** to view running jobs.

To create an Exchange backup job

Notes:

To access newly added databases or storage groups, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

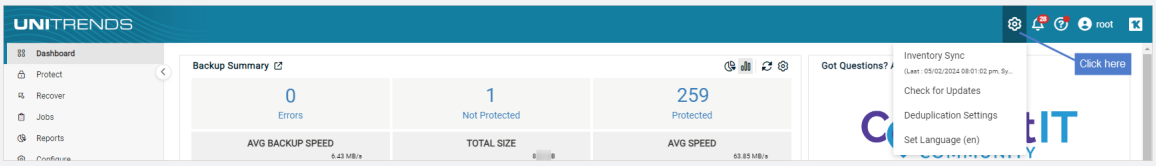

1 Select **Jobs > Active Jobs > Create Job > Backup**.

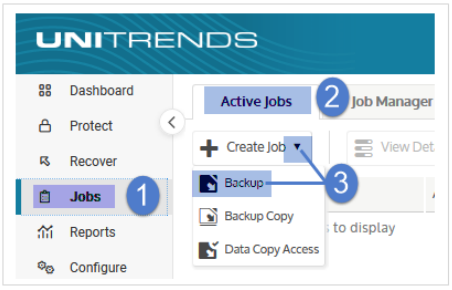

2 Select **Exchange** in the **What do you want to backup?** list.

3 In the Inventory tree, expand the Exchange server and check boxes to select databases or storage groups to protect.

To locate the asset by name, use the **Search** field below.

4 (Optional) Check the **Auto-include new databases** box to automatically add newly discovered databases to the schedule.

5 Click **Next**.

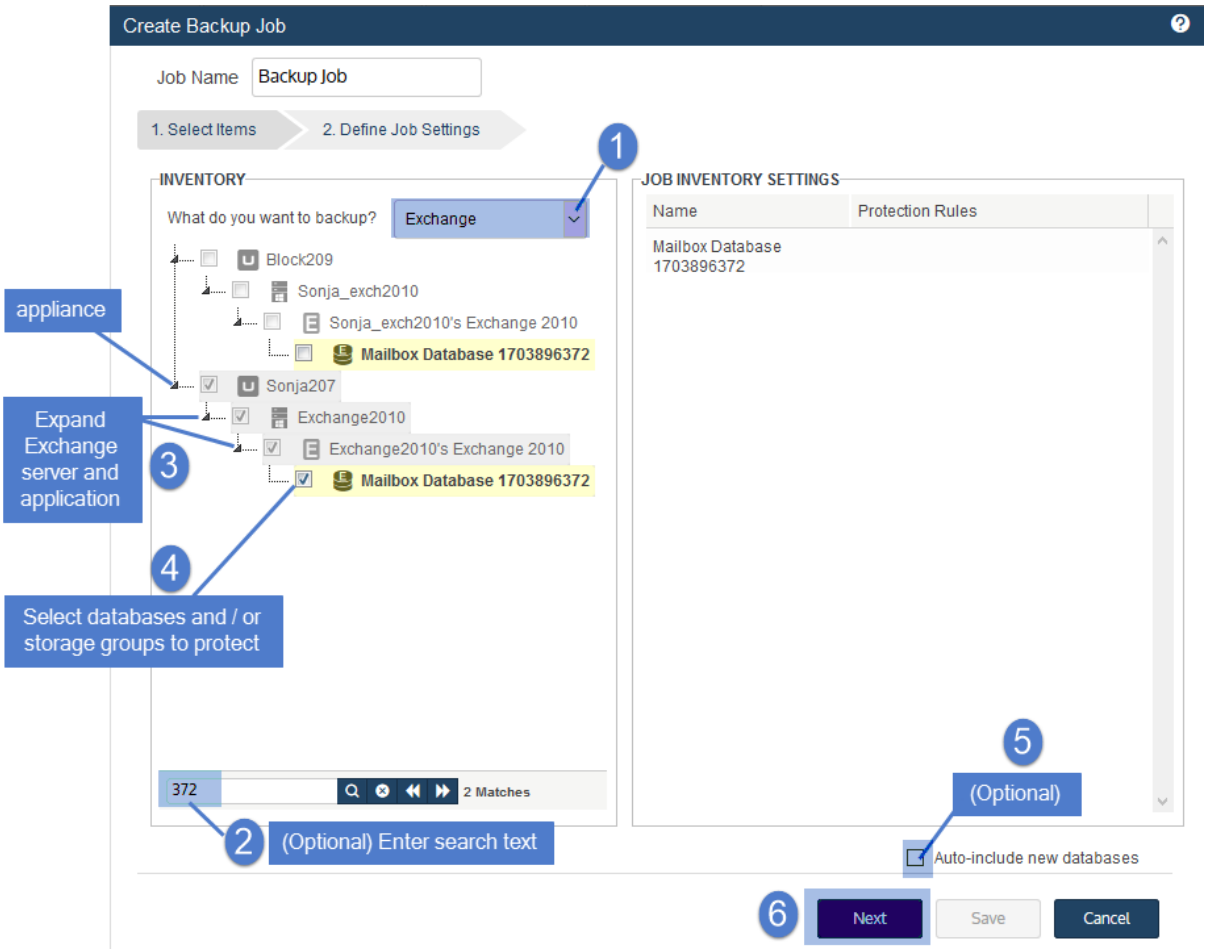

6 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

7 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

8 Click **Save**.

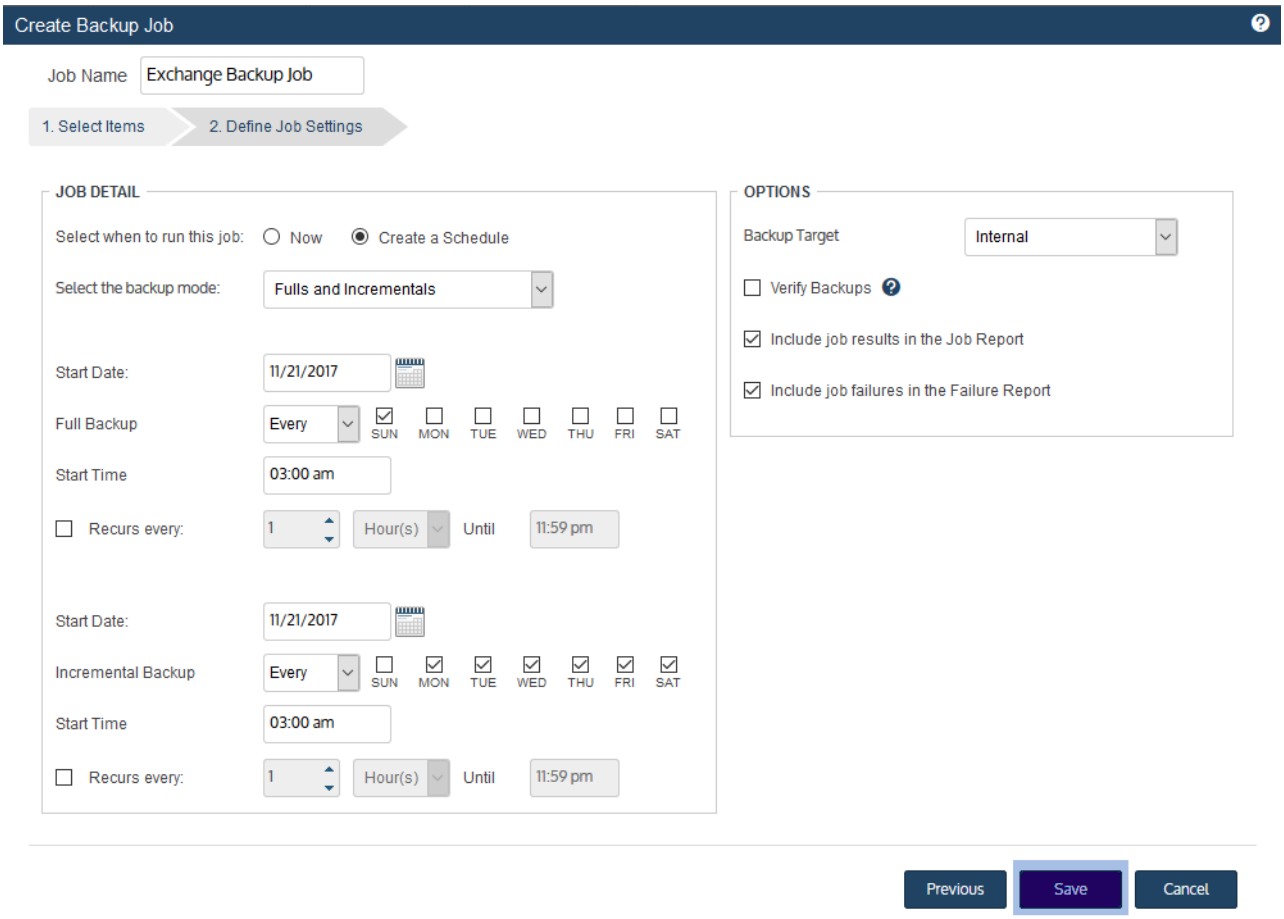

9 Click **OK** to close the Success message.

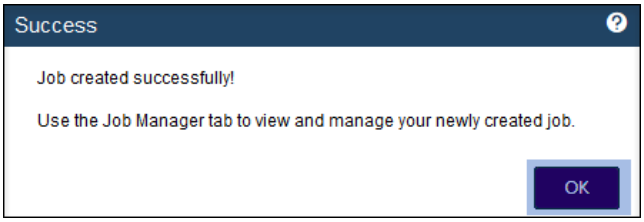

- If you created a schedule, the job runs at the date and times you specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

To create an Oracle backup job

Notes:

To access newly added databases, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

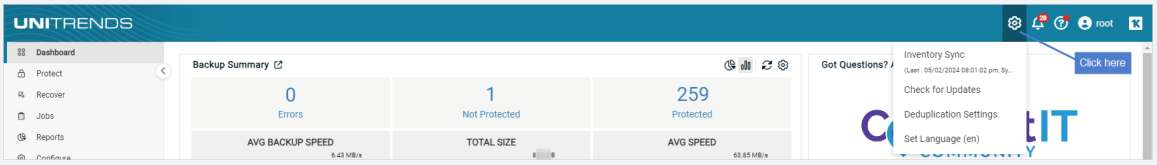

1 Select **Jobs > Active Jobs > Create Job > Backup**.

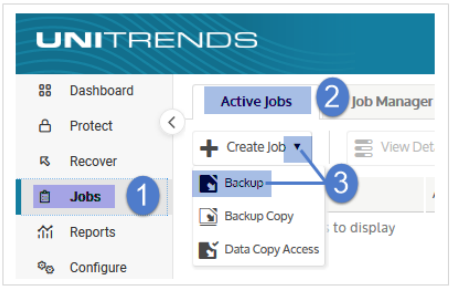

2 Select **Oracle** in the **What do you want to backup?** list.

3 In the Inventory tree, expand the Oracle server and check boxes to select databases to protect.

To locate the asset by name, use the **Search** field below.

> **Note:** If a Samba client is not installed, no databases show as available for backup. The Oracle Dependency from the latest agent release must be installed to protect Oracle data (download from https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads).

4 (Optional) Check the **Auto-include new databases** box to automatically add newly discovered databases to the schedule.

5 Click **Next**.

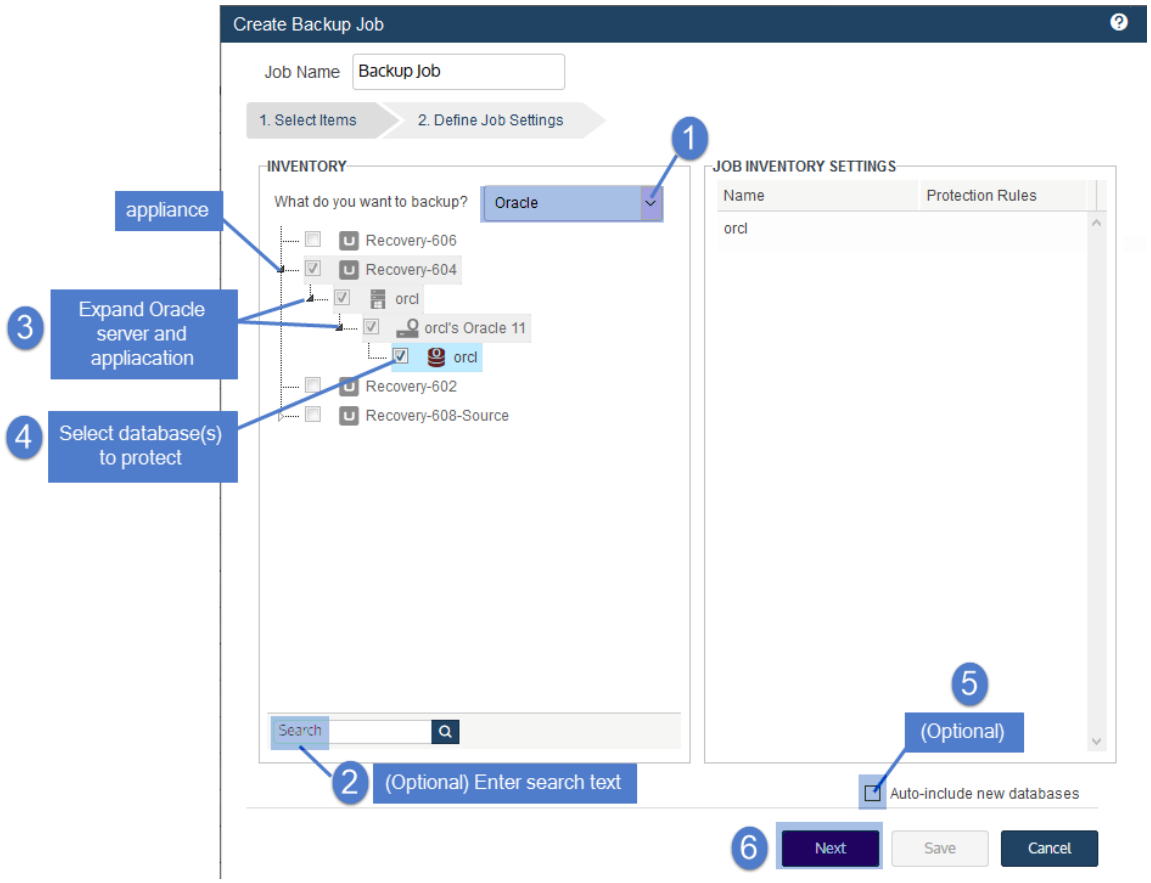

6 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

7 Set remaining Job Details and Options.

> Note: **For Oracle on Linux.** If you are running an incremental forever schedule, you must also exclude Oracle database directories from journal tracking. See Oracle Database: Incremental Forever Schedules on Linux Platforms for details.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

8 Click **Save**.

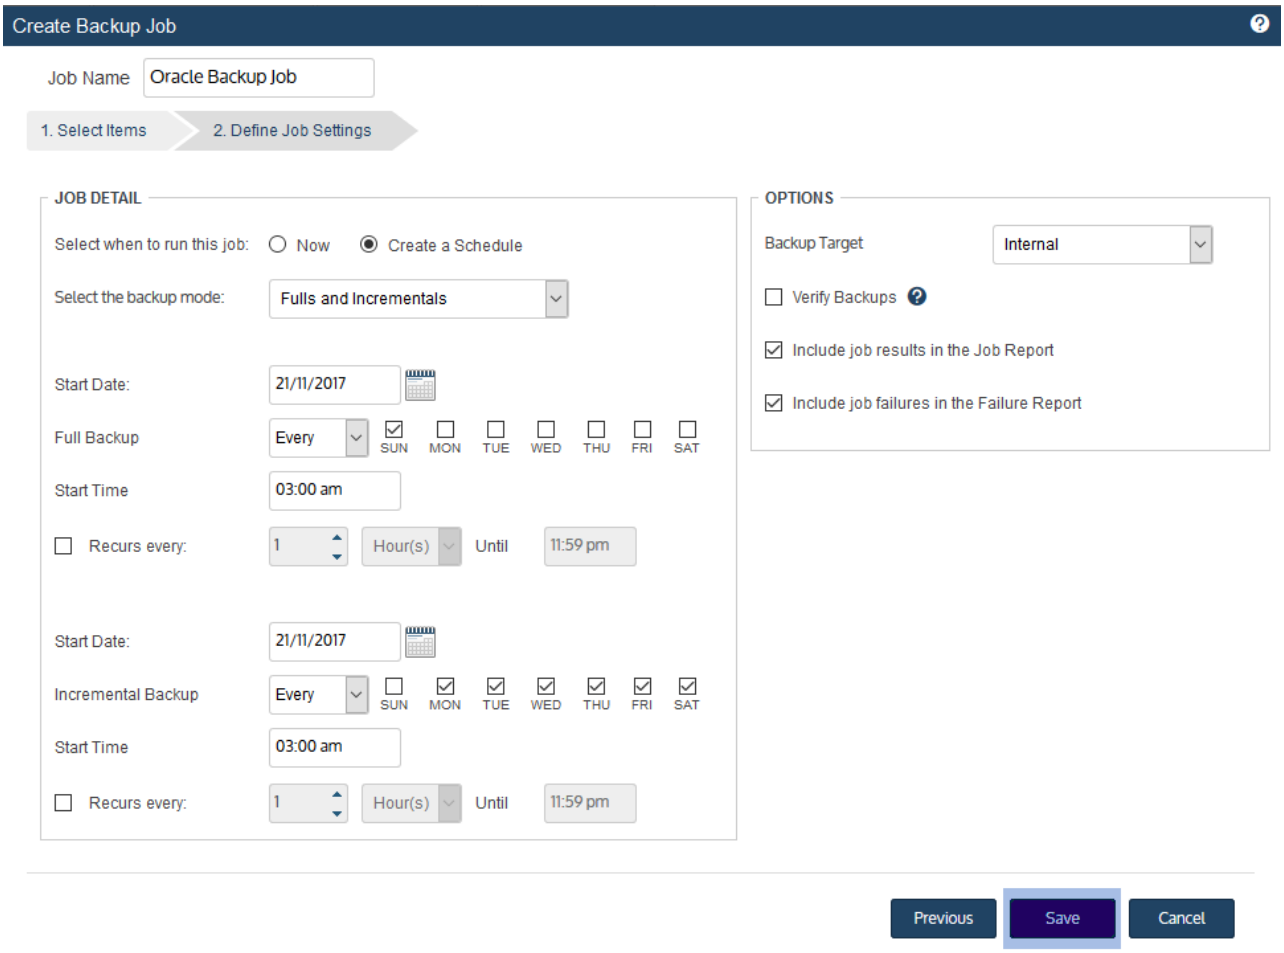

9 Click **OK** to close the Success message.

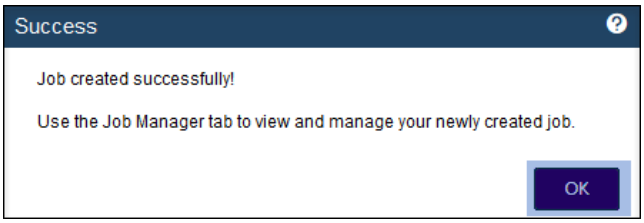

- If you created a schedule, the job runs at the date and times you specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

To create a SQL backup job

This procedure assumes that you have added to the appliance each SQL asset you wish to protect.

Note: Because clustered instances and availability groups can move between server nodes, you must add them as separate assets. If you have a clustered SQL environment and have not added these assets, see "Protecting SQL clusters and availability groups" on page 748 for instructions.

In the Create Backup Job dialog, each SQL instance or availability group displays under its host asset:

- For SQL server assets, hosted non-clustered instances display.
- For SQL cluster assets, the clustered database instance displays.
- For availability group assets, the availability group displays.
- You expand the asset to view and select hosted instances, availability groups, and/or databases to protect.
- To check the SQL asset type, hover over the asset. *SQL Cluster* displays for clustered instance assets, *SQL Availability Group* displays for availability group assets.

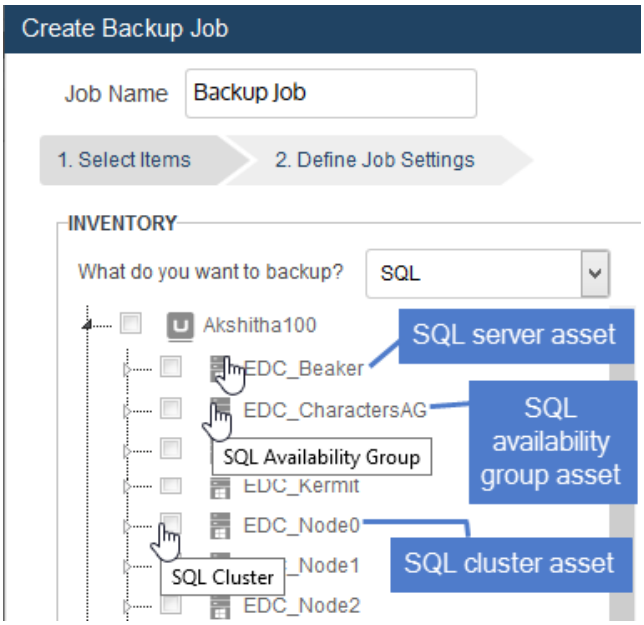

Steps for creating a SQL backup job

Notes:

To access newly added databases, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

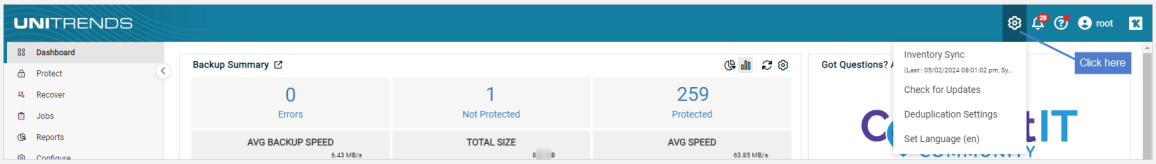

1. Select **Jobs > Active Jobs > Create Job > Backup**.

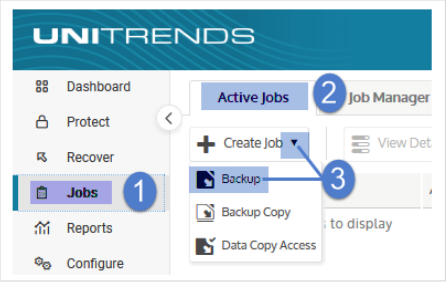

2 Select **SQL** in the **What do you want to backup?** list.

3 In the Inventory tree, expand the SQL host and check boxes to select databases to protect. Selected databases display in the Job Inventory Settings area.

- To locate an asset by name, use the **Search** field below.
- To view individual databases, expand the SQL host and any instances and availability groups.
- To quickly select multiple databases, click an instance or availability group checkbox.
- To select one database, click its checkbox.

4 (Optional) Repeat step 3 to include databases on other host assets.

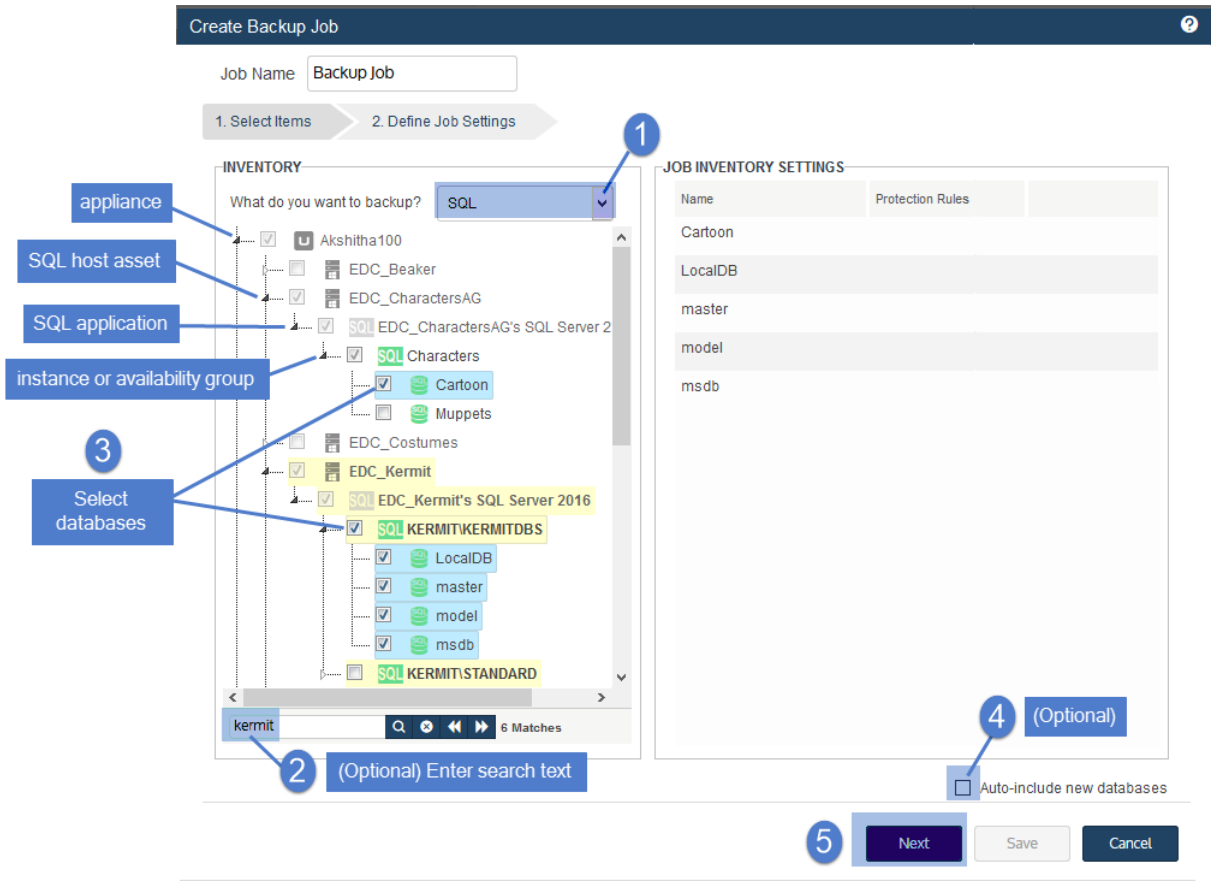

5 (Optional) Check the **Auto-include new databases** box to automatically add newly discovered databases to the schedule.

6 Click **Next**.

7 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

8 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

9 Click **Save**.

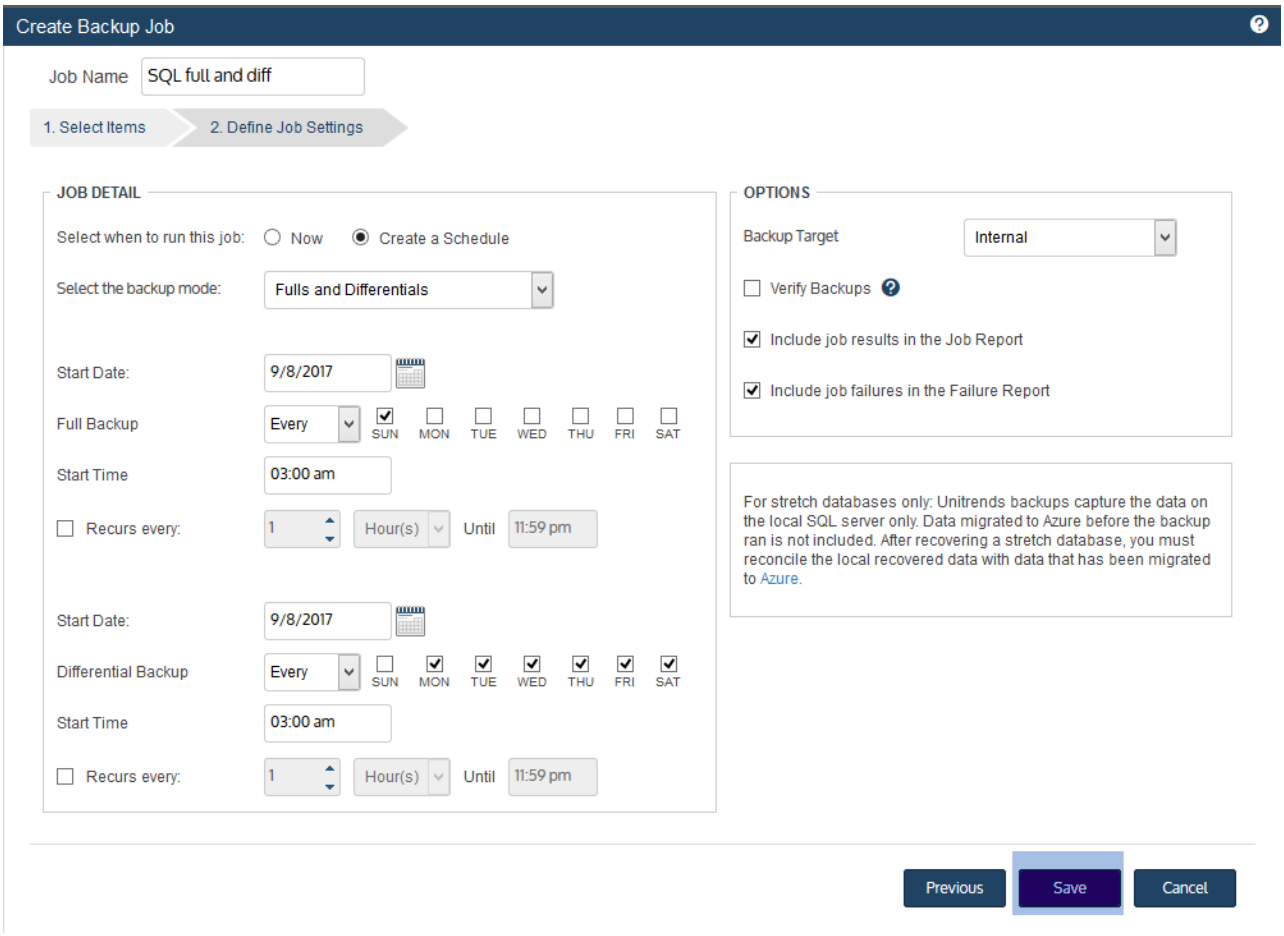

10 Click **OK** to close the Success message.

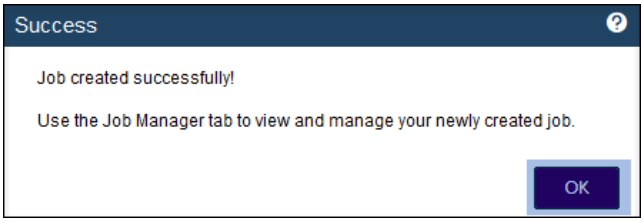

- If you created a schedule, the job runs at the date and times you specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

To create a SharePoint backup job

Notes:

To access a newly installed or newly started SharePoint farm, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

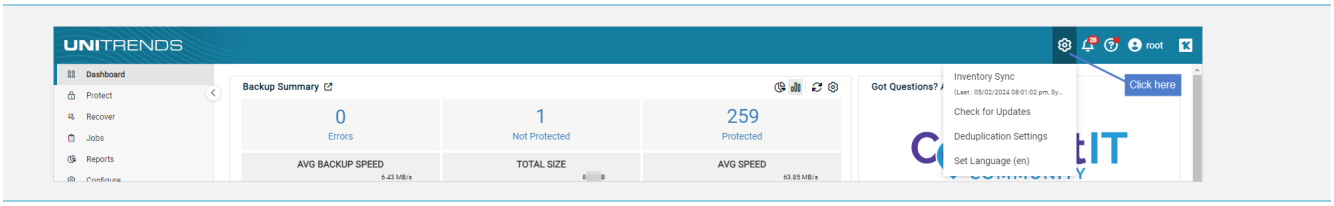

1 Select **Jobs** > **Active Jobs** > **Create Job** > **Backup**.

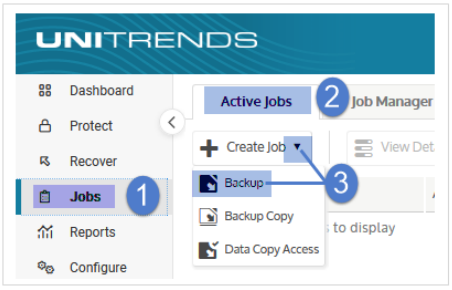

2 Select **SharePoint** in the **What do you want to backup?** list.

3 In the Inventory tree, expand the SharePoint server and check the box to select the farm to protect.

To locate the asset by name, use the **Search** field below.

4 Click **Next**.

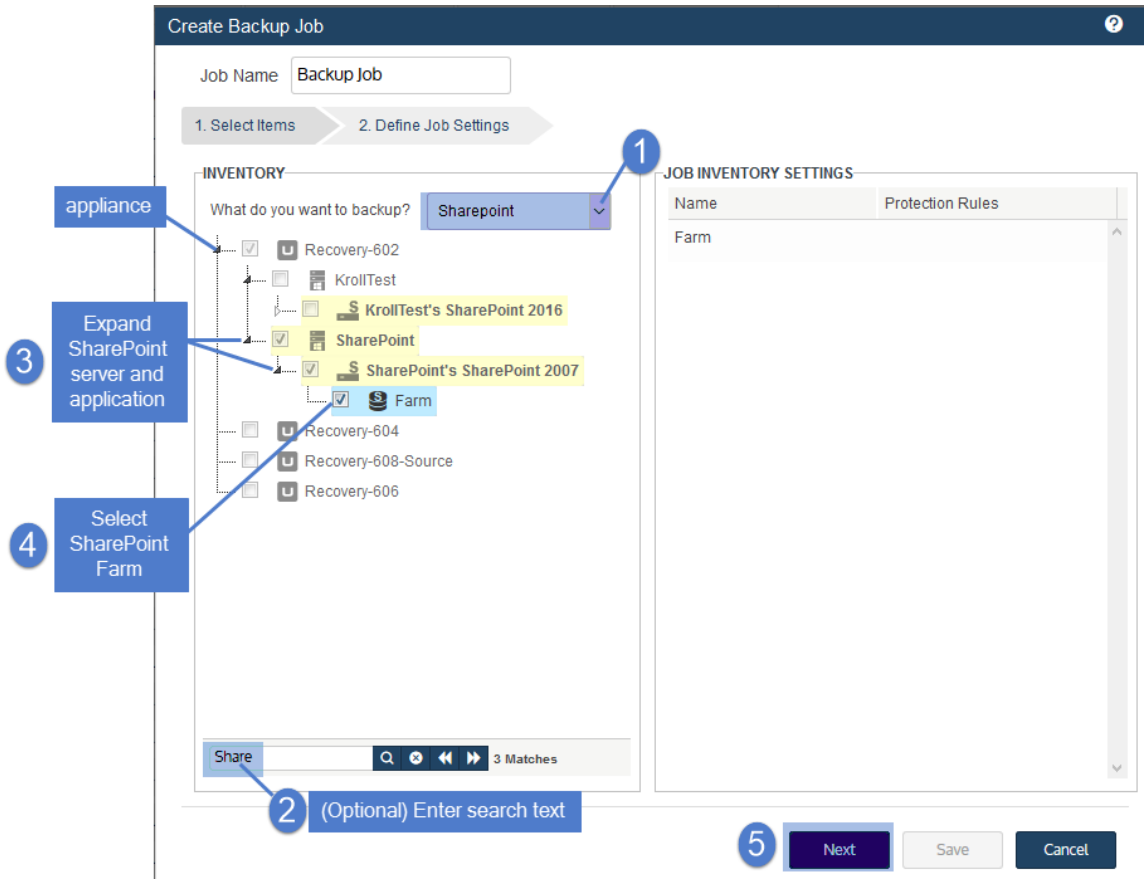

5 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

6 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

7 Click **Save**.

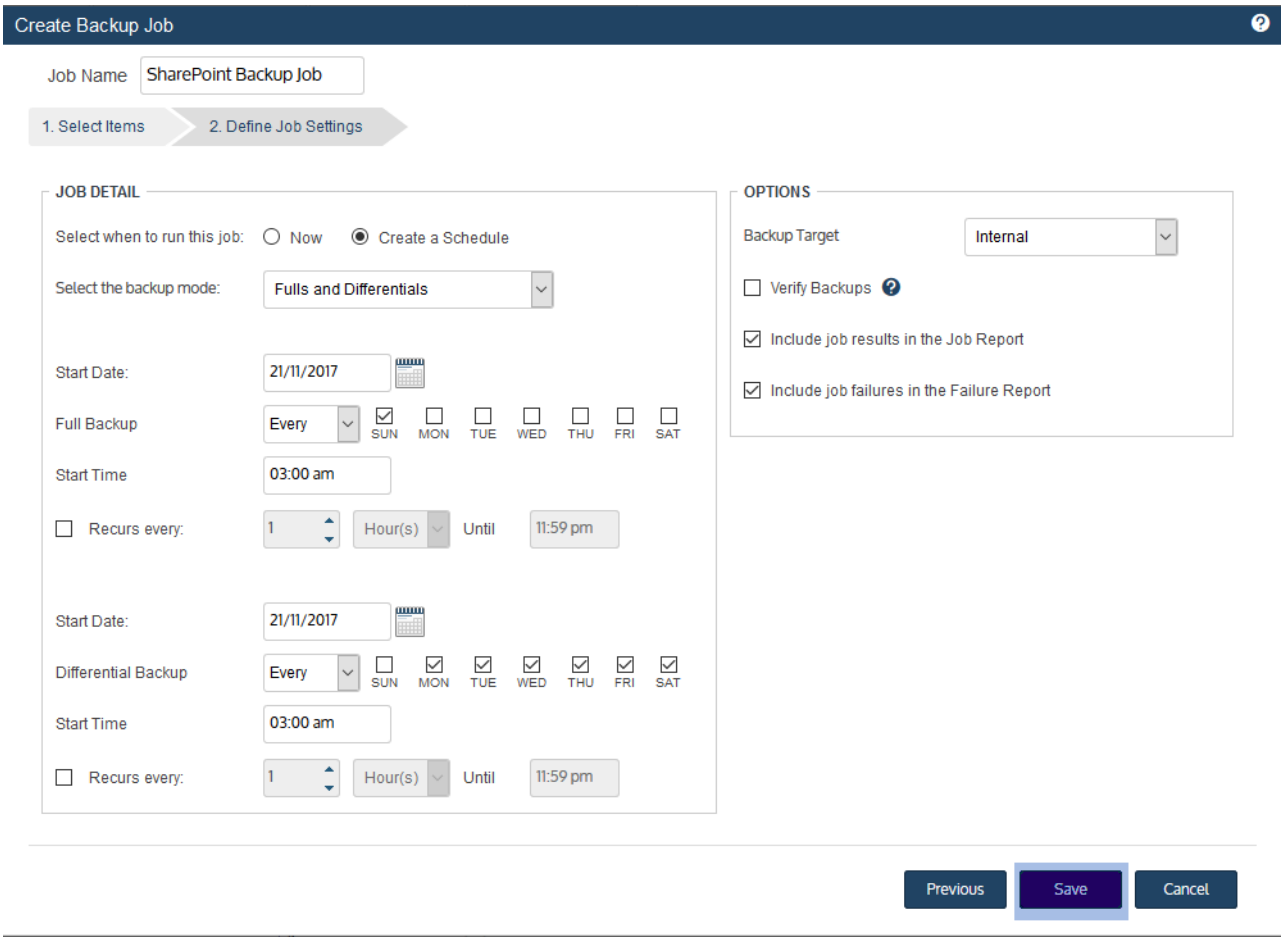

8 Click **OK** to close the Success message.

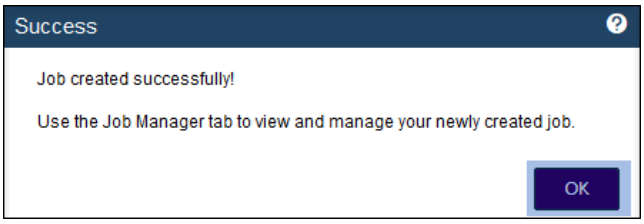

- If you created a schedule, the job runs at the date and times you specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

To create a UCS service profile backup job

1 Select **Jobs > Active Jobs > Create Job > Backup**.

2 Select **Cisco UCS** in the **What do you want to backup?** list.

3 In the Inventory tree, click to select the UCS asset.

4 Click **Next**.

5 Select **Now** or **Create a Schedule** to specify when you want this job to run. If you create a schedule, enter a unique job name.

6 Set remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, you can choose the **Custom** mode to create a custom backup calendar. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

7 Click **Save**.

8 Click **OK** to close the success message.

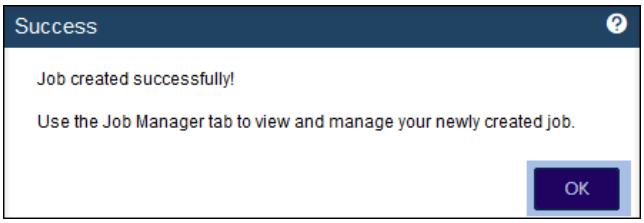

- If you created a schedule, the job runs at the date and times you specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running job.

Creating backup copy jobs

Backup copies are duplicates of your backups and are stored on a secondary target.

See the following topics to create backup copy jobs:

- "Preparing to create a backup copy job" on page 491
- "Selecting assets for backup copy" on page 492
- "Backup copy job procedures" on page 496

Note: For Windows, Linux, VMware, Hyper-V, and AHV assets, you can create SLA policies instead of creating individual jobs. For details on how SLA policies work, see "Methods for scheduling jobs" on page 428.

Preparing to create a backup copy job

Use these procedures to manually create backup copy jobs. Before creating jobs, you must first add the target to your backup appliance (see "Backup copy targets" on page 214). For an overview of the types of targets you can use, see "Backup copies" on page 101. Unitrends also recommends that you review the following information to develop the best protection strategy for your environment: "Preparing for backups" on page 425 and "About creating backup and backup copy jobs" on page 426.

The following considerations apply to backup copies:

- Backups that complete successfully (green) or with warnings (yellow) are eligible for backup copy jobs. Failed (red) backups are not copied.
- Your backup copy storage can contain only one copy of a given backup. If you attempt to create a copy of a backup that has already been written to the target, it is not written, but other backups in the job are copied. The original backup copy remains intact with the original backup copy date.
- Backup copies stored on external media are known as *cold backup copies*. Cold backup copies reside on cloud storage managed by third-party vendors or on other media, such as eSATA, tape, and NAS devices.
- Backup copies stored in the Unitrends Cloud or on a second Unitrends appliance are known as *hot backup copies*.
- After creating a cold backup copy job, backups for the selected assets are copied to the target when the backup copy job runs. You can choose to copy the last backups or backups within a specified date range.
- After creating a hot backup copy job, subsequent backups for the selected assets are copied to the target as backups complete.
- For cold backup copies, incrementals are not copied directly. Instead, incremental backups are synthesized into differential backups for all assets included in a scheduled backup copy job. These differentials are then copied to the cold backup copy media.
- To ensure successful backup copies, make sure the following reserved directory is available on the backup appliance: */mnt/archive/tmp*.

Selecting assets for backup copy

When creating a backup copy job, you first select the assets whose backups you want to copy. Any physical machine, virtual machine, or application is an asset. The Create Backup Copy Job dialog displays assets in an inventory tree where:

- Each physical (agent-based) asset and virtual host displays as a top-level node under its backup appliance.

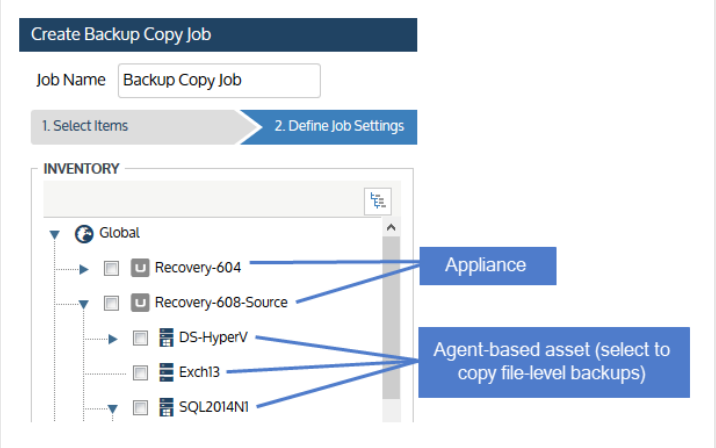

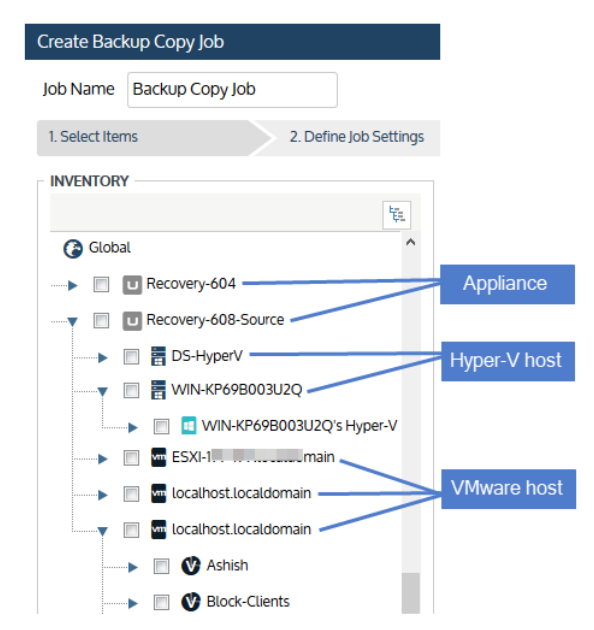

For Windows assets that are eligible for image-level backups, an image-level sub-node displays. Select the primary node to copy file-level backups. Select the image-level sub-node to copy image-level backups.

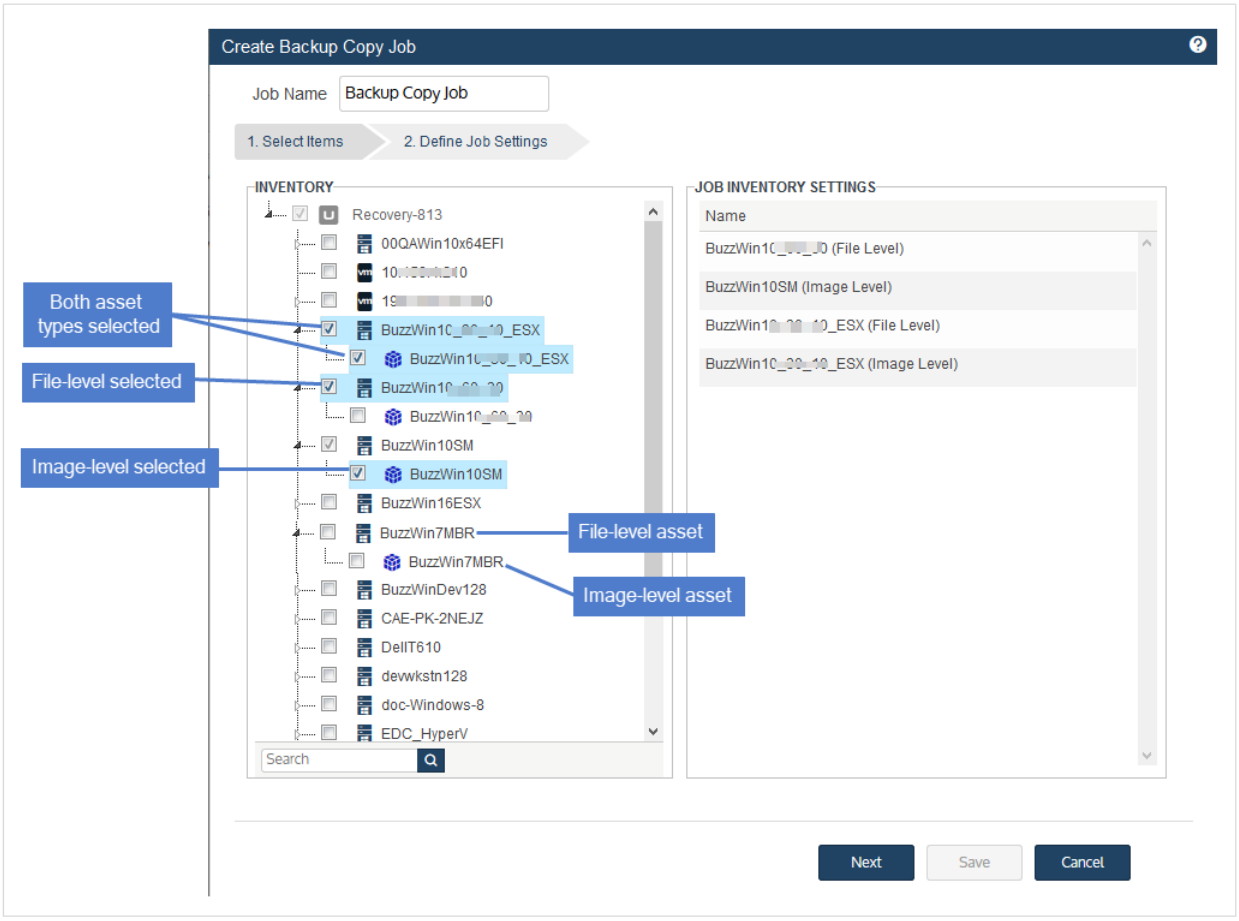

- Each application displays as a sub-node under its physical host asset.

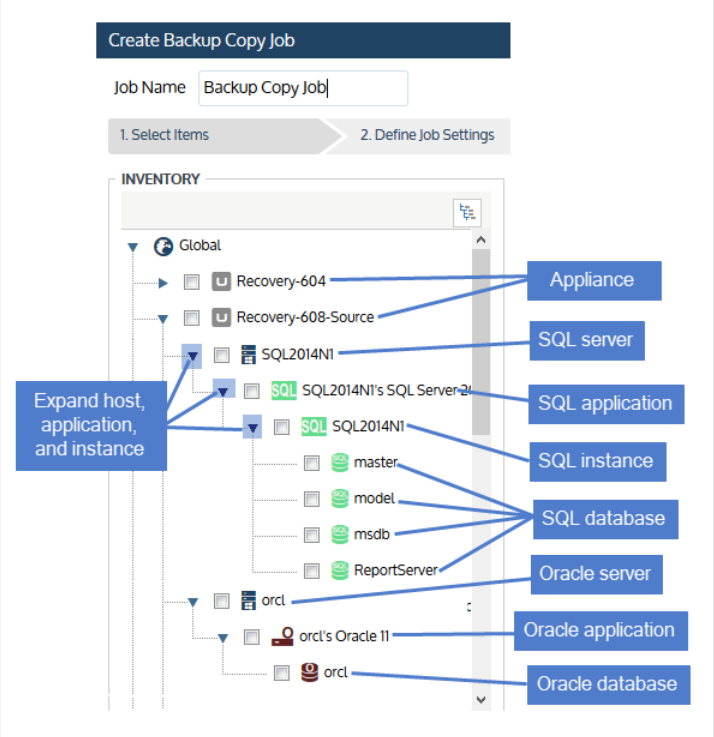

- Each VM displays as a sub-node under its virtual host asset.

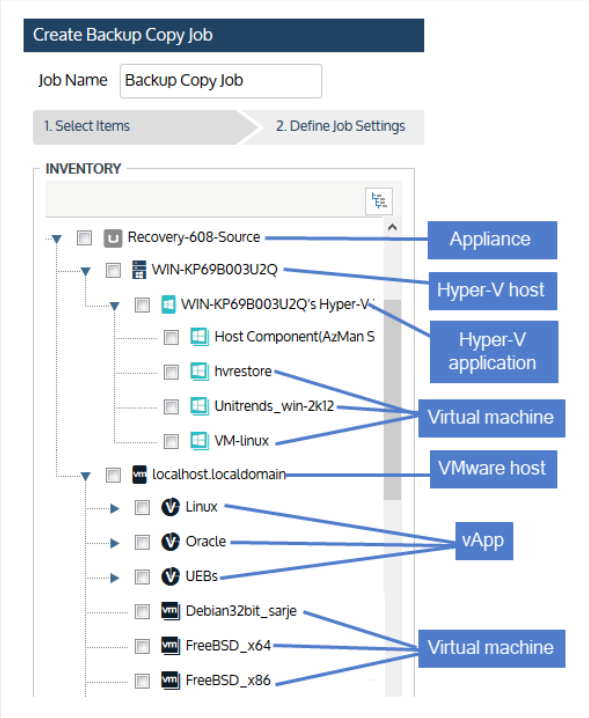

Click to select one or more assets in the inventory tree. The assets you select display in Job Inventory Settings.

Backup copy job procedures

Use these procedures to create backup copy jobs:

- "To create a backup copy job for a Unitrends appliance target" on page 496
- "To copy a full backup to a hot backup copy target on-demand" on page 499
- "To create a backup copy job for a Unitrends Cloud target" on page 501
- "To create a backup copy job for a third-party cloud target" on page 503
- "To create a backup copy job for an attached disk target" on page 508
- "To create a backup copy job for a NAS target" on page 513
- "To create a backup copy job for a SAN target" on page 518
- "To create a backup copy job for an eSATA or USB target" on page 523
- "To create a backup copy job for a tape target" on page 528

To create a backup copy job for a Unitrends appliance target

Use this procedure to copy backups to another Unitrends appliance.

1. Log in to the source backup appliance.
2. Click **Jobs > Active Jobs > Create job > Backup Copy**.

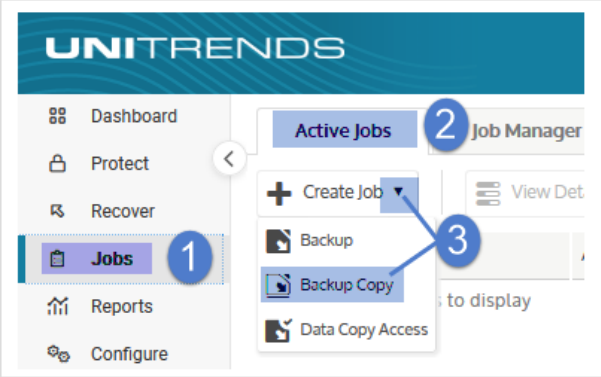

3. Enter a unique **Job Name**.
4. In the Inventory tree, check boxes to select assets whose backups will be copied.
   - To locate an asset by name, use the **Search** field below.
   - Expand the tree as necessary to select VMs and applications.
   - Select a virtual host to select all of its VM assets.
   - Select an application instance to select all of its databases or storage groups.
5. Click **Next**.

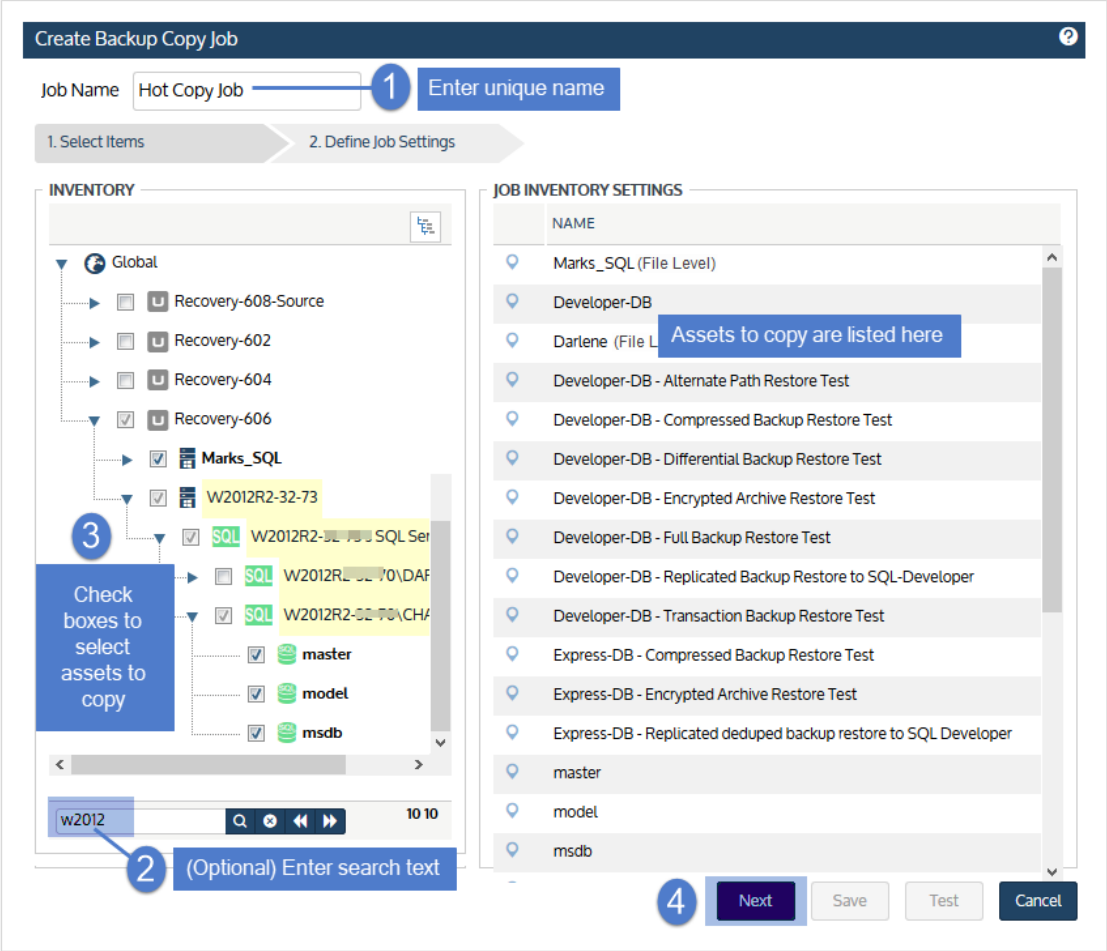

6 In the Backup Copy Type list, select **Backup Copy (Hot)**.

7 In the Backup Copy Target list, select one or more target appliances. (Press **Ctrl** or **Shift** to select more than one appliance.)

Note: A given asset can be included in only one job per backup copy target. If you attempt to add an asset to a second job for the same backup copy target, this message displays:

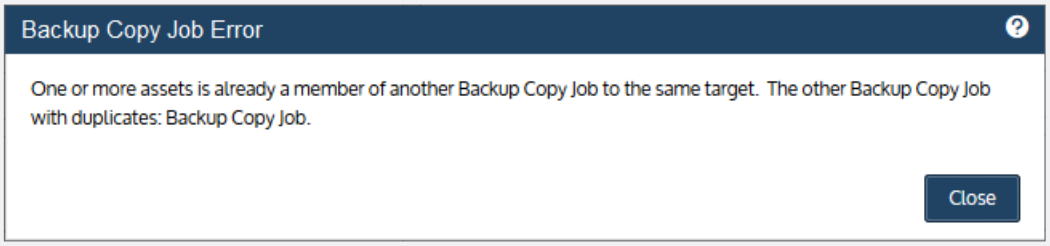

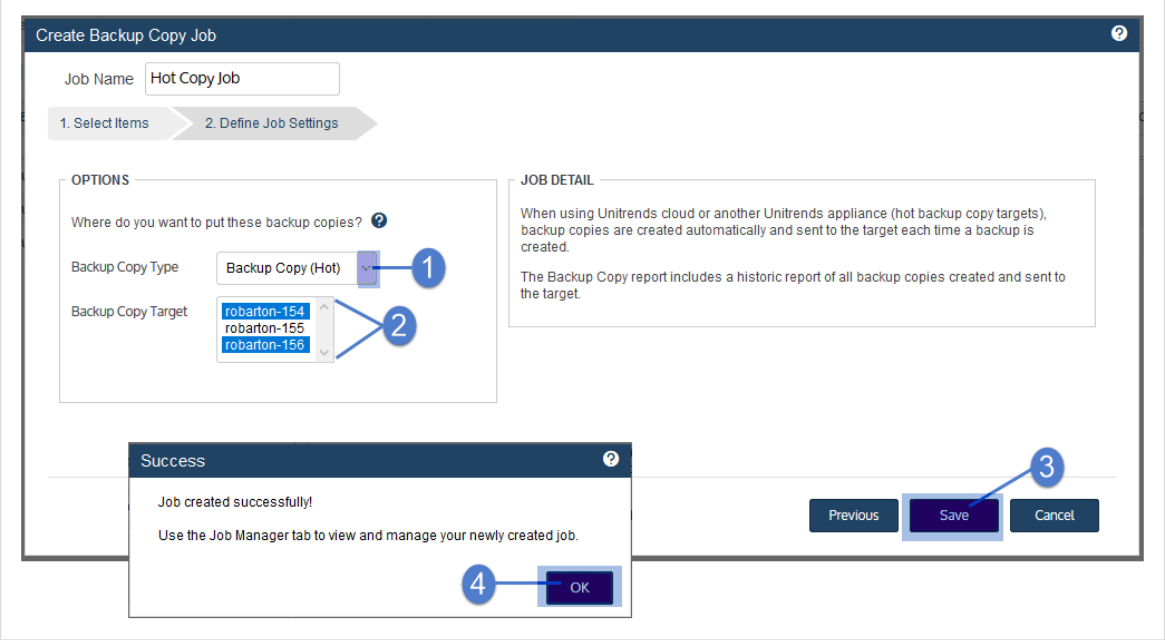

The job is created and backups are copied according to the queue scheme that was configured for the source backup appliance (see "To tune connection options for a Unitrends Cloud or Unitrends appliance target" on page 265).

To copy a full backup to a hot backup copy target on-demand

Use this procedure to manually copy a successful full backup to the Unitrends Cloud or to a Unitrends appliance target. This procedure adds the backup copy job to the Active Jobs queue if these conditions are met:

- A hot backup copy target has been added to the backup appliance.
- The backup has not been copied to the target.
- The backup copy job is not in the Active Jobs queue.
- The source backup appliance is copying to only one hot backup copy target. If the appliance has been configured with multiple hot backup copy targets, the **Add to Hot Backup Copy Queue** button does not display in the backup log and this procedure is not supported.

1. Click **Jobs > Recent Jobs**.
2. Select the full backup and click **View Log**.

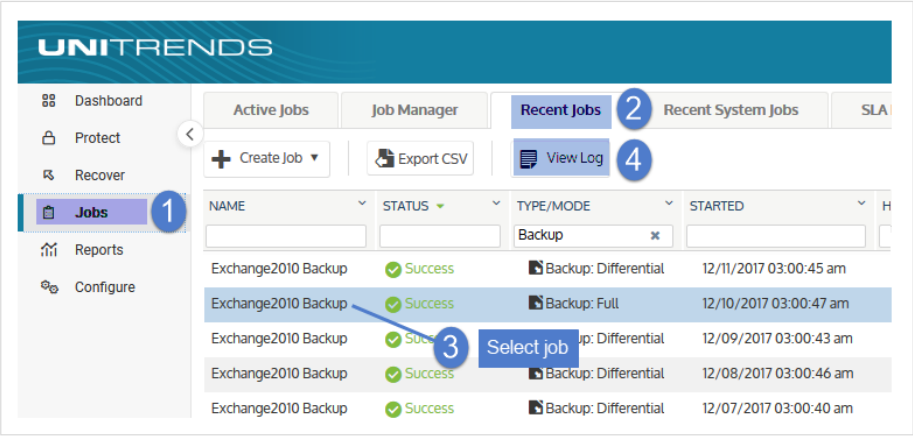

3 Click **Add to Hot Backup Copy Queue**. The backup copy job is added to the queue.

Note: The **Add to Hot Backup Copy Queue** button does not display if hot backup copy is not supported for the backup or if the source appliance is configured with multiple hot backup copy targets.

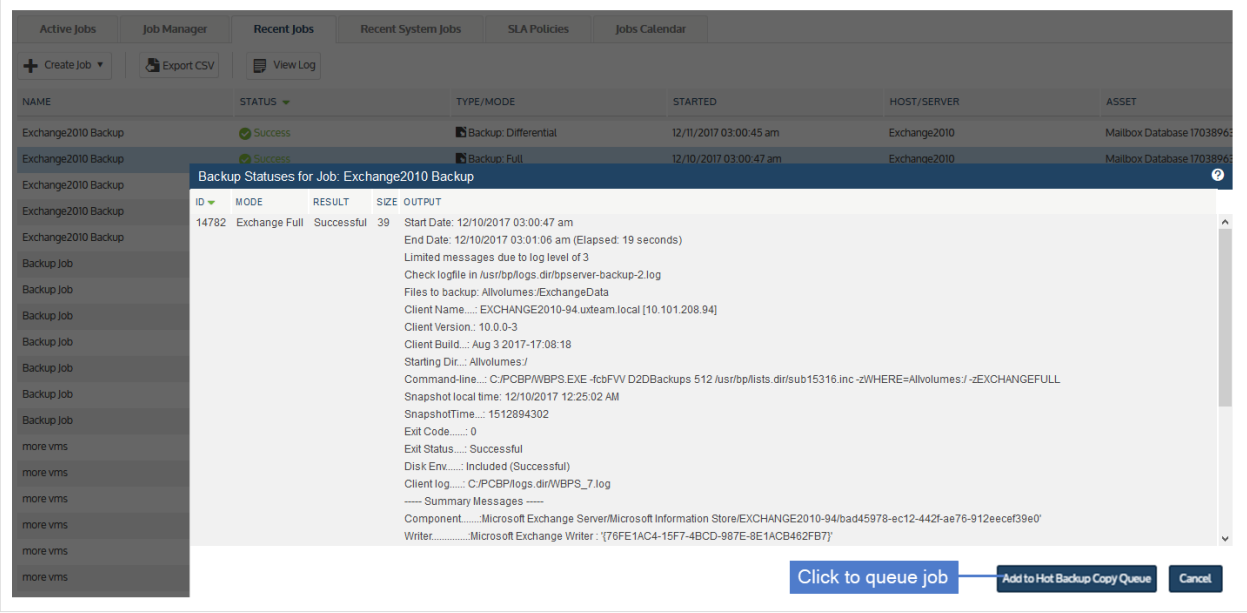

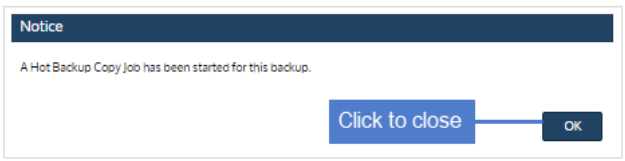

To create a backup copy job for a Unitrends Cloud target

Use this procedure to copy backups to the Unitrends Cloud.

Note: A given asset can be included in only one job per backup copy target.

1 Log in to the source backup appliance.

2 Click **Jobs > Active Jobs > Create job > Backup Copy**.

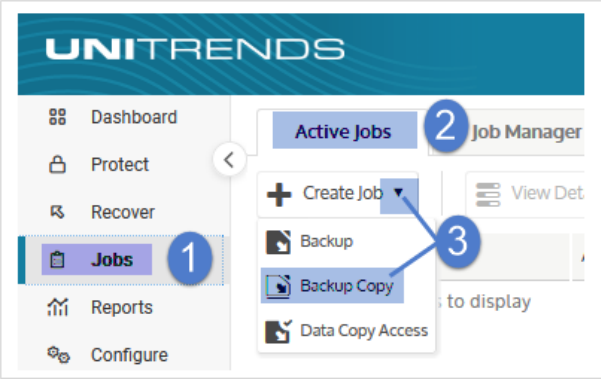

3 Enter a unique **Job Name**.

4 In the Inventory tree, check boxes to select assets whose backups will be copied.

- To locate an asset by name, use the **Search** field below.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select an application instance to select all of its databases or storage groups.

5 Click **Next**.

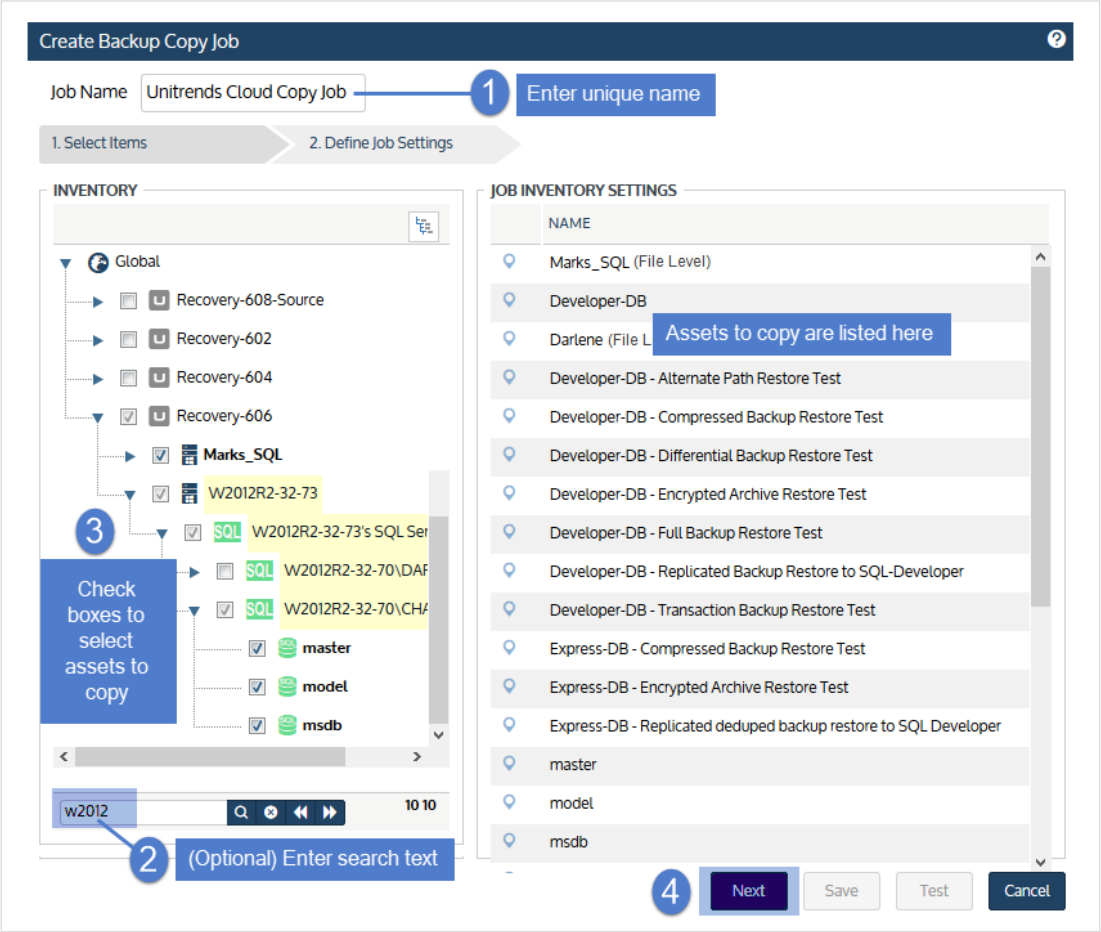

6 Select **Unitrends Cloud** in the **Backup Copy Target** list.

7 Click **Save**.

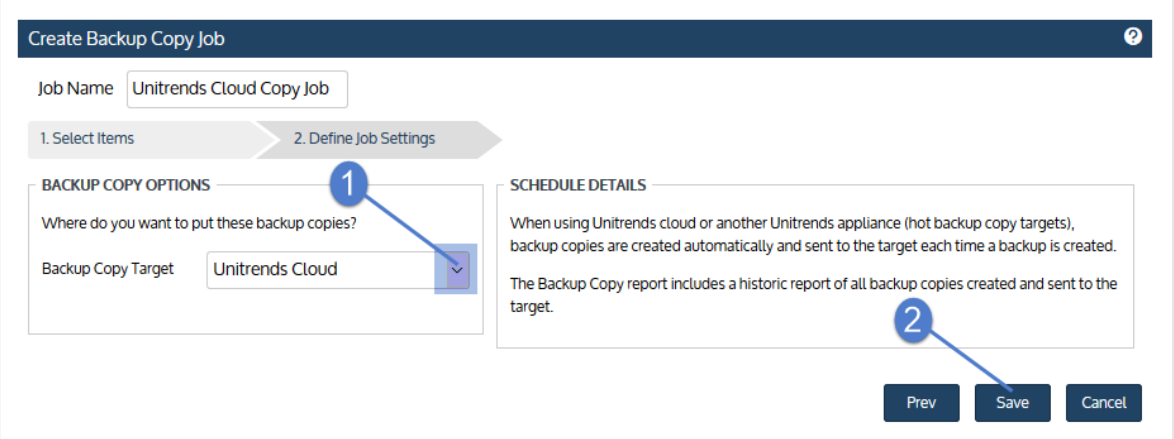

The job is created and backups are copied according to the queue scheme that was configured for the source backup appliance (see "To tune connection options for a Unitrends Cloud or Unitrends appliance target" on page 265).

To create a backup copy job for a third-party cloud target

Use this procedure to copy backups to a Google, Amazon, Rackspace, or Azure Blob cloud target.

IMPORTANT! If you do not have a storage threshold defined for the target, there is no limit to the amount of data the job will copy to the cloud (regardless of the storage threshold setting you define in the job). To define a threshold for the target, see "To view or edit a backup copy target" on page 261.

1 Click **Jobs > Active Jobs > Create job > Backup Copy**.

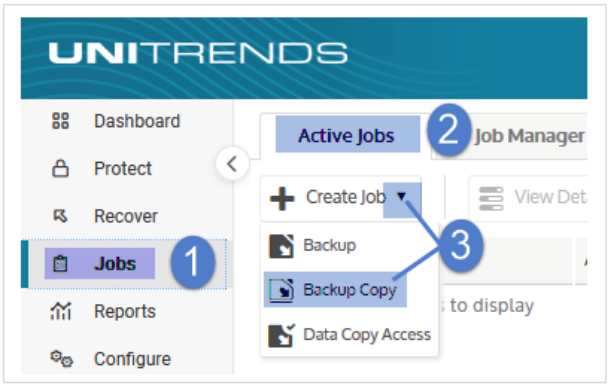

2 Enter a unique **Job Name**.

3 In the Inventory tree, check boxes to select assets whose backups will be copied.

- To locate an asset by name, use the **Search** field below.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.

- Select an application instance to select all of its databases or storage groups.

4 Click **Next**.

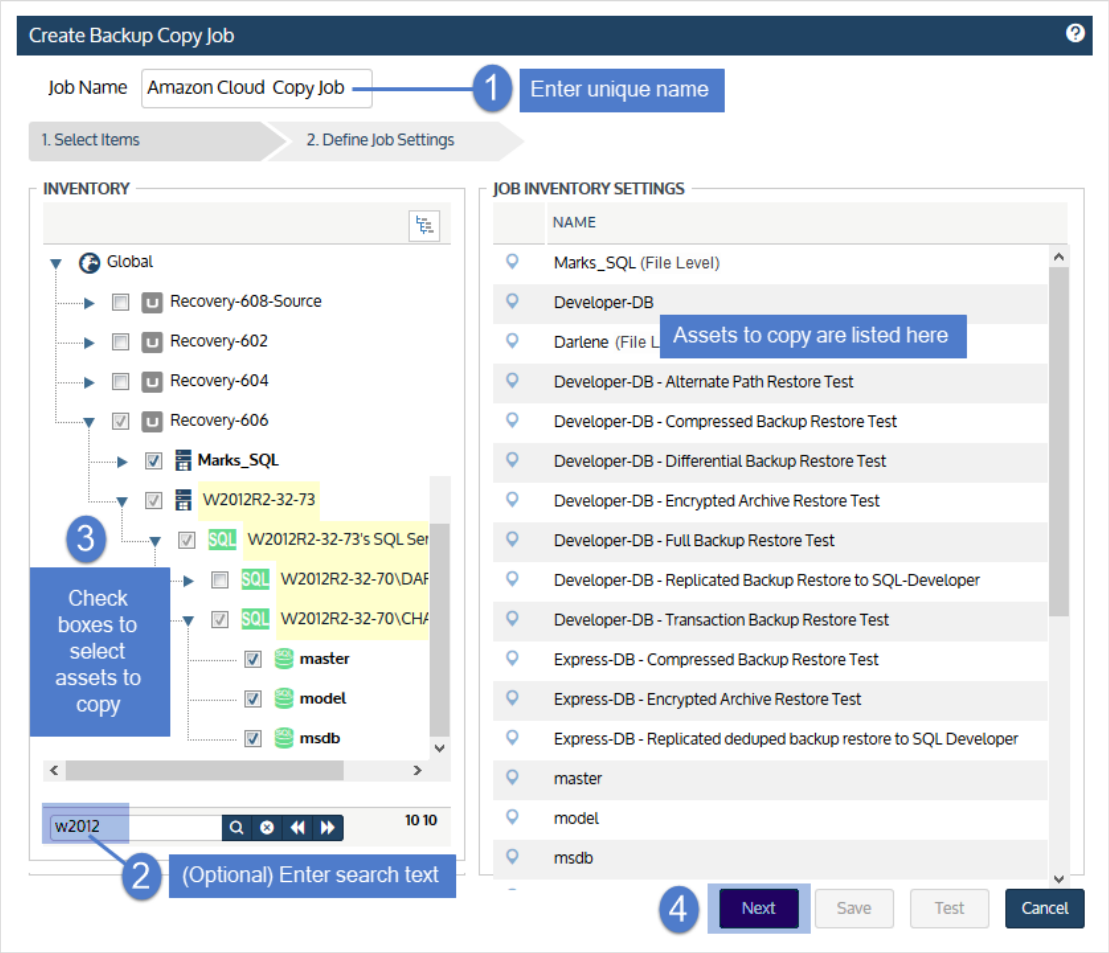

5 Select the cloud target in the **Backup Copy Target** list.

6 Set remaining Backup Copy Options, Schedule Details, and Date Range.

For descriptions of each setting, see "Backup copy job settings" below.

7 (Optional) Click **Test** to see the estimated size of the job and whether the target has enough space available for the new copies. Click **OK** to close the test results Notice.

8 Click **Save**.

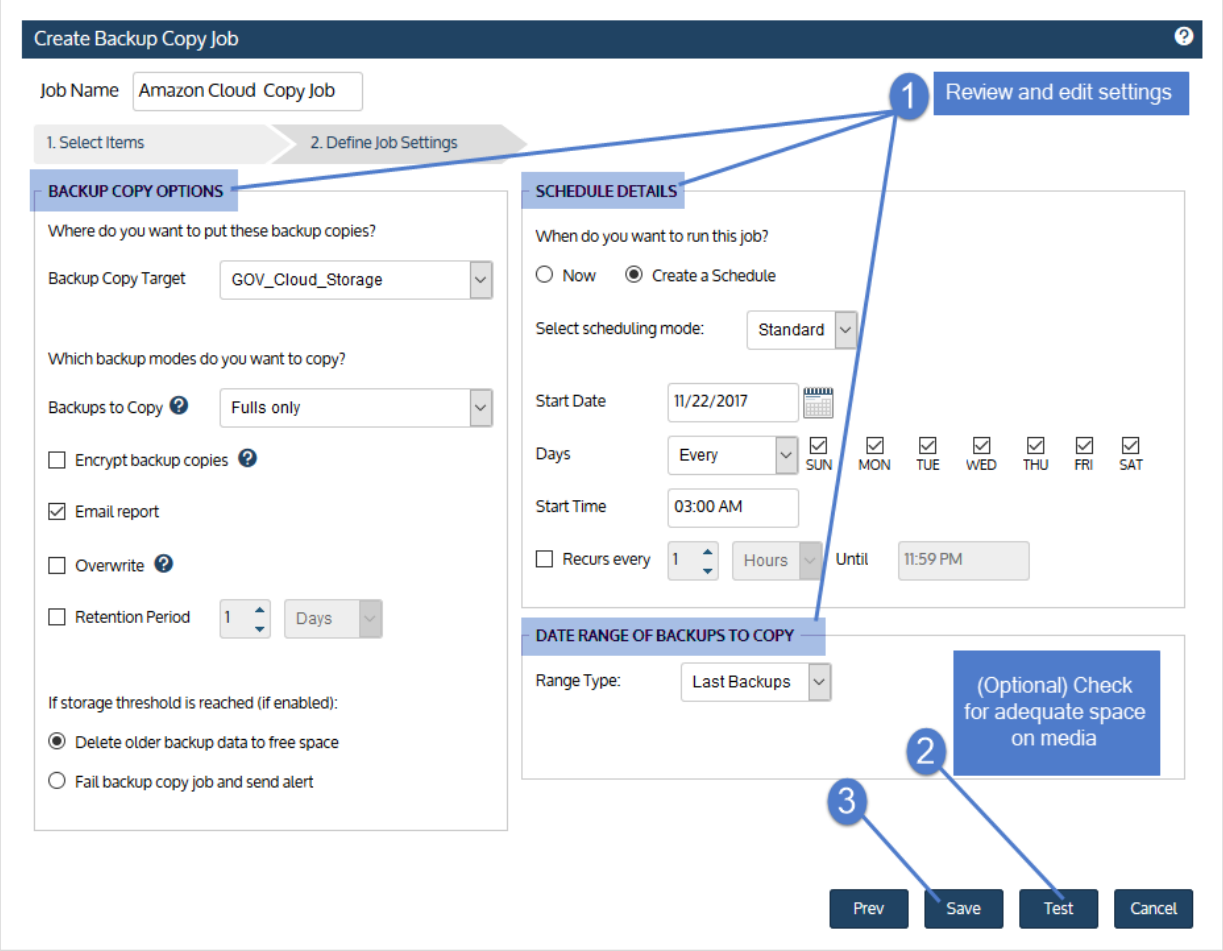

The job is created and will run at the date and times specified.

Backup copy job settings

Item	Description
Job Name	Name of the backup copy job. If creating a schedule, you must enter a unique name.
Backup Copy Target	Select the target where backups will be copied. The list contains all backup targets that have been added to the appliance. For details on adding a target, see "Backup copy targets" on page 214.
Backups to Copy	Select **Fulls only** to copy only full backups or **All Modes** to copy backups of any mode. Backups that complete successfully (green) or with warnings (yellow) are eligible for backup copy jobs. Failed (red) backups are not copied.

Item	Description
Encrypt backup copies	Check this box to encrypt cold backup copies. (Encryption must also be configured on the appliance. For details, see "Encryption" on page 155.) If the backup copy target device is configured for encryption, copies are encrypted regardless of this setting.
Email report	Check this box to email a report when this job completes.
Overwrite	Use with the Retention Period option. Check the **Overwrite** box to overwrite backup copies that are older than the specified Retention Period. If you do not use the Overwrite and Retention Period options, copy jobs fail if there is insufficient space available on the backup copy target. Overwrite does the following: • If all copies on the backup copy target have exceeded their retention, overwrite deletes them and replaces them with the new copies. • Copies on the backup copy target are overwritten only if all have exceeded their retention settings. • All copies are overwritten regardless of available space. • The job fails and nothing is written to the target if overwrite cannot create adequate space for the entire job. • The default retention period is 0 days. If you have not specified another retention period, all existing copies are overwritten each time the backup copy job runs. • If you have selected both the Overwrite and *Delete older backup data to free space*, the Overwrite option takes precedence (Overwrite is applied instead of *Delete older backup data to free space*.)
Retention Period	Use with the *Overwrite* or *Delete older backup data to free space* option. Check the **Retention Period** box to specify the length of time a copy is retained before it can be deleted or overwritten. To define the retention period, enter a number and select Days, Weeks, Months, or Years. Example: enter **2** and select **Weeks** to retain copies for at least 2 weeks. If you do not use the Retention Period option, copy jobs fail if there is insufficient space available on the backup copy target. Modifying the retention period does not change the retention period of existing cold copies. The new setting is applied to subsequent backup copies only.
If storage threshold is reached	Determines whether to delete copies or fail the job if the target does not have sufficient space.

Item	Description
	IMPORTANT! For Google, Amazon, Rackspace, or Azure Blob cloud targets, you must set a storage threshold to limit the amount of data that can be copied. If no threshold is set, the target will always have sufficient space for the copy job and these options do not apply. For details on setting a threshold, see "To view or edit a backup copy target" on page 261.
Delete older backup data to free space	Select this option to remove older copies to make room for new ones. • Copies are deleted only if there is not enough space for the copy job. • Copies held by a retention policy cannot be deleted. • If deleting eligible copies cannot free adequate space for the entire job, the job fails and nothing is written to the target. • If you have selected both *Overwrite* and *Delete older backup data to free space*, overwrite takes precedence (*Overwrite* is applied instead of *Delete older backup data to free space*.)
Fail backup copy job and send alert	Select this option to fail the backup copy job if there is insufficient free space on the target.
Schedule Details	Select **Now** to run the job immediately or select **Create a Schedule** and define additional settings.
Select scheduling mode	In most cases, the **Standard** scheduling mode can be used to create the schedule. If you need more granularity, choose the **Custom** mode to create a custom calendar. For details, see "Using the Custom scheduling mode in the Create Backup Copy Job dialog".
Start Date	Date when the schedule will run for the first time.
Days	Days when the schedule will run.
Start Time	Time of day when the schedule will run. (This is also the Start Time used by the *Recurs every* option.)
Recurs every / Until	Use to run the schedule every *N* hours until the specified time. Check the **Recurs every** box, select a frequency (number of hours), and specify an end time.

Item	Description
Date Range	Choose to copy the most recent backups or specify a date range of backups to copy.

To create a backup copy job for an attached disk target

1 Click **Jobs** > **Active Jobs** > **Create job** > **Backup Copy**.

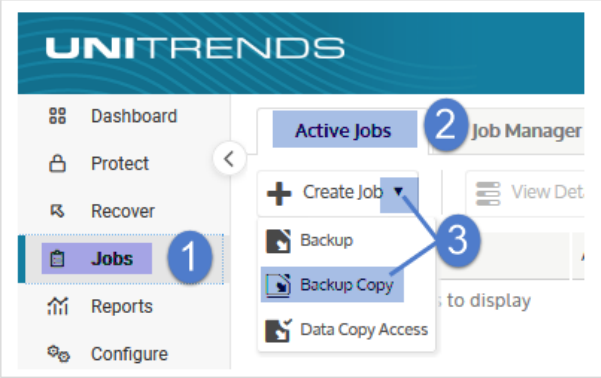

2 Enter a unique **Job Name**.

3 In the Inventory tree, check boxes to select assets whose backups will be copied.

- To locate an asset by name, use the **Search** field below.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select an application instance to select all of its databases or storage groups.

4 Click **Next**.

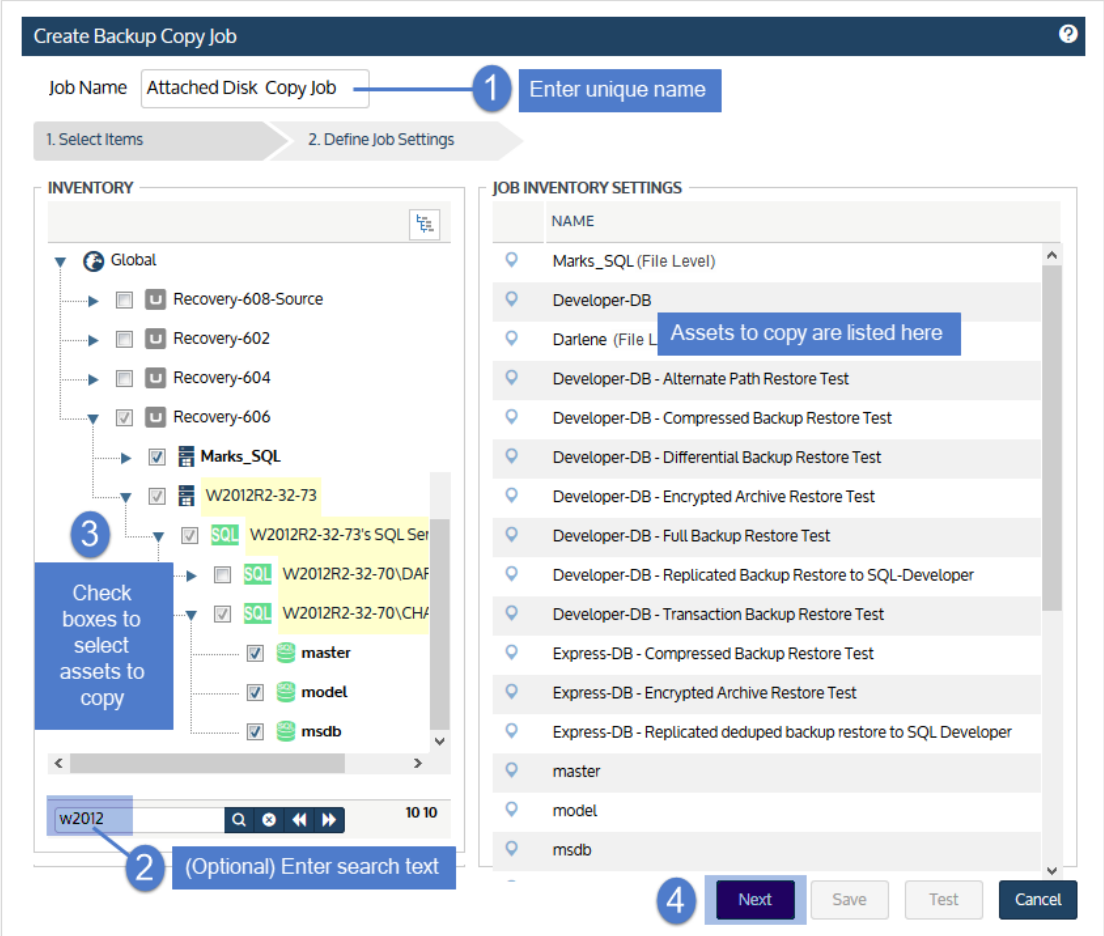

5 Select the disk target in the **Backup Copy Target** list.

6 Set remaining Backup Copy Options, Schedule Details, and Date Range.

For descriptions of each setting, see "Backup copy job settings" below.

7 (Optional) Click **Test** to see the estimated size of the job and whether the target has enough space available for the new copies. Click **OK** to close the test results Notice.

8 Click **Save**.

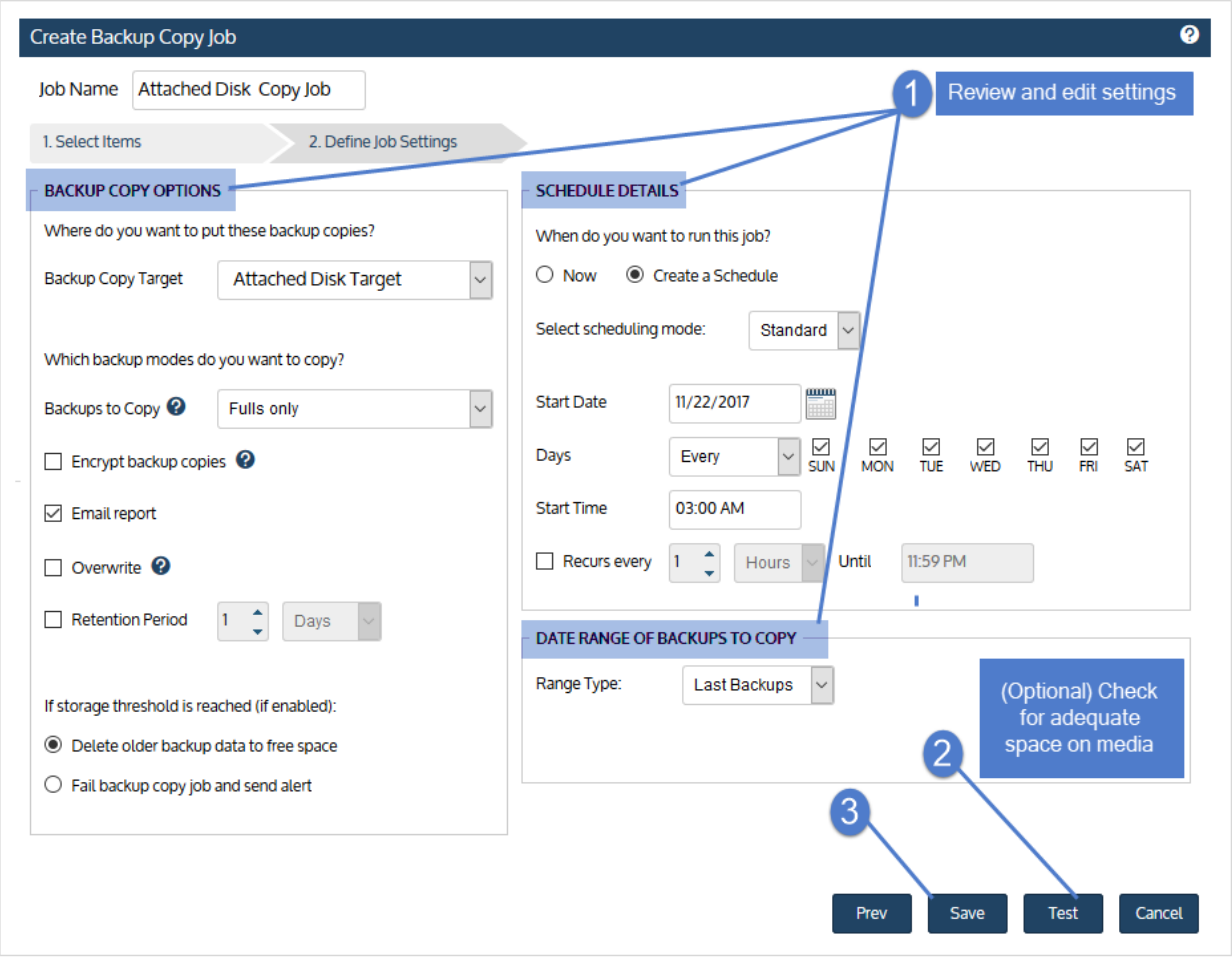

The job is created and will run at the date and times specified.

Backup copy job settings

Item	Description
Job Name	Name of the backup copy job. If creating a schedule, you must enter a unique name.
Backup Copy Target	Select the target where backups will be copied. The list contains all backup targets that have been added to the appliance. For details on adding a target, see "Backup copy targets" on page 214.
Backups to Copy	Select **Fulls only** to copy only full backups or **All Modes** to copy backups of any mode. Backups that complete successfully (green) or with warnings (yellow) are eligible for

Item	Description
	backup copy jobs. Failed (red) backups are not copied.
Encrypt backup copies	Check this box to encrypt cold backup copies. (Encryption must also be configured on the appliance. For details, see "Encryption" on page 155.) If the backup copy target device is configured for encryption, copies are encrypted regardless of this setting.
Email report	Check this box to email a report when this job completes.
Overwrite	Use with the Retention Period option. Check the **Overwrite** box to overwrite backup copies that are older than the specified Retention Period. If you do not use the Overwrite and Retention Period options, copy jobs fail if there is insufficient space available on the backup copy target. Overwrite does the following: • If all copies on the backup copy target have exceeded their retention, overwrite deletes them and replaces them with the new copies. • Copies on the backup copy target are overwritten only if all have exceeded their retention settings. • All copies are overwritten regardless of available space. • The job fails and nothing is written to the target if overwrite cannot create adequate space for the entire job. • The default retention period is 0 days. If you have not specified another retention period, all existing copies are overwritten each time the backup copy job runs. • If you have selected both the Overwrite and *Delete older backup data to free space*, the Overwrite option takes precedence (Overwrite is applied instead of *Delete older backup data to free space*.)
Retention Period	Use with the *Overwrite* or *Delete older backup data to free space* option. Check the **Retention Period** box to specify the length of time a copy is retained before it can be deleted or overwritten. To define the retention period, enter a number and select Days, Weeks, Months, or Years. Example: enter **2** and select **Weeks** to retain copies for at least 2 weeks. If you do not use the Retention Period option, copy jobs fail if there is insufficient space available on the backup copy target. Modifying the retention period does not change the retention period of existing cold copies. The new setting is applied to subsequent backup copies only.
If storage threshold is reached	Determines whether to delete copies or fail the job if the target does not have sufficient space.

Item	Description
	IMPORTANT! For Google, Amazon, Rackspace, or Azure Blob cloud targets, you must set a storage threshold to limit the amount of data that can be copied. If no threshold is set, the target will always have sufficient space for the copy job and these options do not apply. For details on setting a threshold, see "To view or edit a backup copy target" on page 261.
Delete older backup data to free space	Select this option to remove older copies to make room for new ones. • Copies are deleted only if there is not enough space for the copy job. • Copies held by a retention policy cannot be deleted. • If deleting eligible copies cannot free adequate space for the entire job, the job fails and nothing is written to the target. • If you have selected both *Overwrite* and *Delete older backup data to free space*, overwrite takes precedence (*Overwrite* is applied instead of *Delete older backup data to free space*.)
Fail backup copy job and send alert	Select this option to fail the backup copy job if there is insufficient free space on the target.
Schedule Details	Select **Now** to run the job immediately or select **Create a Schedule** and define additional settings.
Select scheduling mode	In most cases, the **Standard** scheduling mode can be used to create the schedule. If you need more granularity, choose the **Custom** mode to create a custom calendar. For details, see "Using the Custom scheduling mode in the Create Backup Copy Job dialog".
Start Date	Date when the schedule will run for the first time.
Days	Days when the schedule will run.
Start Time	Time of day when the schedule will run. (This is also the Start Time used by the *Recurs every* option.)
Recurs every / Until	Use to run the schedule every *N* hours until the specified time. Check the **Recurs every** box, select a frequency (number of hours), and specify an end time.

Item	Description
Date Range	Choose to copy the most recent backups or specify a date range of backups to copy.

To create a backup copy job for a NAS target

1 Click **Jobs** > **Active Jobs** > **Create job** > **Backup Copy**.

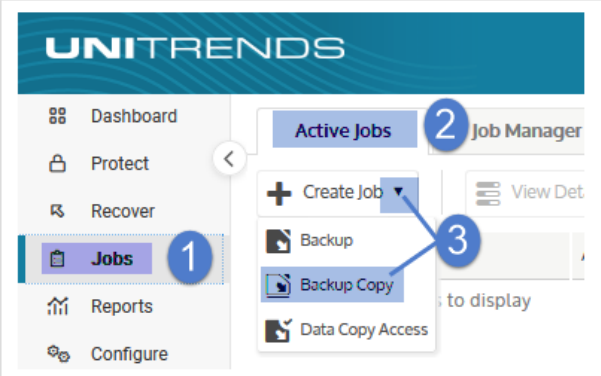

2 Enter a unique **Job Name**.

3 In the Inventory tree, check boxes to select assets whose backups will be copied.

- To locate an asset by name, use the **Search** field below.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select an application instance to select all of its databases or storage groups.

4 Click **Next**.

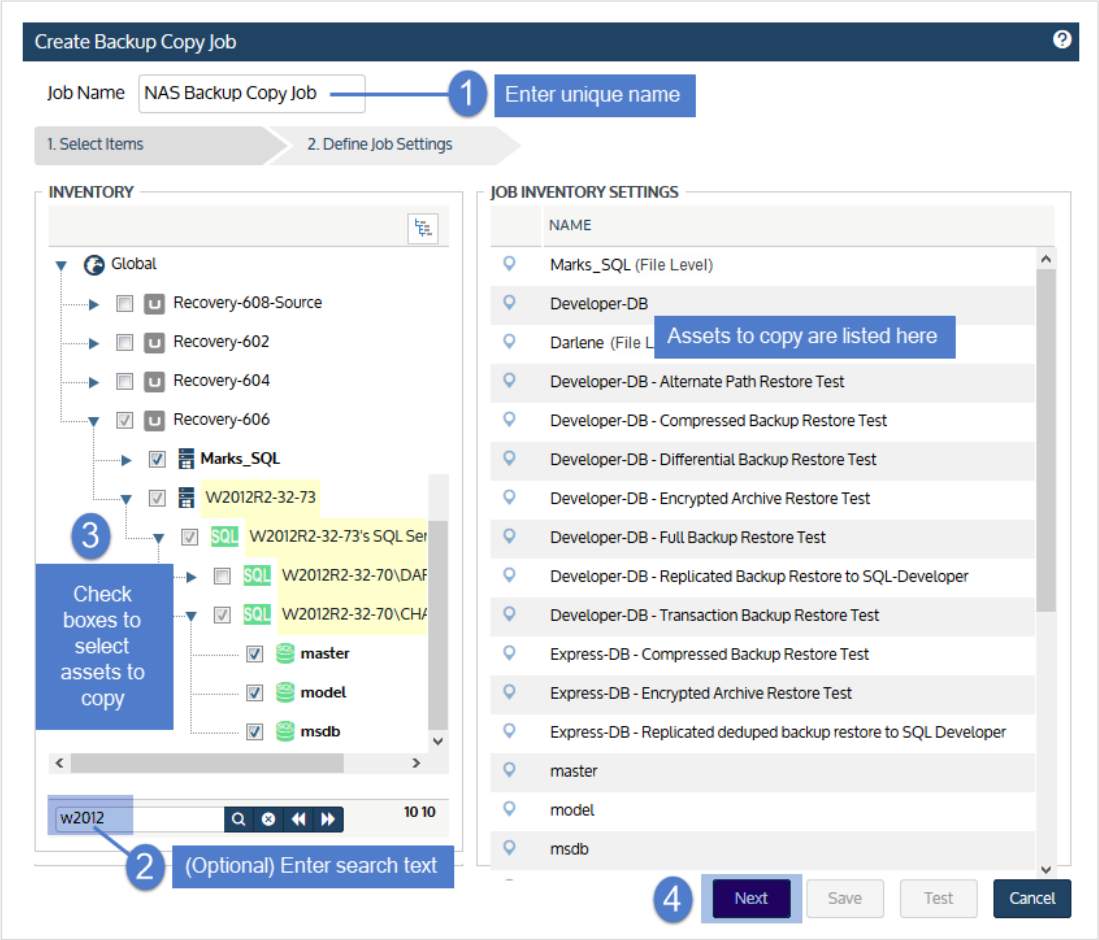

5 Select the NAS target in the **Backup Copy Target** list.

6 Set remaining Backup Copy Options, Schedule Details, and Date Range.

For descriptions of each setting, see "Backup copy job settings" below.

7 (Optional) Click **Test** to see the estimated size of the job and whether the target has enough space available for the new copies. Click **OK** to close the test results Notice.

8 Click **Save**.

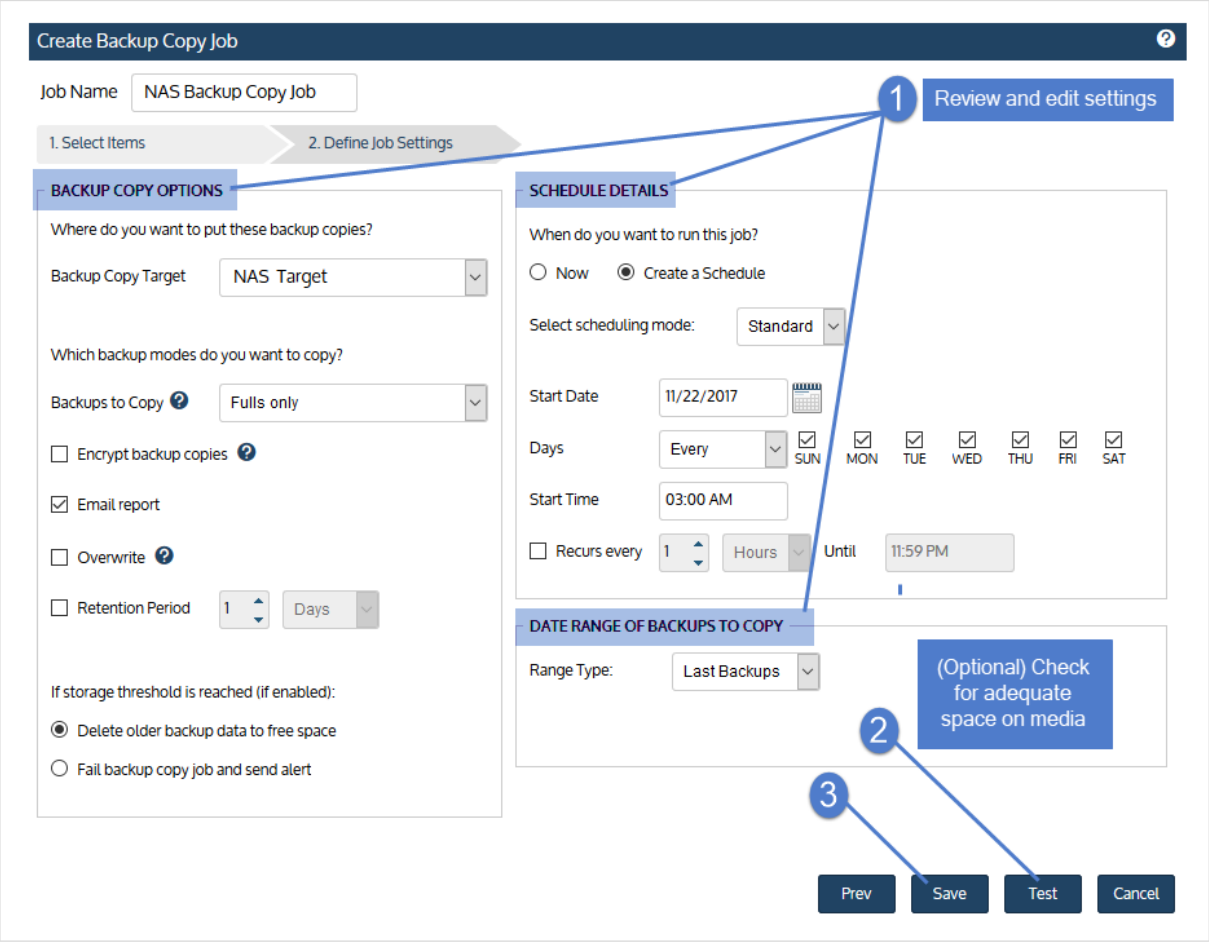

The job is created and will run at the date and times specified.

Backup copy job settings

Item	Description
Job Name	Name of the backup copy job. If creating a schedule, you must enter a unique name.
Backup Copy Target	Select the target where backups will be copied. The list contains all backup targets that have been added to the appliance. For details on adding a target, see "Backup copy targets" on page 214.
Backups to Copy	Select **Fulls only** to copy only full backups or **All Modes** to copy backups of any mode. Backups that complete successfully (green) or with warnings (yellow) are eligible for backup copy jobs. Failed (red) backups are not copied.

Item	Description
Encrypt backup copies	Check this box to encrypt cold backup copies. (Encryption must also be configured on the appliance. For details, see "Encryption" on page 155.) If the backup copy target device is configured for encryption, copies are encrypted regardless of this setting.
Email report	Check this box to email a report when this job completes.
Overwrite	Use with the Retention Period option. Check the **Overwrite** box to overwrite backup copies that are older than the specified Retention Period. If you do not use the Overwrite and Retention Period options, copy jobs fail if there is insufficient space available on the backup copy target. Overwrite does the following: • If all copies on the backup copy target have exceeded their retention, overwrite deletes them and replaces them with the new copies. • Copies on the backup copy target are overwritten only if all have exceeded their retention settings. • All copies are overwritten regardless of available space. • The job fails and nothing is written to the target if overwrite cannot create adequate space for the entire job. • The default retention period is 0 days. If you have not specified another retention period, all existing copies are overwritten each time the backup copy job runs. • If you have selected both the Overwrite and *Delete older backup data to free space*, the Overwrite option takes precedence (Overwrite is applied instead of *Delete older backup data to free space*.)
Retention Period	Use with the *Overwrite* or *Delete older backup data to free space* option. Check the **Retention Period** box to specify the length of time a copy is retained before it can be deleted or overwritten. To define the retention period, enter a number and select Days, Weeks, Months, or Years. Example: enter **2** and select **Weeks** to retain copies for at least 2 weeks. If you do not use the Retention Period option, copy jobs fail if there is insufficient space available on the backup copy target. Modifying the retention period does not change the retention period of existing cold copies. The new setting is applied to subsequent backup copies only.
If storage threshold is reached	Determines whether to delete copies or fail the job if the target does not have sufficient space.

Item	Description
	IMPORTANT! For Google, Amazon, Rackspace, or Azure Blob cloud targets, you must set a storage threshold to limit the amount of data that can be copied. If no threshold is set, the target will always have sufficient space for the copy job and these options do not apply. For details on setting a threshold, see "To view or edit a backup copy target" on page 261.
Delete older backup data to free space	Select this option to remove older copies to make room for new ones. • Copies are deleted only if there is not enough space for the copy job. • Copies held by a retention policy cannot be deleted. • If deleting eligible copies cannot free adequate space for the entire job, the job fails and nothing is written to the target. • If you have selected both *Overwrite* and *Delete older backup data to free space*, overwrite takes precedence (*Overwrite* is applied instead of *Delete older backup data to free space*.)
Fail backup copy job and send alert	Select this option to fail the backup copy job if there is insufficient free space on the target.
Schedule Details	Select **Now** to run the job immediately or select **Create a Schedule** and define additional settings.
Select scheduling mode	In most cases, the **Standard** scheduling mode can be used to create the schedule. If you need more granularity, choose the **Custom** mode to create a custom calendar. For details, see "Using the Custom scheduling mode in the Create Backup Copy Job dialog".
Start Date	Date when the schedule will run for the first time.
Days	Days when the schedule will run.
Start Time	Time of day when the schedule will run. (This is also the Start Time used by the *Recurs every* option.)
Recurs every / Until	Use to run the schedule every *N* hours until the specified time. Check the **Recurs every** box, select a frequency (number of hours), and specify an end time.

Item	Description
Date Range	Choose to copy the most recent backups or specify a date range of backups to copy.

To create a backup copy job for a SAN target

1 Click **Jobs** > **Active Jobs** > **Create job** > **Backup Copy**.

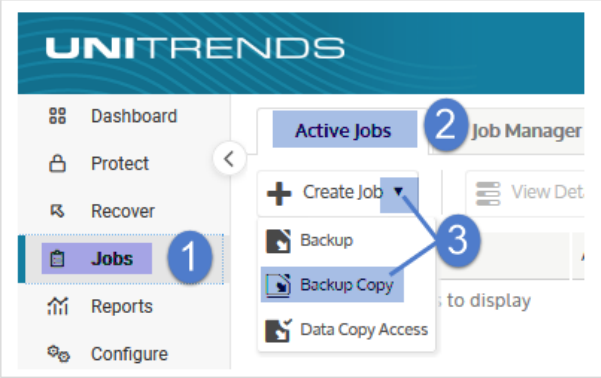

2 Enter a unique **Job Name**.

3 In the Inventory tree, check boxes to select assets whose backups will be copied.

- To locate an asset by name, use the **Search** field below.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select an application instance to select all of its databases or storage groups.

4 Click **Next**.

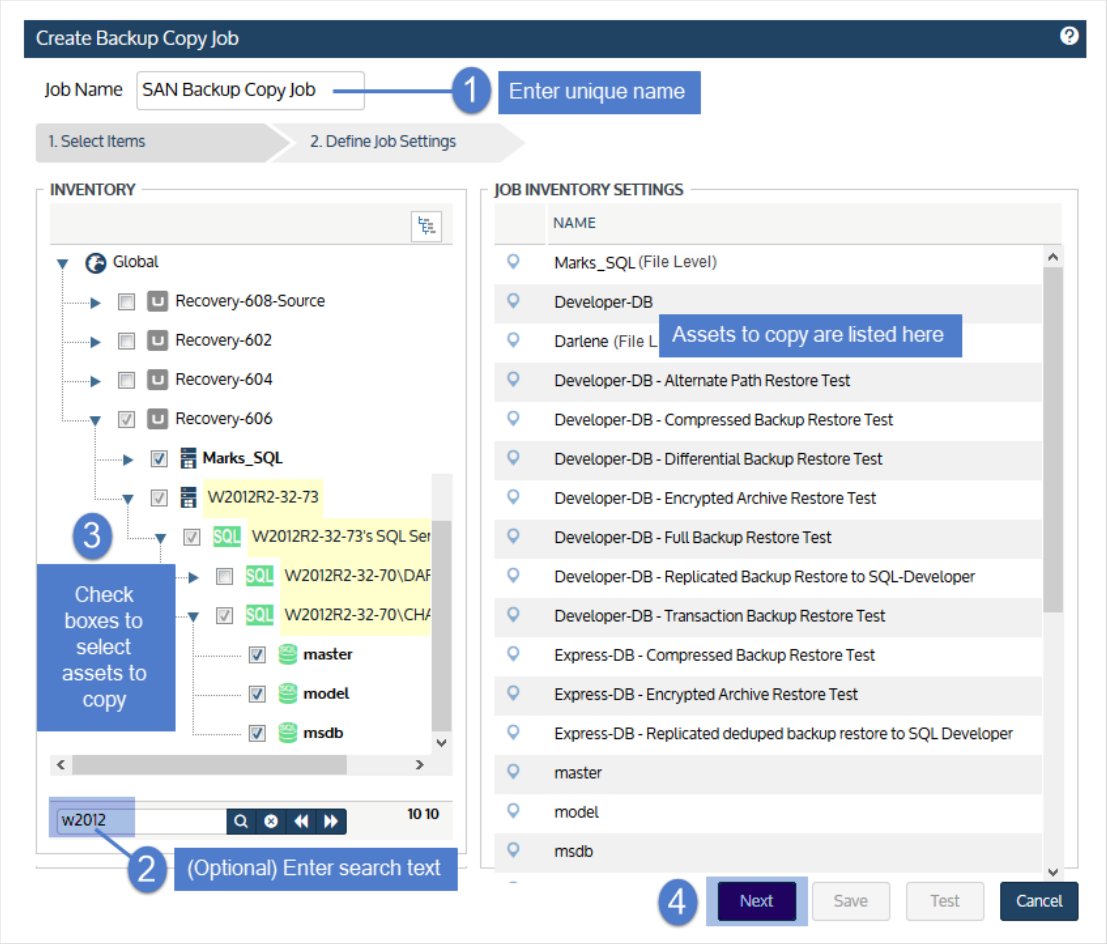

5 Select the SAN target in the **Backup Copy Target** list.

6 Set remaining Backup Copy Options, Schedule Details, and Date Range.

For descriptions of each setting, see "Backup copy job settings" below.

7 (Optional) Click **Test** to see the estimated size of the job and whether the target has enough space available for the new copies. Click **OK** to close the test results Notice.

8 Click **Save**.

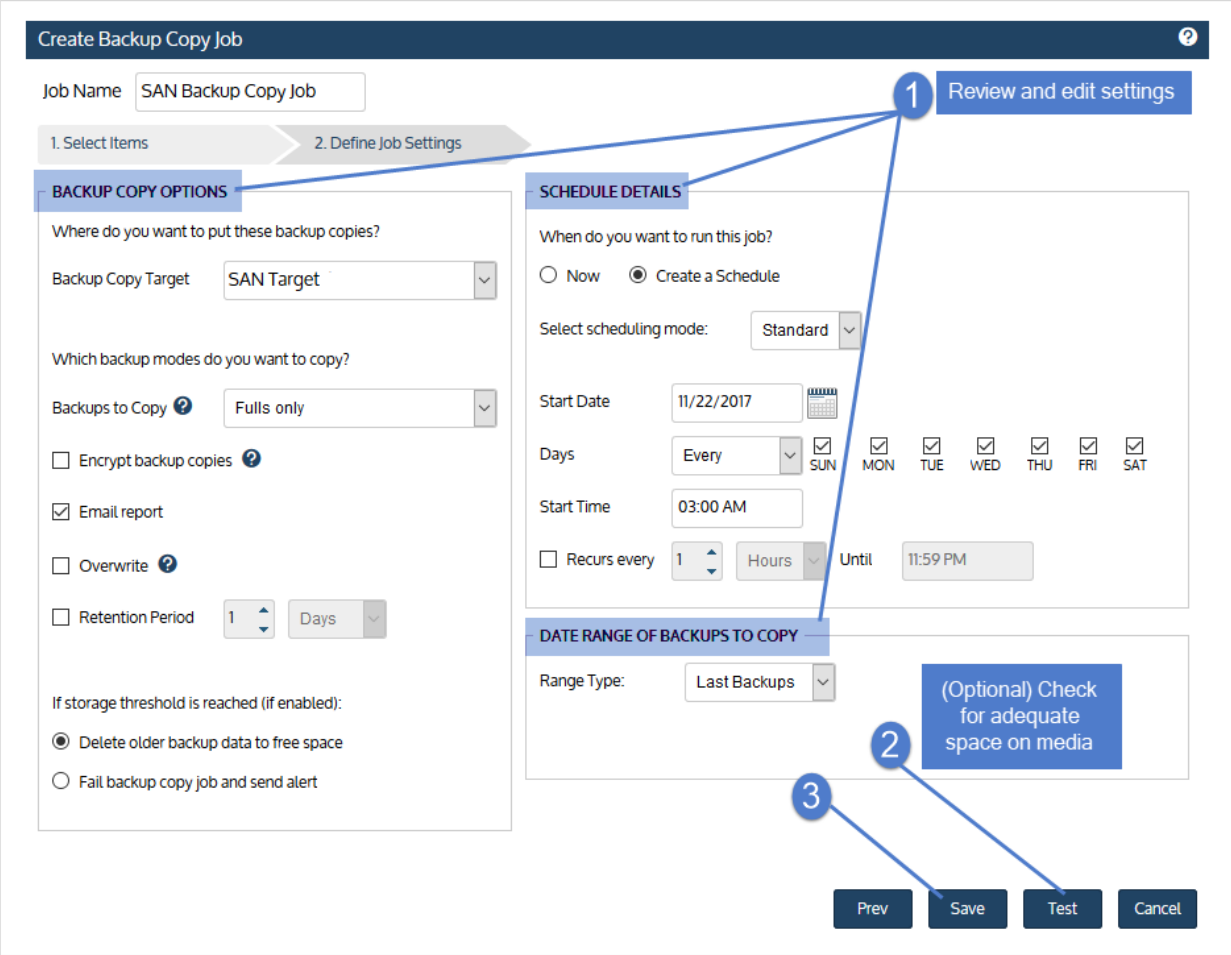

The job is created and will run at the date and times specified.

Backup copy job settings

Item	Description
Job Name	Name of the backup copy job. If creating a schedule, you must enter a unique name.
Backup Copy Target	Select the target where backups will be copied. The list contains all backup targets that have been added to the appliance. For details on adding a target, see "Backup copy targets" on page 214.
Backups to Copy	Select **Fulls only** to copy only full backups or **All Modes** to copy backups of any mode. Backups that complete successfully (green) or with warnings (yellow) are eligible for

Item	Description
	backup copy jobs. Failed (red) backups are not copied.
Encrypt backup copies	Check this box to encrypt cold backup copies. (Encryption must also be configured on the appliance. For details, see "Encryption" on page 155.) If the backup copy target device is configured for encryption, copies are encrypted regardless of this setting.
Email report	Check this box to email a report when this job completes.
Overwrite	Use with the Retention Period option. Check the **Overwrite** box to overwrite backup copies that are older than the specified Retention Period. If you do not use the Overwrite and Retention Period options, copy jobs fail if there is insufficient space available on the backup copy target. Overwrite does the following: • If all copies on the backup copy target have exceeded their retention, overwrite deletes them and replaces them with the new copies. • Copies on the backup copy target are overwritten only if all have exceeded their retention settings. • All copies are overwritten regardless of available space. • The job fails and nothing is written to the target if overwrite cannot create adequate space for the entire job. • The default retention period is 0 days. If you have not specified another retention period, all existing copies are overwritten each time the backup copy job runs. • If you have selected both the Overwrite and *Delete older backup data to free space*, the Overwrite option takes precedence (Overwrite is applied instead of *Delete older backup data to free space*.)
Retention Period	Use with the *Overwrite* or *Delete older backup data to free space* option. Check the **Retention Period** box to specify the length of time a copy is retained before it can be deleted or overwritten. To define the retention period, enter a number and select Days, Weeks, Months, or Years. Example: enter **2** and select **Weeks** to retain copies for at least 2 weeks. If you do not use the Retention Period option, copy jobs fail if there is insufficient space available on the backup copy target. Modifying the retention period does not change the retention period of existing cold copies. The new setting is applied to subsequent backup copies only.
If storage threshold is reached	Determines whether to delete copies or fail the job if the target does not have sufficient space.

Item	Description
	IMPORTANT! For Google, Amazon, Rackspace, or Azure Blob cloud targets, you must set a storage threshold to limit the amount of data that can be copied. If no threshold is set, the target will always have sufficient space for the copy job and these options do not apply. For details on setting a threshold, see "To view or edit a backup copy target" on page 261.
Delete older backup data to free space	Select this option to remove older copies to make room for new ones. • Copies are deleted only if there is not enough space for the copy job. • Copies held by a retention policy cannot be deleted. • If deleting eligible copies cannot free adequate space for the entire job, the job fails and nothing is written to the target. • If you have selected both *Overwrite* and *Delete older backup data to free space*, overwrite takes precedence (*Overwrite* is applied instead of *Delete older backup data to free space*.)
Fail backup copy job and send alert	Select this option to fail the backup copy job if there is insufficient free space on the target.
Schedule Details	Select **Now** to run the job immediately or select **Create a Schedule** and define additional settings.
Select scheduling mode	In most cases, the **Standard** scheduling mode can be used to create the schedule. If you need more granularity, choose the **Custom** mode to create a custom calendar. For details, see "Using the Custom scheduling mode in the Create Backup Copy Job dialog".
Start Date	Date when the schedule will run for the first time.
Days	Days when the schedule will run.
Start Time	Time of day when the schedule will run. (This is also the Start Time used by the *Recurs every* option.)
Recurs every / Until	Use to run the schedule every *N* hours until the specified time. Check the **Recurs every** box, select a frequency (number of hours), and specify an end time.

Item	Description
Date Range	Choose to copy the most recent backups or specify a date range of backups to copy.

To create a backup copy job for an eSATA or USB target

1 Click **Jobs** > **Active Jobs** > **Create job** > **Backup Copy**.

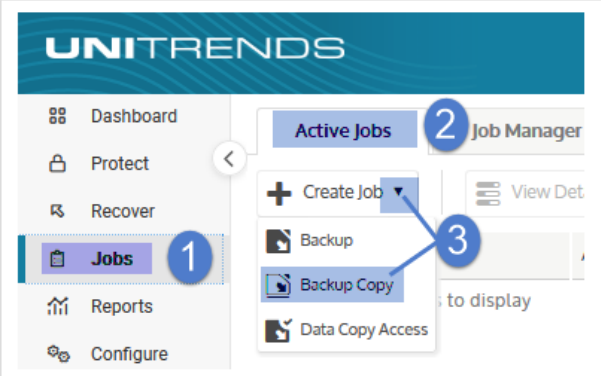

2 Enter a unique **Job Name**.

3 In the Inventory tree, check boxes to select assets whose backups will be copied.

- To locate an asset by name, use the **Search** field below.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select an application instance to select all of its databases or storage groups.

4 Click **Next**.

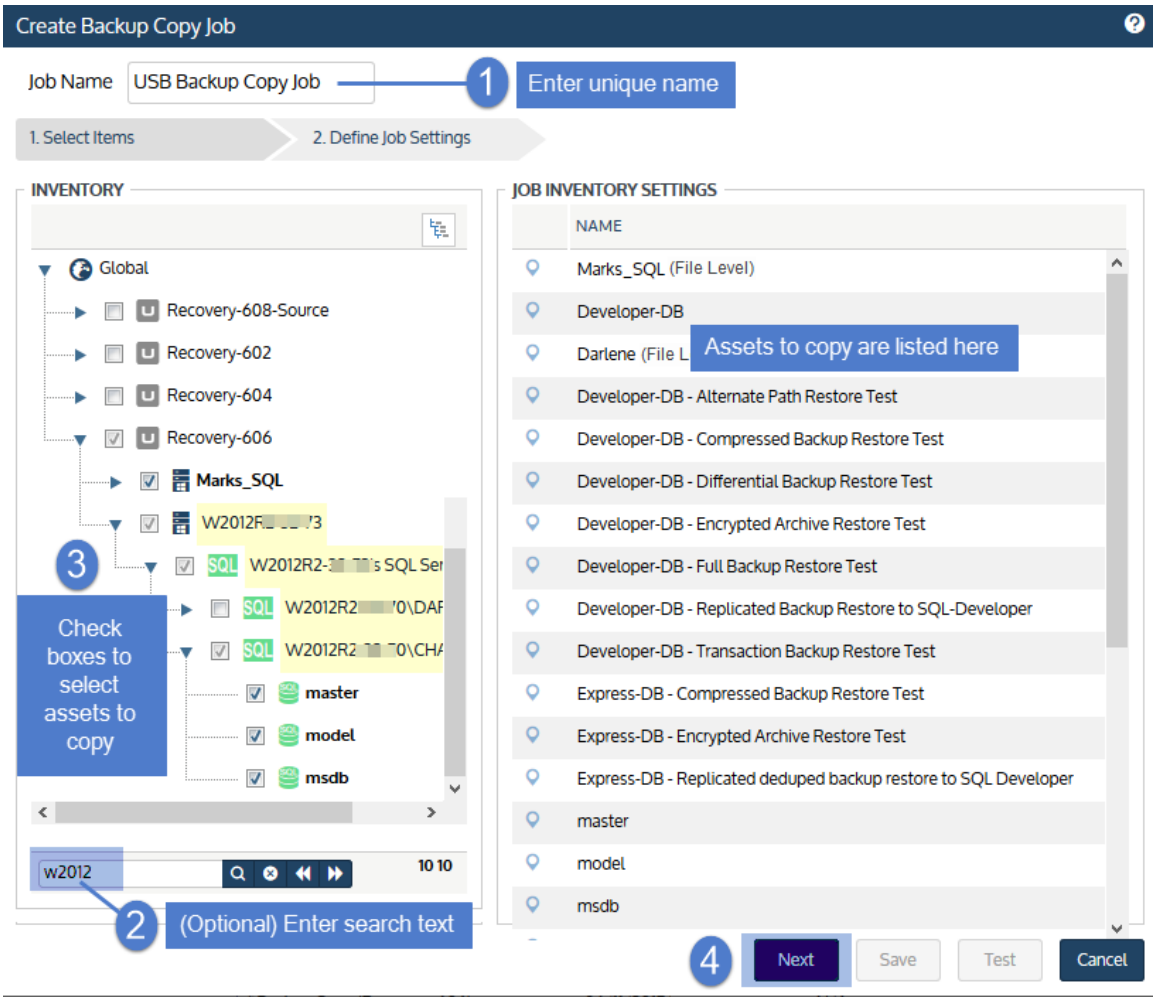

5 Select the disk target in the **Backup Copy Target** list.

6 Set remaining Backup Copy Options, Schedule Details, and Date Range.

For descriptions of each setting, see "Backup copy job settings" below.

7 (Optional) Click **Test** to see the estimated size of the job and whether the target has enough space available for the new copies. Click **OK** to close the test results Notice.

8 Click **Save**.

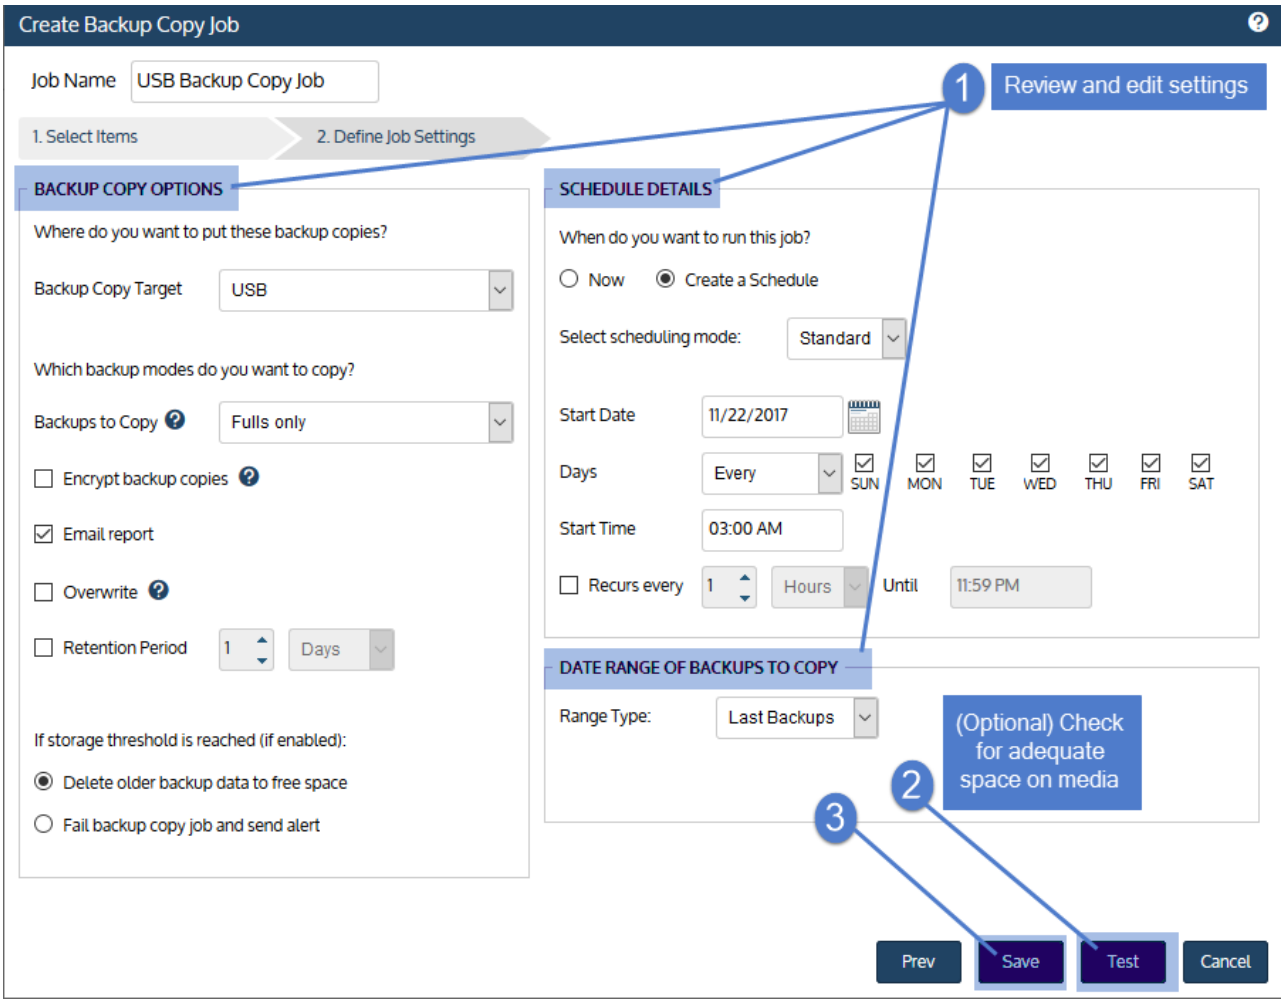

The job is created and will run at the date and times specified.

Backup copy job settings

Item	Description
Job Name	Name of the backup copy job. If creating a schedule, you must enter a unique name.
Backup Copy Target	Select the target where backups will be copied. The list contains all backup targets that have been added to the appliance. For details on adding a target, see "Backup copy targets" on page 214.
Backups to Copy	Select **Fulls only** to copy only full backups or **All Modes** to copy backups of any mode. Backups that complete successfully (green) or with warnings (yellow) are eligible for backup copy jobs. Failed (red) backups are not copied.

Item	Description
Encrypt backup copies	Check this box to encrypt cold backup copies. (Encryption must also be configured on the appliance. For details, see "Encryption" on page 155.) If the backup copy target device is configured for encryption, copies are encrypted regardless of this setting.
Email report	Check this box to email a report when this job completes.
Overwrite	Use with the Retention Period option. Check the **Overwrite** box to overwrite backup copies that are older than the specified Retention Period. If you do not use the Overwrite and Retention Period options, copy jobs fail if there is insufficient space available on the backup copy target. Overwrite does the following: • If all copies on the backup copy target have exceeded their retention, overwrite deletes them and replaces them with the new copies. • Copies on the backup copy target are overwritten only if all have exceeded their retention settings. • All copies are overwritten regardless of available space. • The job fails and nothing is written to the target if overwrite cannot create adequate space for the entire job. • The default retention period is 0 days. If you have not specified another retention period, all existing copies are overwritten each time the backup copy job runs. • If you have selected both the Overwrite and *Delete older backup data to free space*, the Overwrite option takes precedence (Overwrite is applied instead of *Delete older backup data to free space*.)
Retention Period	Use with the *Overwrite* or *Delete older backup data to free space* option. Check the **Retention Period** box to specify the length of time a copy is retained before it can be deleted or overwritten. To define the retention period, enter a number and select Days, Weeks, Months, or Years. Example: enter **2** and select **Weeks** to retain copies for at least 2 weeks. If you do not use the Retention Period option, copy jobs fail if there is insufficient space available on the backup copy target. Modifying the retention period does not change the retention period of existing cold copies. The new setting is applied to subsequent backup copies only.
If storage threshold is reached	Determines whether to delete copies or fail the job if the target does not have sufficient space.

Item	Description
	IMPORTANT! For Google, Amazon, Rackspace, or Azure Blob cloud targets, you must set a storage threshold to limit the amount of data that can be copied. If no threshold is set, the target will always have sufficient space for the copy job and these options do not apply. For details on setting a threshold, see "To view or edit a backup copy target" on page 261.
Delete older backup data to free space	Select this option to remove older copies to make room for new ones. • Copies are deleted only if there is not enough space for the copy job. • Copies held by a retention policy cannot be deleted. • If deleting eligible copies cannot free adequate space for the entire job, the job fails and nothing is written to the target. • If you have selected both *Overwrite* and *Delete older backup data to free space*, overwrite takes precedence (*Overwrite* is applied instead of *Delete older backup data to free space*.)
Fail backup copy job and send alert	Select this option to fail the backup copy job if there is insufficient free space on the target.
Schedule Details	Select **Now** to run the job immediately or select **Create a Schedule** and define additional settings.
Select scheduling mode	In most cases, the **Standard** scheduling mode can be used to create the schedule. If you need more granularity, choose the **Custom** mode to create a custom calendar. For details, see "Using the Custom scheduling mode in the Create Backup Copy Job dialog".
Start Date	Date when the schedule will run for the first time.
Days	Days when the schedule will run.
Start Time	Time of day when the schedule will run. (This is also the Start Time used by the *Recurs every* option.)
Recurs every / Until	Use to run the schedule every *N* hours until the specified time. Check the **Recurs every** box, select a frequency (number of hours), and specify an end time.

Item	Description
Date Range	Choose to copy the most recent backups or specify a date range of backups to copy.

To create a backup copy job for a tape target

1 Click **Jobs** > **Active Jobs** > **Create job** > **Backup Copy**.

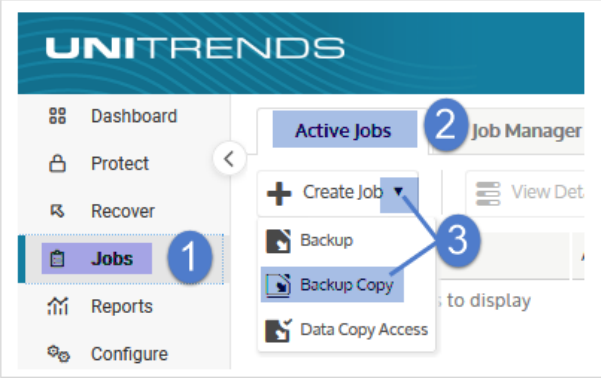

2 Enter a unique **Job Name**.

3 In the Inventory tree, check boxes to select assets whose backups will be copied.

- To locate an asset by name, use the **Search** field below.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select an application instance to select all of its databases or storage groups.

4 Click **Next**.

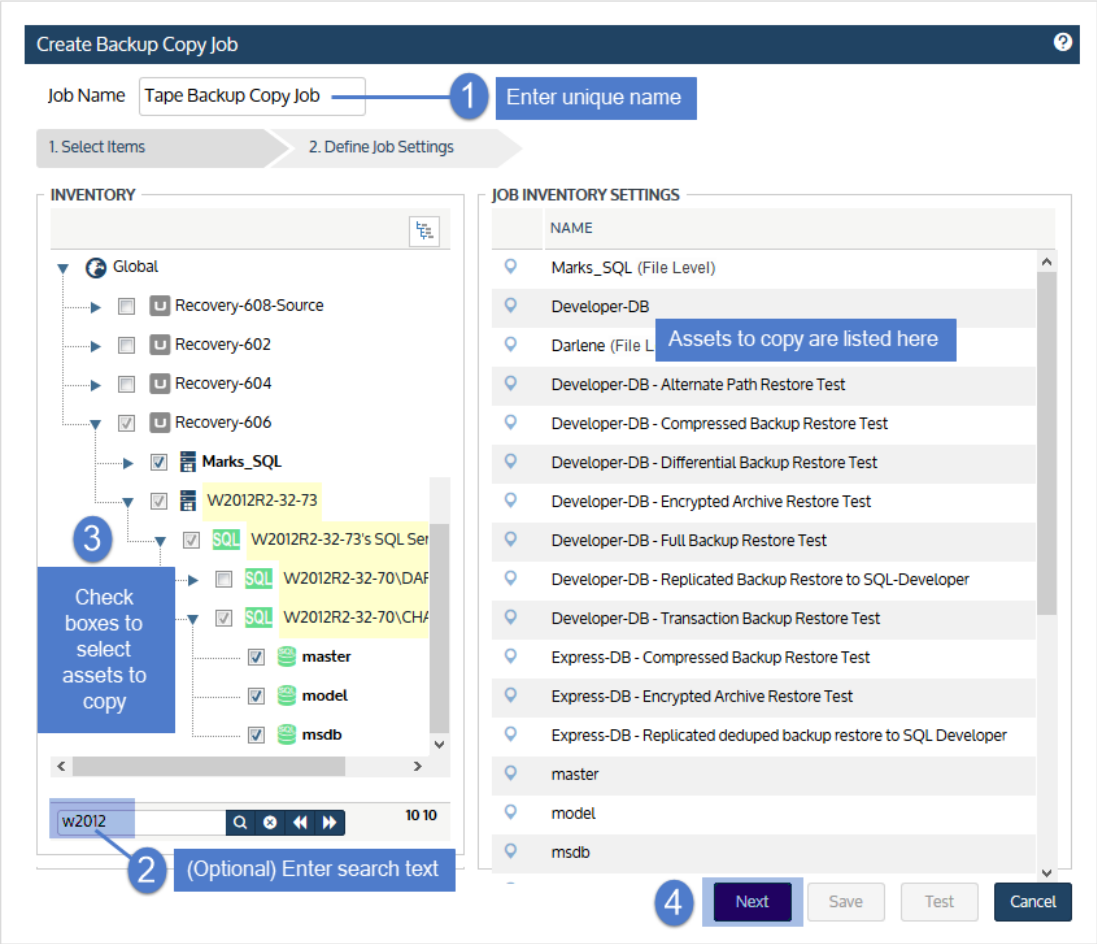

5 Select the tape drive or changer in the **Backup Copy Target** list.

6 Set remaining Backup Copy Options, Schedule Details, and Date Range.

For descriptions of each setting, see "Tape backup copy job settings" below.

7 Click **Save**.

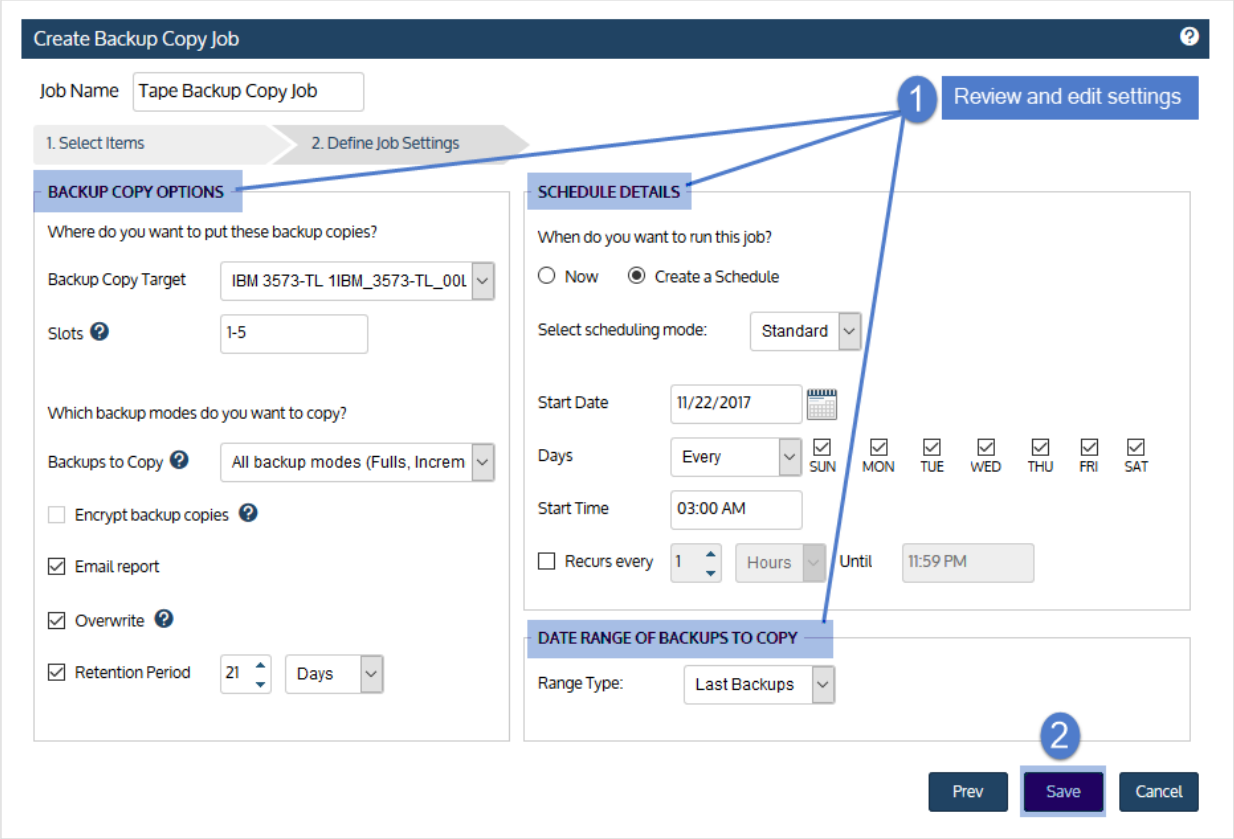

The job is created and will run at the date and times specified.

Tape backup copy job settings

Option	Description
Job Name	Name of the backup copy job. If creating a schedule, you must enter a unique name.
Backup Copy Target	Select the target where backups will be copied. The list contains all backup targets that have been added to the appliance. For details on adding a target, see "Backup copy targets" on page 214.
Slots	Applies to autochangers only. Leave this field empty if your tape device does not have an autochanger. Enter the slots to use when writing backups to tape: • Enter slot numbers using a comma-separated list and/or ranges. Example: 1,2,3,5-8

Option	Description
	• To use all slots that contain available tapes, enter **all** in the Slots field.
Backups to Copy	Select **Fulls only** to copy only full backups or **All Modes** to copy backups of any mode. Backups that complete successfully (green) or with warnings (yellow) are eligible for backup copy jobs. Failed (red) backups are not copied.
Encrypt backup copies	Not used for tape devices. If the tape device is configured for encryption, copies are encrypted regardless of this setting.
Email report	Check this box to email a report when this job completes.
Overwrite	Use with the Retention Period option. Check the **Overwrite** box to overwrite backup copies that are older than the specified Retention Period. If you do not use the Overwrite and Retention Period options, copy jobs fail if there is insufficient space available on the tape(s). Overwrite does the following: • If all copies on the tape media have exceeded their retention, overwrite deletes them and replaces them with the new copies. • Copies on the tape media are overwritten only if all have exceeded their retention settings. • All copies are overwritten regardless of available space. • The job fails and nothing is written to the media if overwrite cannot create adequate space for the entire job. • The default retention period is 0 days. If you have not specified another retention period, all existing copies are overwritten each time the backup copy job runs.
Retention Period	Use with the Overwrite option. Check the **Retention Period** box to specify the length of time a copy is retained before it can be deleted or overwritten. To define the retention period, enter a number and select Days, Weeks, Months, or Years. Example: enter **2** and select **Weeks** to retain copies for at least 2 weeks. If you do not use the Retention Period option, copy jobs fail if there is insufficient space available on the tape(s). Modifying the retention period does not change the retention period of existing cold copies. The new setting is applied to subsequent backup copies only.
Schedule Details	Select **Now** to run the job immediately or select **Create a Schedule** and define additional settings.

Option	Description
Select scheduling mode	In most cases, the **Standard** scheduling mode can be used to create the schedule. If you need more granularity, choose the **Custom** mode to create a custom calendar. For details, see "Using the Custom scheduling mode in the Create Backup Copy Job dialog".
Start Date	Date when the schedule will run for the first time.
Days	Days when the schedule will run.
Start Time	Time of day when the schedule will run. (This is also the Start Time used by the *Recurs every* option.)
Recurs every / Until	Use to run the schedule every *N* hours until the specified time. Check the **Recurs every** box, select a frequency (number of hours), and specify an end time.
Date Range	Choose to copy the most recent backups or specify a date range of backups to copy.

Using the Custom scheduling mode in the Create Backup Copy Job dialog

In most cases, the standard scheduling mode can be used to create backup copy schedules. If you need more granularity, you can opt to use the Custom mode for cold backup copy job schedules.

After adding assets to the cold backup copy job, the Define Job Settings step displays. To use the Custom backup mode:

Note: There are many configuration options in the Define Job Settings step. This procedure describes settings that are required to use the Custom scheduling mode. For a complete description of these settings, see one of the following: "Backup copy job settings" or "Tape backup copy job settings".

1. In the Options area:
   - Select **Backup Copy (Cold)** in the **Backup Copy Type** list.
   - Select your target from the **Backup Copy Target** list.
2. Select **Custom** in the **Select scheduling mode** list.
3. Click the calendar icon.

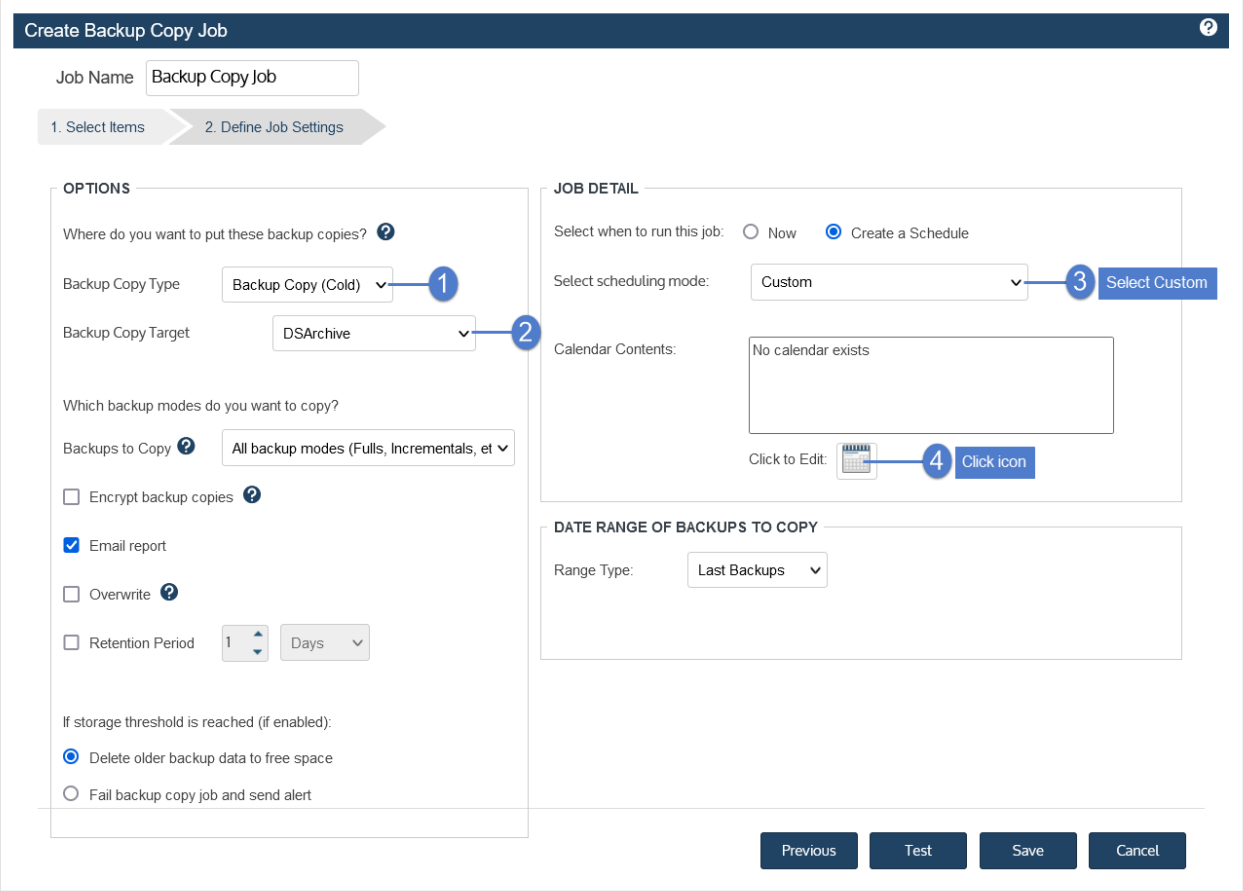

4 In the Calendar dialog, select **Backup Copy** and drag it to a day on the calendar. (You cannot drag to a day in the past.)

5 In the Add Backup Copy dialog, modify settings and click **Save**.

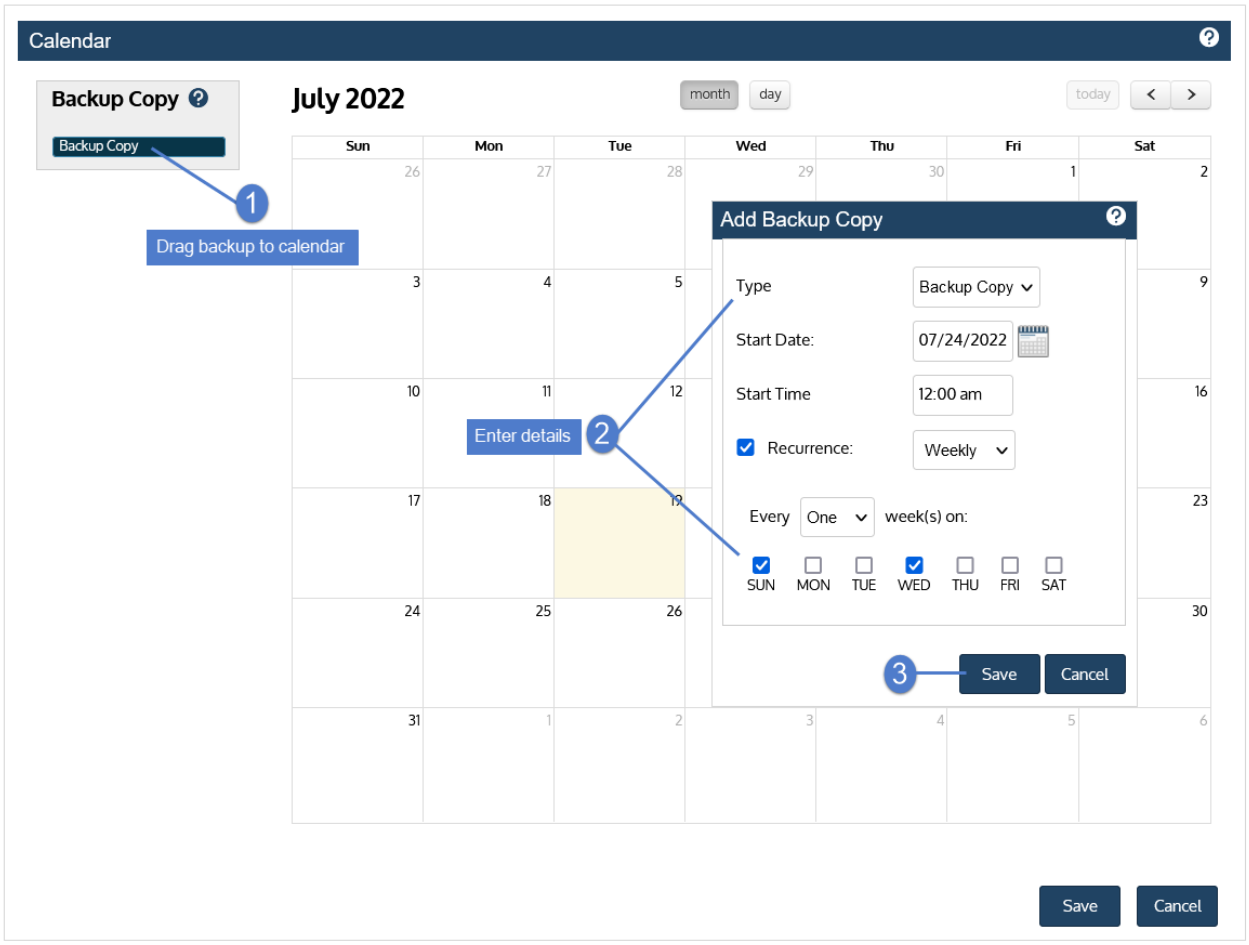

6 Click **Save** to save the settings and close the Calendar dialog.

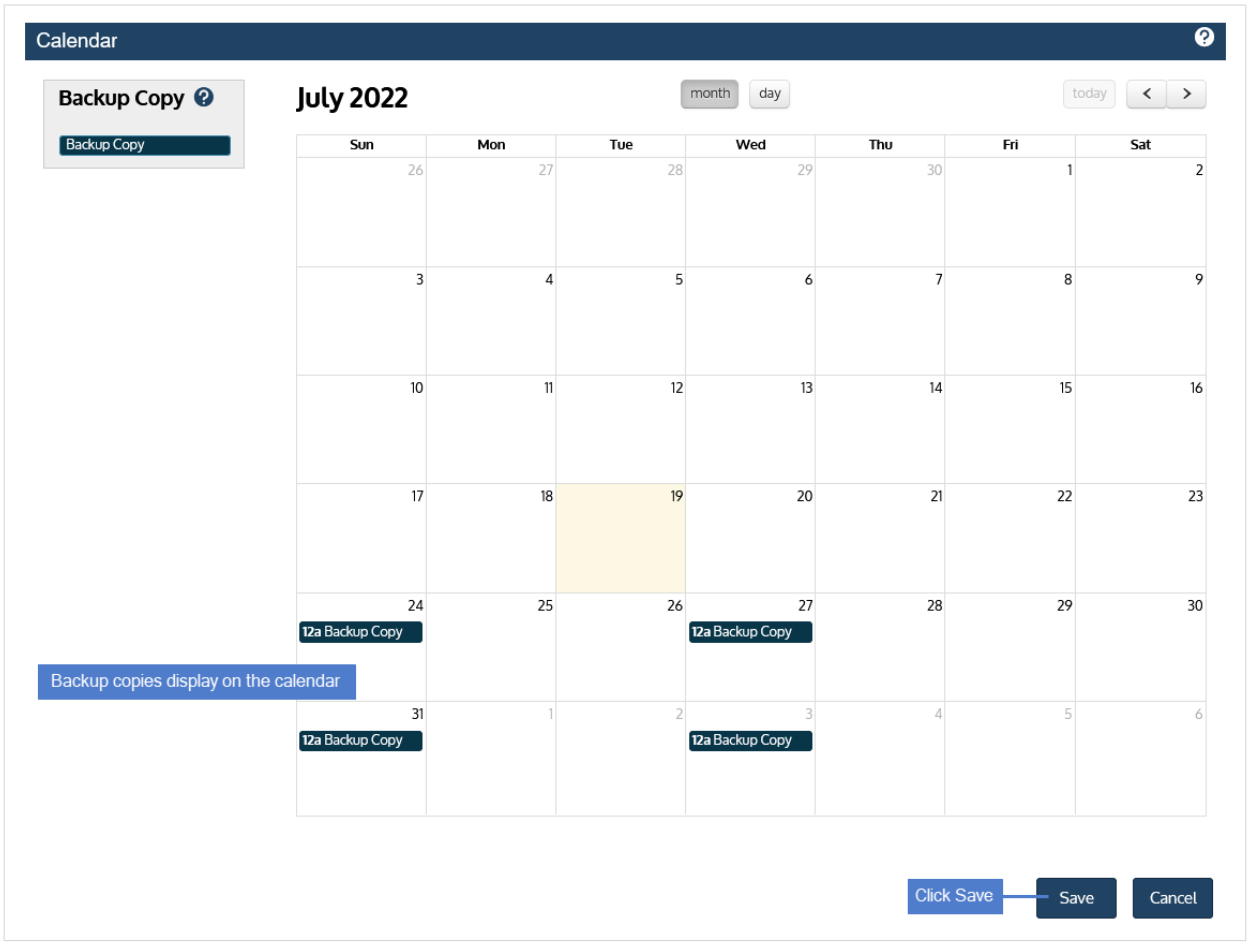

7 Click **Save** to save the schedule.

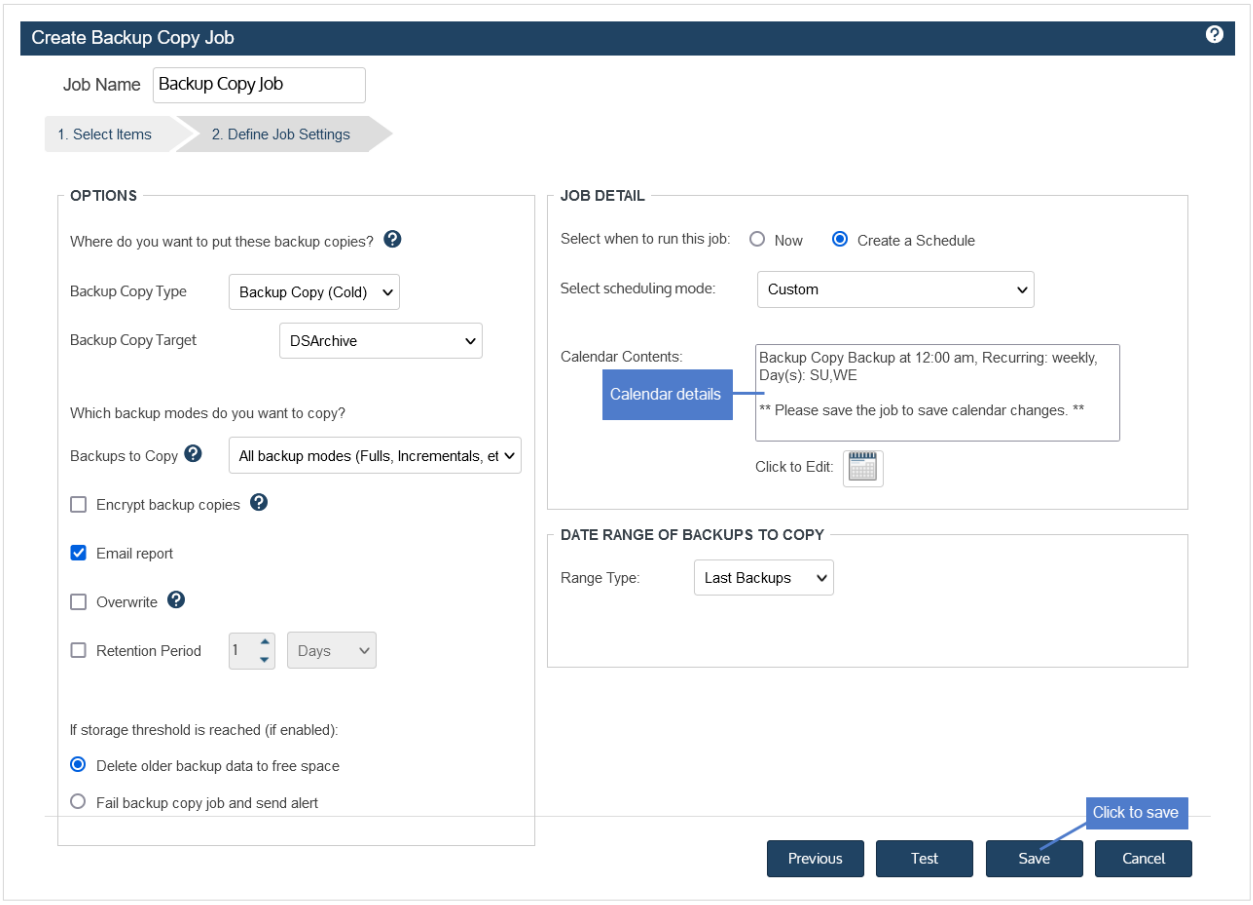

Creating SLA policies

Use the SLA Policy Automation feature to quickly implement a protection strategy that aligns with your business continuity plan. Simply create an SLA policy and the appliance automatically creates the backup and backup copy jobs needed for the RPO and retention settings you specified. For details on how SLA policies work, see "Methods for scheduling jobs" on page 428.

SLA policy requirements

The following requirements and considerations apply to SLA Policy Automation:

- The backup appliance must be running version 10.0 or higher (10.3 or higher for Windows image-level assets).
- The following asset types are supported: file-level Windows, file-level Linux, image-level Windows, and VMware, Hyper-V, and AHV virtual machines. You must manually create job schedules for other asset types.
- You must log in to the backup appliance directly to create SLA policies. You cannot create a policy for a managed appliance by logging in to its manager appliance.
- An asset can be assigned to only one SLA policy.

- A file-level asset can be assigned to one SLA policy and/or to one or more manually created backup schedules.
- An image-level Windows asset can be assigned to one SLA policy and/or to one or more manually created backup schedules.
- A virtual machine asset can be assigned either to one SLA policy or to one manually created backup schedule (to ensure that the VM exists in only one backup schedule).
- SLA policy schedules do not support the auto-include assets option. You must manually create a backup schedule to use this option.
- The hot backup copy option is supported only if a Unitrends appliance or the Unitrends Cloud has been added as a backup copy target. (For details on adding a hot target, see "Backup copy targets" on page 214.)
- The cold backup copy option is supported only if the cold target has been added to the appliance and this target is one of the following types: third-party cloud, NAS, or iSCSI. (For details on adding a cold target, see "Backup copy targets" on page 214.)
- If multiple cold targets exist, the policy copies backups to the one that was added first. To copy to a different cold target, you must manually create a backup copy job instead.
- Do not directly edit job schedules that were created by an SLA policy. Instead, modify the SLA policy itself. The appliance automatically modifies the policy's schedules based on the changes that you make. On the Job Manager tab, SLA policy schedule names display with the prefix *_SLA*, so you can easily distinguish them from manually created schedules.

SLA policy procedures

Use these procedures to create SLA policies:

- "To create an SLA policy for Windows and Linux file-level assets" on page 537
- "To create an SLA policy for Windows image-level assets" on page 545
- "To create an SLA policy for VMware assets" on page 551
- "To create an SLA policy for Hyper-V assets" on page 556
- "To create an SLA policy for AHV assets" on page 559

To create an SLA policy for Windows and Linux file-level assets

Note: A file-level asset can be assigned to one SLA policy and/or to one or more manually created backup schedules.

1 Select **Jobs > SLA Policies > Create Policy**.

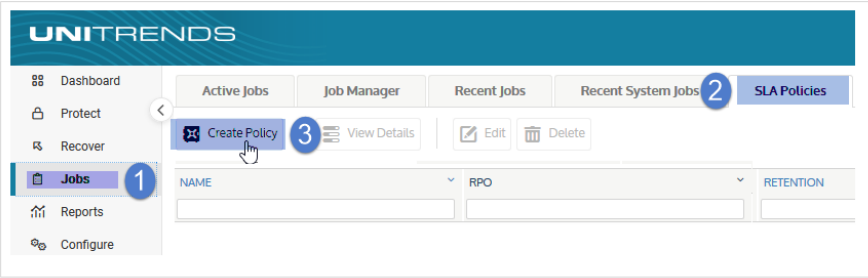

2 Select **File Level** in the **What do you want to protect?** list.

3 In the Inventory tree, check boxes to select the asset(s) you want to protect. Selected assets display in the Job Inventory Settings area.

To locate an asset by name, use the **Search** field below.

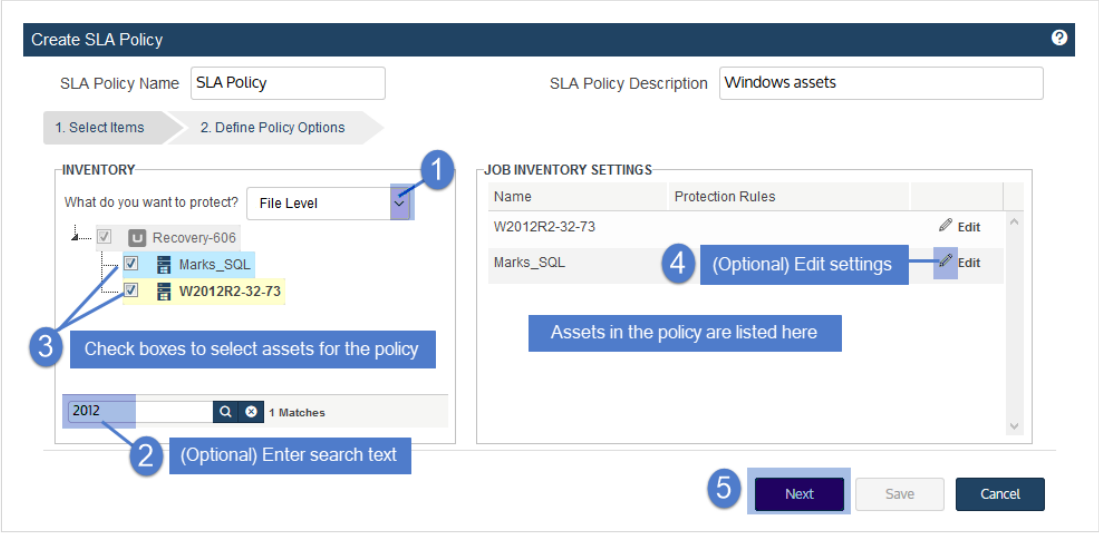

4 (Optional) Select an asset and click **Edit** to apply options, such as data to include or exclude and commands to run pre- and/or post-backup. Click **Save** to retain any changes.

See the following for details and considerations:

File-level setting	Description
General considerations for including or excluding data from an asset's backups	Review the following before specifying data to include or exclude: • When creating a schedule, any inclusions or exclusions you add are applied to jobs run by that schedule only. Inclusions and exclusions are not applied automatically in other cases.

File-level setting	Description
	• Running an on-demand backup of the asset does not automatically apply any inclusions or exclusions specified in the asset's schedule. To run an on-demand backup, do one of the following: – Create a one-time job that has the same inclusions and exclusions as in the schedule. – Manually run the schedule (select the schedule under **Jobs > Job Manager** and click **Run**). – Run a one-time Selective backup (so that a new full is not created). • If you specify both files to include and files to exclude, the inclusion is applied first. Any exclusions are then applied to the subset of included files.
Inclusion tab	Click to specify files, folders, or volumes to include in backups of this asset. • Data that does not meet the criteria you specify here is NOT included in the backup. • Type in the full path (e.g., **C:/Documents**) or **Browse** the asset to specify data to include. (Wildcards are not supported.) • If you are running a full backup and include files or folders in the system drive (typically C:), do *not* check the System State box on the Advanced tab. Full backups fail if system state is excluded. • Run a new full backup upon creating or modifying included files. Example: 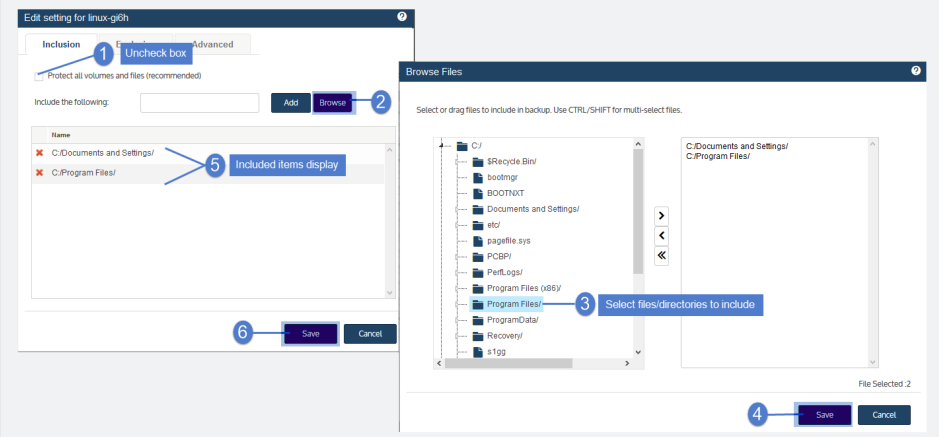

File-level setting	Description
Exclusion tab	Click to specify files, folders, or volumes to exclude from backups of this asset. • Data that does not meet the criteria you specify here IS included in the backup. • To specify files to exclude, do any of the following: – Type in the full path (e.g., **C:/Documents**). – **Browse** the asset. – Enter a selection pattern. Wildcards are supported for Windows assets. Wildcards are not supported for these asset types: Linux, Unix, and NAS. See these rows below for usage examples: "Wildcard * usage", "Wildcard ? usage", and "Multiple wildcards". • Run a new full backup upon creating or modifying excluded files. Example: 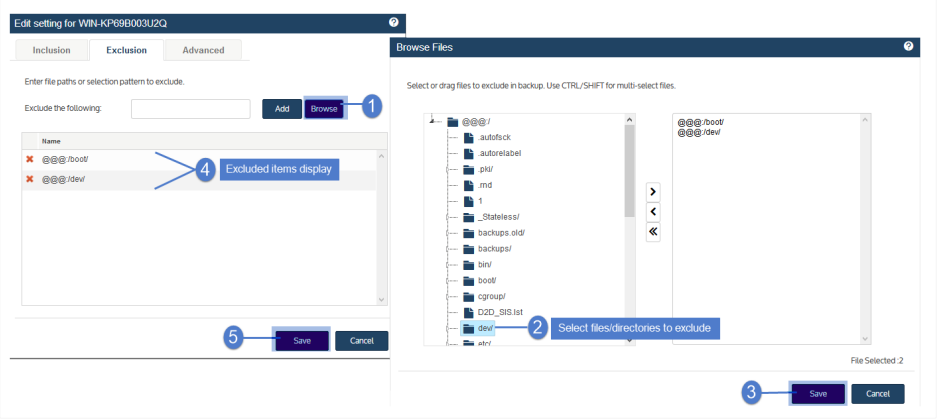
Wildcard * usage	An example of how to exclude all files with zero or more characters that match exclusion pattern: ***.txt** An example of how to exclude directories with zero or more characters and their contents within a specified path that match the exclusion pattern: **C:/windows/sys*** Limitations: • ***folder_abc** cannot be used to exclude all folders that match *folder_abc* on the protected asset. The full path must be provided. • If an entire directory is excluded, the directory name will still appear in the

File-level setting	Description
	backup; however, its contents will be empty. • Multiple wildcard matches like the following are not supported: **C:**\abc.txt**
Wildcard ? usage	An example of how to exclude all files within specified path that matches a single character within exclusion pattern: **C:/PCBP/Lists.dir/pro_client?.spr** An example of how to exclude all directories and their contents within specified path that matches a single character within exclusion pattern: **C:/Programfiles/Case?/** Limitation: If an entire directory is excluded, the directory name itself will still appear in the backup; however its contents will be empty.
Multiple wildcards	An example that uses multiple “?” wildcards and only one * wildcard: **C:/?Log?/*.logs** Limitation: If an entire directory is excluded, the directory name itself will still appear in the backup; however its contents will be empty.
Advanced tab	Use this tab to specify advanced options. See these rows below for details: "Advanced Exclusions", "Command to run Pre-Backup", and "Command to run Post-Backup". Example: 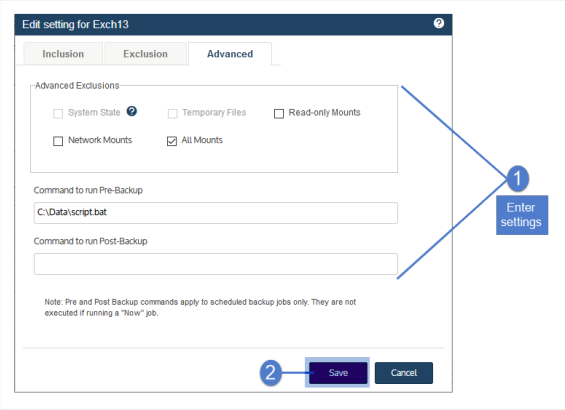
Advanced Exclusions	Check one or more boxes to exclude any of the following: system state, temporary files, read-only mounts, network mounts, or all mounts. Consider the following before applying advanced exclusions:

File-level setting	Description
	• To perform bare metal recovery or use Windows replicas, the following must be included in the backup: system state and all boot and critical system (OS) disks/volumes. If you need these features for the asset, do not specify data to include or exclude unless you are sure these disks/volumes will be included. • If you are running a full backup and have selected files or folders in the system drive (typically C:) on the Inclusion tab, do *not* check the System State box on the Advanced tab. Full backups fail if system state is excluded. • Creating aliases for an asset - Adhere to the following when creating aliases for an asset: – You must include the system state on the asset whose backups contain the boot and critical OS volumes. – You must exclude the system state on the other aliased assets. This approach ensures you can perform bare metal recovery of the asset. – Only one asset can include the system state. Disaster recovery of the asset fails if the system state is not included with the boot and OS volume or if the system state is included on aliased assets that do not include the boot and OS volume. **IMPORTANT!** For Windows assets, the backup must contain the system state, boot disk and any other system critical volumes to use the integrated bare metal recovery and Windows replica features. Be sure one of the aliased assets contains all of these disks to use these features.
Command to run Pre-Backup	To run a command or script on the asset before a scheduled backup starts, enter the full path to the command or script in the **Command to run Pre-Backup** field. For example, *C:\Data\script.bat* or */usr/jsmith/script.sh*. **Note:** Pre-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).
Command to run Post-Backup	To run a command or script on the asset after a scheduled backup completes, enter the full path to the command or script in the **Command to run Post-Backup** field. For example, *C:\Data\script.bat* or */usr/jsmith/script.sh*.

File-level setting	Description
	Note: Post-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).

5 Click **Next**.

6 Define the remaining Policy Options, then click **Save**. See "SLA policy options" on page 562 for descriptions of each option.

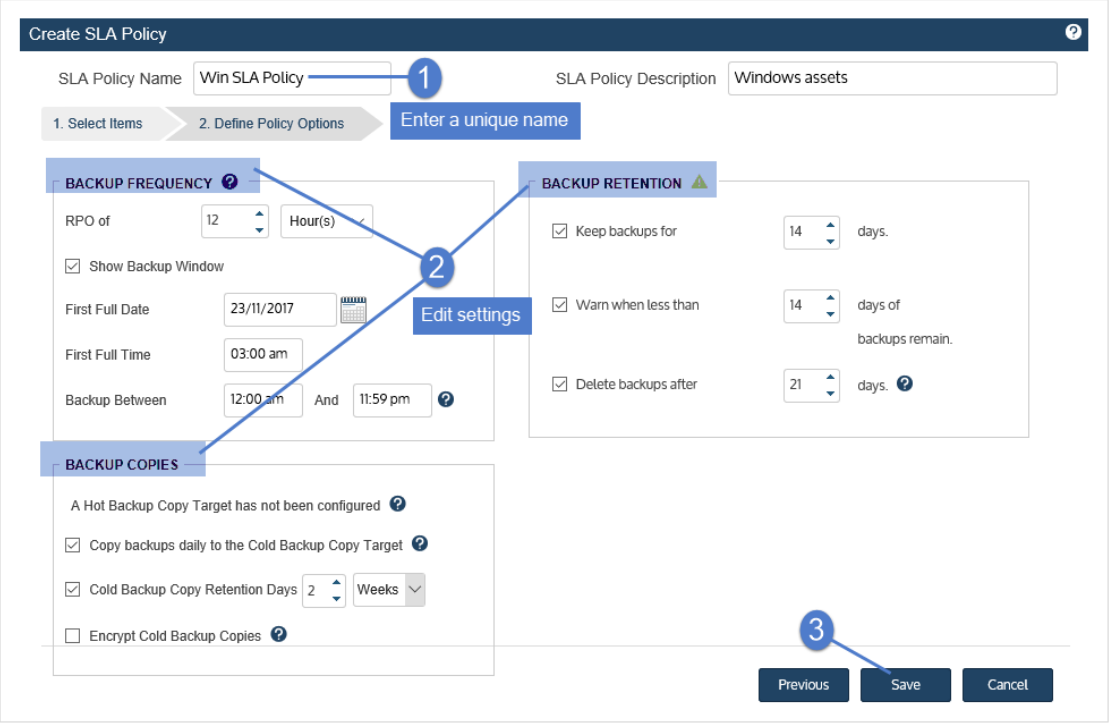

7 The appliance creates the policy and related jobs. Click **Close** to close the status message.

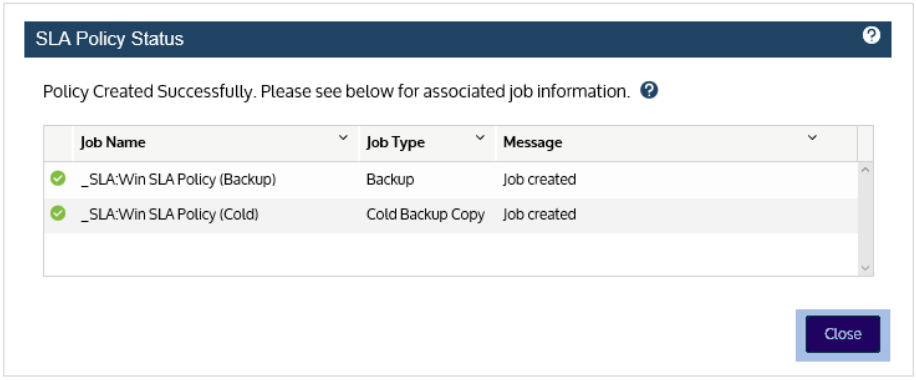

The policy displays on the SLA Policies tab:

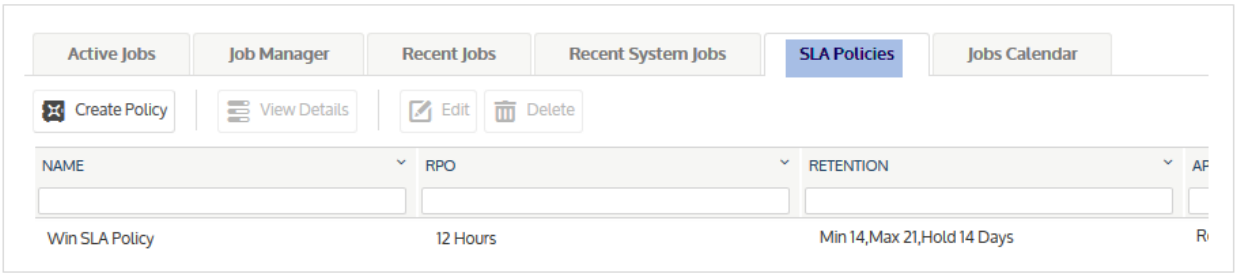

Jobs display on the Job Manager tab and are named with the prefix *_SLA:*

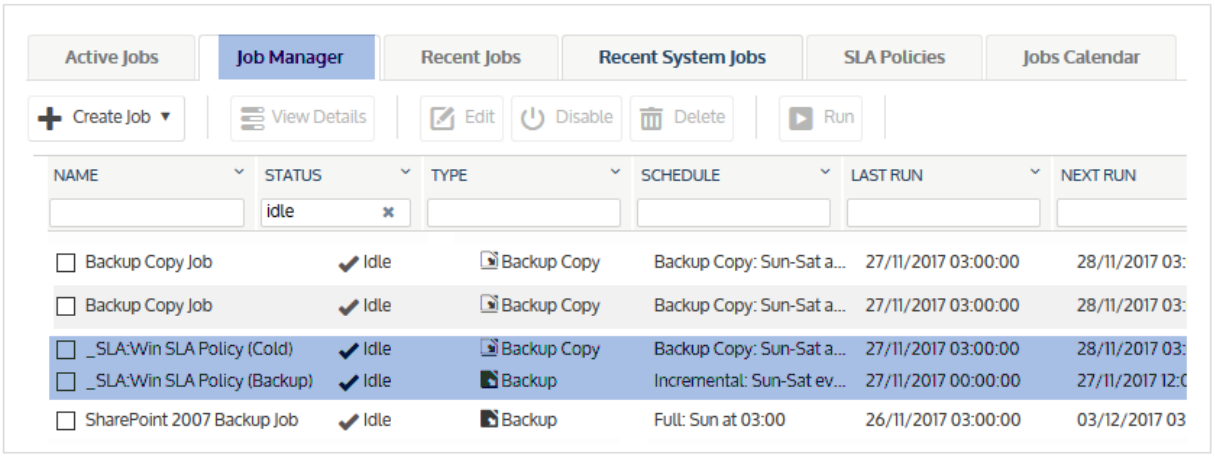

To create an SLA policy for Windows image-level assets

Note: An image-level asset can be assigned to one SLA policy and/or to one or more manually created backup schedules.

1 Select **Jobs > SLA Policies > Create Policy**.

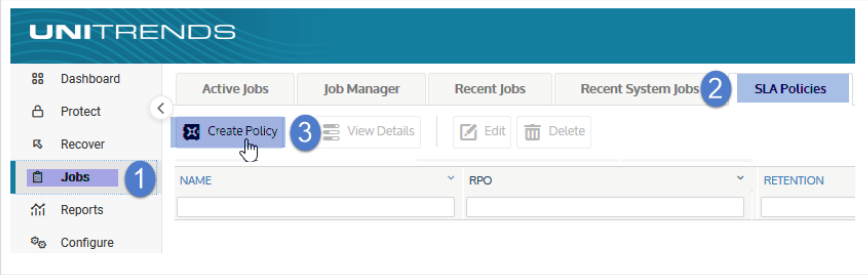

2 Select **Image Level** in the **What do you want to protect?** list.

3 In the Inventory tree, check boxes to select the asset(s) you want to protect. Selected assets display in the Job Inventory Settings area.

To locate an asset by name, use the **Search** field below.

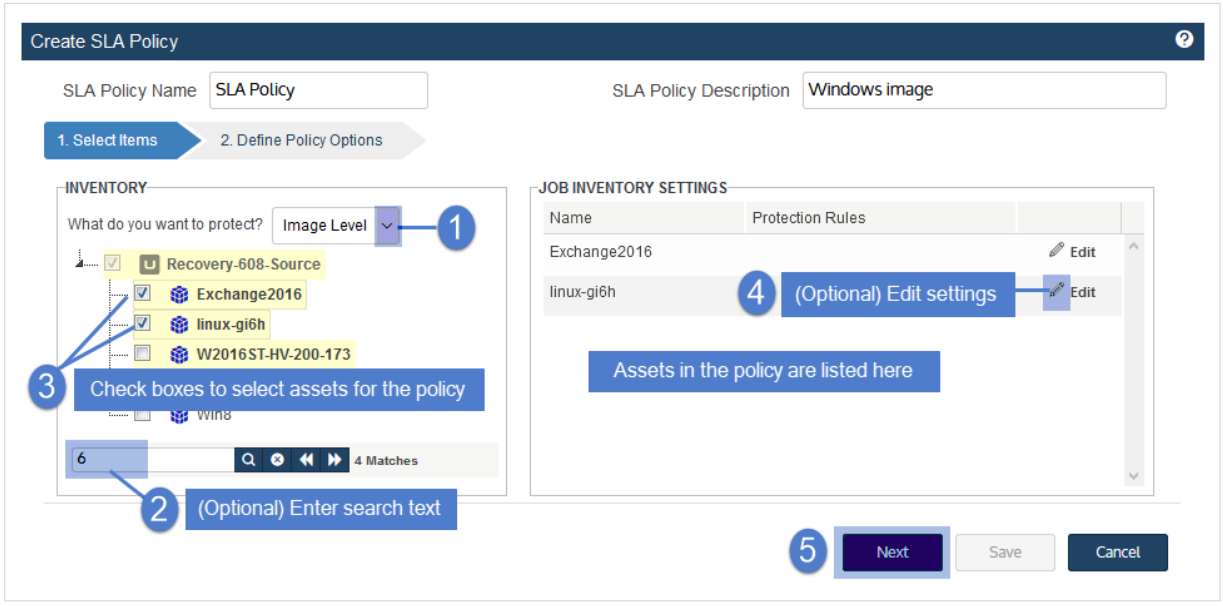

4 (Optional) Select an asset and click **Edit** to apply options, such as data to include or exclude and commands to run pre- and/or post-backup. Click **Save** to retain any changes.

See the following for details and considerations:

Notes:

- Critical system volumes are required for the image-level replicas feature and to recover the entire asset . Use care when omitting volumes from backup.
- When you recover the entire asset, any existing data on the target is overwritten or deleted. Volumes on the target disk that were excluded from backup may also be overwritten. For details, see "Windows unified bare metal recovery" on page 1209.
- To recover a SQL server, the master, model, and msdb system databases must be present in the image-level backup of the Windows asset. (These are included by default. If you want the recovered asset to include a hosted SQL application, use care not to exclude these system databases from the image-level backup.)
- Image-level protection is not supported for read-only disks. You must exclude all volumes on read-only disks from the backup job or run file-level backups. Image-level backups fail if read-only volumes have not been excluded.
- Removable media is automatically excluded from image-level backups. (You do not need to exclude volumes on a read-only disk that resides on removable media.)

See the following for details:

Image-level setting	Description
General considerations	Review the following before specifying volumes to include or exclude: • When creating a schedule, any inclusions or exclusions you add are applied to jobs run by that schedule only. Inclusions and exclusions are not applied automatically in other cases. • Running an on-demand backup of the asset does not automatically apply any inclusions or exclusions specified in the asset's schedule. To run an on-demand backup, do one of the following: – Create a one-time job that has the same inclusions and exclusions as in the schedule. – Manually run the schedule (select the schedule under **Jobs > Job Manager** and click **Run**).
Inclusion tab	Click to specify volumes to include in backups of this asset. • Any volumes that you do not select here are NOT included in the backup. • Type in the volume (e.g., **C:/**) or **Browse** the asset to select volumes to include. • Run a new full backup upon creating or modifying included

Image-level setting	Description
	volumes.
Inclusion example: 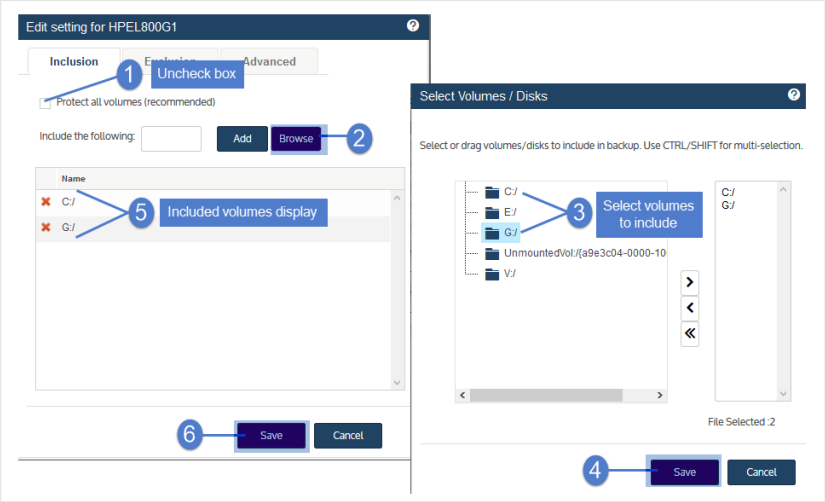	
Exclusion tab	Click to specify volumes to exclude from backups of this asset. • Any volumes that you do not select here are included in the backup. • Type in the volume (e.g., **C:/**) or **Browse** the asset to select volumes to exclude. • Run a new full backup upon creating or modifying excluded volumes.
Exclusion example:	

Image-level setting	Description
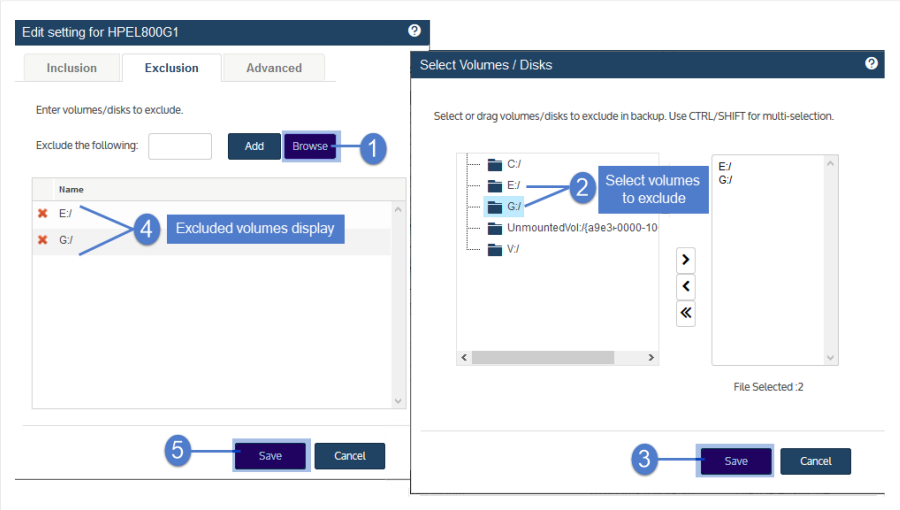	
Advanced tab	Use this tab to specify commands to run before or after the backup job. (Advanced Exclusion options are not supported for image-level backups.) See these rows below for details: "Command to run Pre-Backup" and "Command to run Post-Backup". Example: 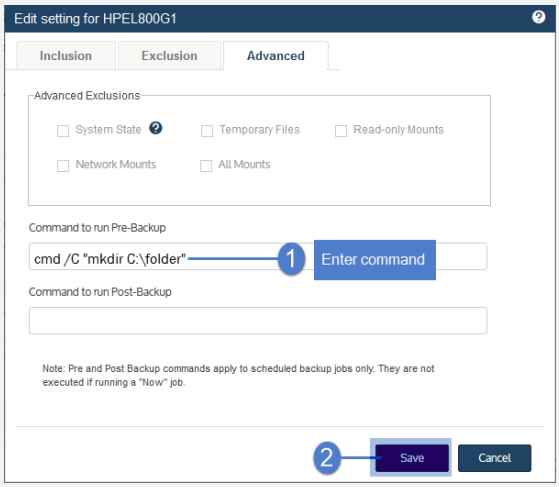
Command to run Pre-Backup	For image-level backup schedules, you cannot enter command syntax as you would inside the Windows Command Prompt. Instead, you must use this format: *cmd /<argument> "<commandSyntax>"*

Image-level setting	Description
	For example: *cmd /C "mkdir C:\folder"* To run a command or script on the asset before a scheduled backup starts, enter the command in the **Command to run Pre-Backup** field. **Note:** Pre-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).
Command to run Post-Backup	For image-level backup schedules, you cannot enter command syntax as you would inside the Windows Command Prompt. Instead, you must use this format: *cmd /<argument> "<commandSyntax>"* For example: *cmd /C "mkdir C:\folder"* To run a command or script on the asset after a scheduled backup completes, enter the command in the **Command to run Post-Backup** field. **Note:** Post-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).

5 Click **Next**.

6 Define the remaining Policy Options, then click **Save**. See "SLA policy options" on page 562 for descriptions of each option.

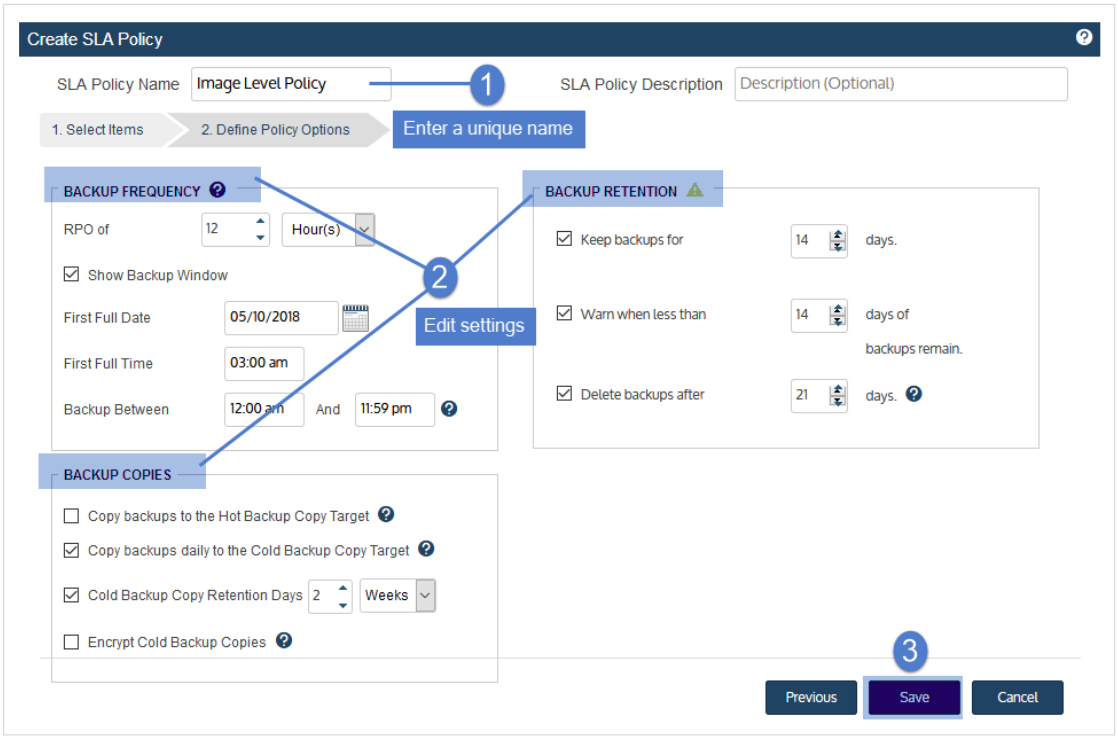

7 The appliance creates the policy and related jobs. Click **Close** to close the status message.

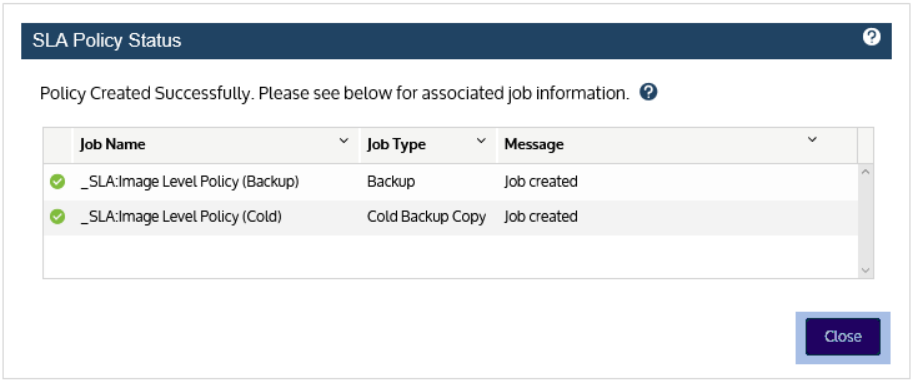

The policy displays on the SLA Policies tab:

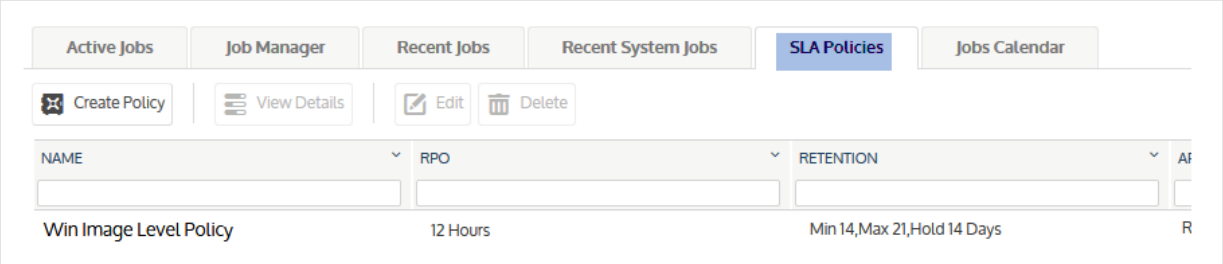

Jobs display on the Job Manager tab and are named with the prefix *_SLA:*

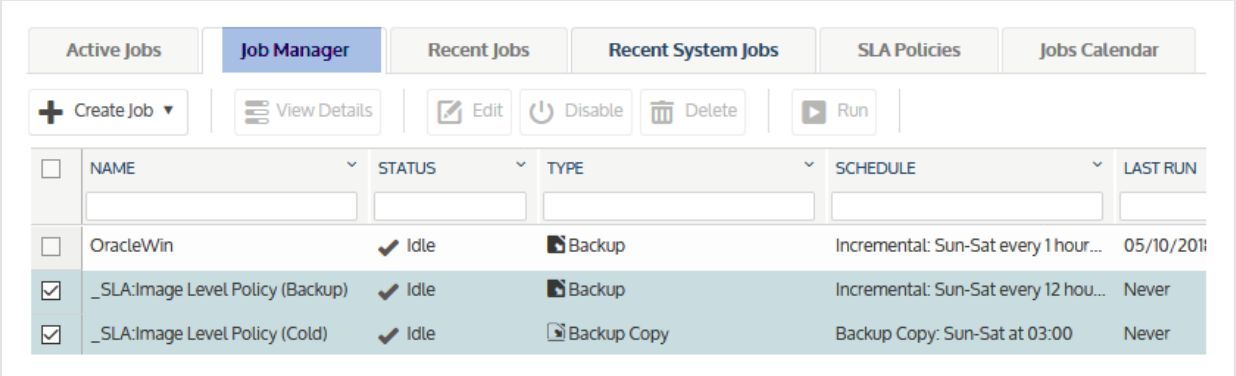

To create an SLA policy for VMware assets

Notes:

- A VMware asset can be assigned either to one SLA policy or to one manually created backup schedule (to ensure that the VM exists in only one backup schedule).
- The policy can contain VMs that are managed by a single vCenter or ESX server.
- A VM can be assigned to only one hot backup copy schedule. The policy does not create a hot backup copy schedule if any of its VMs exist in another hot backup copy schedule.
- To access newly added virtual machines, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

1 Select **Jobs > SLA Policies > Create Policy**.

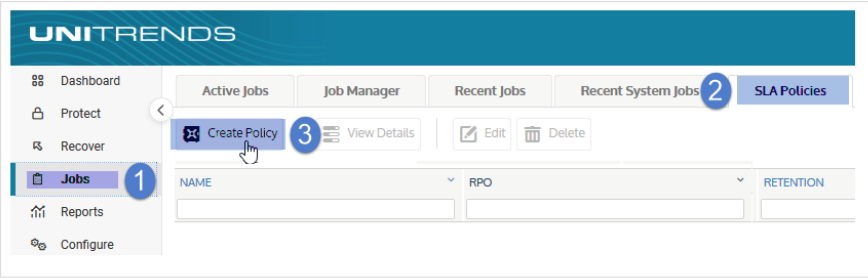

2 Select **VMware** in the **What do you want to protect?** list.

3 In the Inventory tree, expand the virtual host and check boxes to select virtual machines to protect. Selected VMs display in the Job Inventory Settings area.

- To locate an asset by name, use the **Search** field below.
- To protect all VMs, select the virtual host.

4 (Optional) Edit Job Inventory Settings to exclude VM disks from backup:

- Locate the VM in the Job Inventory Settings list.
- Click **Edit** to specify disks to exclude.
- Click **Save** to retain any changes.

> **Note:** To recover the entire virtual machine requires critical system volumes. Use care when omitting disks from backup.

5 Click **Next**.

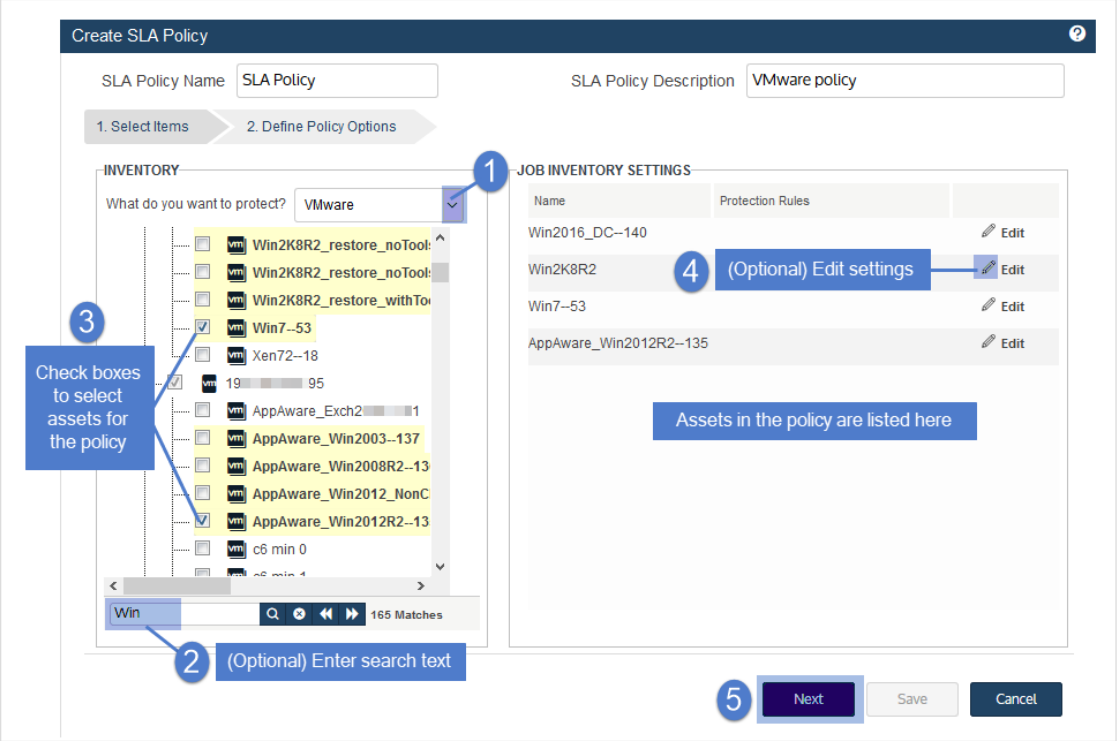

6 Define the remaining Policy Options, then click **Save**. See "SLA policy options" on page 562 for descriptions of each option.

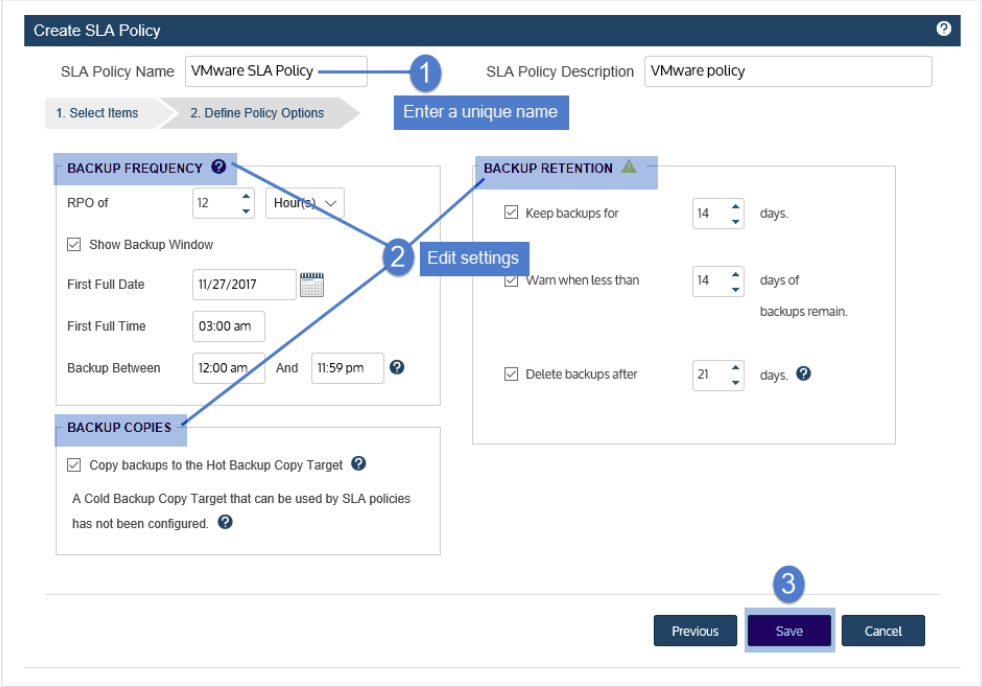

7 The appliance creates the policy and related jobs. Click **Close** to close the status message.

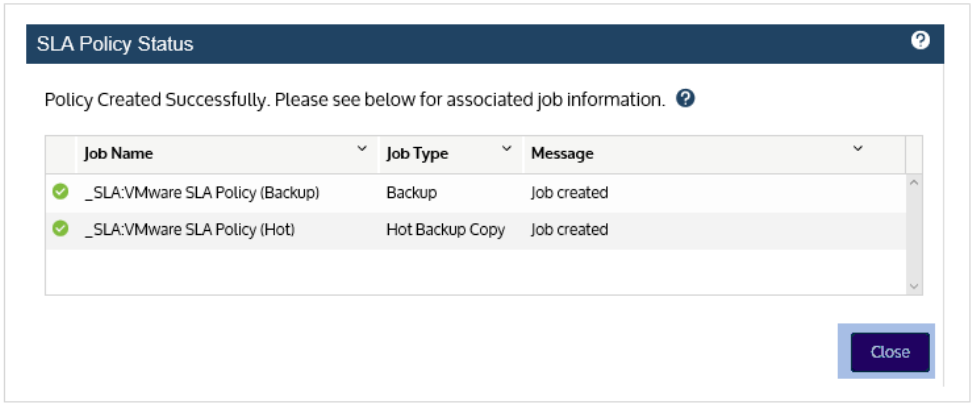

The policy displays on the SLA Policies tab:

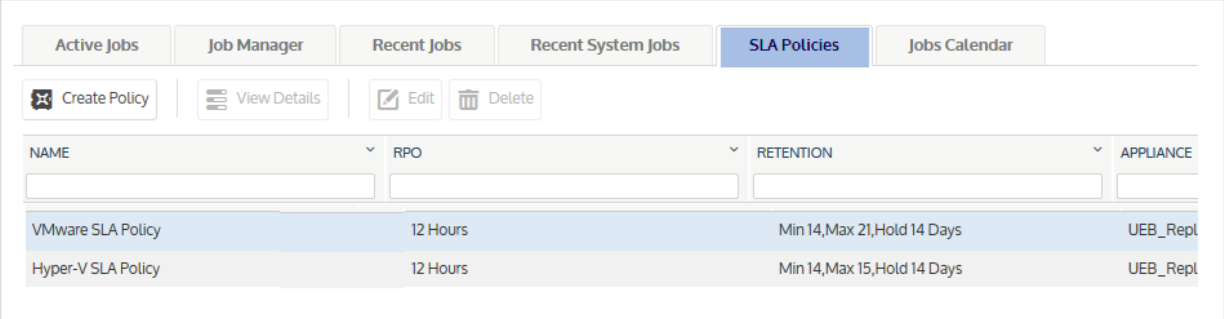

Jobs display on the Job Manager tab and are named with the prefix *_SLA:*

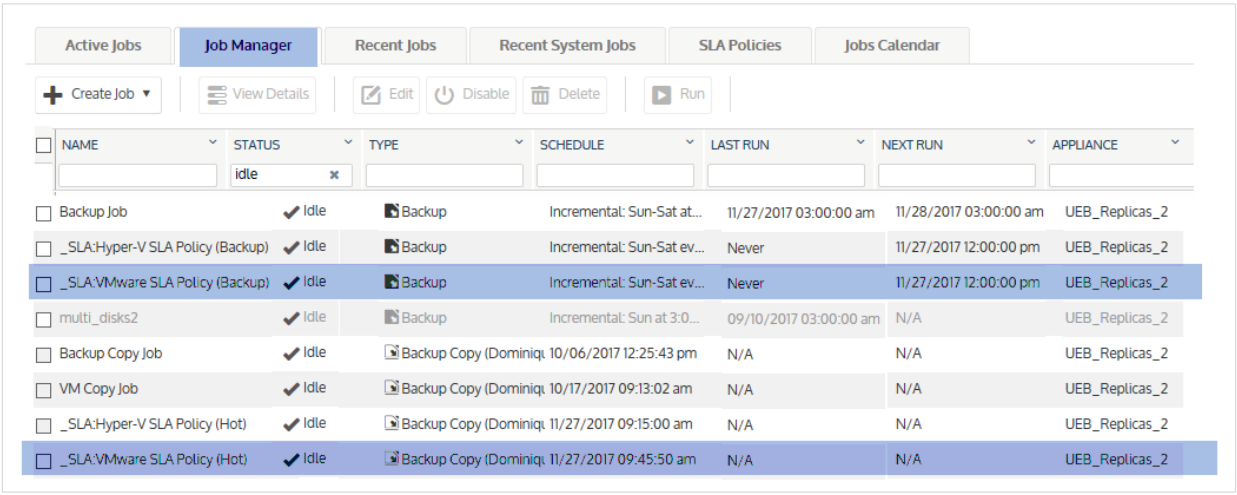

To create an SLA policy for Hyper-V assets

Notes:

- A Hyper-V asset can be assigned either to one SLA policy or to one manually created backup schedule (to ensure that the VM exists in only one backup schedule).
- The policy can contain VMs that are managed by a single Hyper-V server.
- A VM can be assigned to only one hot backup copy schedule. The policy does not create a hot backup copy schedule if any of its VMs exist in another hot backup copy schedule.
- To access newly added virtual machines, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

1 Select **Jobs > SLA Policies > Create Policy**.

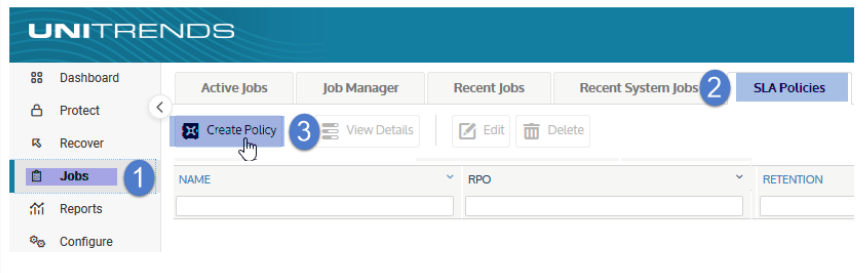

2 Select **Hyper-V** in the **What do you want to protect?** list.

3 In the Inventory tree, expand the Hyper-V server and host application, then check boxes to select virtual machines to protect.

- To locate an asset by name, use the **Search** field below.
- To protect all VMs, select the Hyper-V host.

4 Click **Next**.

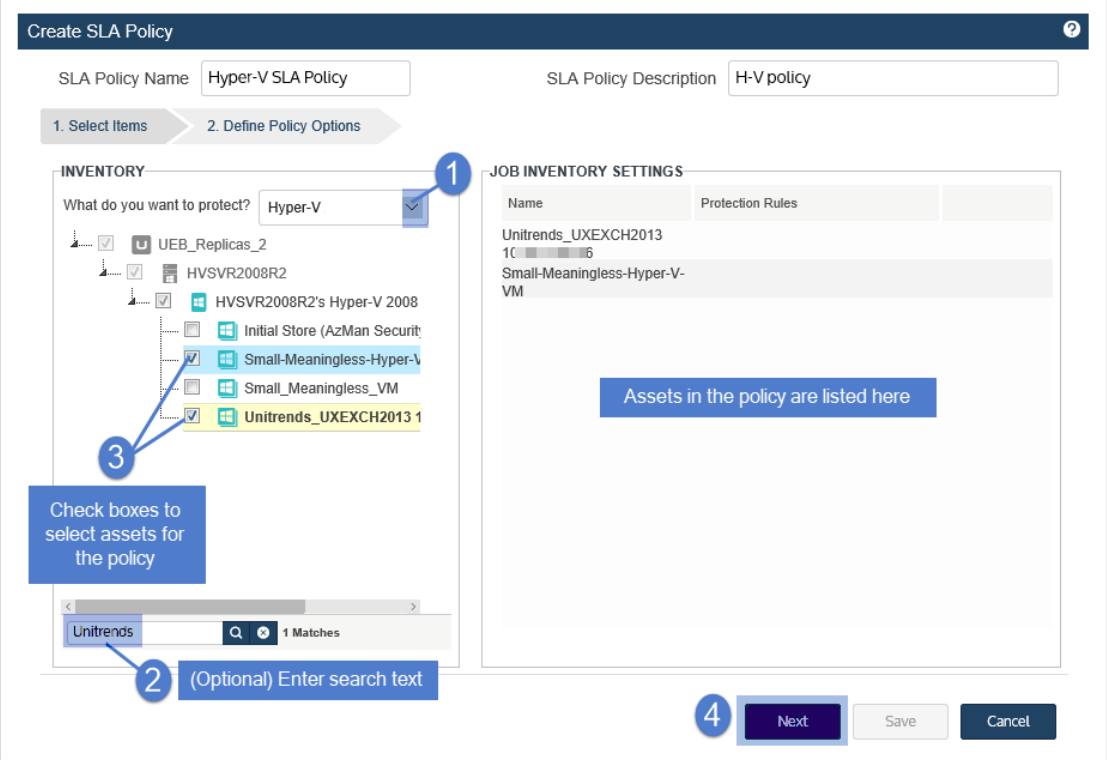

5 Define the remaining Policy Options, then click **Save**. See "SLA policy options" on page 562 for descriptions of each option.

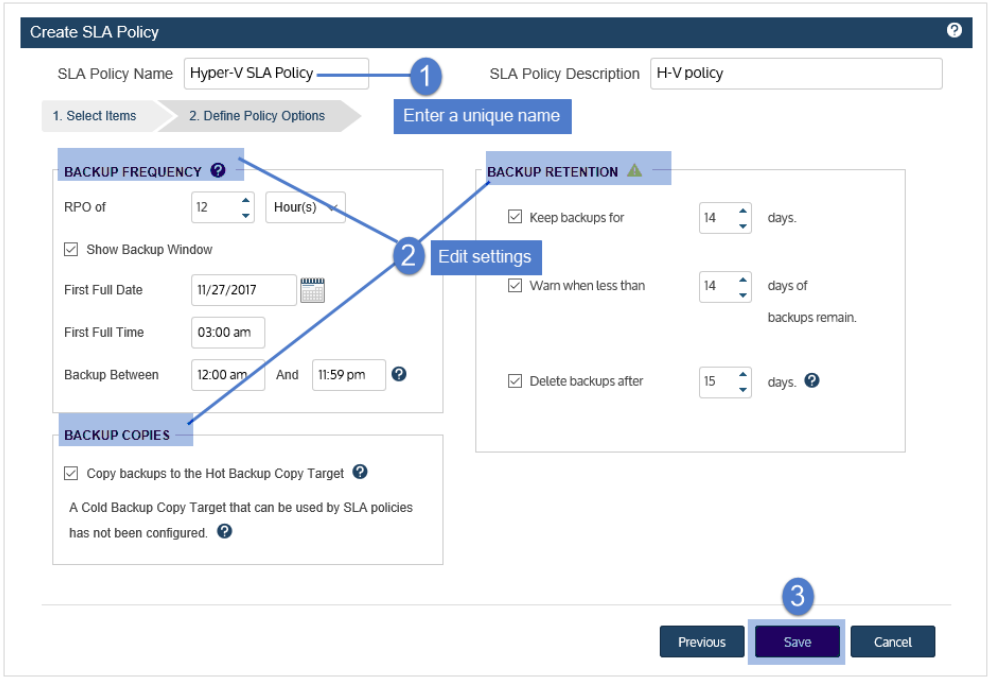

6 The appliance creates the policy and related jobs. Click **Close** to close the status message.

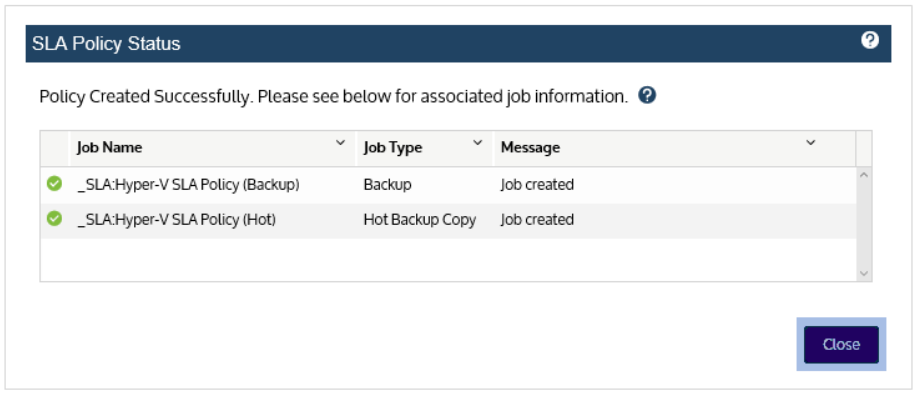

The policy displays on the SLA Policies tab:

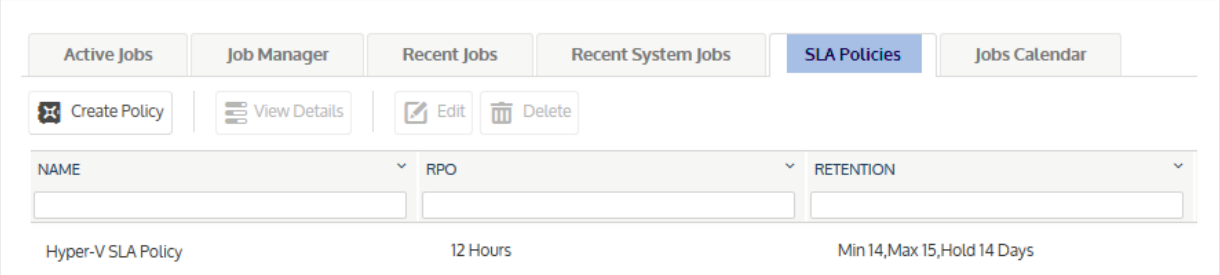

Jobs display on the Job Manager tab and are named with the prefix *_SLA:*

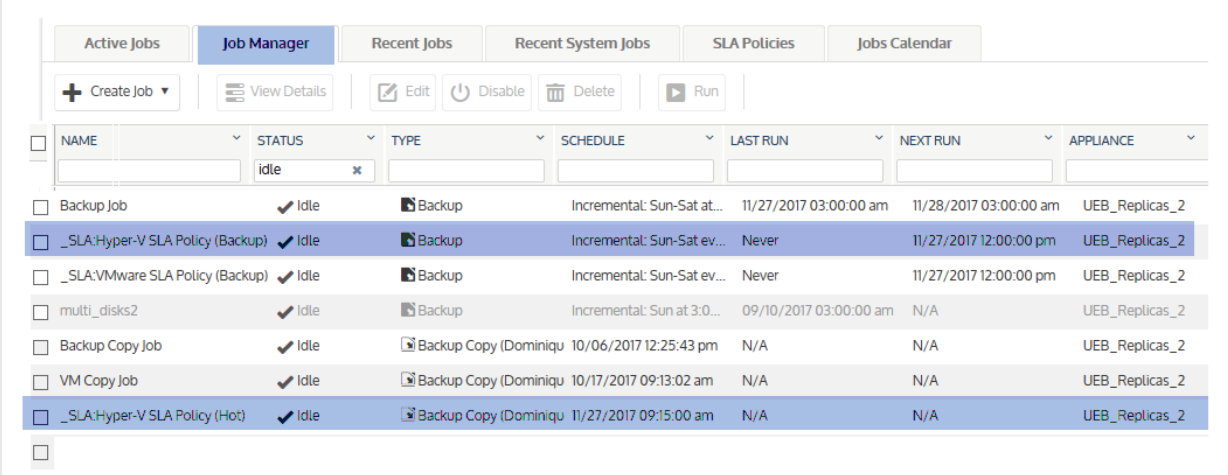

To create an SLA policy for AHV assets

Notes:

- An AHV asset can be assigned either to one SLA policy or to one manually created backup schedule (to ensure that the VM exists in only one backup schedule).
- The policy can contain VMs that are managed by a single AHV cluster.
- A VM can be assigned to only one hot backup copy schedule. The policy does not create a hot backup copy schedule if any of its VMs exist in another hot backup copy schedule.
- To access newly added virtual machines, sync inventory before creating your job by clicking the **Gear** icon in the upper-right of the UI and selecting **Inventory Sync**.

1 Select **Jobs > SLA Policies > Create Policy**.

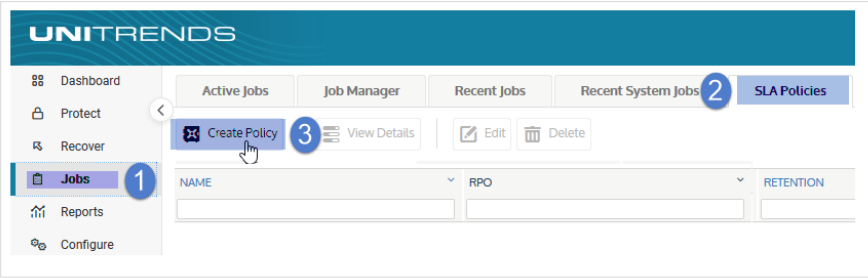

2 Select **Nutanix AHV** in the **What do you want to protect?** list.

3 In the Inventory tree, expand the AHV host cluster and check boxes to select virtual machines to protect. Selected VMs display in the Job Inventory Settings area.

- To locate an asset by name, use the **Search** field below.
- To protect all VMs, select the AHV host.

4 (Optional) Edit Job Inventory Settings to exclude VM disks from backup:

- Locate the VM in the Job Inventory Settings list.
- Click **Edit** to specify disks to exclude.
- Click **Save** to retain any changes.

> **Note:** To recover the entire virtual machine requires critical system volumes. Use care when omitting disks from backup.

5 Click **Next**.

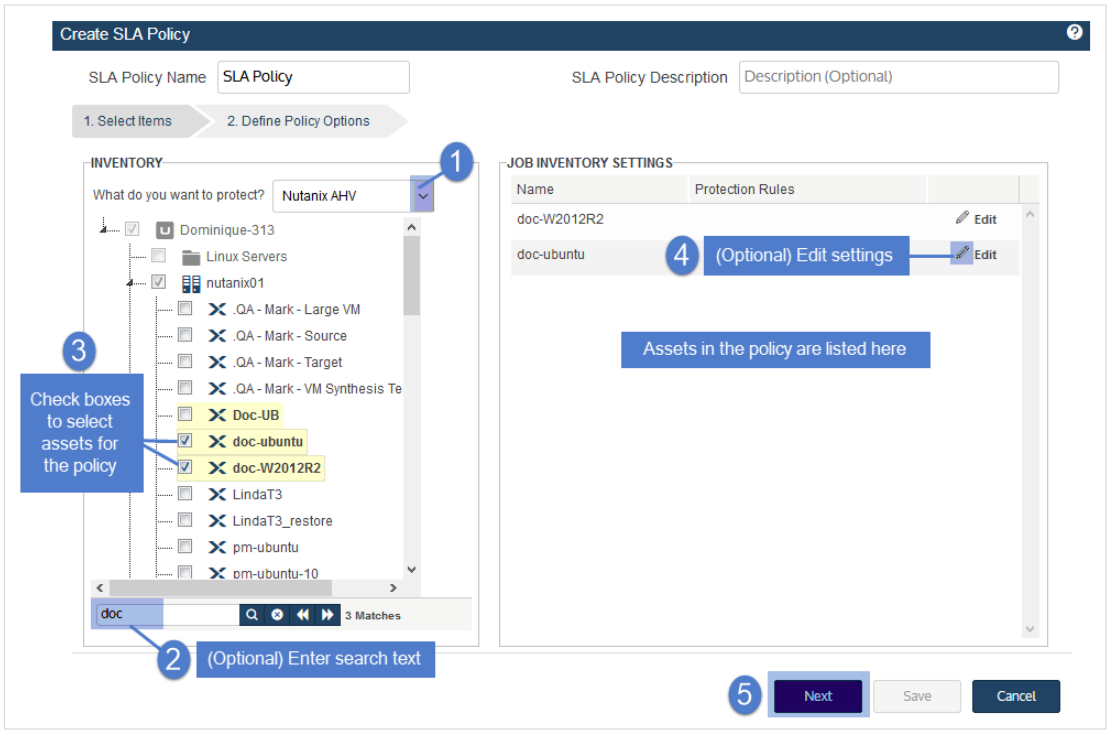

6 Define the remaining Policy Options, then click **Save**. See "SLA policy options" on page 562 for descriptions of each option.

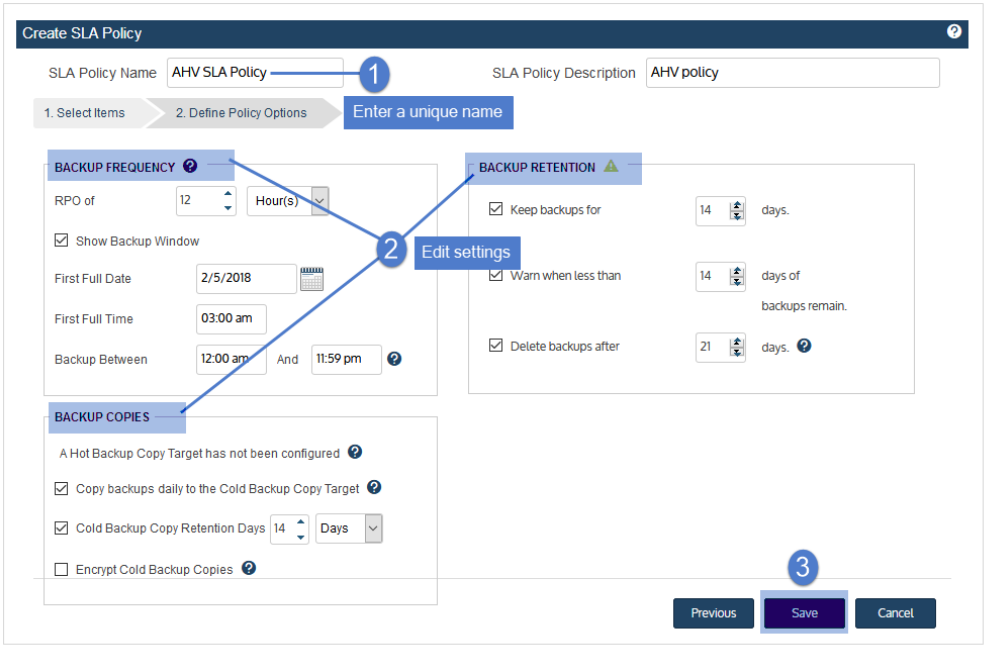

7 The appliance creates the policy and related jobs. Click **Close** to close the status message.

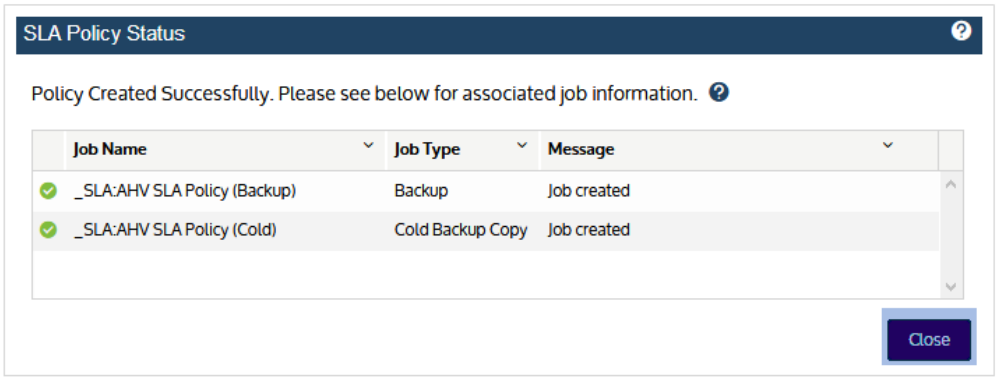

The policy displays on the SLA Policies tab:

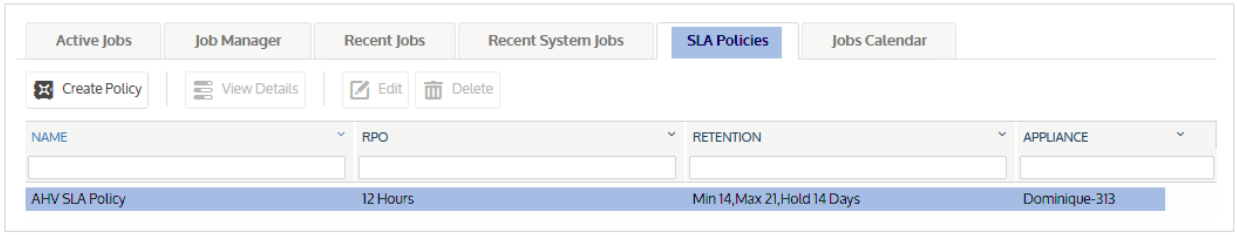

Jobs display on the Job Manager tab and are named with the prefix *_SLA:*

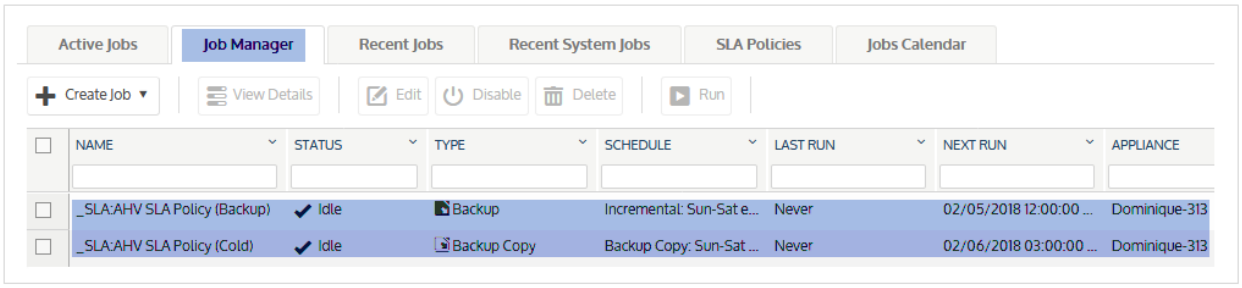

SLA policy options

The following table describes the options used to create an SLA policy.

SLA policy setting	Description
SLA Policy Name	Enter a unique name for the policy.
SLA Policy Description	(Optional) Enter a short description of the policy.
RPO	Recovery Point Objective – The maximum interval of time between backups (the maximum threshold of data loss tolerated by your business continuity plan). Determines how often backups will run.

SLA policy setting	Description
	Enter the number of hours or minutes to define the RPO interval.
Show Backup Window	Check this box to view and/or edit the following: • First Full Date – Date when the policy's first full backups will run. (Applies to assets that do not yet have a successful full backup.) • First Full Time – Time when the policy's first full backups will run. (Applies to assets that do not yet have a successful full backup.) • Backup Between – Hours of the day when backups will be taken.
Copy backups to the Hot Backup Copy Target	Check this box to copy backups to your hot backup copy target. Supported only when a Unitrends appliance or the Unitrends Cloud has been added as a backup copy target. (For details on adding a hot target, see "Backup copy targets" on page 214.)
Copy backups daily to the Cold Backup Copy Target	Check this box to copy backups to your cold backup copy target. Supported only when a cold backup copy target has been added to the backup appliance. Supported for these types of cold targets only: third-party cloud, NAS, or iSCSI. If multiple cold targets exist, the policy copies to the one that was added first. • To copy to a different cold target, manually create a backup copy job instead, as described in "Creating backup copy jobs" on page 491. • To add a cold target to the backup appliance, see "Backup copy targets" on page 214.
Cold Backup Copy Retention Days	Check this box to specify the length of time a copy must be retained before it can be deleted. To define the retention period, enter a number and select Days, Weeks, Months, or Years. For example, enter **2** and select **Weeks** to retain copies for 2 weeks.
Encrypt Cold Backup Copies	Check this box to encrypt cold backup copies. (Encryption must also be configured on the appliance. For details, see "Encryption" on page 155.) **Note:** If the backup copy target device is configured for encryption, copies are encrypted regardless of this setting.
Keep backups for *N* days	Number of days backups must be retained. Backups that are younger than *N* days are not purged for any reason, including at the expense of new, incoming backups.
Warn when less than *N* days of backups remain	Use this option to receive an email notification if this asset has less than *N* days of backups stored on the appliance.
Delete backups after *N* days	Number of days after which the appliance will delete backups.

Managing scheduled jobs

Once you have created scheduled jobs, use these procedures to view, edit, enable, disable, and delete schedules or to run them on-demand:

Notes:

- Do not directly edit or delete a job schedule that was created by an SLA policy. Instead, modify the SLA policy itself (see "Managing SLA policies" on page 589). The appliance automatically modifies the policy's schedules

based on the changes that you make. On the Job Manager tab, SLA policy schedule names display with the prefix *_SLA:*, so you can easily distinguish them from manually created schedules.

- Deleted VMs and applications that had been protected in a schedule display as *unavailable* in the Edit Backup Job and Create Backup Copy Job dialogs. For details, see "To remove a deleted VM or application from a backup schedule" on page 585.
- iSeries schedules are managed by using the dpuconfig console interface. See "iSeries Backups Overview and Procedures" on page 767 for details on working with iSeries schedules.

To view all scheduled jobs

1. Select **Jobs > Job Manager**.
2. The Job Manager tab lists all scheduled jobs.

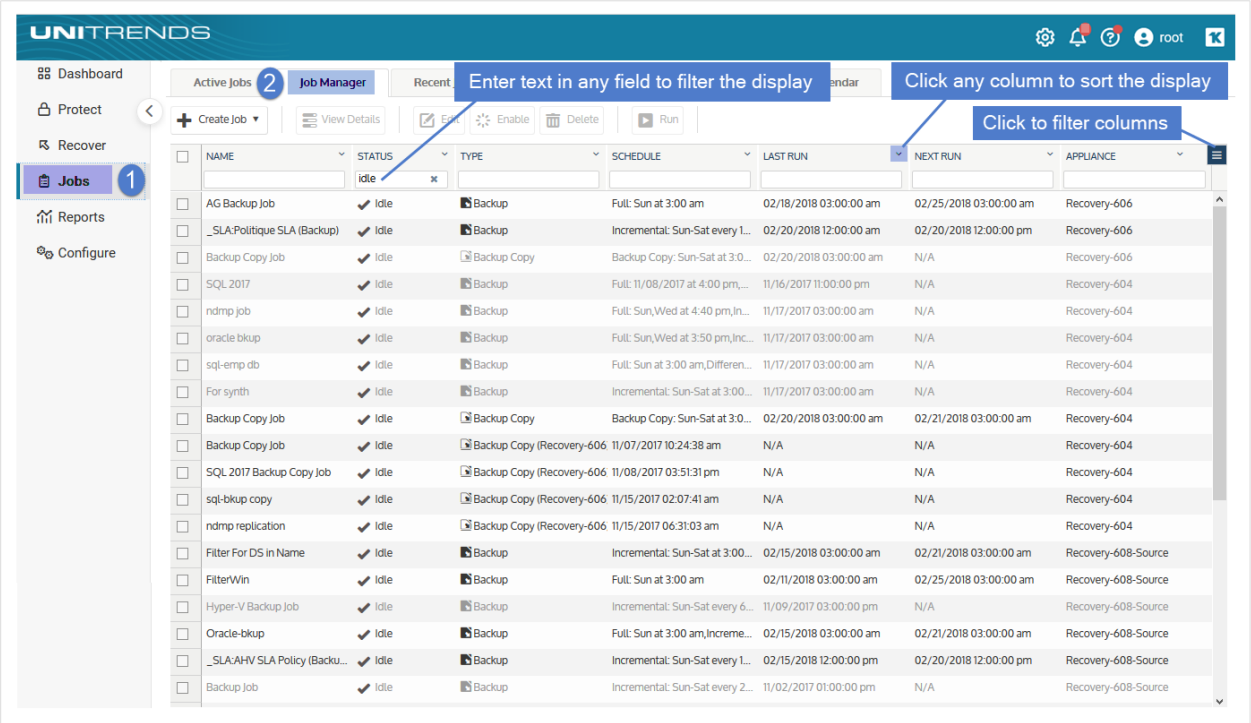

- Click on any column to sort alphabetically (a to z). To reverse the order (z to a), click again.
- Jobs that display in lighter type have been disabled.
- The following information displays for each scheduled job:

Column	Description
Name	The name of the scheduled job. Schedules that were created by SLA policies adhere to the following naming conventions: • SLA policy schedule names begin with the prefix *_SLA:*, so you can easily distinguish them from manually created schedules. • SLA policy schedule names end with one of the following: – *(Backup)* for backup schedules. – *(Hot)* for hot backup copy schedules. – *(Cold)* for cold backup copy schedules.
Status	The current status of the job: • Running - The job is running now.

Column	Description
	• Idle - The job is not running.
Type	Job type: Backup or Backup Copy.
Schedule	Description of the schedule.
Last Run	The date and time the job last ran.
Next Run	The date and time of the next scheduled run.
Appliance	The appliance running the job.

To view job details

1 Click **Jobs > Job Manager.**

2 Select the job and click **View Details.**

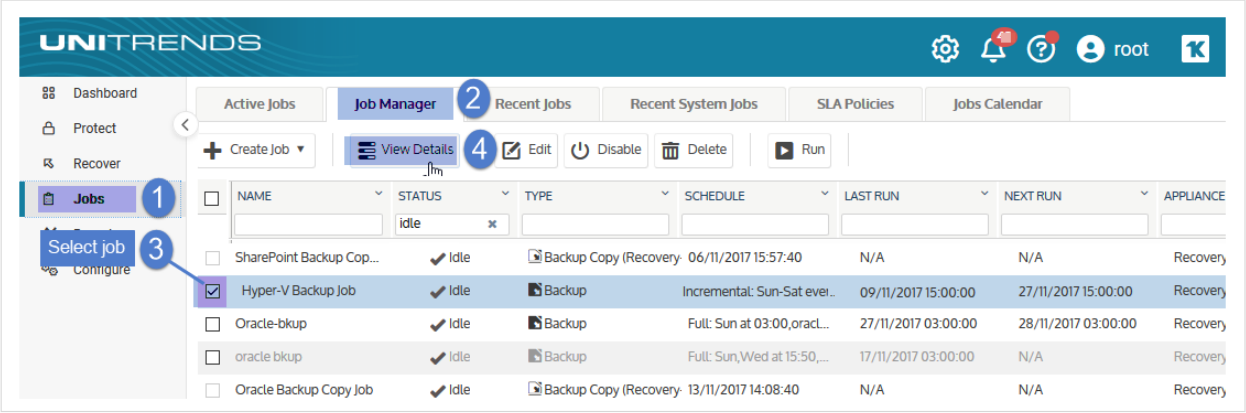

3 Job details display below.

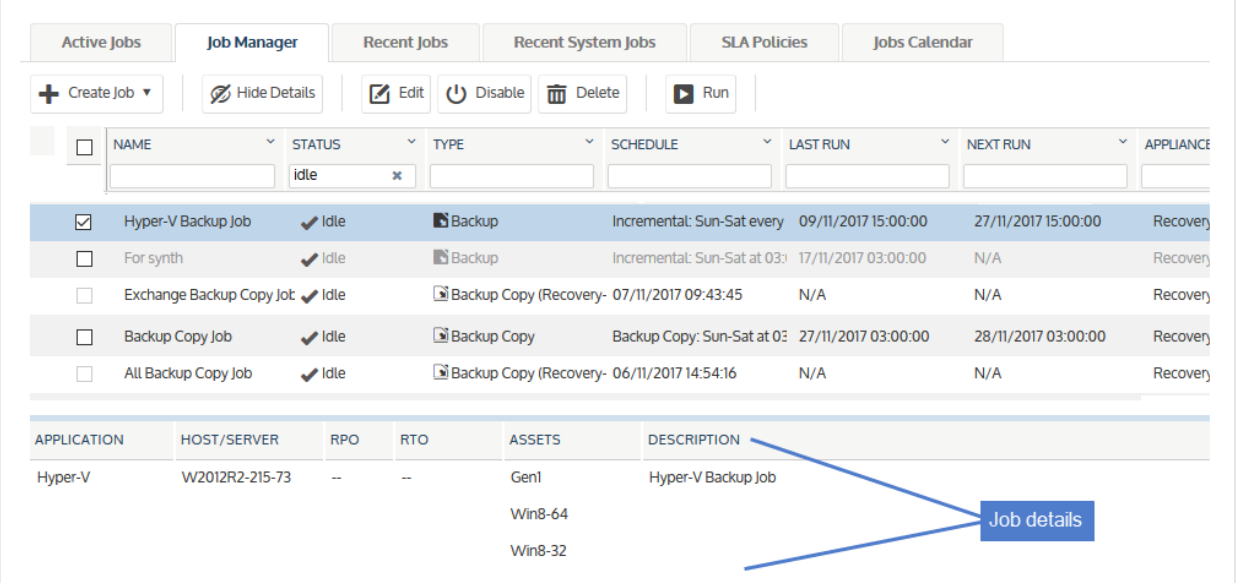

4 Click **Hide Details** to return to the original page view.

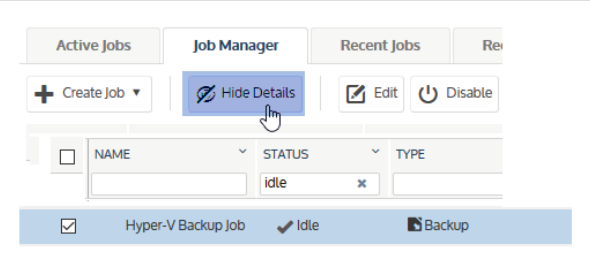

To view or edit a backup job

Notes:

- SLA policies – Do not use this procedure to edit a backup schedule that was created by an SLA policy. Instead, modify the SLA policy itself (see "To view or edit an SLA policy" on page 592). The appliance automatically modifies the policy's schedules based on the changes that you make. On the Job Manager tab, SLA policy schedule names display with the prefix *_SLA:*, so you can easily distinguish them from manually created schedules.
- SQL clusters that use a Distributed Transaction Coordinator (DTC) – There is a known issue where the DTC's IP address is assigned to the SQL instance on the Unitrends appliance. To edit the schedule, change the IP of the SQL instance to the DTC's IP. After editing the schedule, change the IP of the SQL instance back to its original address.

1 Click **Jobs > Job Manager**.

2 Select the job and click **Edit**.

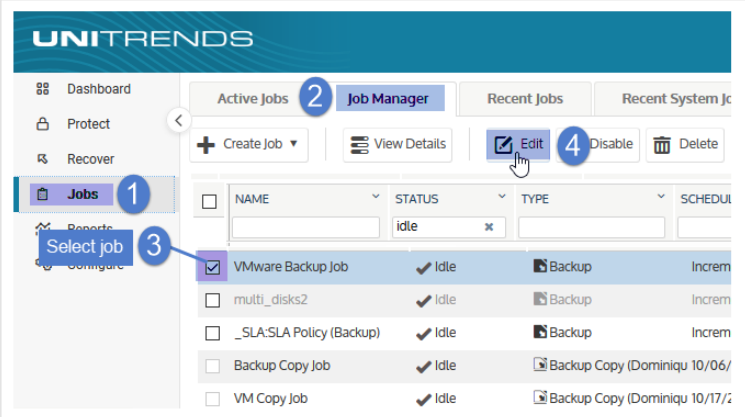

3 (Optional) In the Inventory tree, check boxes to add or remove asset(s) from the list of assets protected by this schedule.

4 (Optional) Modify the auto-include new VMs or databases setting (if applicable).

5 (Optional) Select an asset in the Job Inventory Settings area and click **Edit** to modify options. Click **Save** to retain any changes.

Notes:

- A new full backup is required after modifying the data that is included in or excluded from an asset's backups. See these topics below for details: "Considerations for VMware, AHV, and XenServer backups", "Considerations for file-level backups", and "Considerations for Windows image-level backups" on page 575.
- Editing backup options is not supported for all asset types. If you do not see the Edit option, edits are not supported.

6 Click **Next**.

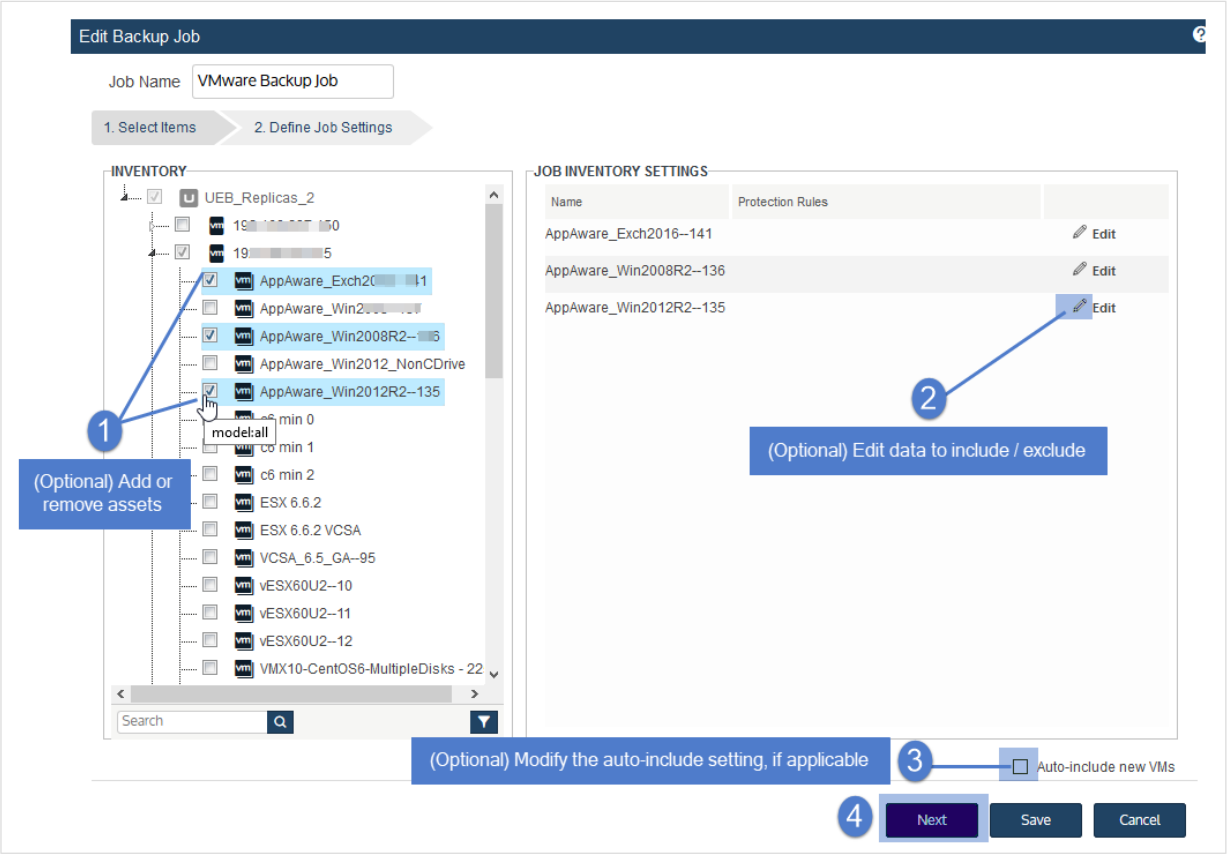

7 (Optional) Modify remaining Job Details and Options.

In most cases, the standard backup modes can be used to create the schedule. If you need more granularity, choose the **Custom** mode. For details, see "Using the Custom backup mode in the Create Backup Job dialog" on page 445.

8 Click **Save**. The schedule is updated with your changes.

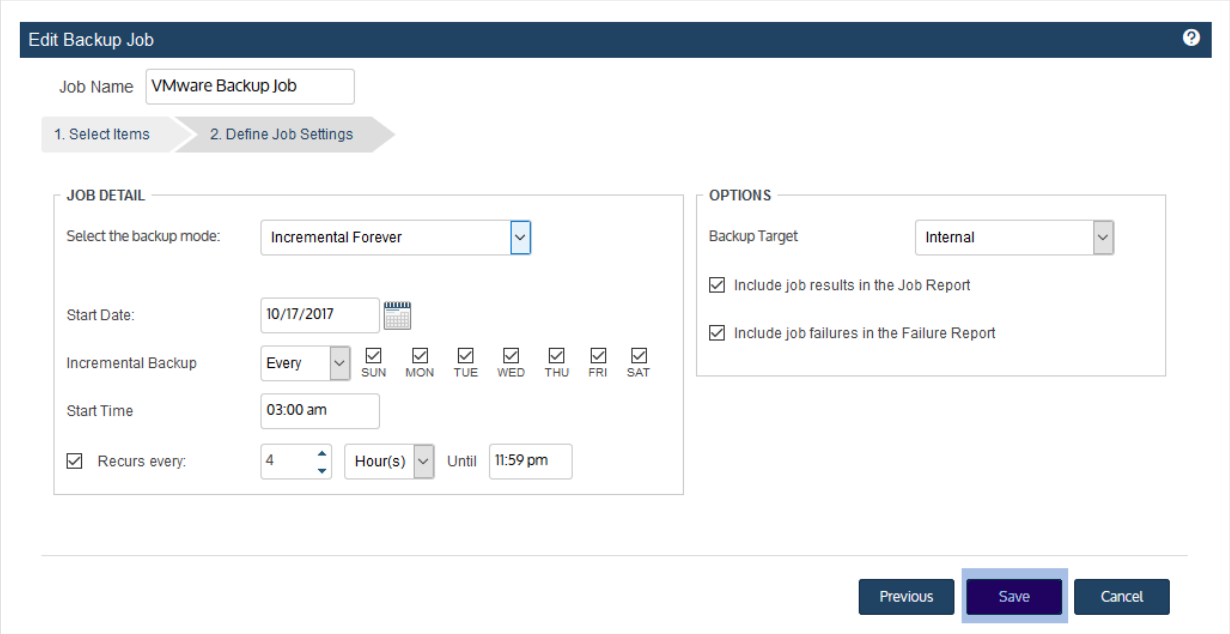

Considerations for VMware, AHV, and XenServer backups

For VMware, AHV, and XenServer, you can specify disks to exclude from backup. To exclude a disk, click to uncheck its checkbox.

Note: To recover the entire virtual machine requires critical system volumes. Use care when omitting disks from backup.

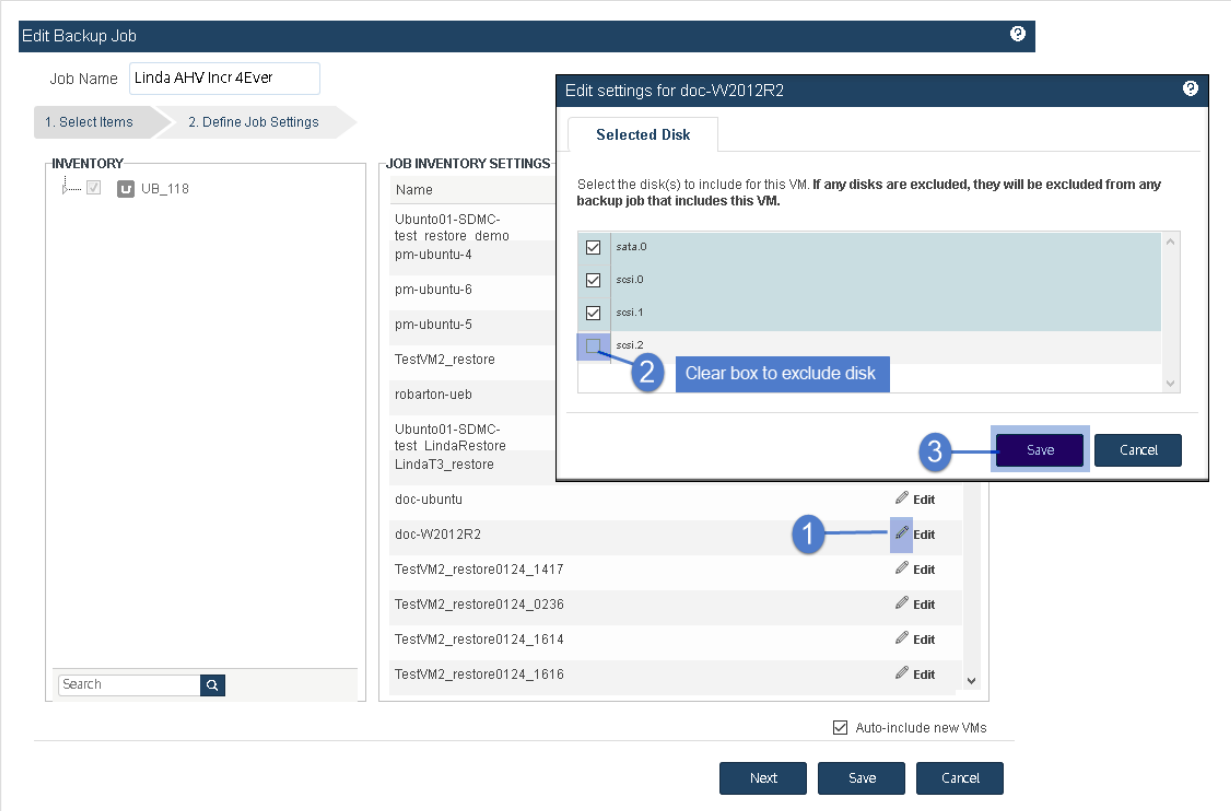

Considerations for file-level backups

File-level setting	Description
General considerations for including or excluding data from an asset's backups	Review the following before specifying data to include or exclude: • When creating a schedule, any inclusions or exclusions you add are applied to jobs run by that schedule only. Inclusions and exclusions are not applied automatically in other cases. • Running an on-demand backup of the asset does not automatically apply any inclusions or exclusions specified in the asset's schedule. To run an on-demand backup, do one of the following: – Create a one-time job that has the same inclusions and exclusions as in the schedule. – Manually run the schedule (select the schedule under **Jobs > Job Manager** and click **Run**). – Run a one-time Selective backup (so that a new full is not created).

File-level setting	Description
	• If you specify both files to include and files to exclude, the inclusion is applied first. Any exclusions are then applied to the subset of included files.
Inclusion tab	Click to specify files, folders, or volumes to include in backups of this asset. • Data that does not meet the criteria you specify here is NOT included in the backup. • Type in the full path (e.g., **C:/Documents**) or **Browse** the asset to specify data to include. (Wildcards are not supported.) • If you are running a full backup and include files or folders in the system drive (typically C:), do *not* check the System State box on the Advanced tab. Full backups fail if system state is excluded. • Run a new full backup upon creating or modifying included files. Example: 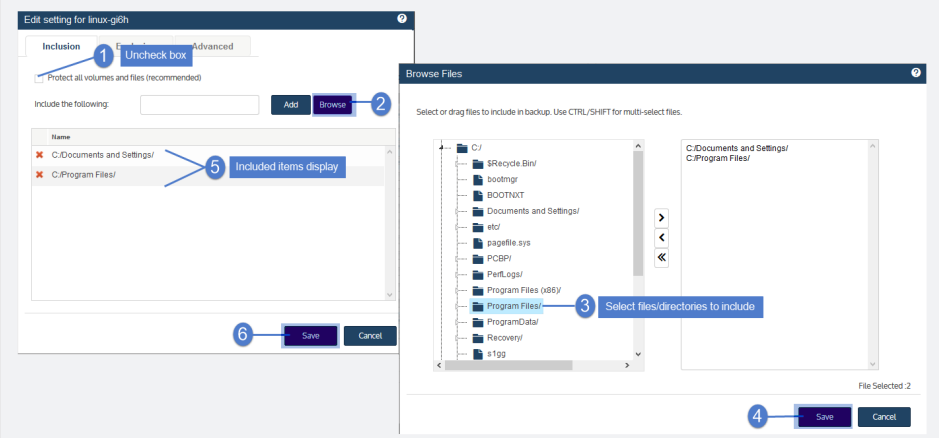
Exclusion tab	Click to specify files, folders, or volumes to exclude from backups of this asset. • Data that does not meet the criteria you specify here IS included in the backup. • To specify files to exclude, do any of the following: – Type in the full path (e.g., **C:/Documents**). – **Browse** the asset. – Enter a selection pattern. Wildcards are supported for Windows assets.

File-level setting	Description
	Wildcards are not supported for these asset types: Linux, Unix, and NAS. See these rows below for usage examples: "Wildcard * usage", "Wildcard ? usage", and "Multiple wildcards". • Run a new full backup upon creating or modifying excluded files. Example: 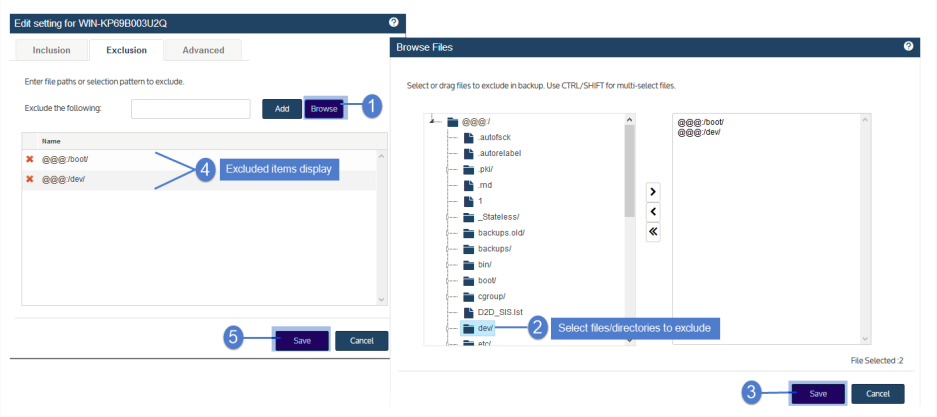
Wildcard * usage	An example of how to exclude all files with zero or more characters that match exclusion pattern: ***.txt** An example of how to exclude directories with zero or more characters and their contents within a specified path that match the exclusion pattern: **C:/windows/sys*** Limitations: • ***folder_abc** cannot be used to exclude all folders that match *folder_abc* on the protected asset. The full path must be provided. • If an entire directory is excluded, the directory name will still appear in the backup; however, its contents will be empty. • Multiple wildcard matches like the following are not supported: **C:**\abc.txt**
Wildcard ? usage	An example of how to exclude all files within specified path that matches a single character within exclusion pattern: **C:/PCBP/Lists.dir/pro_client?.spr** An example of how to exclude all directories and their contents within specified path that matches a single character within exclusion pattern:

File-level setting	Description
	C:/Programfiles/Case?/ Limitation: If an entire directory is excluded, the directory name itself will still appear in the backup; however its contents will be empty.
Multiple wildcards	An example that uses multiple “?” wildcards and only one * wildcard: **C:/?Log?/*.logs** Limitation: If an entire directory is excluded, the directory name itself will still appear in the backup; however its contents will be empty.
Advanced tab	Use this tab to specify advanced options. See these rows below for details: "Advanced Exclusions", "Command to run Pre-Backup", and "Command to run Post-Backup". Example: 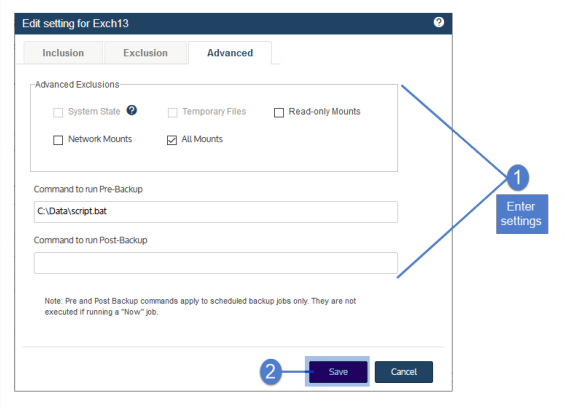
Advanced Exclusions	Check one or more boxes to exclude any of the following: system state, temporary files, read-only mounts, network mounts, or all mounts. Consider the following before applying advanced exclusions: • To perform bare metal recovery or use Windows replicas, the following must be included in the backup: system state and all boot and critical system (OS) disks/volumes. If you need these features for the asset, do not specify data to include or exclude unless you are sure these disks/volumes will be included. • If you are running a full backup and have selected files or folders in the system drive (typically C:) on the Inclusion tab, do *not* check the System State box on the Advanced tab. Full backups fail if system state is excluded. • Creating aliases for an asset - Adhere to the following when creating aliases

File-level setting	Description
	for an asset: – You must include the system state on the asset whose backups contain the boot and critical OS volumes. – You must exclude the system state on the other aliased assets. This approach ensures you can perform bare metal recovery of the asset. – Only one asset can include the system state. Disaster recovery of the asset fails if the system state is not included with the boot and OS volume or if the system state is included on aliased assets that do not include the boot and OS volume. **IMPORTANT!** For Windows assets, the backup must contain the system state, boot disk and any other system critical volumes to use the integrated bare metal recovery and Windows replica features. Be sure one of the aliased assets contains all of these disks to use these features.
Command to run Pre-Backup	To run a command or script on the asset before a scheduled backup starts, enter the full path to the command or script in the **Command to run Pre-Backup** field. For example, *C:\Data\script.bat* or */usr/jsmith/script.sh*. **Note:** Pre-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).
Command to run Post-Backup	To run a command or script on the asset after a scheduled backup completes, enter the full path to the command or script in the **Command to run Post-Backup** field. For example, *C:\Data\script.bat* or */usr/jsmith/script.sh*. **Note:** Post-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).

Considerations for Windows image-level backups

Notes:

- Critical system volumes are required for the image-level replicas feature and to recover the entire asset . Use care when omitting volumes from backup.
- When you recover the entire asset, any existing data on the target is overwritten or deleted. Volumes on the target disk that were excluded from backup may also be overwritten. For details, see "Windows unified bare metal recovery" on page 1209.

- To recover a SQL server, the master, model, and msdb system databases must be present in the image-level backup of the Windows asset. (These are included by default. If you want the recovered asset to include a hosted SQL application, use care not to exclude these system databases from the image-level backup.)
- Image-level protection is not supported for read-only disks. You must exclude all volumes on read-only disks from the backup job or run file-level backups. Image-level backups fail if read-only volumes have not been excluded.
- Removable media is automatically excluded from image-level backups. (You do not need to exclude volumes on a read-only disk that resides on removable media.)

See the following for details:

Image-level setting	Description
General considerations	Review the following before specifying volumes to include or exclude: • When creating a schedule, any inclusions or exclusions you add are applied to jobs run by that schedule only. Inclusions and exclusions are not applied automatically in other cases. • Running an on-demand backup of the asset does not automatically apply any inclusions or exclusions specified in the asset's schedule. To run an on-demand backup, do one of the following: – Create a one-time job that has the same inclusions and exclusions as in the schedule. – Manually run the schedule (select the schedule under **Jobs > Job Manager** and click **Run**).
Inclusion tab	Click to specify volumes to include in backups of this asset. • Any volumes that you do not select here are NOT included in the backup. • Type in the volume (e.g., **C:/**) or **Browse** the asset to select volumes to include. • Run a new full backup upon creating or modifying included volumes.
Inclusion example:	

Image-level setting	Description
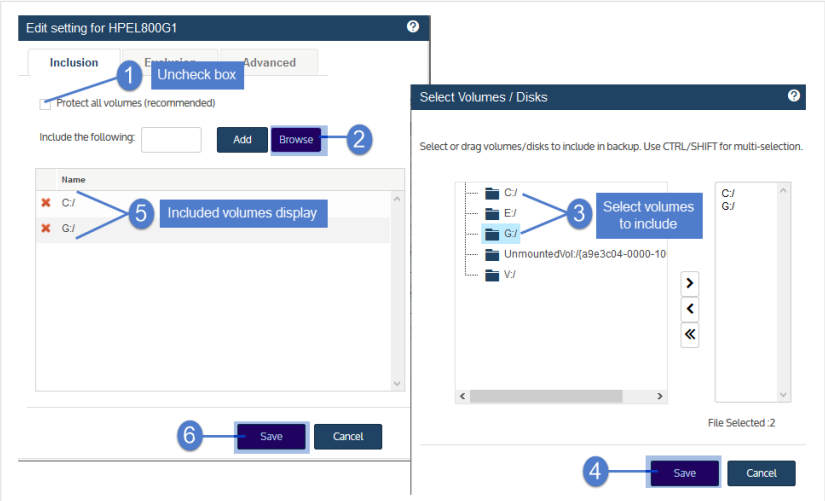	
Exclusion tab	Click to specify volumes to exclude from backups of this asset. • Any volumes that you do not select here are included in the backup. • Type in the volume (e.g., **C:/**) or **Browse** the asset to select volumes to exclude. • Run a new full backup upon creating or modifying excluded volumes.
Exclusion example:	

Image-level setting	Description
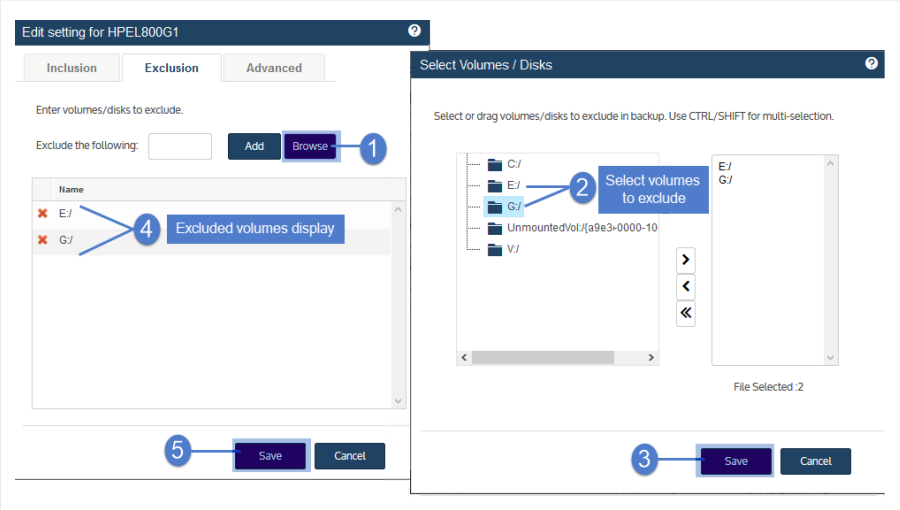	
Advanced tab	Use this tab to specify commands to run before or after the backup job. (Advanced Exclusion options are not supported for image-level backups.) See these rows below for details: "Command to run Pre-Backup" and "Command to run Post-Backup". Example: 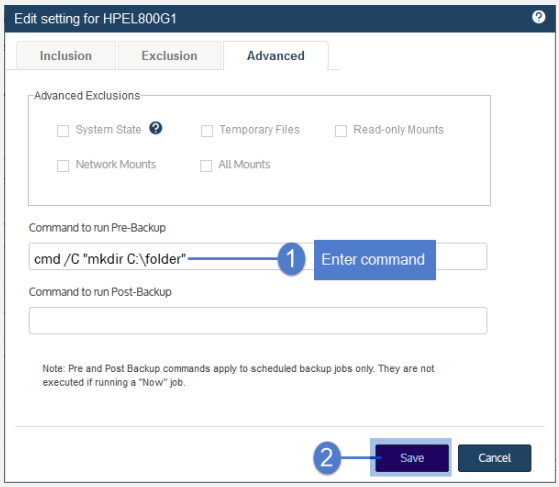
Command to run Pre-Backup	For image-level backup schedules, you cannot enter command syntax as you would inside the Windows Command Prompt. Instead, you must use this format: *cmd /<argument> "<commandSyntax>"*

Image-level setting	Description
	For example: *cmd /C "mkdir C:\folder"* To run a command or script on the asset before a scheduled backup starts, enter the command in the **Command to run Pre-Backup** field. **Note:** Pre-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).
Command to run Post-Backup	For image-level backup schedules, you cannot enter command syntax as you would inside the Windows Command Prompt. Instead, you must use this format: *cmd /<argument> "<commandSyntax>"* For example: *cmd /C "mkdir C:\folder"* To run a command or script on the asset after a scheduled backup completes, enter the command in the **Command to run Post-Backup** field. **Note:** Post-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).

To view or edit a backup copy job

Note: Do not use this procedure to edit a backup copy schedule that was created by an SLA policy. Instead, modify the SLA policy itself (see "To view or edit an SLA policy" on page 592). The appliance automatically modifies the policy's schedules based on the changes that you make. On the Job Manager tab, SLA policy schedule names display with the prefix *_SLA:*, so you can easily distinguish them from manually created schedules.

1. Click **Jobs > Job Manager**.
2. Select the job and click **Edit**.

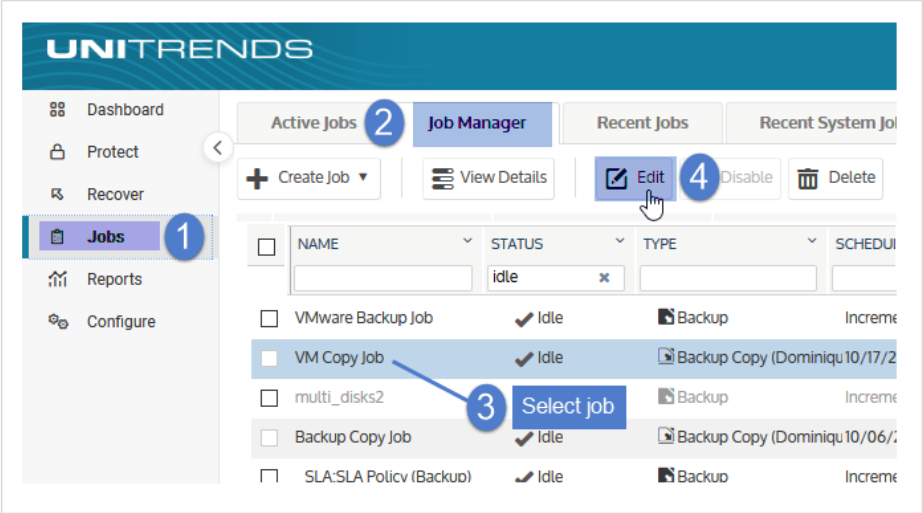

3 (Optional) Modify the job name and/or check boxes to add or remove asset(s) from the list of assets protected by this schedule. Click **Next**.

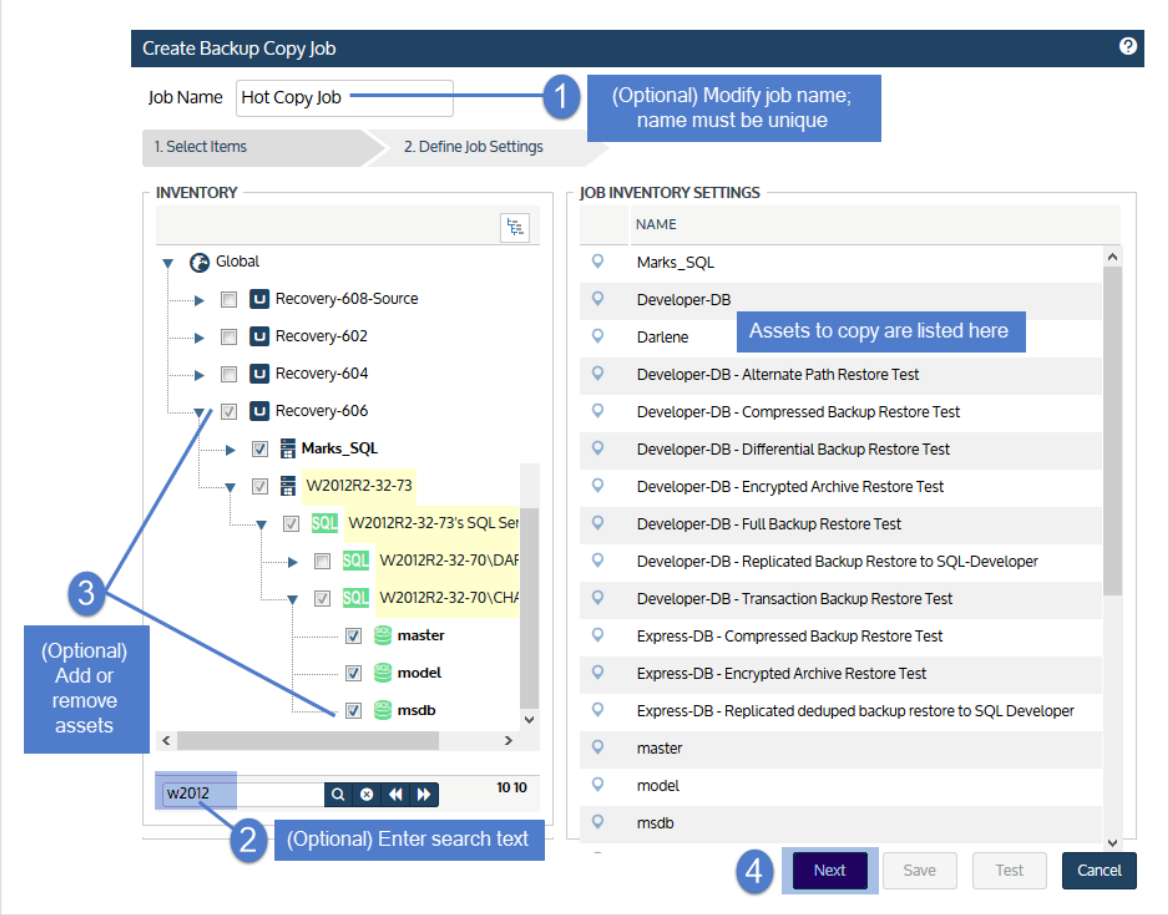

4 (Optional) Modify Job Settings. For details, see "Edit backup copy schedule examples".

Note: For Google, Amazon, AWS, and Rackspace cloud targets, reducing the storage threshold to a value less than the amount of space currently used by backup copies results in data being deleted the next time the job runs (reducing the amount of data in the cloud to meet the new threshold setting). For more information, see "Managing the amount of data copied to a third-party cloud target" on page 246.

5 Click **Save**. The schedule is updated with your changes.

Edit backup copy schedule examples

Example settings for hot backup copy schedule:

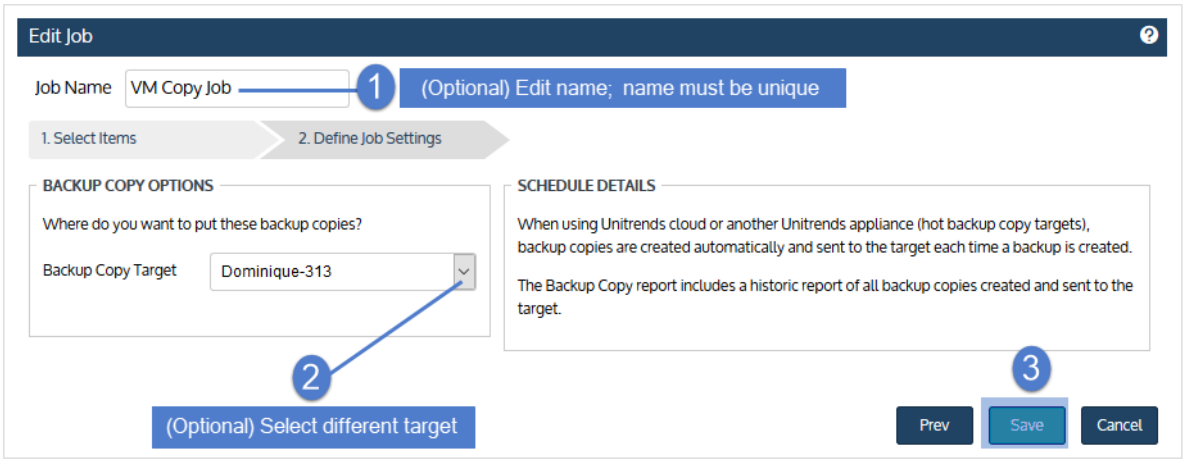

Example settings for cold backup copy schedule:

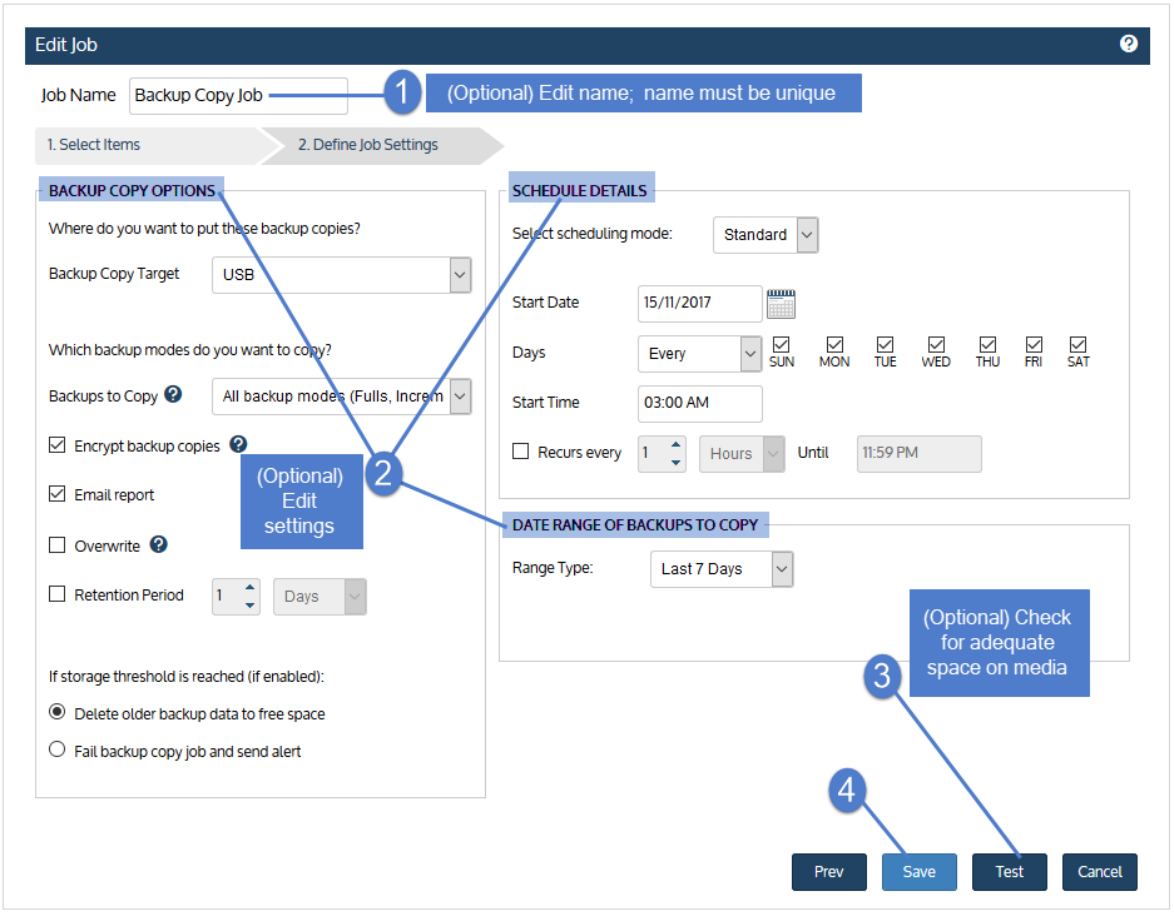

To enable or disable a job

Notes:

- Jobs are enabled by default. When a job is disabled, none of its scheduled backups or backup copies run. Be aware that disabling jobs can leave assets unprotected.
- Hot backup copy jobs cannot be disabled. Either delete the job, edit it to remove assets, or suspend backup copies from the source appliance as described in "To suspend hot backup copies" on page 266.

1. Click **Jobs > Job Manager**.
2. Select the job.
3. Click either **Enable** / **Disable**. Disabled jobs display in light gray text.

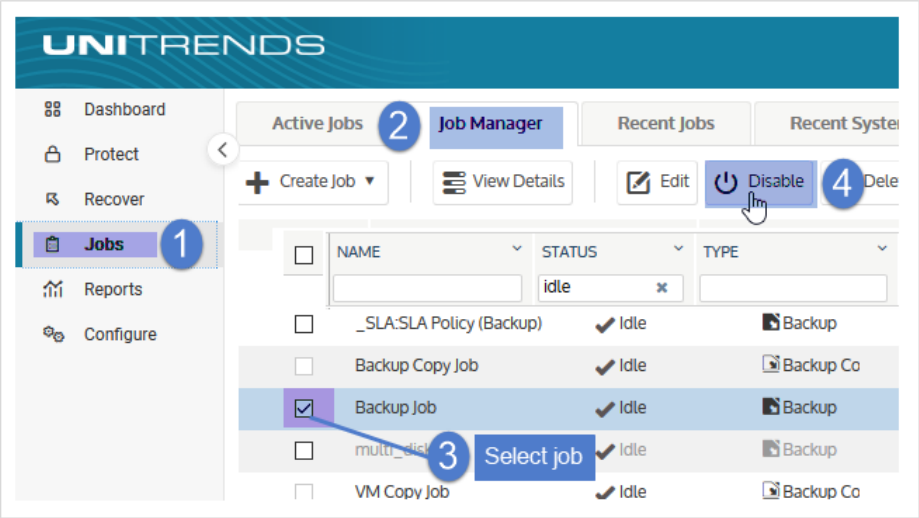

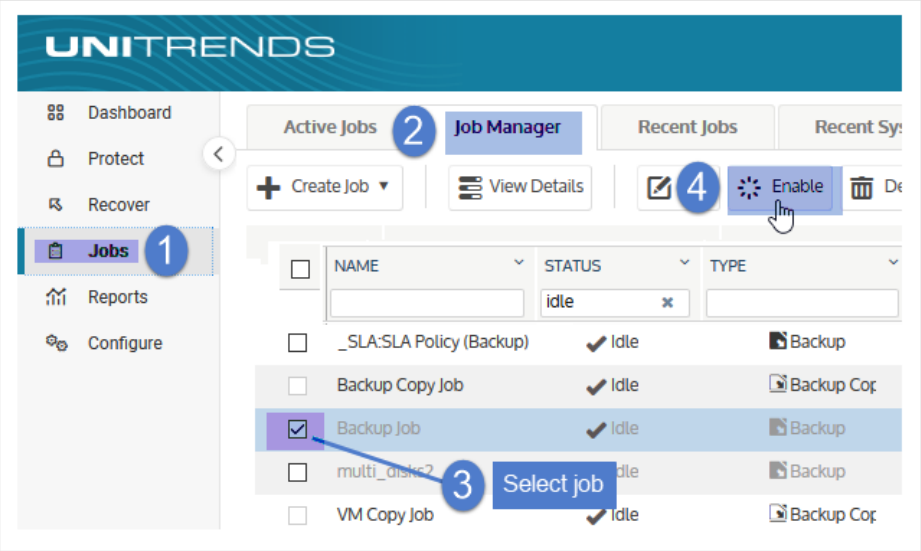

To delete a job

Note: Do not use this procedure to delete a job schedule that was created by an SLA policy. Instead, delete the SLA policy itself (see "To delete an SLA policy" on page 604). On the Job Manager tab, SLA policy schedule names display with the prefix *_SLA:*, so you can easily distinguish them from manually created schedules.

1. Click **Jobs > Job Manager**.
2. Select the job and click **Delete**.

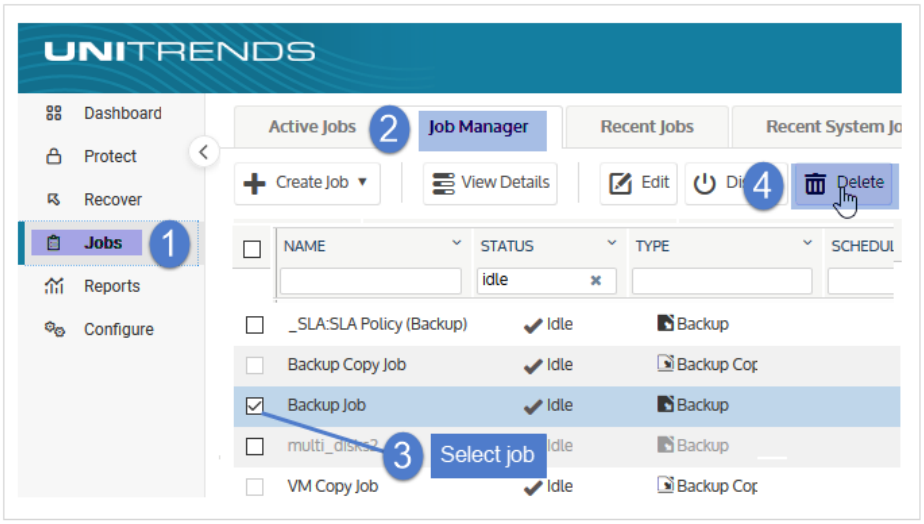

3. Click **Confirm** to delete the job.

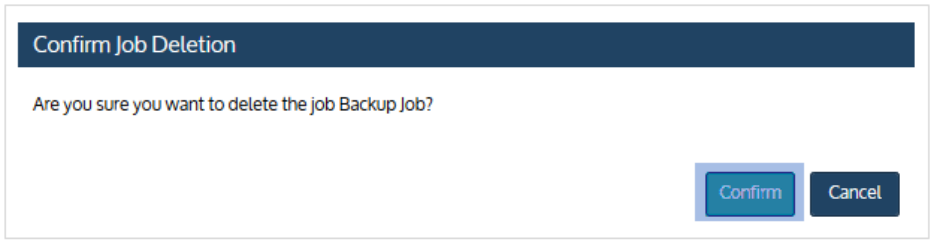

To run a scheduled job on-demand

1. Click **Jobs > Job Manager**.
2. Select the job.
3. Click **Run**. The job queues immediately.

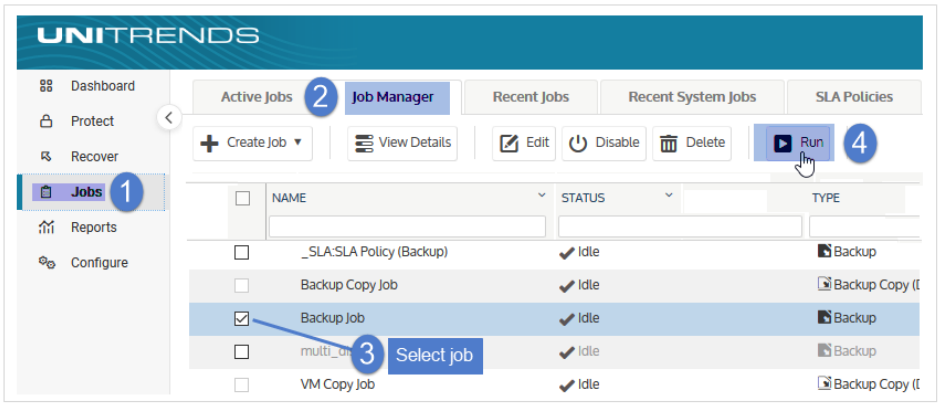

4. To monitor, pause, or cancel the job, go to **Jobs > Active Jobs**.

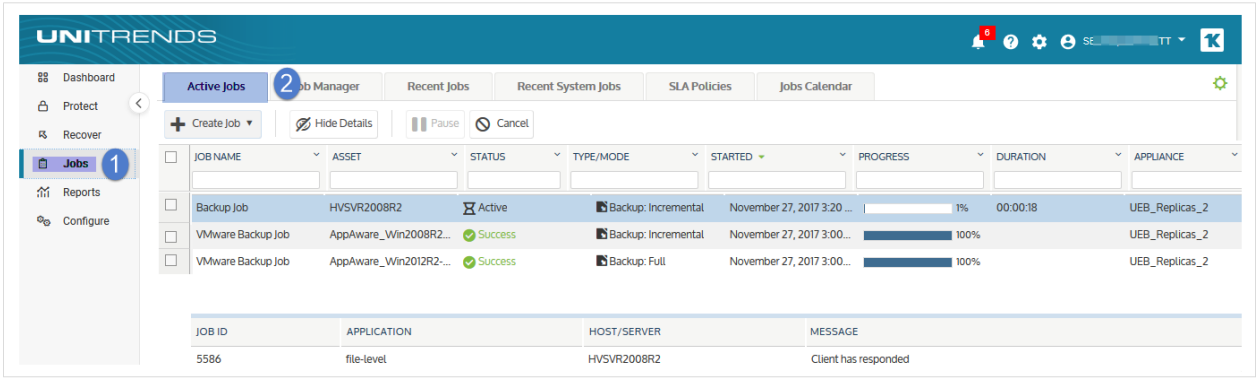

To remove a deleted VM or application from a backup schedule

If a VM or application asset is included in a backup schedule but is no longer available in your environment, alert messages and backup failures occur:

- If the VM or application will be back in service, you can ignore these alerts and failures. The issue will resolve itself once the asset is back online.
- If the VM or application will no longer be used in your environment, use this procedure to remove it from any backup schedules. Once the asset has been removed from all schedules, related alerts are automatically dismissed (and no subsequent backup failures occur because the appliance is no longer attempting to protect the decommissioned asset).

Use these steps to remove the asset from each applicable backup schedule:

1 Click **Jobs > Job Manager**.

2 Select the job and click **Edit**.

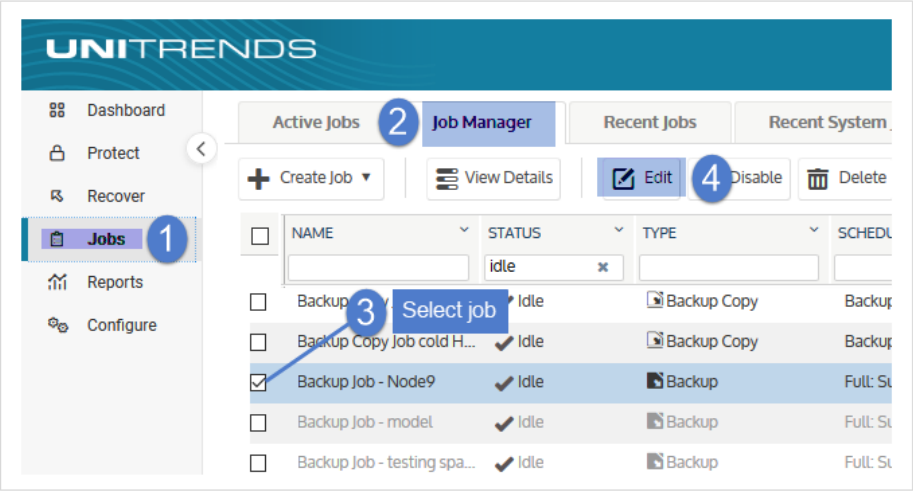

3 Select the unavailable VM or application. Click **Confirm** to remove it from the schedule.

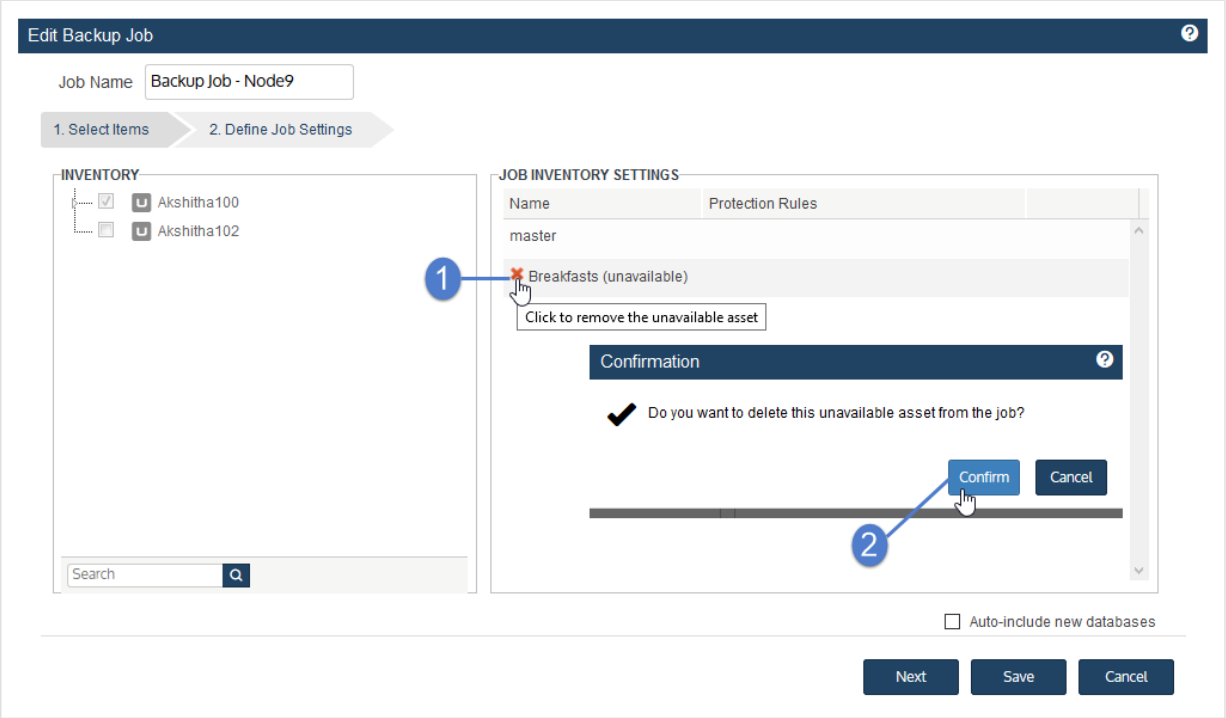

4 Click **Save** to retain the changes.

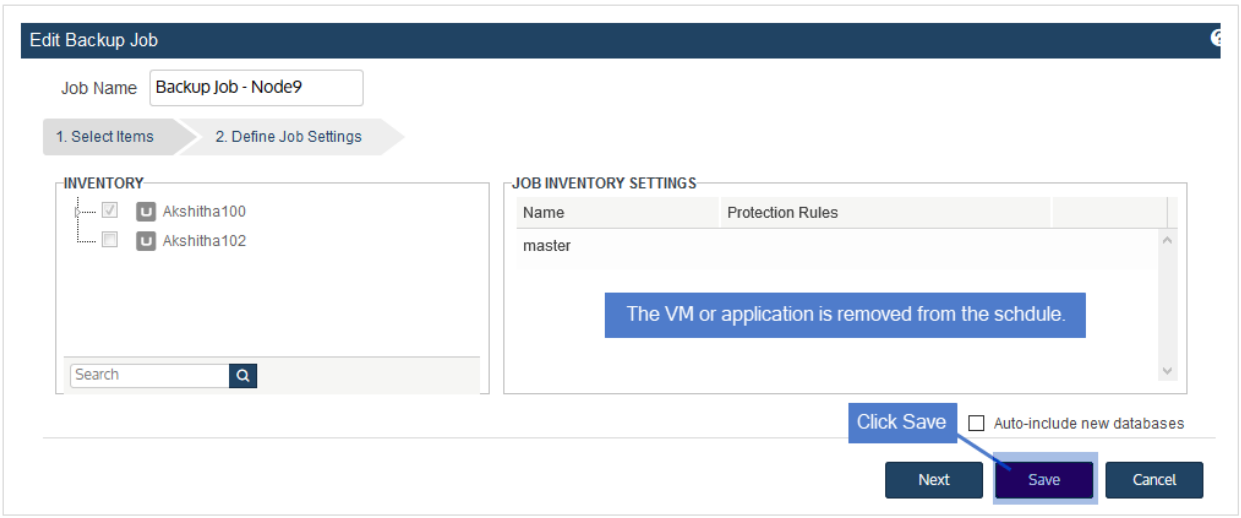

To copy a full backup to a hot backup copy target on-demand

Use this procedure to manually copy a successful full backup to the Unitrends Cloud or to a Unitrends appliance target. This procedure adds the backup copy job to the Active Jobs queue if these conditions are met:

- A hot backup copy target has been added to the backup appliance.

- The backup has not been copied to the target.
- The backup copy job is not in the Active Jobs queue.
- The source backup appliance is copying to only one hot backup copy target. If the appliance has been configured with multiple hot backup copy targets, the **Add to Hot Backup Copy Queue** button does not display in the backup log and this procedure is not supported.

1 Click **Jobs > Recent Jobs**.

2 Select the full backup and click **View Log**.

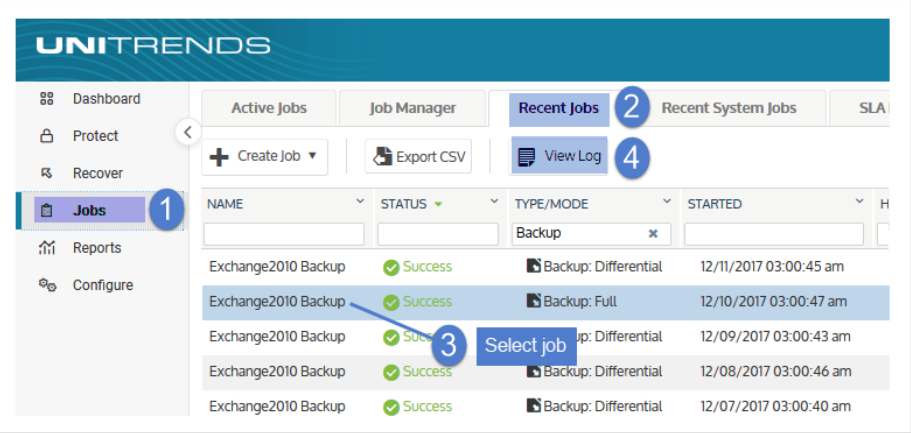

3 Click **Add to Hot Backup Copy Queue**. The backup copy job is added to the queue.

Note: The **Add to Hot Backup Copy Queue** button does not display if hot backup copy is not supported for the backup or if the source appliance is configured with multiple hot backup copy targets.

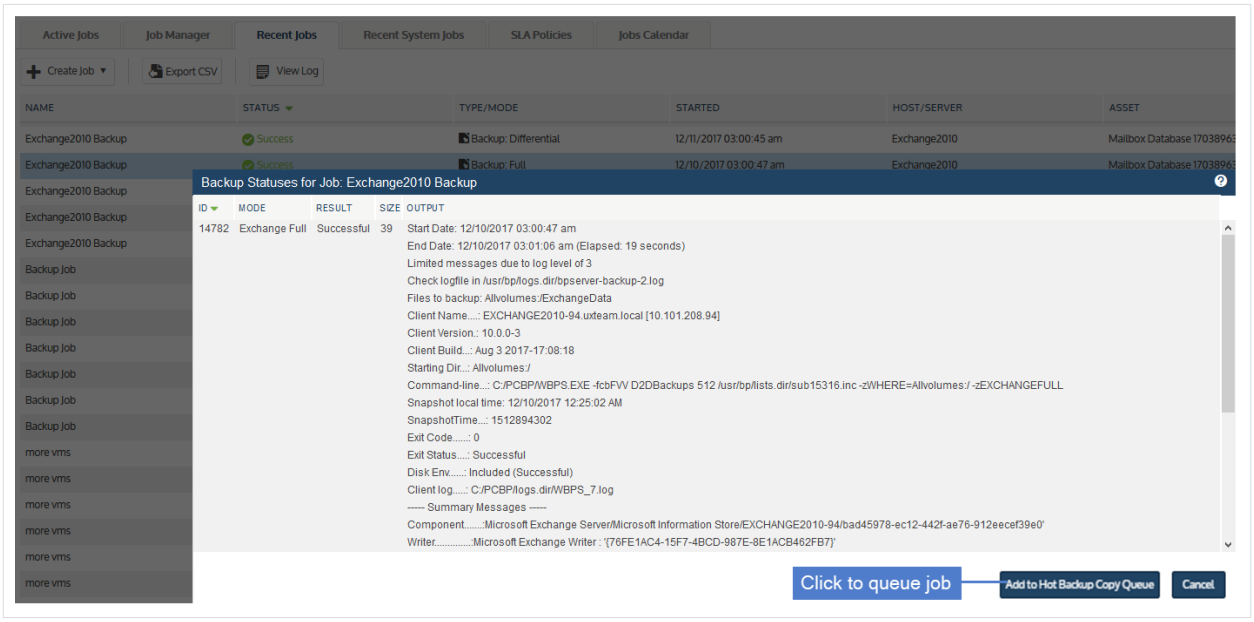

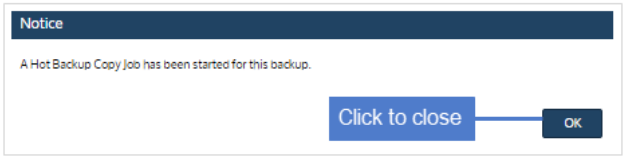

Managing SLA policies

Use these procedures to view, edit, and delete SLA policies:

- "To view all SLA policies" on page 589
- "To view details of an SLA policy" on page 590
- "To view or edit an SLA policy" on page 592
- "To delete an SLA policy" on page 604

To view all SLA policies

1. Select **Jobs > SLA Policies**.
2. The SLA Policies tab lists all policies.

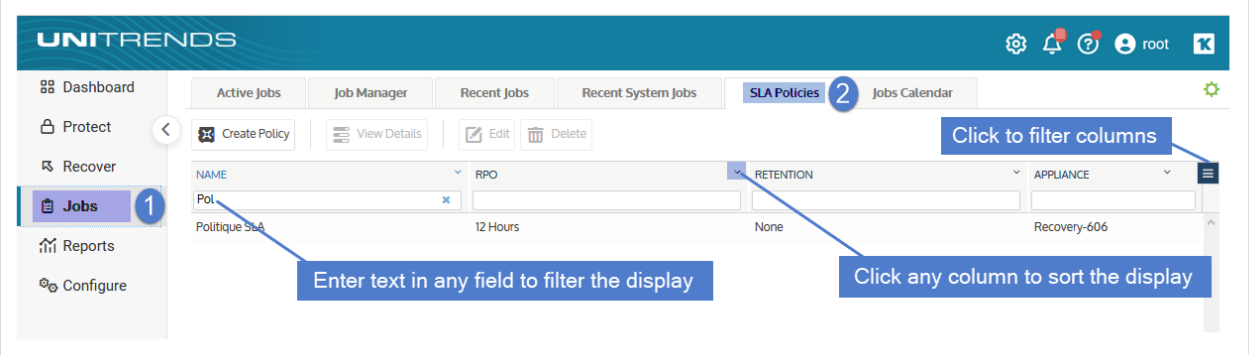

- Click on any column to sort alphabetically (a to z). To reverse the order (z to a), click again.
- The following information displays for each policy:

Column	Description
Name	The name of the SLA policy.
RPO	Recovery Point Objective – The maximum interval of time between backups (the maximum threshold of data loss tolerated by your business continuity plan).
Retention	Settings that determine how long backups must be retained on the appliance: • Min – Email notification threshold, in days. An email notification is sent if assets have less than *N* days of backups stored on the appliance. • Max – Number of days after which the appliance will delete backups. • Hold – Number of days backups must be retained. Backups that are younger than *N* days are not purged for any reason, including at the expense of new, incoming backups.
Appliance	The appliance to which this policy applies.

To view details of an SLA policy

1 Select **Jobs > SLA Policies**.

2 Select the policy in the list and click **View Details**.

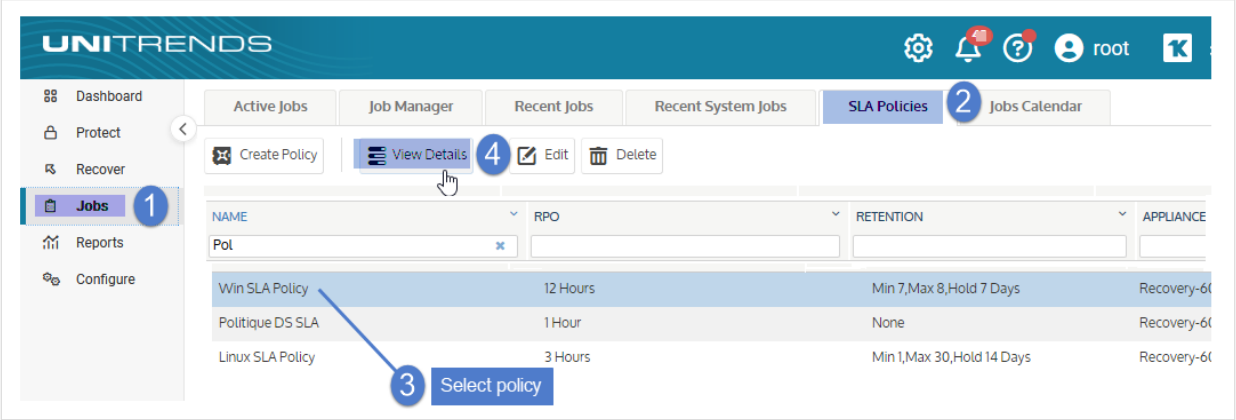

3 Policy details display below.

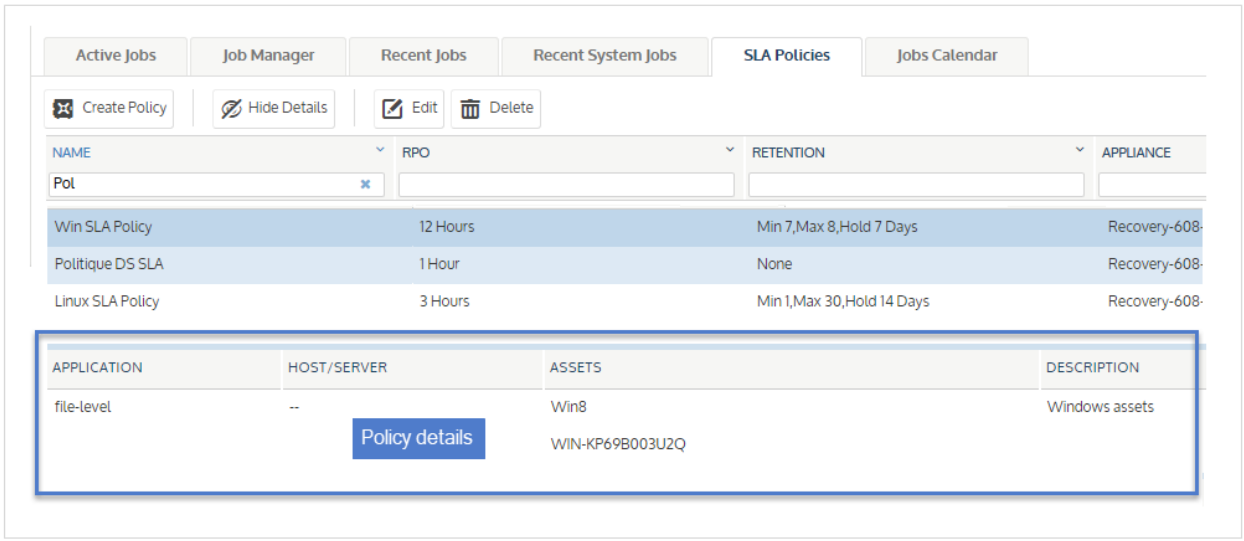

4 Click **Hide Details** to return to the original page view.

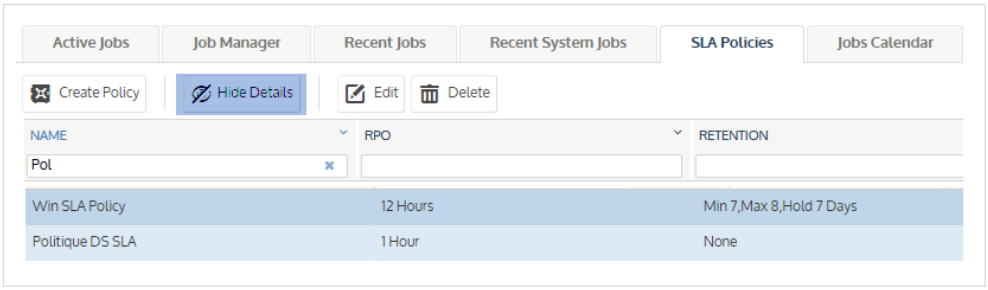

To view or edit an SLA policy

1 Select **Jobs > SLA Policies**.

2 Select the policy in the list and click **Edit**.

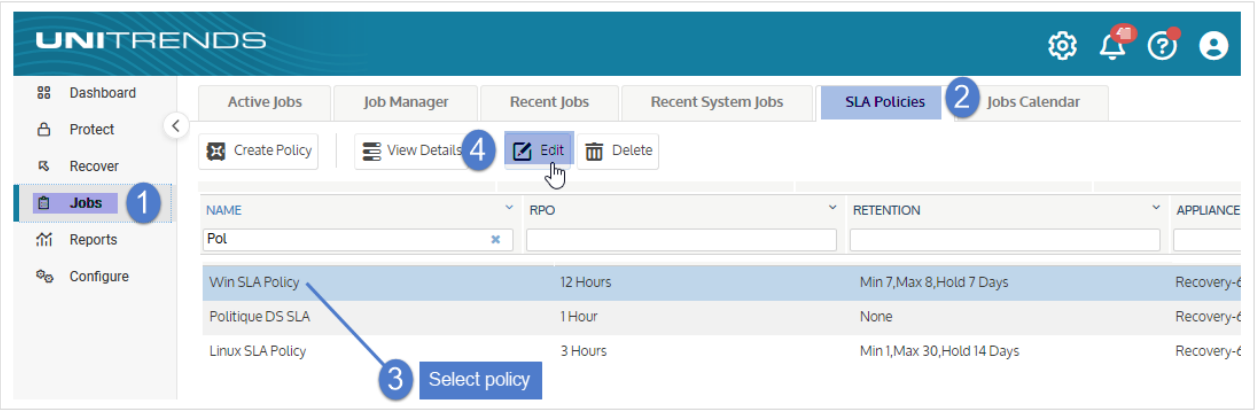

3 (Optional) In the Inventory tree, check boxes to add or remove asset(s) from the list of assets protected by this policy.

4 (Optional) Select an asset in the Job Inventory Settings area and click **Edit** to modify options. Click **Save** to retain any changes.

> **Note:** A new full backup is required after modifying the data that is included in or excluded from an asset's backups. See these topics below for details: "Considerations for VMware and AHV backups", "Considerations for Windows and Linux file-level backups", and "Considerations for Windows image-level backups" on page 598.

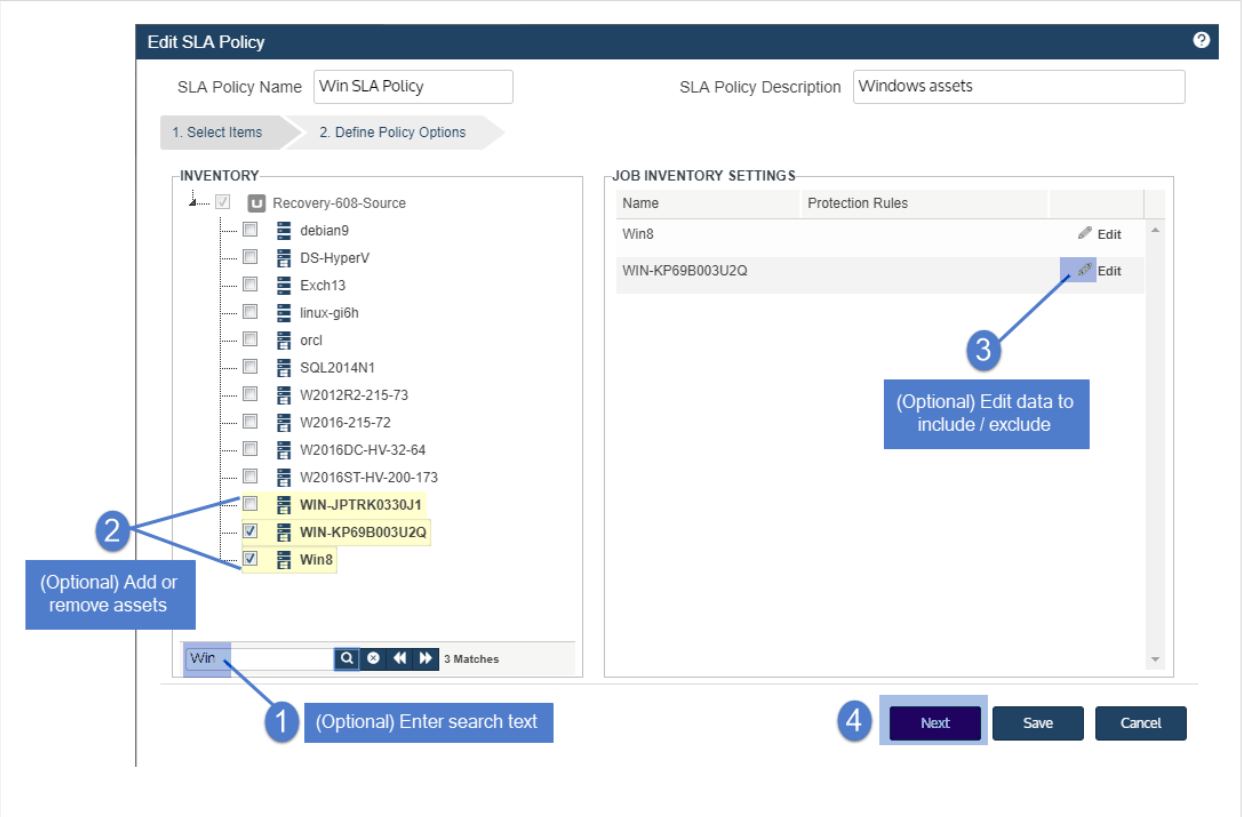

Considerations for VMware and AHV backups

For VMware and AHV, you can specify disks to exclude from backup. To exclude a disk, click to uncheck its checkbox.

Note: To recover the entire virtual machine requires critical system volumes. Use care when omitting disks from backup.

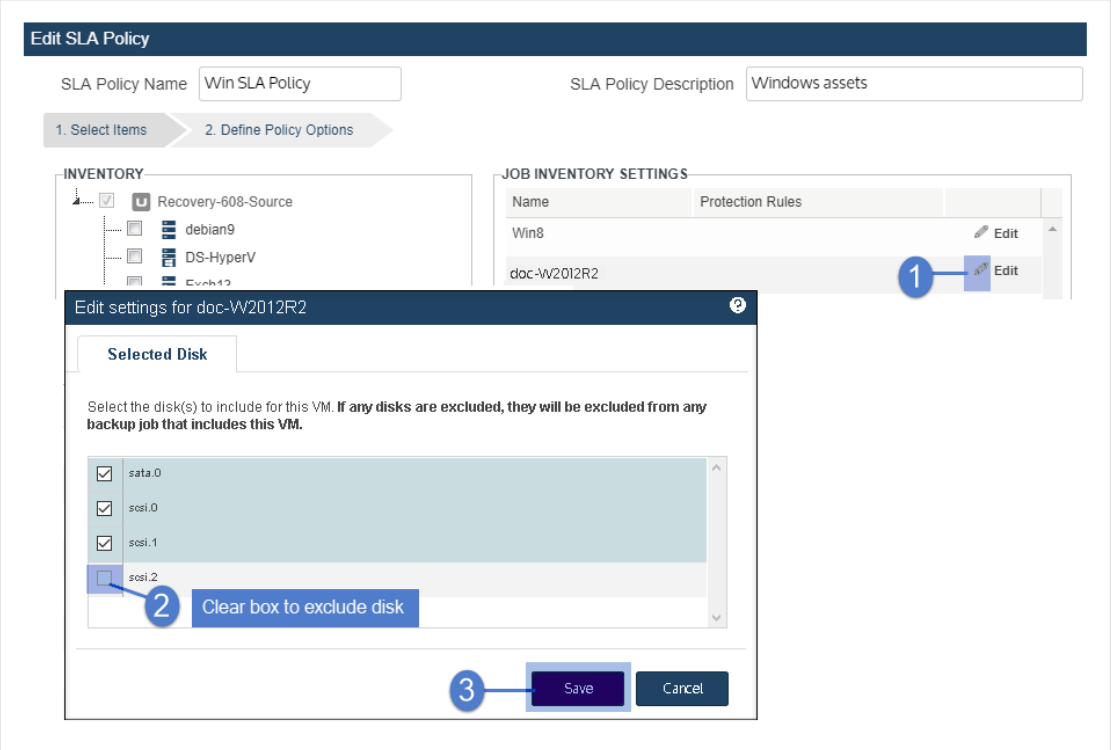

Considerations for Windows and Linux file-level backups

File-level setting	Description
General considerations for including or excluding data from an asset's backups	Review the following before specifying data to include or exclude: • When creating a schedule, any inclusions or exclusions you add are applied to jobs run by that schedule only. Inclusions and exclusions are not applied automatically in other cases. • Running an on-demand backup of the asset does not automatically apply any inclusions or exclusions specified in the asset's schedule. To run an on-demand backup, do one of the following: – Create a one-time job that has the same inclusions and exclusions as in the schedule. – Manually run the schedule (select the schedule under **Jobs > Job Manager** and click **Run**). – Run a one-time Selective backup (so that a new full is not created). • If you specify both files to include and files to exclude, the inclusion is

File-level setting	Description
	applied first. Any exclusions are then applied to the subset of included files.
Inclusion tab	Click to specify files, folders, or volumes to include in backups of this asset. • Data that does not meet the criteria you specify here is NOT included in the backup. • Type in the full path (e.g., **C:/Documents**) or **Browse** the asset to specify data to include. (Wildcards are not supported.) • If you are running a full backup and include files or folders in the system drive (typically C:), do *not* check the System State box on the Advanced tab. Full backups fail if system state is excluded. • Run a new full backup upon creating or modifying included files. Example: 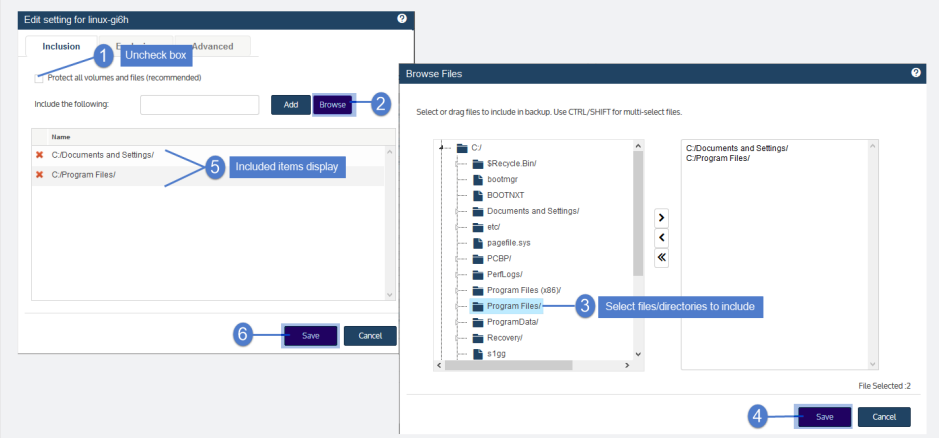
Exclusion tab	Click to specify files, folders, or volumes to exclude from backups of this asset. • Data that does not meet the criteria you specify here IS included in the backup. • To specify files to exclude, do any of the following: – Type in the full path (e.g., **C:/Documents**). – **Browse** the asset. – Enter a selection pattern. Wildcards are supported for Windows assets. Wildcards are not supported for these asset types: Linux, Unix, and NAS. See these rows below for usage examples: "Wildcard * usage", "Wildcard

File-level setting	Description
	? usage", and "Multiple wildcards". • Run a new full backup upon creating or modifying excluded files. Example: 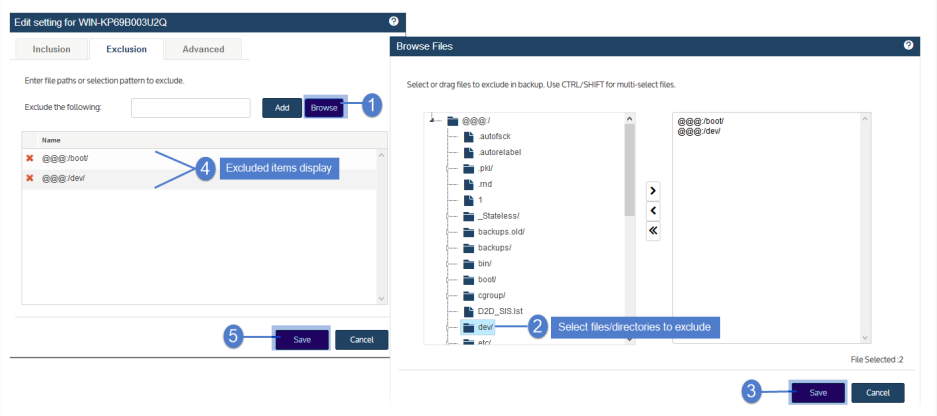
Wildcard * usage	An example of how to exclude all files with zero or more characters that match exclusion pattern: ***.txt** An example of how to exclude directories with zero or more characters and their contents within a specified path that match the exclusion pattern: **C:/windows/sys*** Limitations: • ***folder_abc** cannot be used to exclude all folders that match *folder_abc* on the protected asset. The full path must be provided. • If an entire directory is excluded, the directory name will still appear in the backup; however, its contents will be empty. • Multiple wildcard matches like the following are not supported: **C:**\abc.txt**
Wildcard ? usage	An example of how to exclude all files within specified path that matches a single character within exclusion pattern: **C:/PCBP/Lists.dir/pro_client?.spr** An example of how to exclude all directories and their contents within specified path that matches a single character within exclusion pattern: **C:/Programfiles/Case?/**

File-level setting	Description
	Limitation: If an entire directory is excluded, the directory name itself will still appear in the backup; however its contents will be empty.
Multiple wildcards	An example that uses multiple "?" wildcards and only one * wildcard: **C:/?Log?/*.logs** Limitation: If an entire directory is excluded, the directory name itself will still appear in the backup; however its contents will be empty.
Advanced tab	Use this tab to specify advanced options. See these rows below for details: "Advanced Exclusions", "Command to run Pre-Backup", and "Command to run Post-Backup". Example: 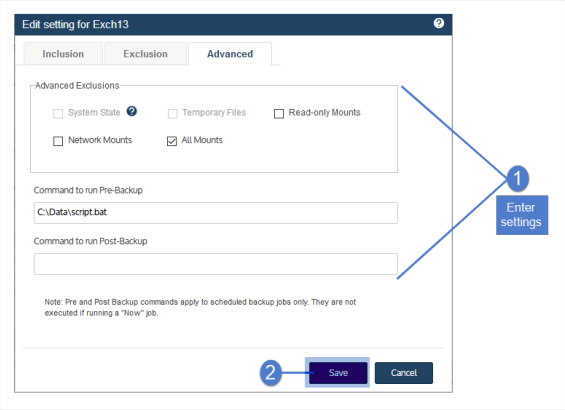
Advanced Exclusions	Check one or more boxes to exclude any of the following: system state, temporary files, read-only mounts, network mounts, or all mounts. Consider the following before applying advanced exclusions: • To perform bare metal recovery or use Windows replicas, the following must be included in the backup: system state and all boot and critical system (OS) disks/volumes. If you need these features for the asset, do not specify data to include or exclude unless you are sure these disks/volumes will be included. • If you are running a full backup and have selected files or folders in the system drive (typically C:) on the Inclusion tab, do *not* check the System State box on the Advanced tab. Full backups fail if system state is excluded. • Creating aliases for an asset - Adhere to the following when creating aliases for an asset:

File-level setting	Description
	– You must include the system state on the asset whose backups contain the boot and critical OS volumes. – You must exclude the system state on the other aliased assets. This approach ensures you can perform bare metal recovery of the asset. – Only one asset can include the system state. Disaster recovery of the asset fails if the system state is not included with the boot and OS volume or if the system state is included on aliased assets that do not include the boot and OS volume. **IMPORTANT!** For Windows assets, the backup must contain the system state, boot disk and any other system critical volumes to use the integrated bare metal recovery and Windows replica features. Be sure one of the aliased assets contains all of these disks to use these features.
Command to run Pre-Backup	To run a command or script on the asset before a scheduled backup starts, enter the full path to the command or script in the **Command to run Pre-Backup** field. For example, *C:\Data\script.bat* or */usr/jsmith/script.sh*. **Note:** Pre-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).
Command to run Post-Backup	To run a command or script on the asset after a scheduled backup completes, enter the full path to the command or script in the **Command to run Post-Backup** field. For example, *C:\Data\script.bat* or */usr/jsmith/script.sh*. **Note:** Post-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).

Considerations for Windows image-level backups

Notes:

- Critical system volumes are required for the image-level replicas feature and to recover the entire asset . Use care when omitting volumes from backup.
- When you recover the entire asset, any existing data on the target is overwritten or deleted. Volumes on the target disk that were excluded from backup may also be overwritten. For details, see "Windows unified bare metal recovery" on page 1209.

- To recover a SQL server, the master, model, and msdb system databases must be present in the image-level backup of the Windows asset. (These are included by default. If you want the recovered asset to include a hosted SQL application, use care not to exclude these system databases from the image-level backup.)
- Image-level protection is not supported for read-only disks. You must exclude all volumes on read-only disks from the backup job or run file-level backups. Image-level backups fail if read-only volumes have not been excluded.
- Removable media is automatically excluded from image-level backups. (You do not need to exclude volumes on a read-only disk that resides on removable media.)

See the following for details:

Image-level setting	Description
General considerations	Review the following before specifying volumes to include or exclude: • When creating a schedule, any inclusions or exclusions you add are applied to jobs run by that schedule only. Inclusions and exclusions are not applied automatically in other cases. • Running an on-demand backup of the asset does not automatically apply any inclusions or exclusions specified in the asset's schedule. To run an on-demand backup, do one of the following: – Create a one-time job that has the same inclusions and exclusions as in the schedule. – Manually run the schedule (select the schedule under **Jobs > Job Manager** and click **Run**).
Inclusion tab	Click to specify volumes to include in backups of this asset. • Any volumes that you do not select here are NOT included in the backup. • Type in the volume (e.g., **C:/**) or **Browse** the asset to select volumes to include. • Run a new full backup upon creating or modifying included volumes.
Inclusion example:	

Image-level setting	Description
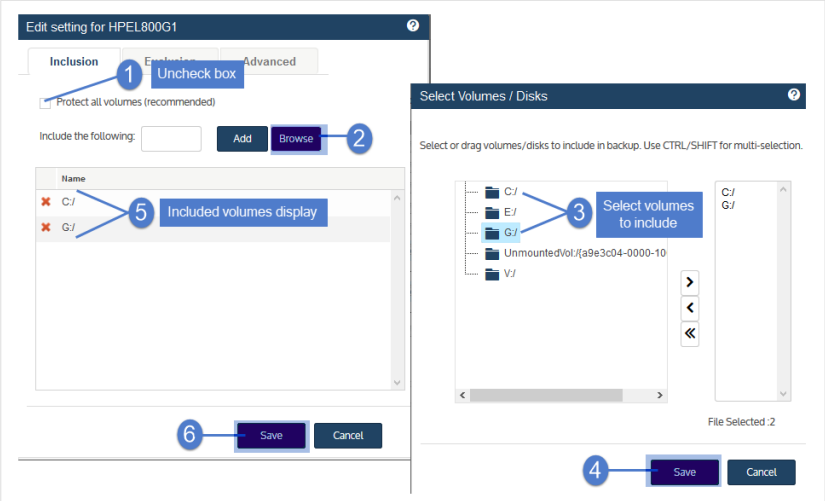	
Exclusion tab	Click to specify volumes to exclude from backups of this asset. • Any volumes that you do not select here are included in the backup. • Type in the volume (e.g., **C:/**) or **Browse** the asset to select volumes to exclude. • Run a new full backup upon creating or modifying excluded volumes.
Exclusion example:	

Image-level setting	Description
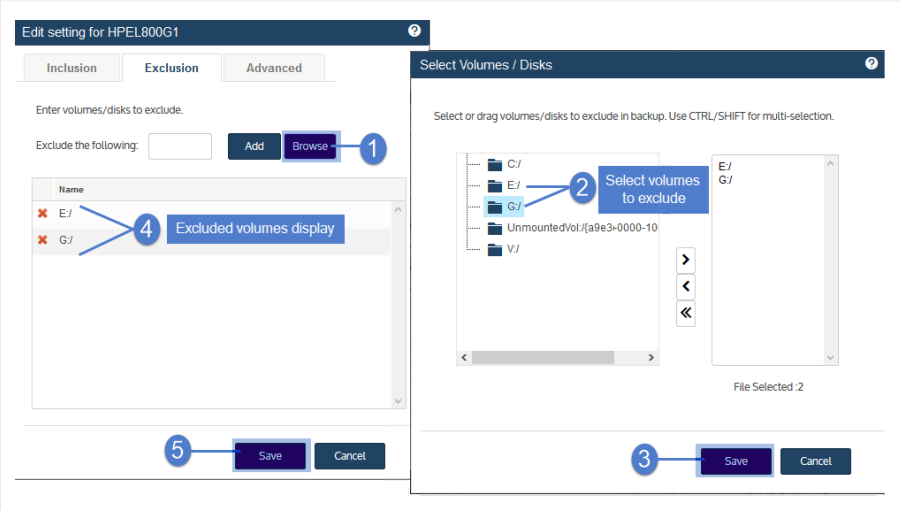	
Advanced tab	Use this tab to specify commands to run before or after the backup job. (Advanced Exclusion options are not supported for image-level backups.) See these rows below for details: "Command to run Pre-Backup" and "Command to run Post-Backup". Example: 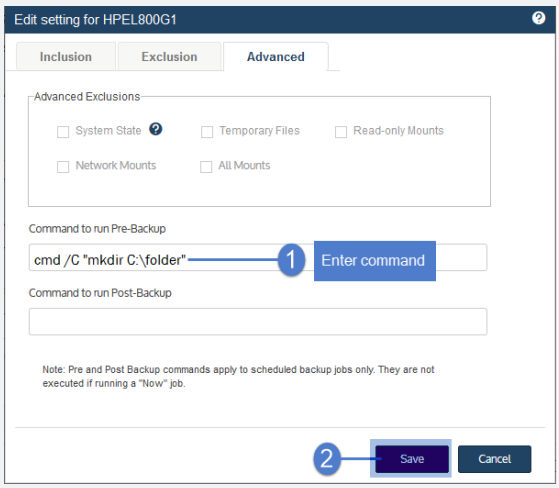
Command to run Pre-Backup	For image-level backup schedules, you cannot enter command syntax as you would inside the Windows Command Prompt. Instead, you must use this format: *cmd /<argument> "<commandSyntax>"* For example: *cmd /C "mkdir C:\folder"*

Image-level setting	Description
	To run a command or script on the asset before a scheduled backup starts, enter the command in the **Command to run Pre-Backup** field. **Note:** Pre-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).
Command to run Post-Backup	For image-level backup schedules, you cannot enter command syntax as you would inside the Windows Command Prompt. Instead, you must use this format: *cmd /<argument> "<commandSyntax>"* For example: *cmd /C "mkdir C:\folder"* To run a command or script on the asset after a scheduled backup completes, enter the command in the **Command to run Post-Backup** field. **Note:** Post-backup commands are used for scheduled jobs only. These commands are not applied to on-demand jobs (jobs run by selecting the **Now** option).

5 Click **Next**.

6 (Optional) Modify Policy Options.

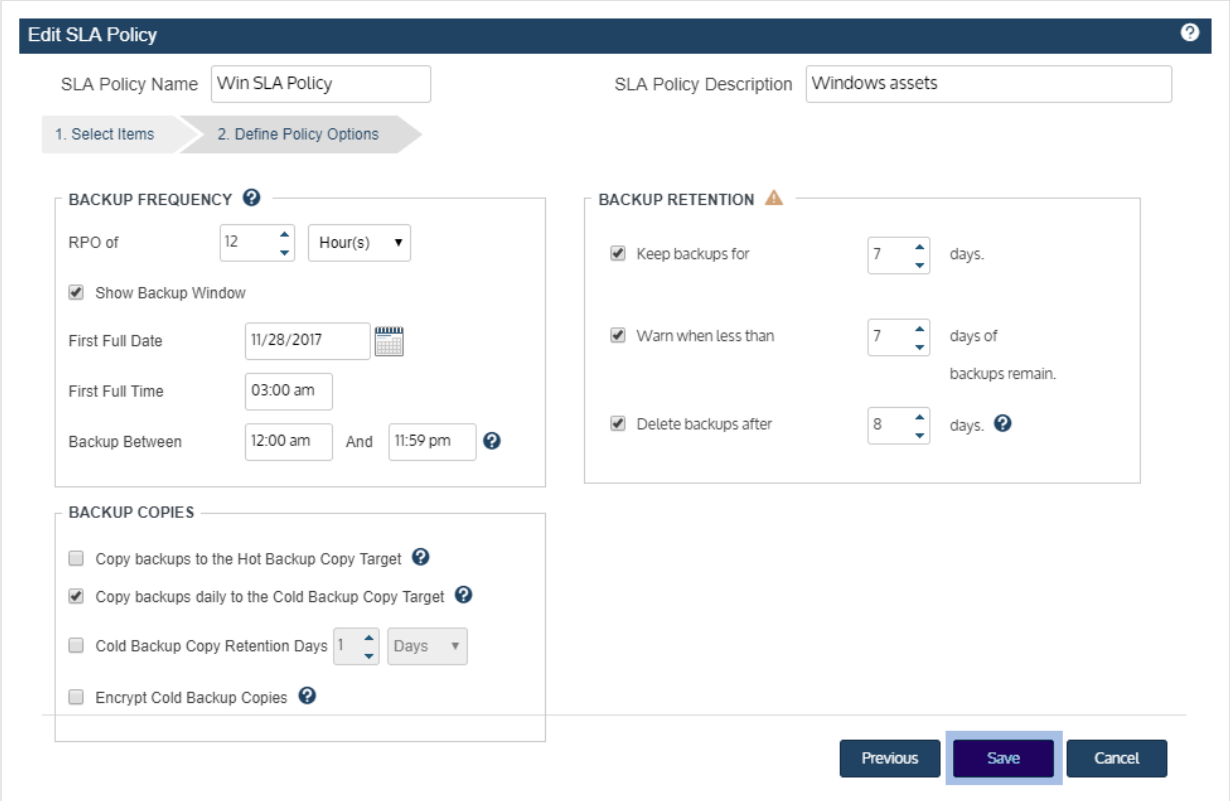

See the following table for details:

SLA policy setting	Description
SLA Policy Name	Enter a unique name for the policy.
SLA Policy Description	(Optional) Enter a short description of the policy.
RPO	Recovery Point Objective – The maximum interval of time between backups (the maximum threshold of data loss tolerated by your business continuity plan). Determines how often backups will run. Enter the number of hours or minutes to define the RPO interval.
Show Backup Window	Check this box to view and/or edit the following: • First Full Date – Date when the policy's first full backups will run. (Applies to assets that do not yet have a successful full backup.) • First Full Time – Time when the policy's first full backups will run. (Applies to assets that do not yet have a successful full backup.) • Backup Between – Hours of the day when backups will be taken.
Copy backups to the Hot Backup Copy Target	Check this box to copy backups to your hot backup copy target. Supported only when a Unitrends appliance or the Unitrends Cloud has been added as a

SLA policy setting	Description
	backup copy target. (For details on adding a hot target, see "Backup copy targets" on page 214.)
Copy backups daily to the Cold Backup Copy Target	Check this box to copy backups to your cold backup copy target. Supported only when a cold backup copy target has been added to the backup appliance. Supported for these types of cold targets only: third-party cloud, NAS, or iSCSI. If multiple cold targets exist, the policy copies to the one that was added first. • To copy to a different cold target, manually create a backup copy job instead, as described in "Creating backup copy jobs" on page 491. • To add a cold target to the backup appliance, see "Backup copy targets" on page 214.
Cold Backup Copy Retention Days	Check this box to specify the length of time a copy must be retained before it can be deleted. To define the retention period, enter a number and select Days, Weeks, Months, or Years. For example, enter **2** and select **Weeks** to retain copies for 2 weeks.
Encrypt Cold Backup Copies	Check this box to encrypt cold backup copies. (Encryption must also be configured on the appliance. For details, see "Encryption" on page 155.) **Note:** If the backup copy target device is configured for encryption, copies are encrypted regardless of this setting.
Keep backups for *N* days	Number of days backups must be retained. Backups that are younger than *N* days are not purged for any reason, including at the expense of new, incoming backups.
Warn when less than *N* days of backups remain	Use this option to receive an email notification if this asset has less than *N* days of backups stored on the appliance.
Delete backups after *N* days	Number of days after which the appliance will delete backups.

7 Do one of the following:

- If you modified the policy settings, click **Save**. The policy and related jobs are updated with your changes.
- If you did not make any changes, click **Cancel** to close the Edit SLA Policy dialog.

To delete an SLA policy

When deleting a policy, you can opt to remove the associated backup and backup copy schedules or opt to retain these schedules.

1 Select **Jobs > SLA Policies**.

2 Select the policy in the list and click **Delete**.

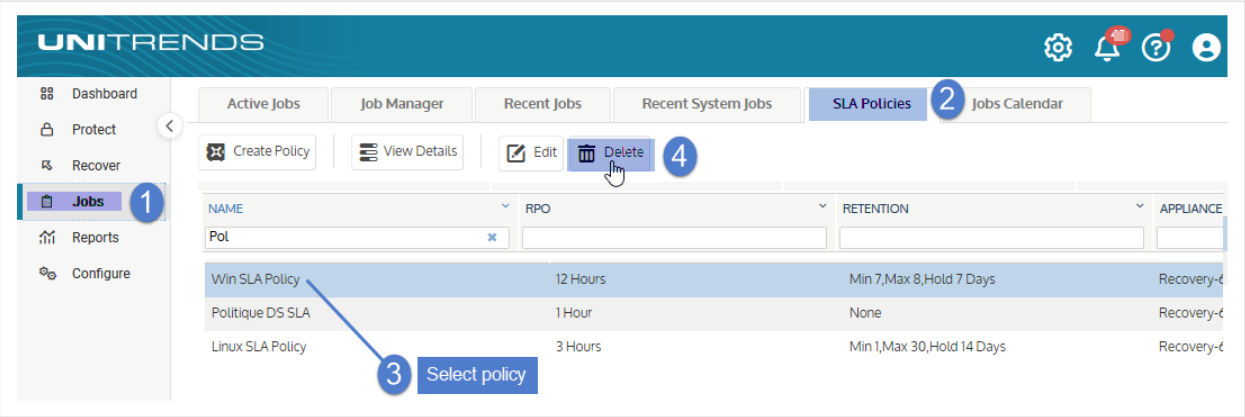

3 Select one of the following options:

- Delete the SLA Policy **and** its related Backup and Backup Copy Jobs
- Delete the SLA Policy only, **not** its related Backup and Backup Copy Jobs

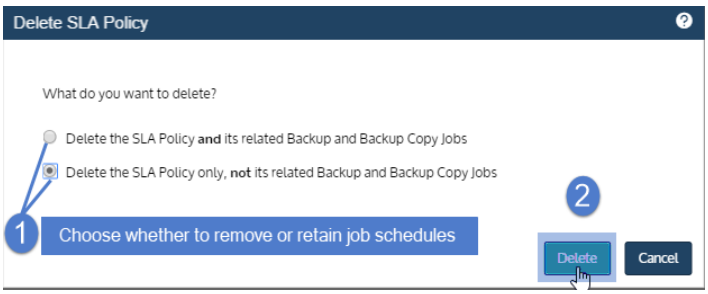

4 Click **Delete** to delete the policy.

5 If you opted to retain job schedules, the jobs can now be managed independently (since they are no longer associated with the SLA policy). On the Job Manager tab, edit the job names to remove the *_SLA* prefix.

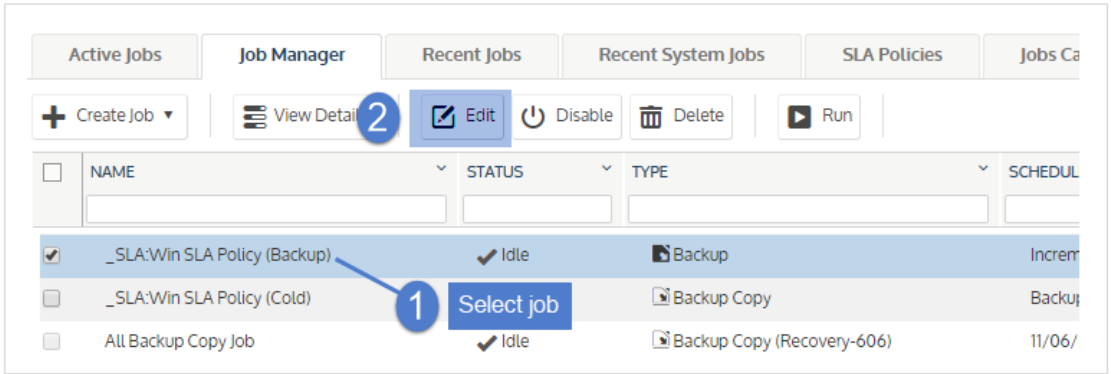

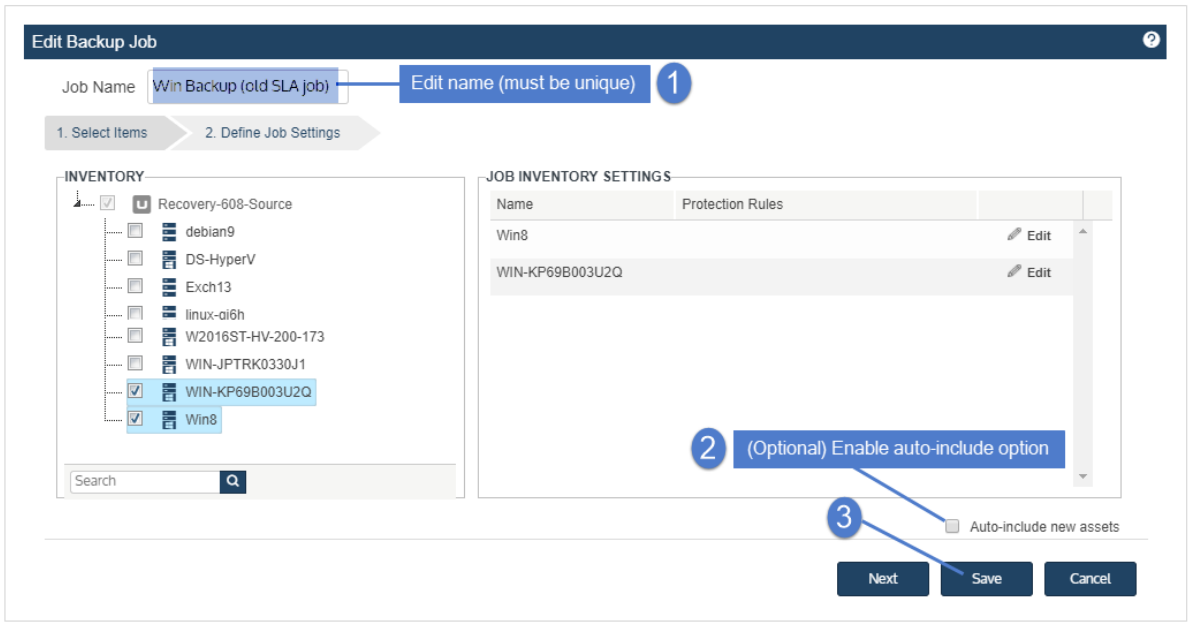
Edit Backup Job
Job Name
Win Backup (old SLA job)
Edit name (must be unique)
1
1. Select Items
2. Define Job Settings
INVENTORY
Recovery-608-Source
debian9
DS-HyperV
Exch13
linux-ai6h
W2016ST-HV-200-173
WIN-JPTRK0330J1
WIN-KP69B003U2Q
Win8
Search
JOB INVENTORY SETTINGS
Name
Protection Rules
Win8
Edit
WIN-KP69B003U2Q
Edit
2
(Optional) Enable auto-include option
3
Auto-include new assets
Next
Save
Cancel

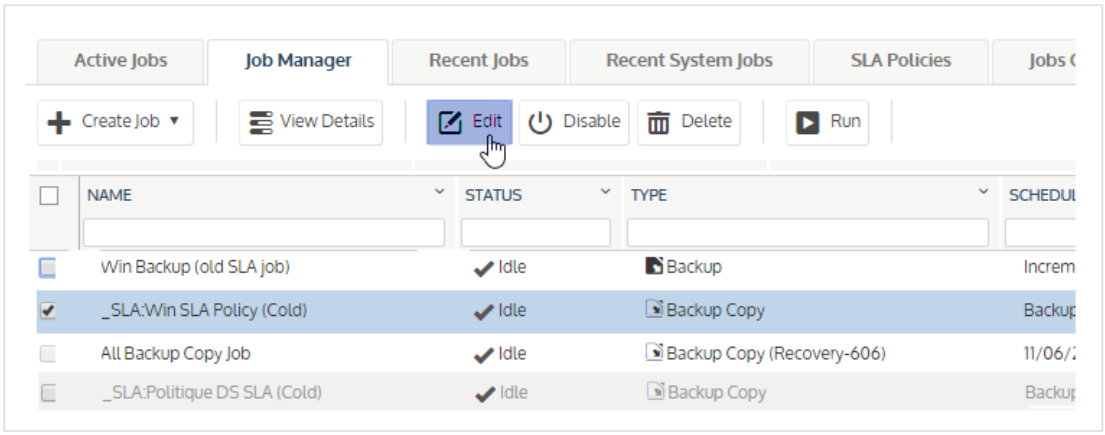
Active Jobs
Job Manager
Recent Jobs
Recent System Jobs
SLA Policies
Create Job
View Details
Edit
Disable
Delete
Run
NAME
STATUS
TYPE
SCHEDUL
Win Backup (old SLA job)
Idle
Backup
Increm
_SLA:Win SLA Policy (Cold)
Idle
Backup Copy
Backup
All Backup Copy Job
Idle
Backup Copy (Recovery-606)
_SLA:Politique DS SLA (Cold)
Idle
Backup Copy

Managing active jobs

The Active Jobs tab on the Jobs page provides a real-time listing of all jobs currently running and all jobs queued to run.

For backup copy jobs, use the Active Jobs tab to see currently running jobs and see the following for additional backup copy management options:

- For cold backup copy targets (eSATA, USB, tape, third-party cloud, attached disk, NAS, and SAN), see "Backup Copy – Cold Targets tile" on page 45 for an at-a-glance view of backup copy performance.
- For Unitrends appliance and Unitrends Cloud backup copy targets, see "Backup Copy – Hot Targets tile" on page 44 for an at-a-glance view of backup copy performance.

Notes: If you are logged in to a Unitrends appliance target, you can access additional information:

- See the "Backup Copy – Target tile" on page 46 for status and performance information about backups that were copied to the target appliance within the last seven days
- See the "Copied Assets tab" on page 88 to view, edit, and remove assets whose backups are being copied to this target appliance.

Use these procedures to manage active jobs:

Note: Monitor active iSeries jobs from the dpuconfig console interface instead. See "iSeries Backups Overview and Procedures" on page 767 for details.

To view all active jobs

1 Select **Jobs > Active Jobs**.

2 All running and queued jobs display in a list on the Active Jobs tab.

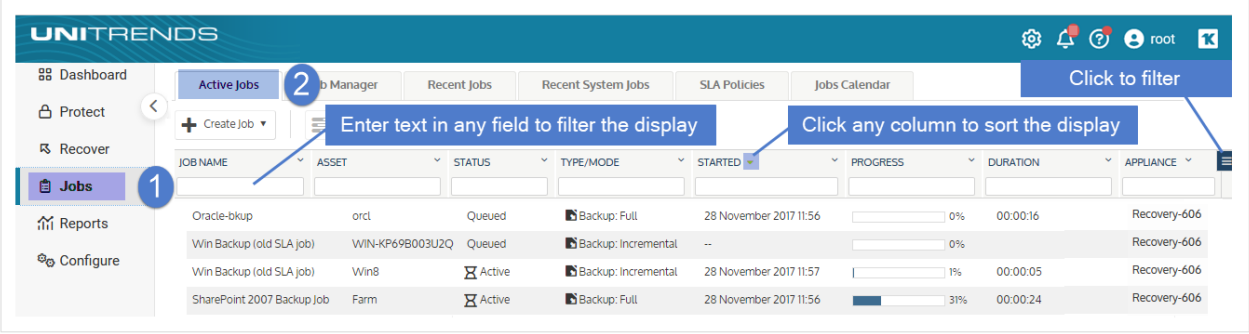

- Click on any column to sort alphabetically (a to z). To reverse the order (z to a), click again.
- The following information is given for each job:

Column	Description
Job Name	The name of the active or queued job. Schedules that were created by SLA policies adhere to the following naming conventions: • SLA policy schedule names begin with the prefix *_SLA:*, so you can easily distinguish them from manually created schedules. • SLA policy schedule names end with one of the following: – *(Backup)* for backup schedules. – *(Hot)* for hot backup copy schedules. – *(Cold)* for cold backup copy schedules.
Asset	The name of the asset whose data is being backed up, copied, or recovered.
Status	The current status of the job: • Active - The job is running now. • Queued - The job is queued and will run as soon as resources become available. • Paused - The job is paused. • Cancelled - An instance of this job was cancelled. • Successful - An instance of this job completed successfully. • Warning - An instance of this job completed with warnings. • Error - An instance of this job encountered an error and could not complete.
Type/Mode	Job type (Backup, Backup Copy, Import, or Restore) and mode (Full, Incremental, Differential, Selective, or Bare Metal). **Note:** VM replicas and Windows replicas are kept up-to-date by applying backups of the original asset as they run. The appliance applies backups by running a restore job. Replica restores display on the Active Jobs tab as type *restore: VM Replica* or *restore: Windows replica*.

Column	Description
Started	The date and time the job began.
Progress	A graphic bar representing the completed percentage of current job's progress.
Duration	Amount of time elapsed since the job started runing.
Appliance	The appliance running the job.
Comment	Job comment.
Est Finish	Estimated date and time at which the job will finish.
Application	Application type: • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013. • For agent-based backups, contains *file-level* or *image-level*. • For System Metadata contains system information, such as appliance configuration and settings. System Metadata is copied to the target during backup copy jobs if changes to this system information are detected.
ID	System-generated ID number assigned to the job.

To view job details

1 Select **Jobs > Active Jobs**.

2 Select the job and click **View Details**.

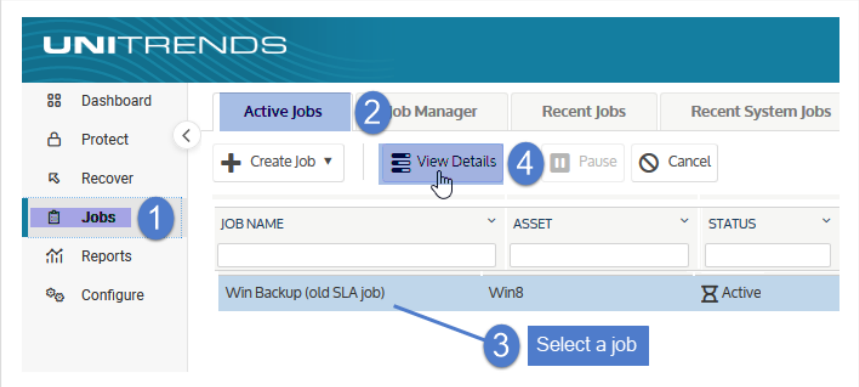

3 Job details display below.

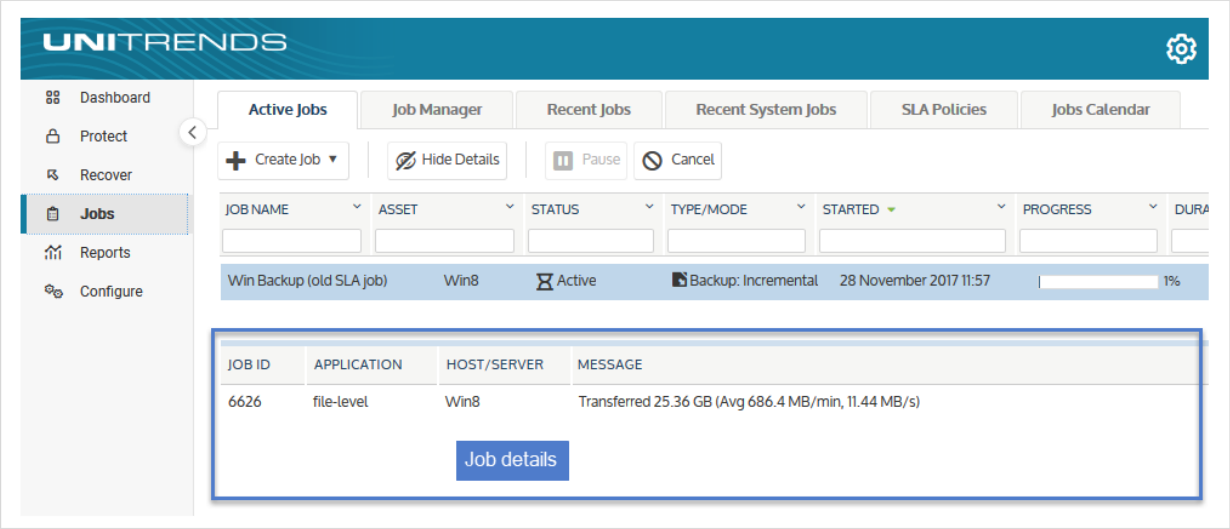

Column	Description
Job ID	A system-generated ID number assigned to the job.
Application	Application type: • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013. • For agent-based backups, contains *file level* or *image level*.

Column	Description
	• For System Metadata contains system information, such as appliance configuration and settings. System Metadata is copied to the target during backup copy jobs if changes to this system information are detected.
Host/Server	The name of the virtual host or physical server.
Message	Any system-generated message produced during the job.

4 Click **Hide Details** to return to the original page view.

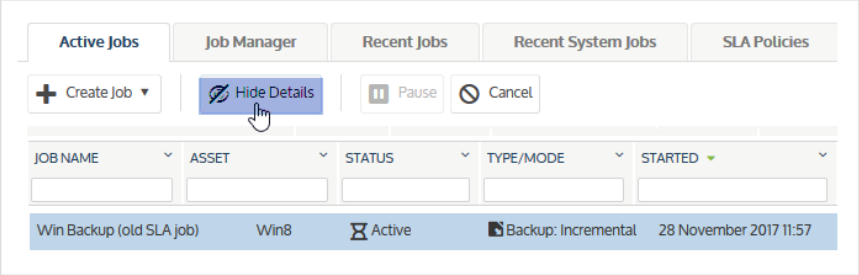

To pause a job

Use this procedure to pause queued jobs. Pausing queued jobs can prove useful if you want to push other jobs to the top of the queue. (You cannot pause a job that is currently running.)

1 Select **Jobs > Active Jobs**.

2 Select a queued job in the list.

3 Click **Pause**.

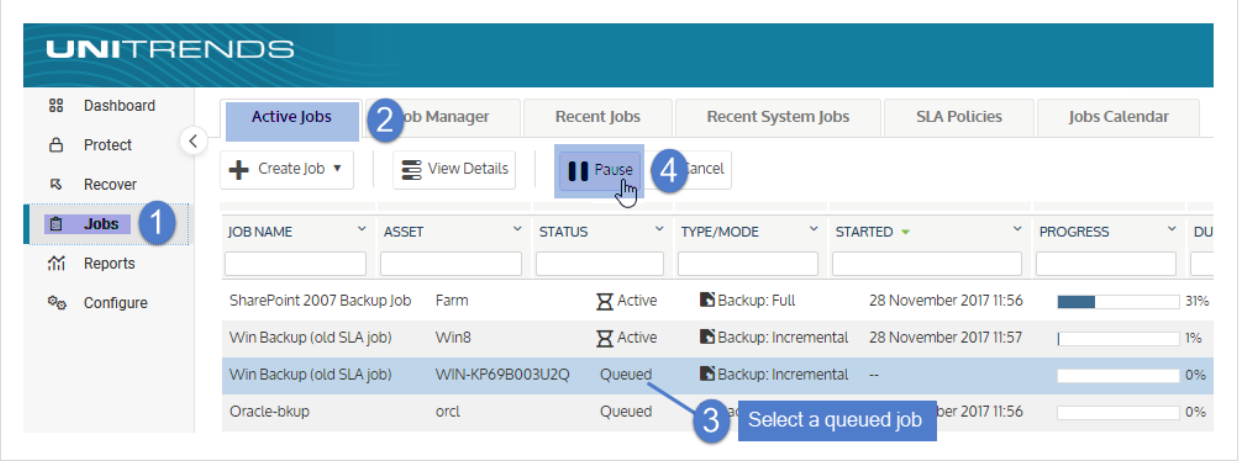

4 The job is paused.

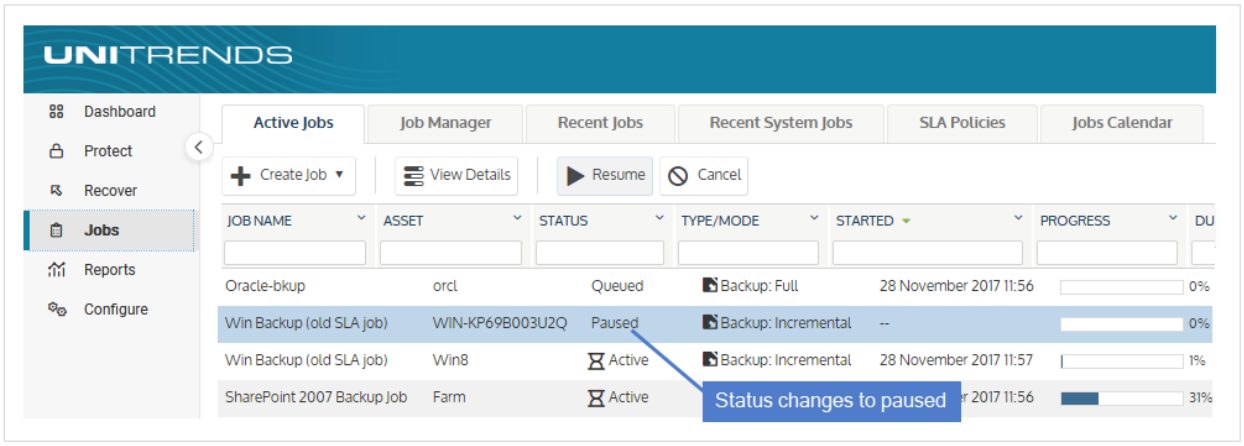

To resume a job

1 Select **Jobs** > **Active Jobs**.

2 Select a paused job in the list.

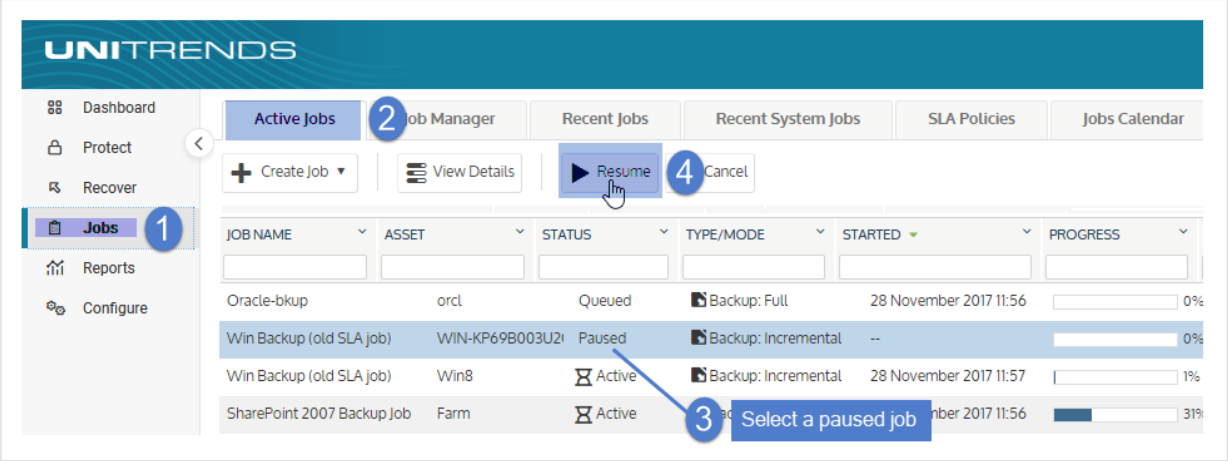

3 Click **Resume** to queue the job.

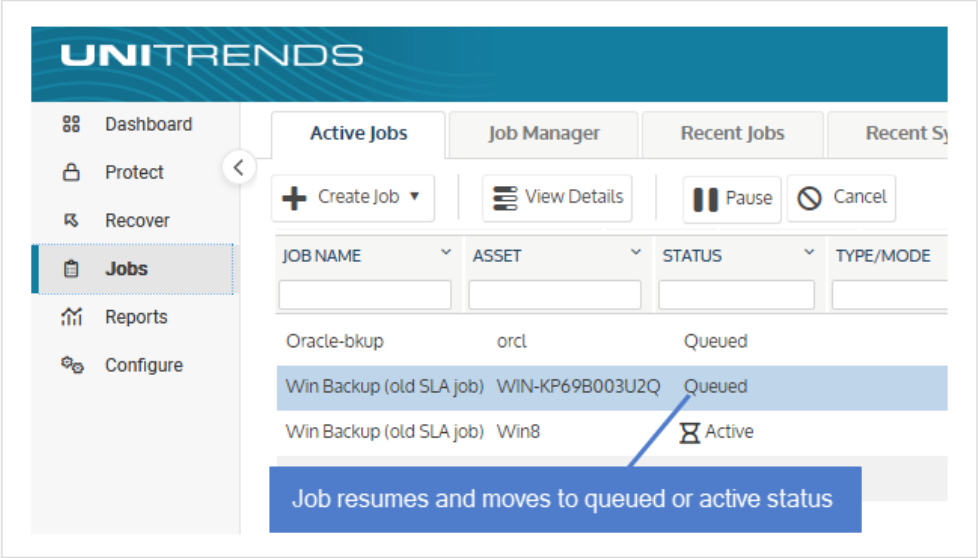

To cancel an active job

Notes:

Do not use this procedure for these job types:

- Hot backup copy jobs – These jobs run to copy backups to the Unitrends Cloud or to another Unitrends appliance.

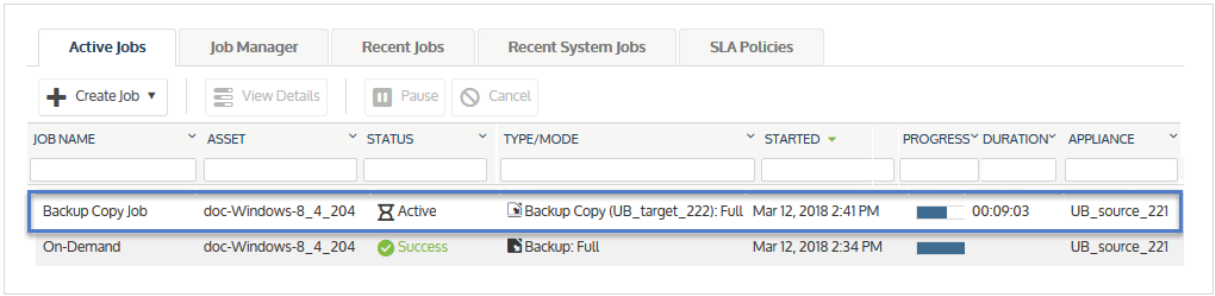

If you cancel a hot copy job by using the Cancel button on the Active Jobs page, the appliance automatically creates a new job to replace the one you canceled. To temporarily stop copying backups, suspend backup copies instead (as described in "To suspend hot backup copies" on page 266). Use the procedure "To resume hot backup copies" on page 268 to start sending hot copies again. Note that copies of all backups that ran while copies were suspended will be sent to the hot backup copy target once you resume. You cannot skip copying a specific backup.

- Replica restore jobs – These jobs run to apply backups to Windows or VM replicas.

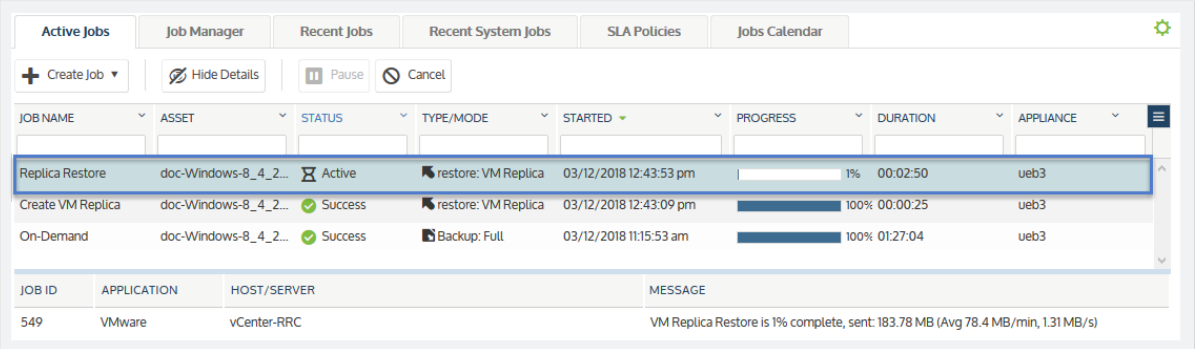

  - Windows replica restore jobs – If you cancel a Windows replica restore job by using the Cancel button on the Active Jobs page, the appliance automatically creates a new job to replace the one you canceled. To temporarily stop applying backups to a Windows replica, bring the replica into audit mode instead (as described in "Auditing a Windows replica" on page 1015). Use the procedure "To exit audit mode" on page 1020 to start applying backups again.
  - VM replica restore jobs – If you cancel a VM replica restore job by using the Cancel button on the Active Jobs page, the replica may become invalid and need to be recreated (for details, see "VM replica modes" on page 902). To temporarily stop applying backups to a VM replica, bring the replica into audit mode instead (as described in "Auditing a VM replica" on page 893). Use the procedure "To exit audit mode" on page 895 to start applying backups again.

Note that all backups that ran while the replica was in audit mode will be applied to the replica upon exiting audit mode. You cannot skip applying a specific backup to a replica.

1. Select **Jobs > Active Jobs**.
2. Select a job in the list.

3 Click **Cancel**.

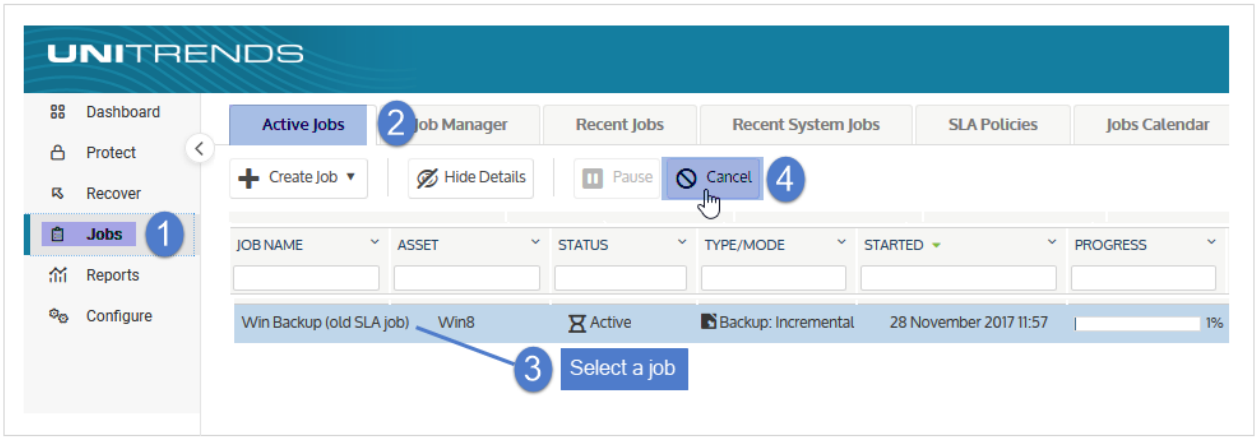

4 The job is canceled and its status changes to *Error*.

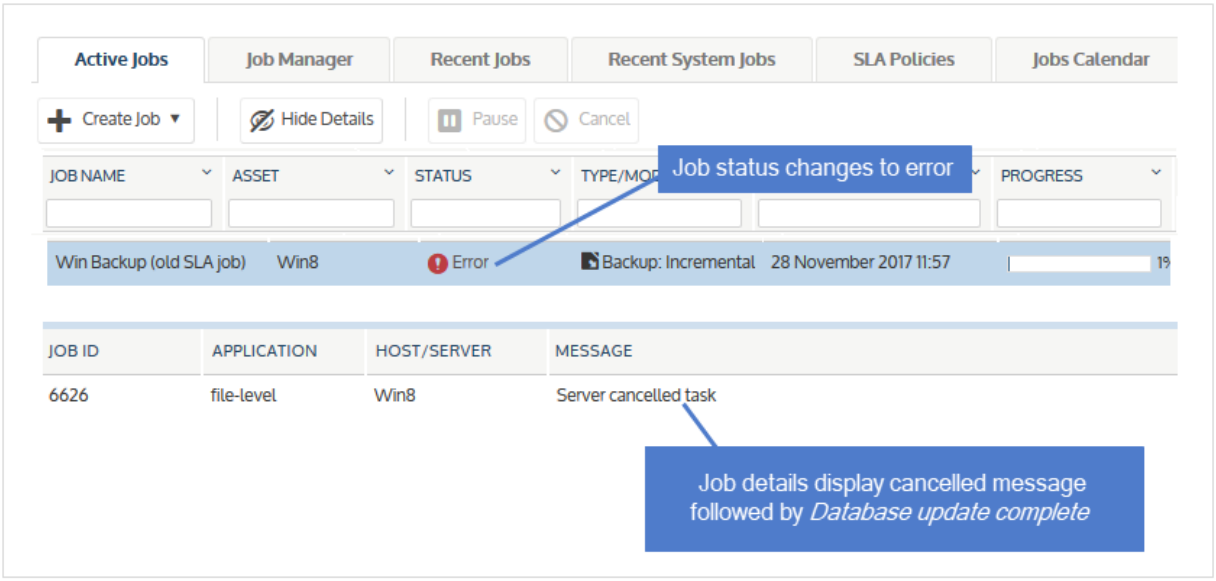

Viewing recent jobs

The Recent Jobs tab captures the results of job activity over the last 7 days.

Notes: See these additional resources while working with recent jobs:

- After a backup or backup copy job completes, the resulting backup or copy displays in the Backup Catalog on the **Recover > Backup Catalog** page. For details on the Backup Catalog and your recovery options, see "Backup Catalog tab" on page 60 and the "Recovery Overview" chapter.

- Reports provide additional detail about the jobs that ran in the last 7 days, as well as information about older jobs. For details, see "Reports" on page 1307.

To view recent jobs

1. Select **Jobs > Recent Jobs**.
2. Jobs that ran over the last 7 days display in a list on the Recent Jobs tab.

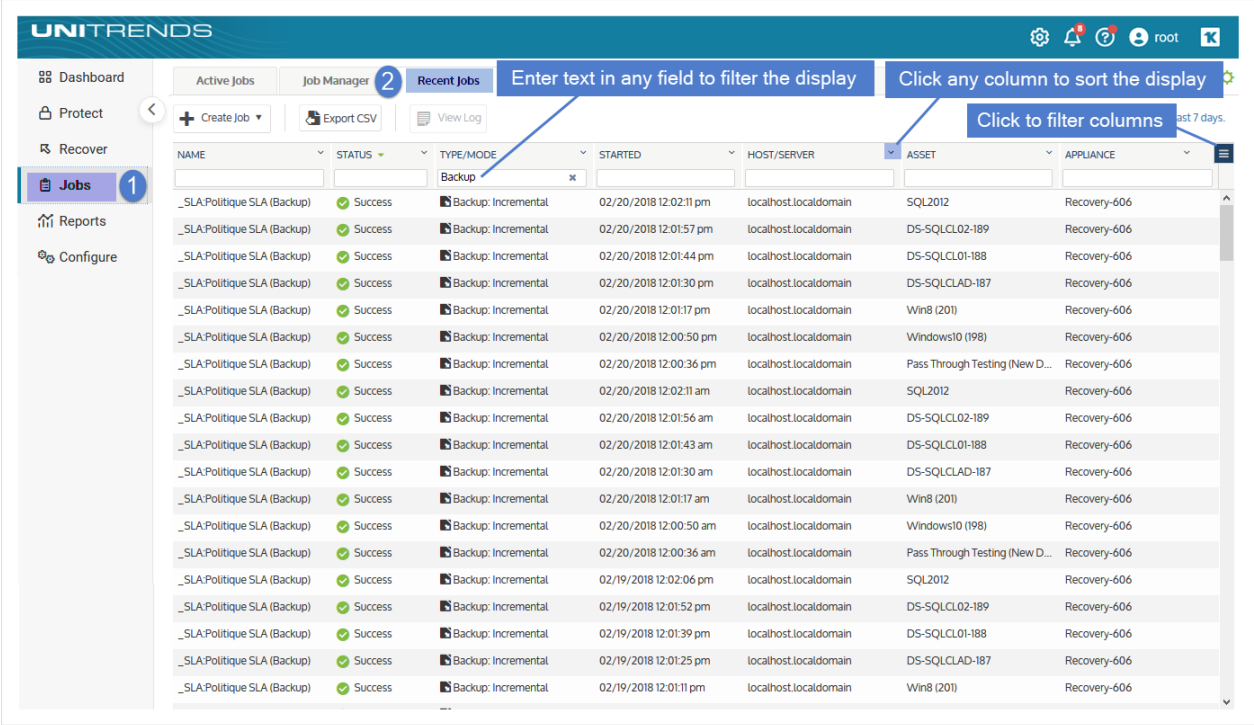

- The following information is given for each job:

Column	Description
Name	The name of the job.
Status	The final status of the job. Each job receives a color code, based on its status: • Green jobs completed successfully. • Yellow jobs completed with a warning. • Red jobs completed with an error.
Type/Mode	The job type (Backup, Backup Copy, or Restore) and mode (Full, Incremental,

Column	Description
	Differential, Selective, or Bare Metal).
Started	The date and time the job began.
Host/Server	The name of the virtual host or physical server.
Asset	Asset for which the job ran.
Appliance	The appliance that ran the job.
Application	Application type: • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013. • For agent-based backups, contains *file-level* or *image-level*. • For System Metadata contains system information, such as appliance configuration and settings. System Metadata is copied to the target during backup copy jobs if changes to this system information are detected.
Size (MB)	Data size, in megabytes.
Files	Number of files, if applicable.
ID	System-generated ID number assigned to the job.

- Click on any column to sort alphabetically (a to z). To reverse the order (z to a), click again.
- Click **Export CSV** to export recent jobs data to a CSV (Excel) file. (To view and/or export a list of files in a file-level backup, see "Backup History report" on page 1320.)

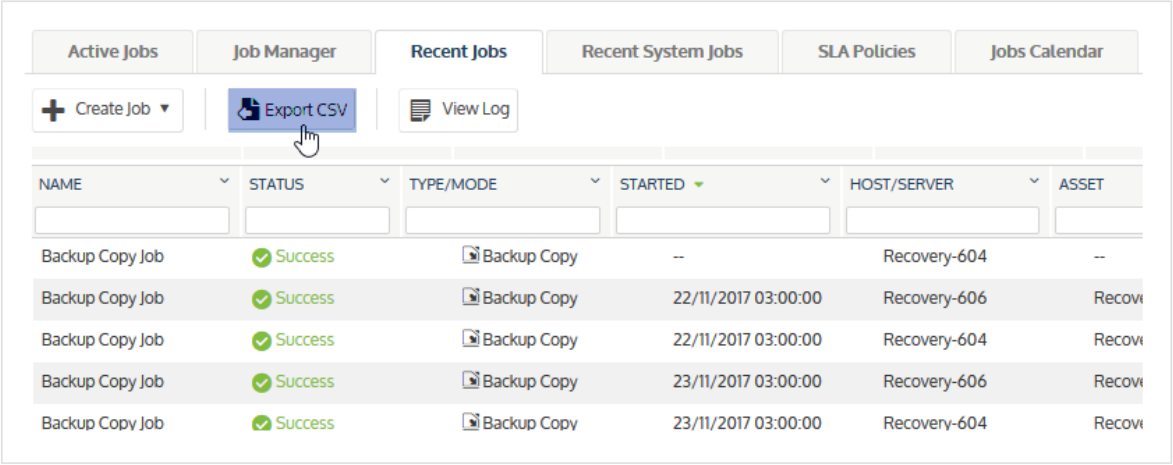
Active Jobs
Job Manager
Recent Jobs
Recent System Jobs
SLA Policies
Jobs Calendar
Create Job
Export CSV
View Log
NAME
STATUS
TYPE/MODE
STARTED
HOST/SERVER
ASSET
Backup Copy Job Success Backup Copy -- Recovery-604 --
Backup Copy Job Success Backup Copy 22/11/2017 03:00:00 Recovery-606
Backup Copy Job Success Backup Copy 22/11/2017 03:00:00 Recovery-604
Backup Copy Job Success Backup Copy 23/11/2017 03:00:00 Recovery-606
Backup Copy Job Success Backup Copy 23/11/2017 03:00:00 Recovery-604

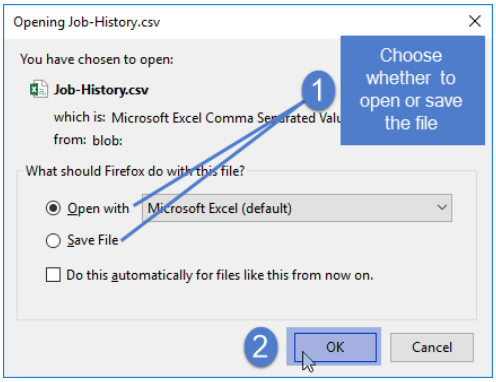
Opening Job-History.csv
You have chosen to open:
Job-History.csv
which is: Microsoft Excel Comma Separated Valu
from: blob:
What should Firefox do with this file?
Open with
Microsoft Excel (default)
Save File
Do this automatically for files like this from now on.
OK
Cancel
1
2
Choose whether to open or save the file

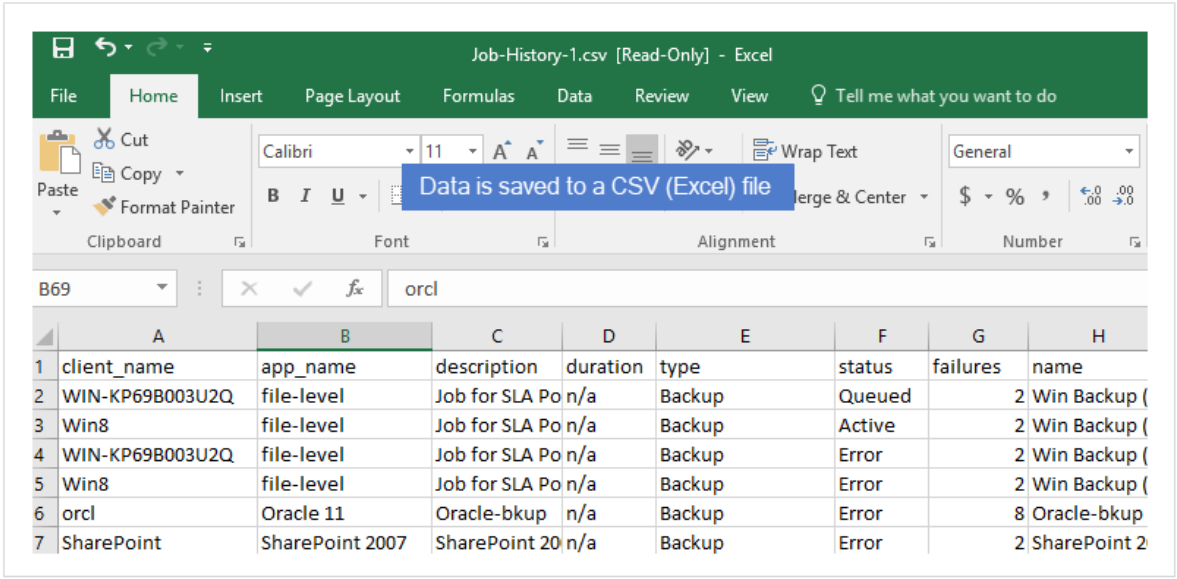
Job-History-1.csv [Read-Only] - Excel
File
Home
Insert
Page Layout
Formulas
Data
Review
View
Tell me what you want to do
Cut
Copy
Format Painter
Paste
Clipboard
Calibri
11
Wrap Text
General
Font
Alignment
Number
Data is saved to a CSV (Excel) file
B69
orcl
client_name app_name description duration type status failures name
WIN-KP69B003U2Q file-level Job for SLA Po n/a Backup Queued 2 Win Backup (
Win8 file-level Job for SLA Po n/a Backup Active 2 Win Backup (
WIN-KP69B003U2Q file-level Job for SLA Po n/a Backup Error 2 Win Backup (
Win8 file-level Job for SLA Po n/a Backup Error 2 Win Backup (
orcl Oracle 11 Oracle-bkup n/a Backup Error 8 Oracle-bkup
SharePoint SharePoint 2007 SharePoint 20 n/a Backup Error 2 SharePoint 2

- Select a job and click **View Log** to view its log file.

Note: Linux and NAS assets – Due to a know issue, Exclusion Lists do not display in the log file. (Any Inclusion Lists do display in the log.)

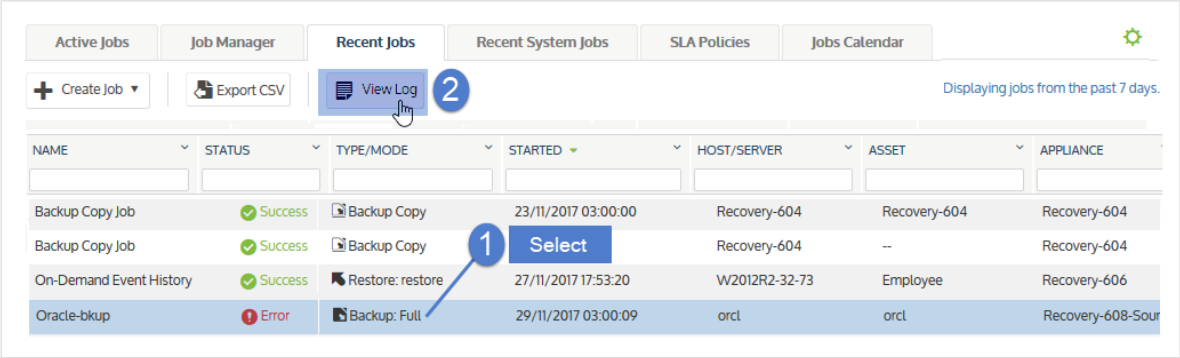

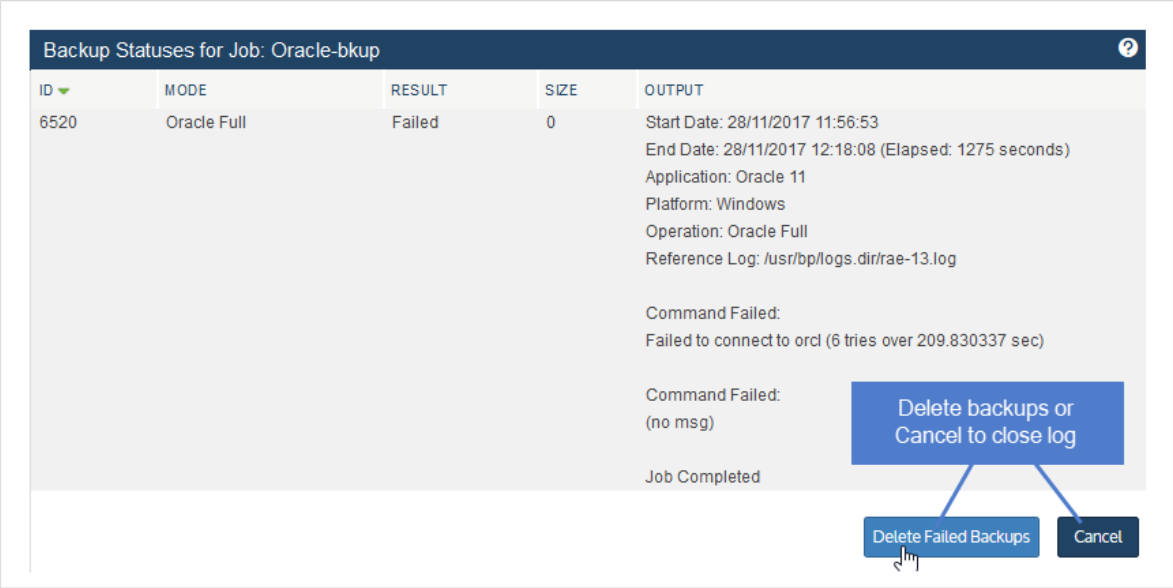

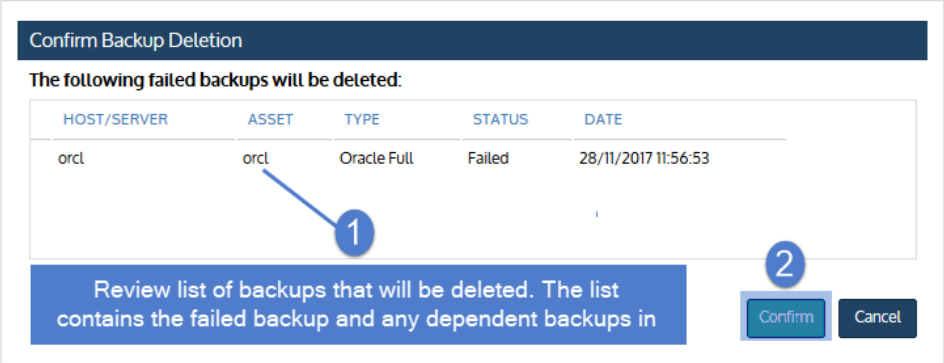

Viewing system jobs

The Recent System Jobs tab captures and displays the results of *system-level* job activity over the last seven days. If more than 250 jobs have run over the last 7 days, the results are limited to the last 250 jobs.

To view recent system jobs

1 Select **Jobs > Recent System Jobs**.

2 Jobs that ran over the last 7 days display in a list on the Recent System Jobs tab.

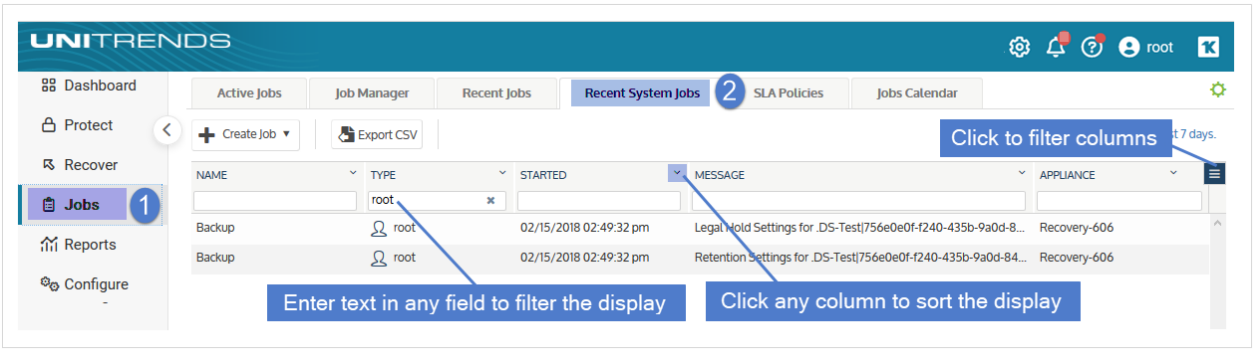

- The following information is given for each job:

Column	Description
Name	The name of the job.
Type	The type of user account that initiated the job, *System* or *Root*.
Started	The date and time the job began.

Column	Description
Message	Displays system-generated messages related to recent system jobs.
Appliance	The appliance that ran the job.

- Click on any column to sort alphabetically (a to z). To reverse the order (z to a), click again.
- Click **Export CSV** to export system jobs data to a CSV (Excel) file.

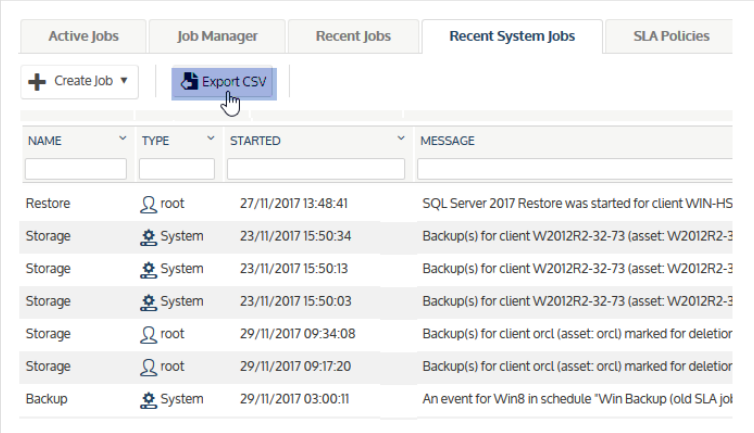

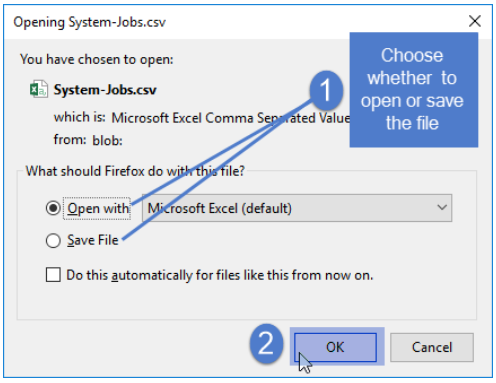

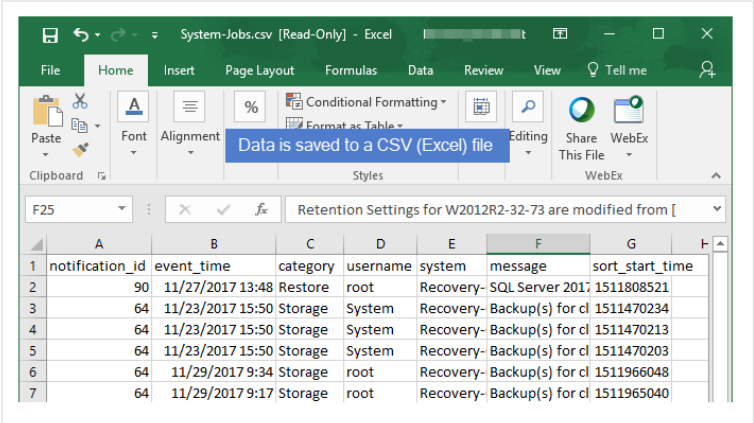

Viewing job details

In some cases, you may want to view job details to gather information for troubleshooting or to verify that the job ran as expected. For example, you can check details of a host-level backup to see which quiesce method or transport mode was used.

Use these procedures to view the details of backup, backup copy, and recovery jobs:

- "To view details of a recent job"
- "To view details of an older job" on page 623
- "To view details from the Backup Catalog tab" on page 626

To view details of a recent job

Use this procedure to quickly view details of a job that ran in the last 7 days.

1. Click **Jobs > Recent Jobs**.
2. Select the job and click **View Log**.

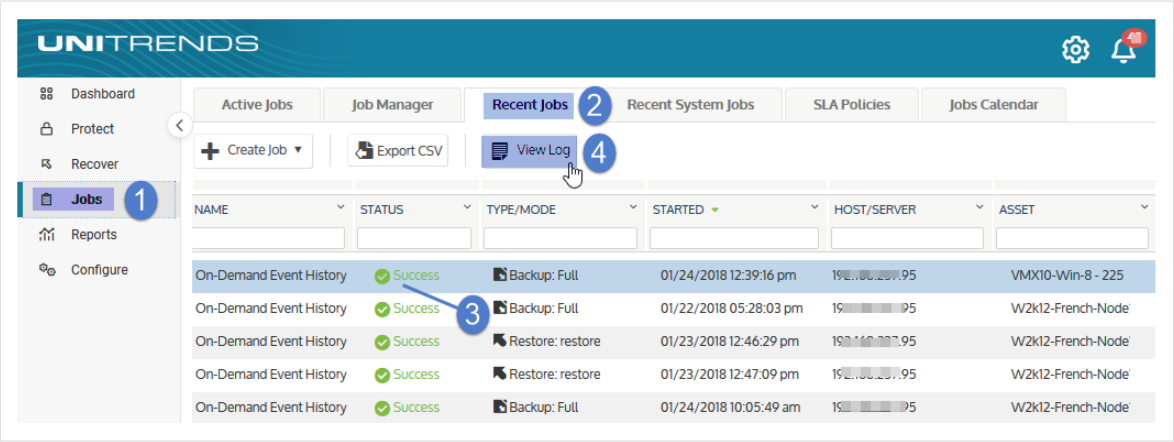

3 Details display in the job status dialog.

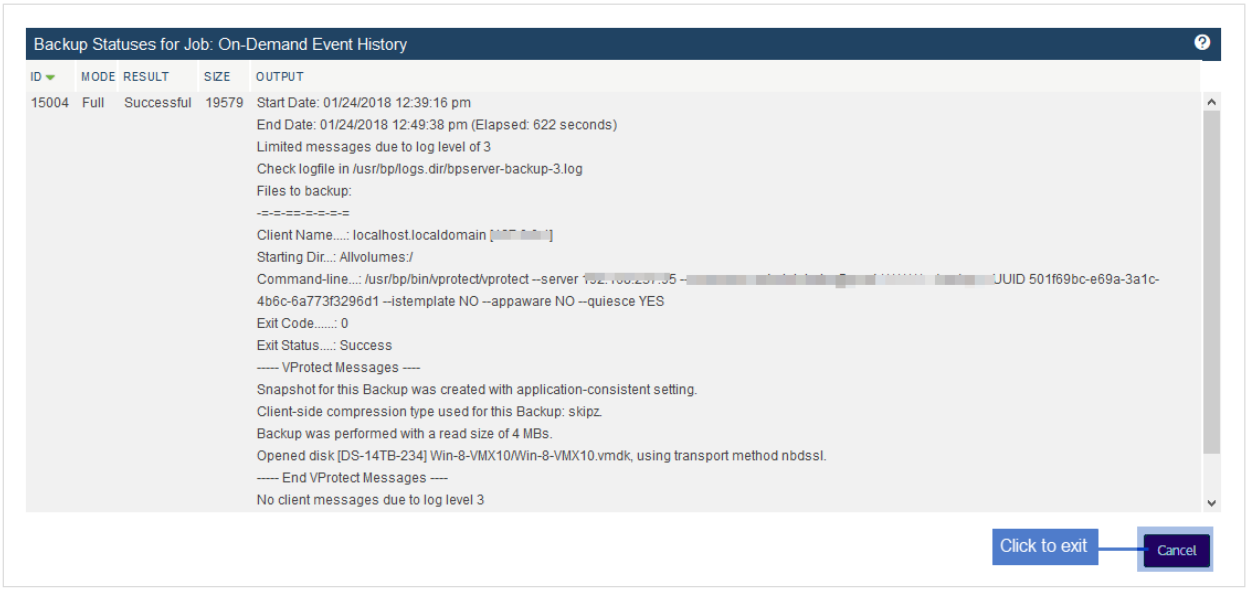

To view details of an older job

Use this procedure to view details of a job that ran more than 7 days ago.

1 Access the applicable job report:

- For backup jobs, click **Reports > Backup > Backup History**.
- For recovery jobs, click **Reports > Recover > Recovery History**.
- For hot copy jobs, click **Reports > Backup Copy > Backup Copy – Hot Targets**.
- For cold copy jobs, click **Reports > Backup Copy > Backup Copy – Cold Targets**.

- For copy data management jobs, click **Reports > Recover > Recovery Assurance**.

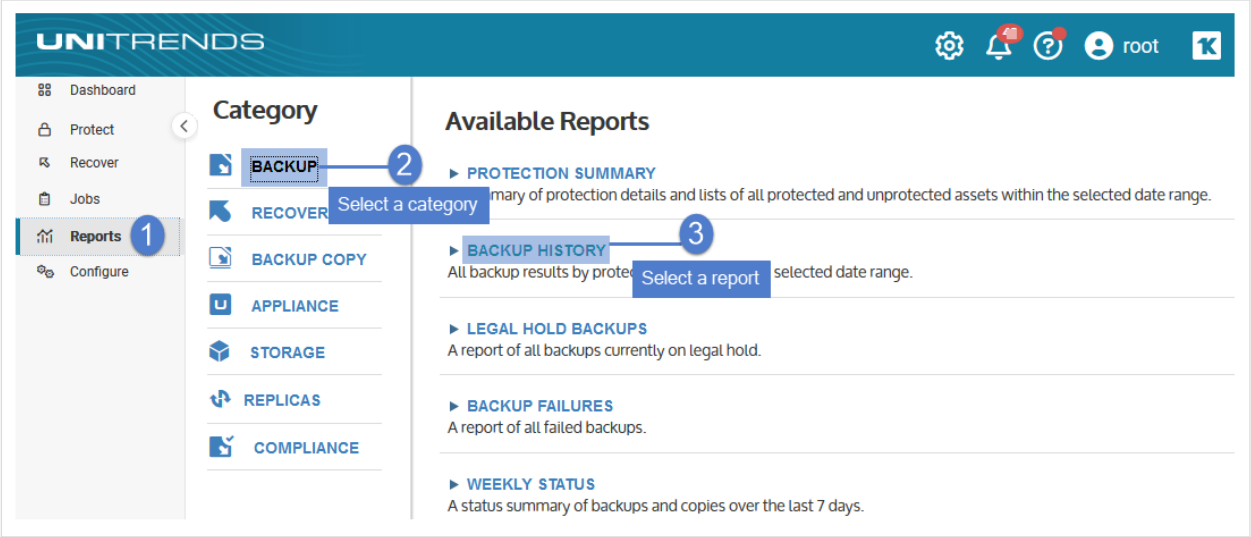

2 Run the report:

- Enter a date range.
- Click **Generate Report**.

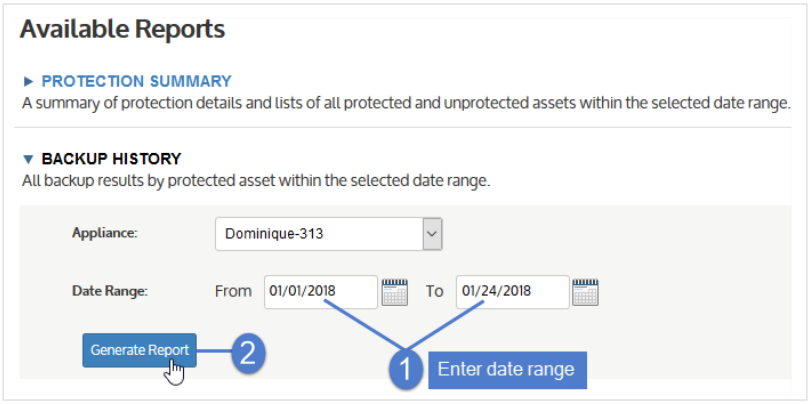

3 Click the job in the list. Details display in the job status dialog.

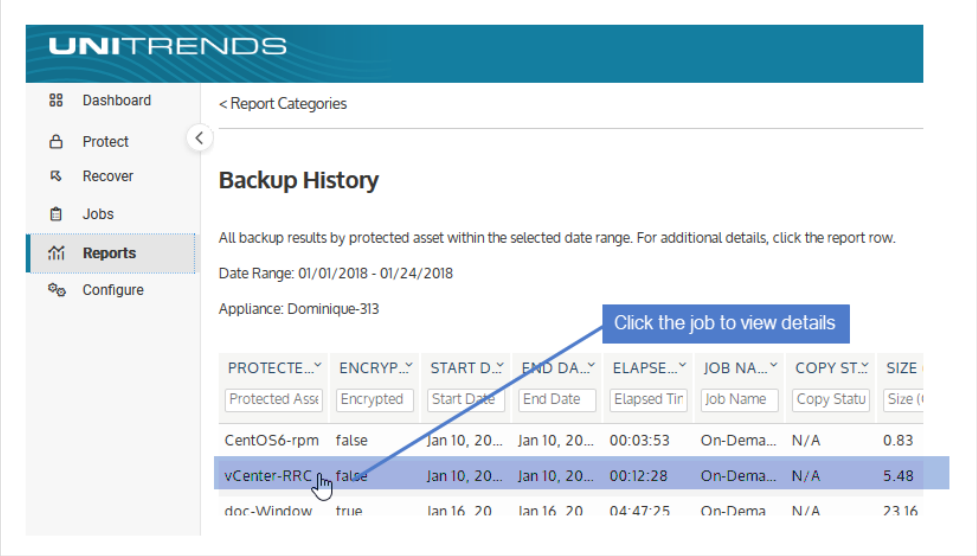
UNITRENDS
Dashboard
Protect
Recover
Jobs
Reports
Configure
< Report Categories
Backup History
All backup results by protected asset within the selected date range. For additional details, click the report row.
Date Range: 01/01/2018 - 01/24/2018
Appliance: Dominique-313
Click the job to view details
PROTECTE...
ENCRYP...
START D...
END DA...
ELAPSE...
JOB NA...
COPY ST...
SIZE
Protected Ass
Encrypted
Start Date
End Date
Elapsed Tir
Job Name
Copy Statu
CentOS6-rpm false Jan 10, 20... Jan 10, 20... 00:03:53 On-Dema... N/A 0.83
vCenter-RRC false Jan 10, 20... Jan 10, 20... 00:12:28 On-Dema... N/A 5.48
doc-Window true Jan 16, 20 Jan 16, 20 04:47:25 On-Dema N/A 23.16

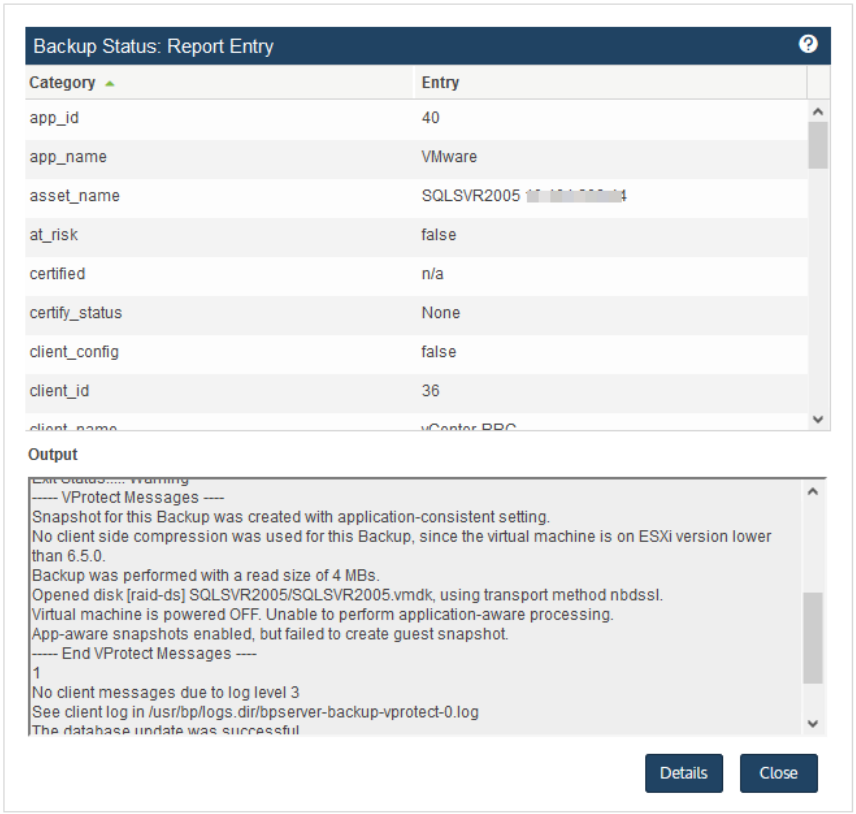
Backup Status: Report Entry
Category
Entry
app_id 40
app_name VMware
asset_name SQLSVR2005
at_risk false
certified n/a
certify_status None
client_config false
client_id 36
Output
----- VProtect Messages -----
Snapshot for this Backup was created with application-consistent setting.
No client side compression was used for this Backup, since the virtual machine is on ESXi version lower than 6.5.0.
Backup was performed with a read size of 4 MBs.
Opened disk [raid-ds] SQLSVR2005/SQLSVR2005.vmdk, using transport method nbdssl.
Virtual machine is powered OFF. Unable to perform application-aware processing.
App-aware snapshots enabled, but failed to create guest snapshot.
----- End VProtect Messages -----
1
No client messages due to log level 3
See client log in /usr/bp/logs.dir/bpserver-backup-vprotect-0.log
Details
Close

To view details from the Backup Catalog tab

While working in the Backup Catalog, you can quickly view job details by selecting jobs and clicking **View Details**. For imported backups, details include the date when the original backup ran and the date the backup was imported.

Example:

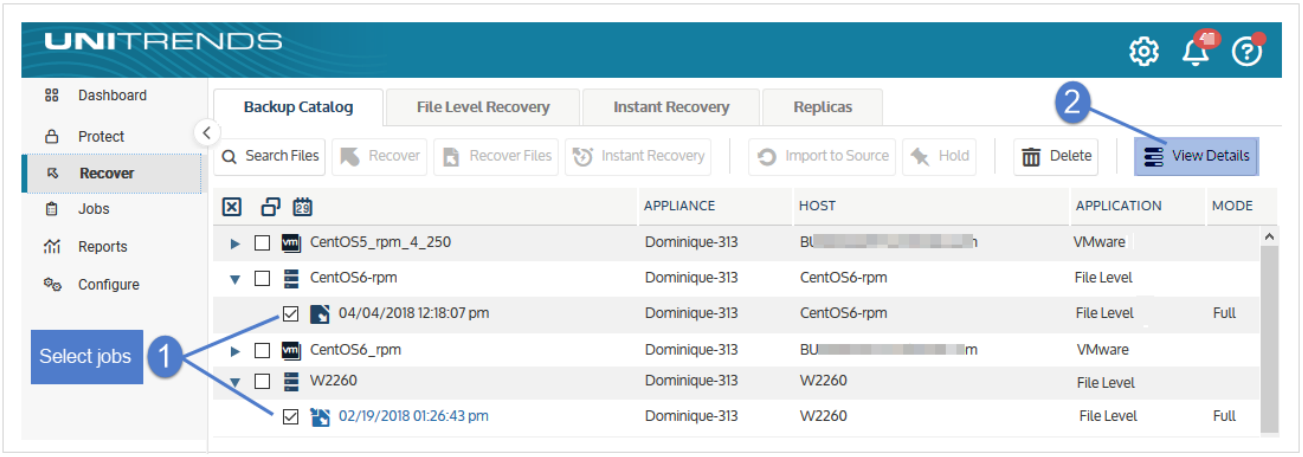

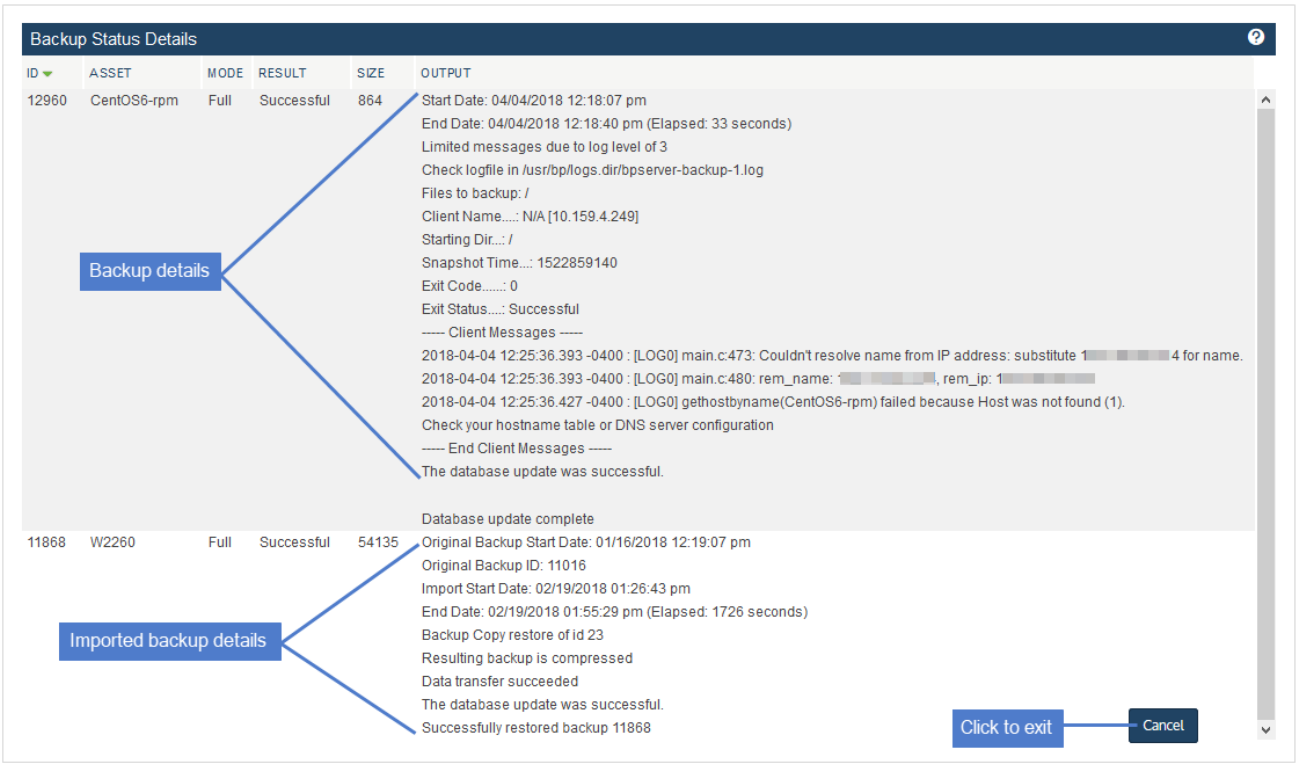

Deleting backups and backup copies

If necessary, you can manually delete backups, imported backup copies, and hot backup copies from the appliance. Deleting a full backup or backup copy also deletes any associated incrementals and differentials in the backup group. (For details, see "Backup groups" on page 98.)

Notes:

- Backups and backup copies are logically deleted immediately and no longer display in the UI. They are physically deleted later by the appliance's purge process. Space reclaimed by removing backups and copies happens when the purge job completes.
- Deleting a single cold backup copy from the target media is not supported. Instead you must erase or prepare the target, which removes all backup copies stored on the media. For details, see one of the following:
  - "To prepare tapes for use with an autochanger device" on page 273 for tape autochanger targets.
  - "To initialize and erase cold backup copy media" on page 271 for all other cold backup copy targets, including tape drive devices that do not have an autochanger.

You can delete backups by using the Backup Browser or the Backup Catalog. The Backup Browser provides appliance-level search capability, advanced search options, and faster filtering performance. The Backup Catalog enables you to browse by asset only. With the Backup Browser you can quickly search for and delete all failed backups. This is much easier than looking for failures under each asset in the Backup Catalog.

To delete imported backups and backup copies, you must use the Backup Catalog.

See the following procedures for details:

- "To delete backups by using the Backup Browser"
- "To delete backups and imported backup copies by using the Backup Catalog" on page 630
- "To delete hot backup copies"

To delete backups by using the Backup Browser

Use this procedure to manually delete backups.

Notes:

- Backups that have been placed on hold cannot be manually deleted. You must first remove the hold status. (See "To unhold backups by using the Backup Browser" on page 640 for details.)
- Backup copies and imported backups cannot be viewed in the Backup Browser. To delete a backup copy or imported backup, use the Backup Catalog instead.

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog** and click **Backup Browser**.

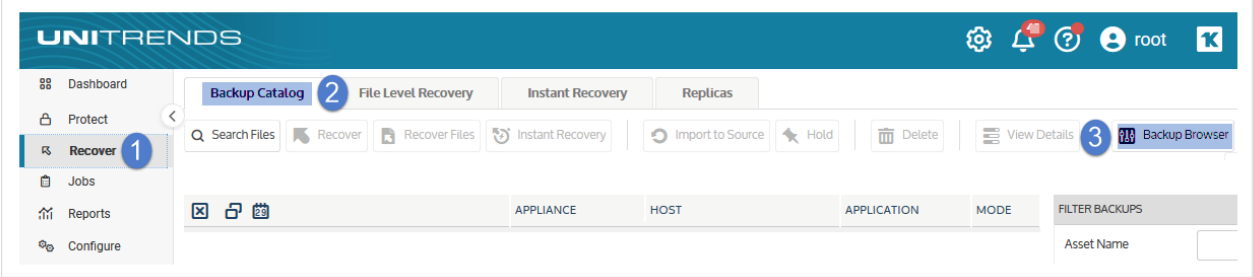

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display:

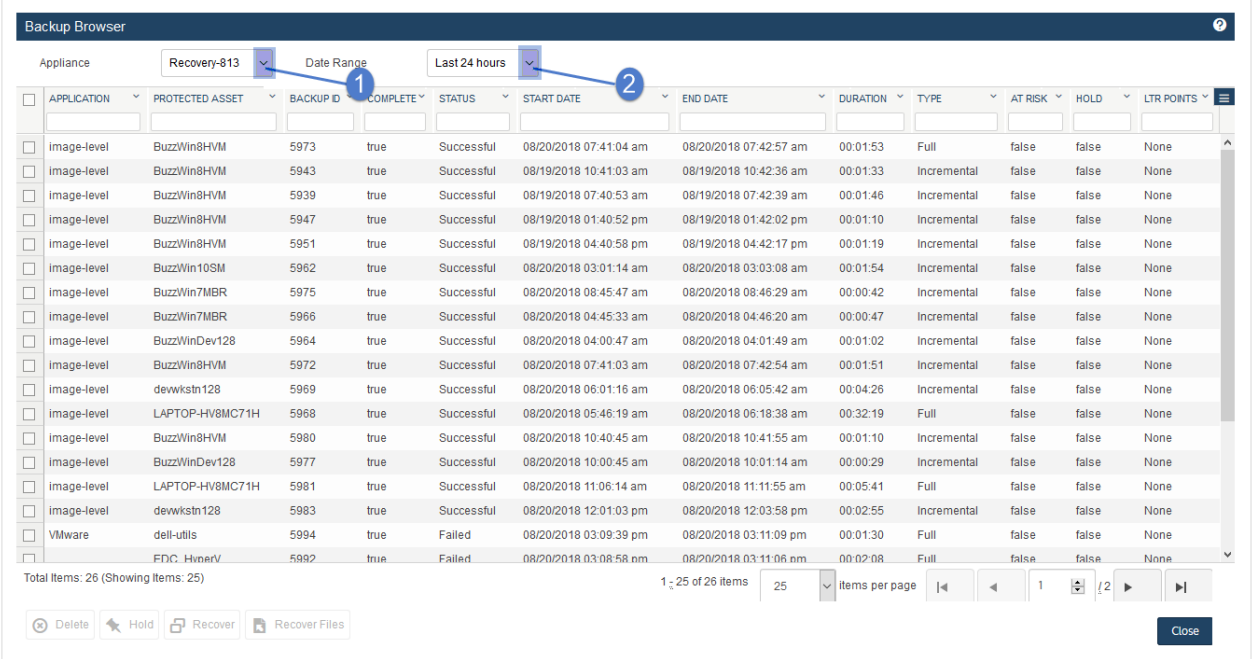

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Check boxes to select the backups to delete. (To select all, check the box at the top of the column.)

6 Click **Delete**.

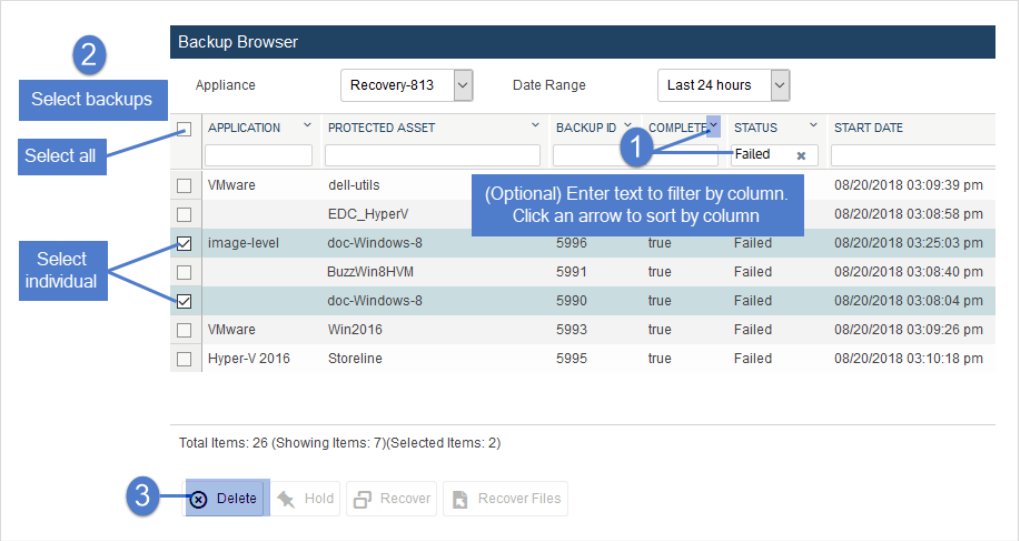

7 Review the backups that will be deleted. Click **Delete Backups** (or **Cancel** to exit without deleting).

Note: Deleting a full backup or backup copy also deletes any associated incrementals and differentials in the backup group.

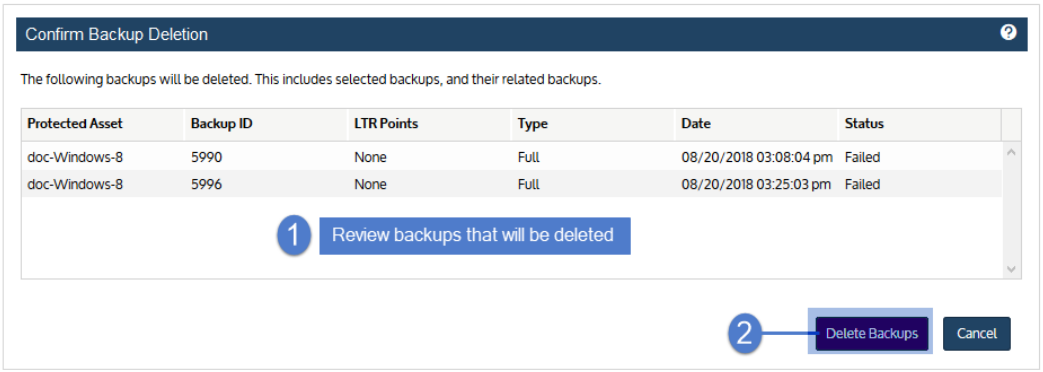

8 Backups are deleted.

Note: Backups are physically deleted later by the appliance's purge process. Space reclaimed by removing backups happens when the purge job completes.

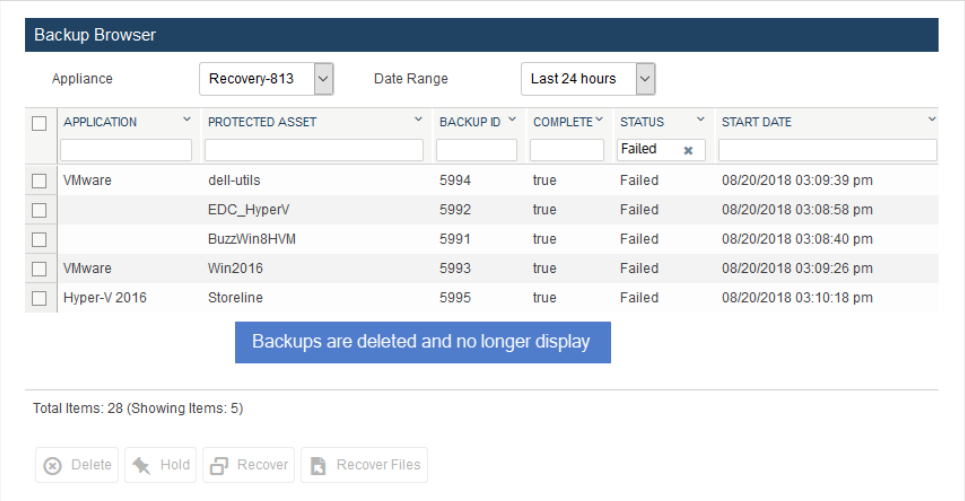

To delete backups and imported backup copies by using the Backup Catalog

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog**.

   Protected assets display. (Optional) Use the Filter Backups fields to filter the display. For details, see "Backup Catalog tab" on page 60.
3. Click to expand assets to view their backups and imported backup copies.
4. Check boxes to select the backups and imported backup copies to delete.
5. Click **Delete**.

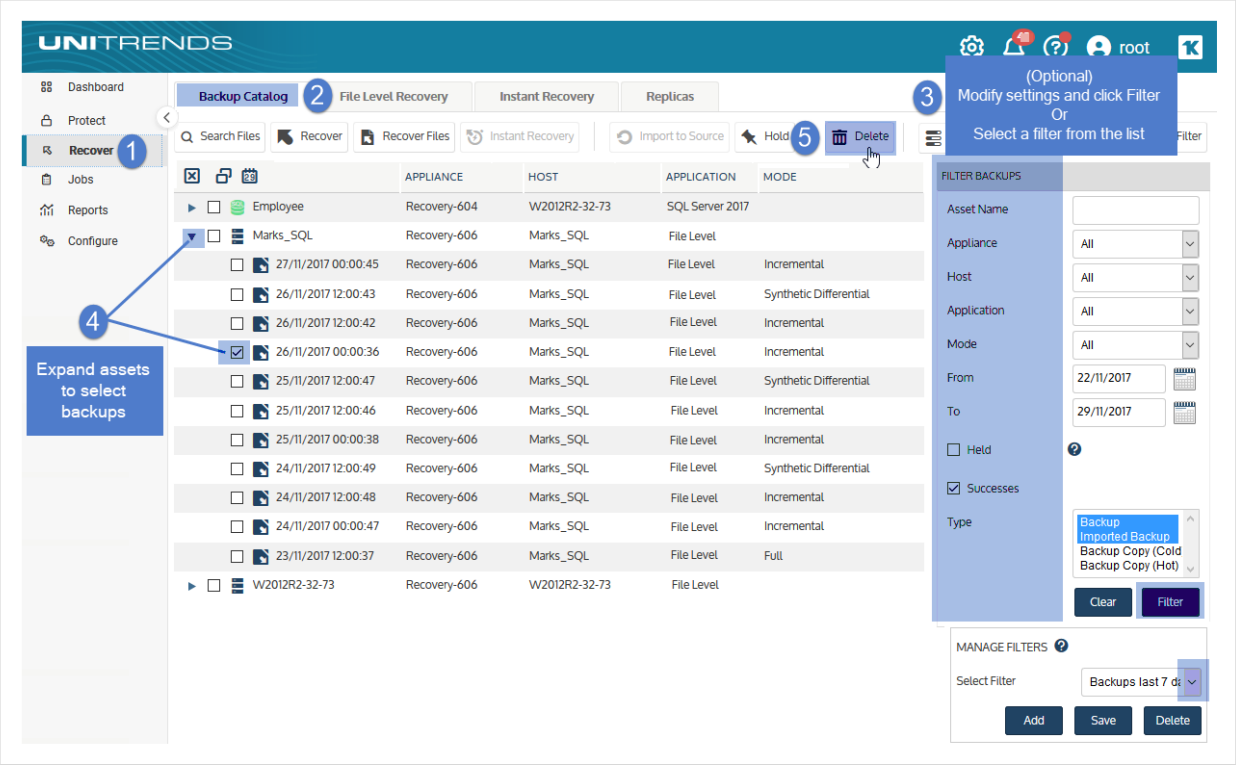

6 Review the list of backups that will be deleted.

7 Click **Confirm** to delete the backups.

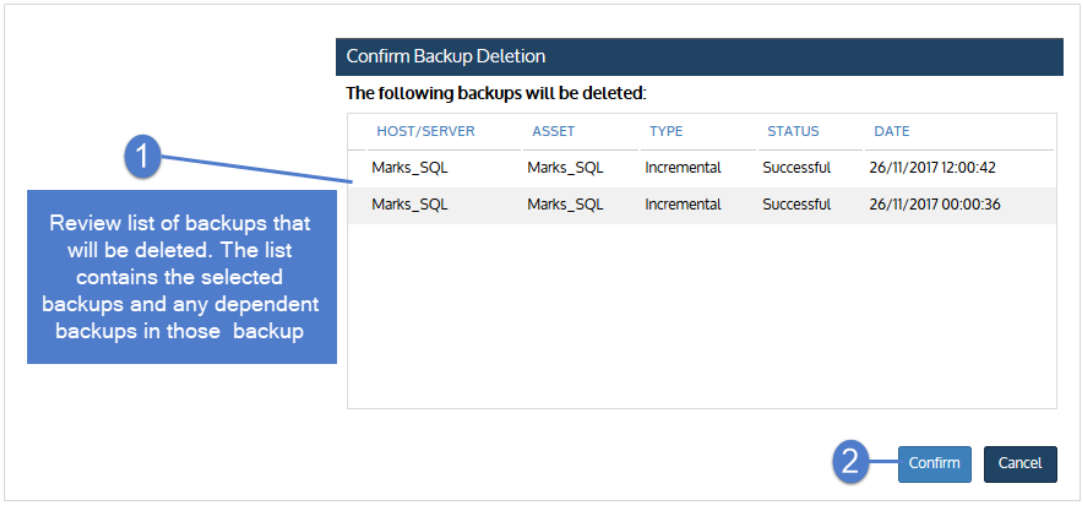

Backups are deleted and no longer display in the Backup Catalog. Note that the backups are logically deleted immediately and physically deleted later by the appliance's purge process. Space reclaimed by removing the backups happens when the purge job completes.

To delete hot backup copies

1 Log in to either the backup appliance or the backup copy target appliance.

2 Select **Recover > Backup Catalog**.

3 In the Filter Backups area to the right, select **Backup Copy (Hot)** in the Type list.

(Optional) Enter other filter options. For details, see "Backup Catalog tab" on page 60.

4 Click **Filter**.

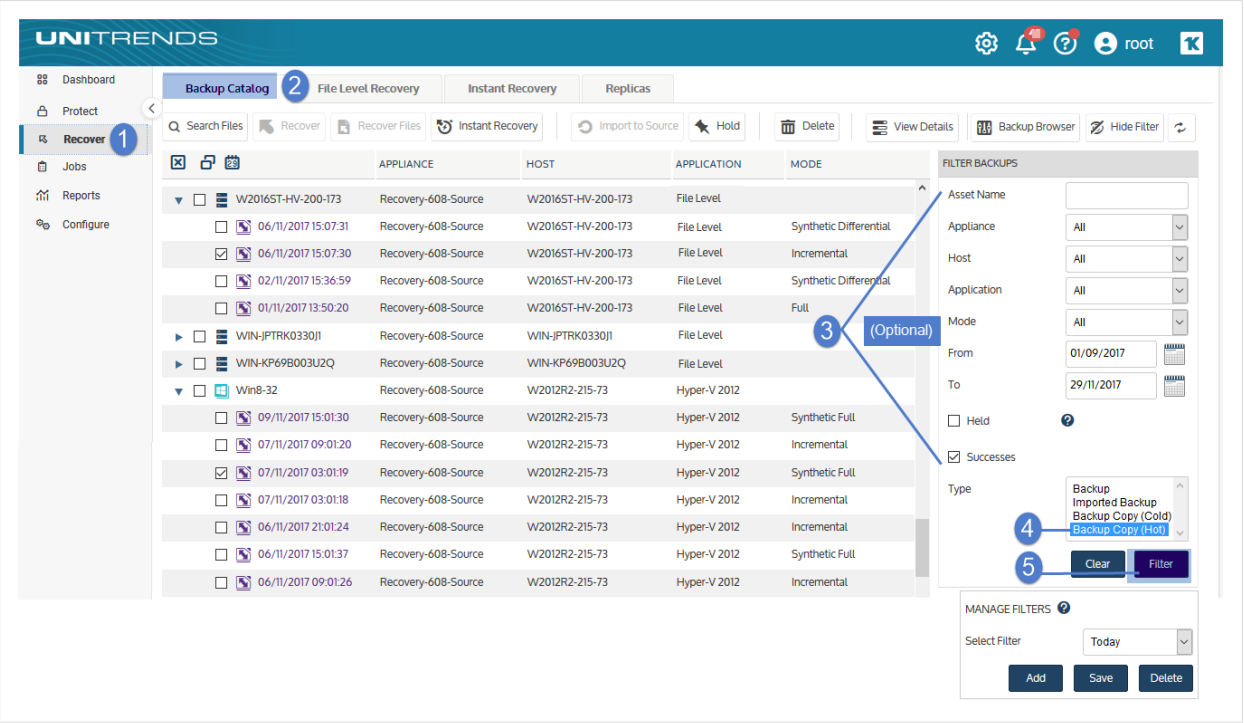

- Assets with backup copies meeting the criteria you specified display in the Backup Catalog list. The source appliance where the backup originated displays in the Appliance column.
- If the appliance is being used as both a backup appliance and a backup copy target, the catalog lists the hot copies stored on this appliance and any backups that were copied from this appliance to another hot backup copy target. (The hot backup copy target could be another appliance or the Unitrends Cloud).
- (Optional) Hover over a backup copy icon to determine whether the copy resides on this appliance or on the remote backup copy target.

  The description *Backup Copy (Hot) on Target* displays for copies that are stored on a remote target:

The description *Backup Copy (Hot)* displays for copies that are stored on the local appliance:

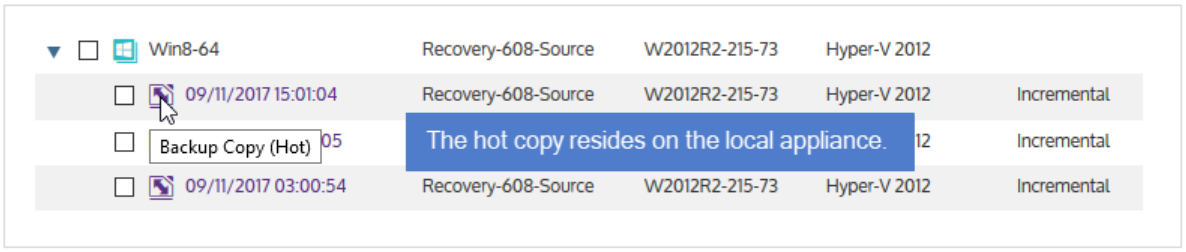

5 Expand assets to view backup copies. Check boxes to select the backup copies to delete.

6 Click **Delete**

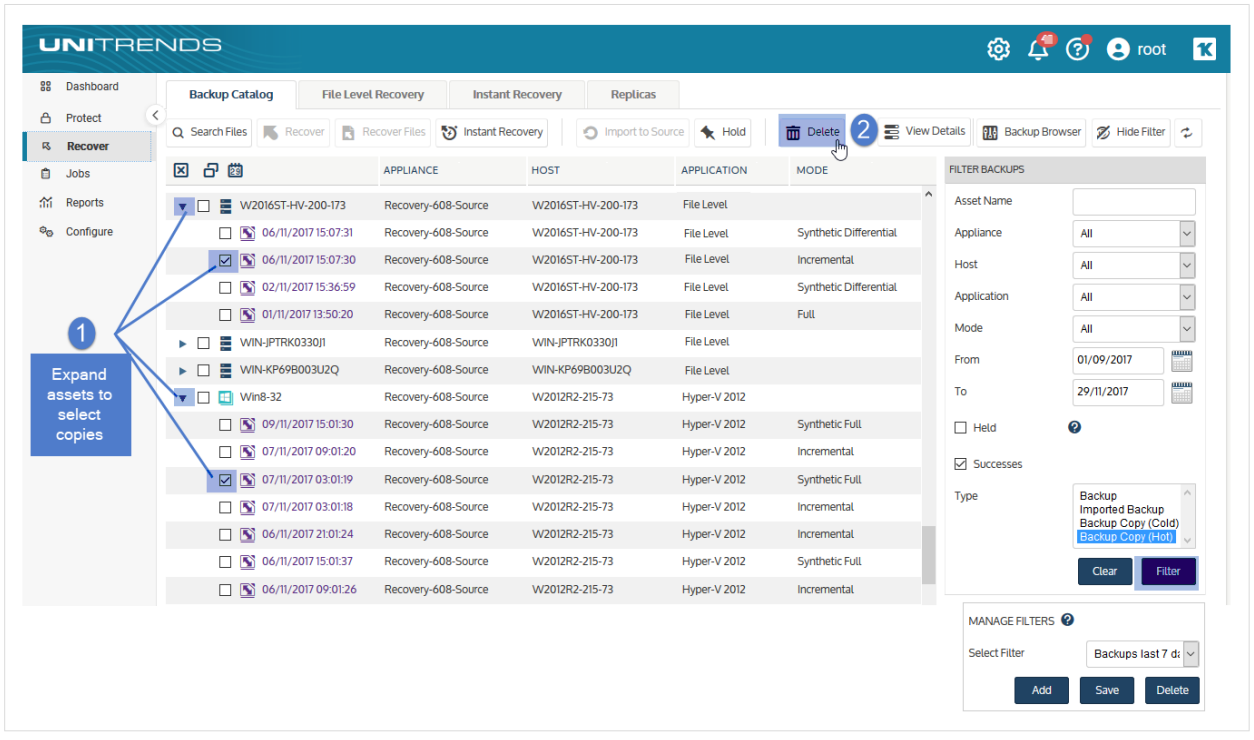

7 Review the list of backup copies that will be deleted.

8 Click **Confirm** to delete the backup copies.

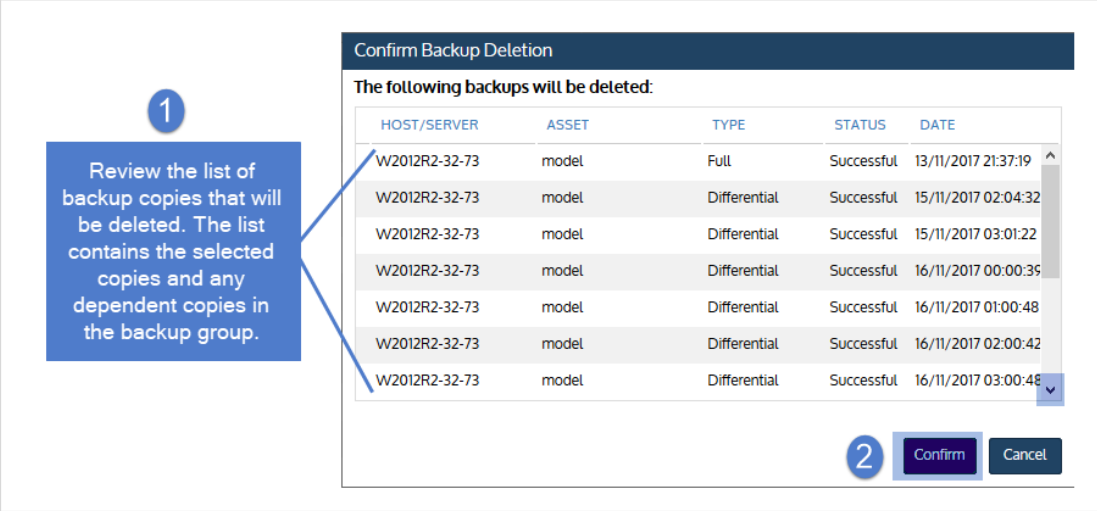

Backup copies are deleted and no longer display in the Backup Catalog on both the source and target appliances. Note that the copies are logically deleted immediately and physically deleted later by the system's purge process. Space reclaimed by removing the copies happens when the purge job completes.

Placing backups on hold

To create space for new backups, Unitrends appliances periodically purge older backups. You can control how long backups remain on the appliance by creating retention policies or by manually placing individual backups on hold. Backups that have been placed on hold cannot be purged until the hold has been removed manually. New backups fail if an appliance cannot purge older backups to create sufficient space.

Notes:

- Holds applied to individual backups override the asset's retention policy. Once the backup is no longer on hold, it is eligible for deletion by the policy.
- To hold an asset's backups temporarily, you can use a long-term retention policy. The policy deletes a backup after the specified retention time period (30 days by default). For details, see "Managing retention with long-term data management" on page 328.

You can manually place backups on hold by using the Backup Browser or the Backup Catalog. The Backup Browser provides appliance-level search capability, advanced search options, and faster filtering performance. The Backup Catalog enables you to browse by asset only.

See these procedures to manually hold and unhold backups:

- "To place backups on hold by using the Backup Catalog"
- "To unhold backups by using the Backup Catalog" on page 636
- "To place backups on hold by using the Backup Browser" on page 637

- "To unhold backups by using the Backup Browser" on page 640

To place backups on hold by using the Backup Catalog

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog**.

   Protected assets display. (Optional) Use the Filter Backups fields to filter the display. For details, see "Backup Catalog tab" on page 60.
3. Click to expand assets to view their backups.
4. Check boxes to select backups.
5. Click **Hold**.

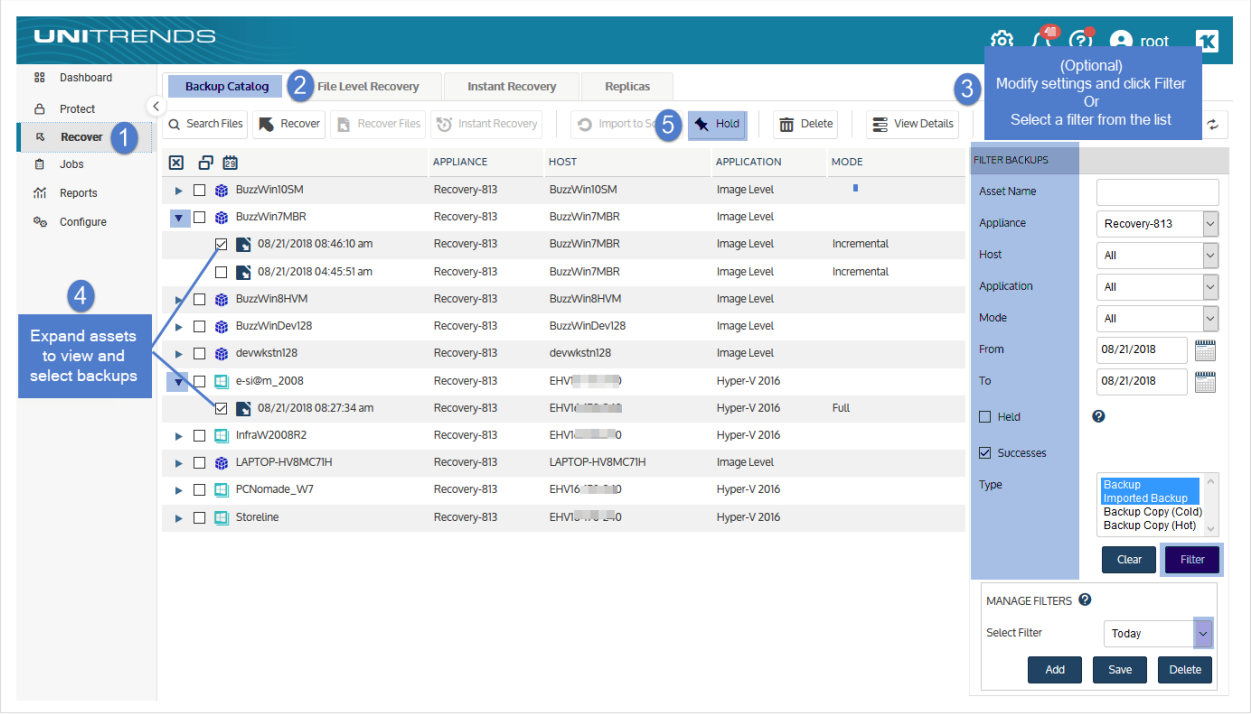

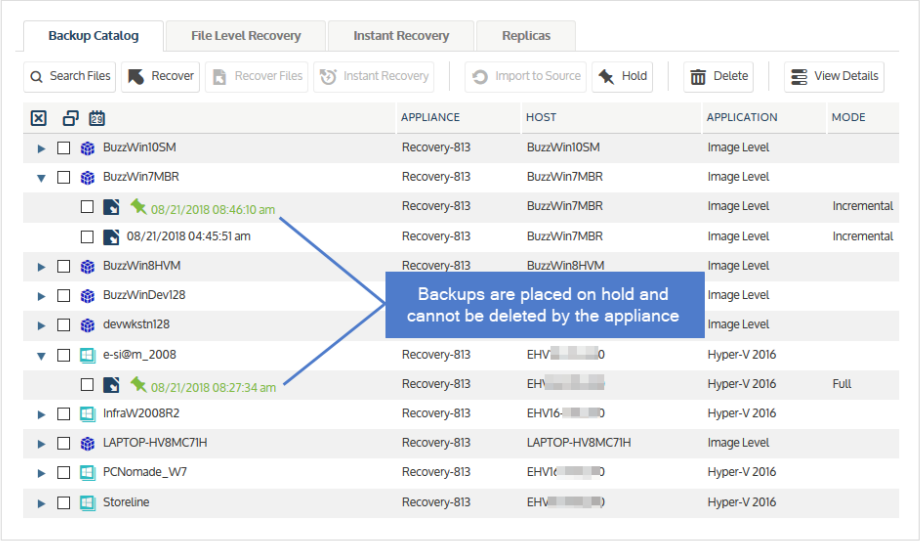

To unhold backups by using the Backup Catalog

Note: The asset's retention policy is applied to the backup upon removing the hold.

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog**.

   Protected assets display. (Optional) Use the Filter Backups fields to filter the display. For details, see "Backup Catalog tab" on page 60.
3. Click to expand assets to view their backups.
4. Check boxes to select held backups.
5. Click **Unhold**.

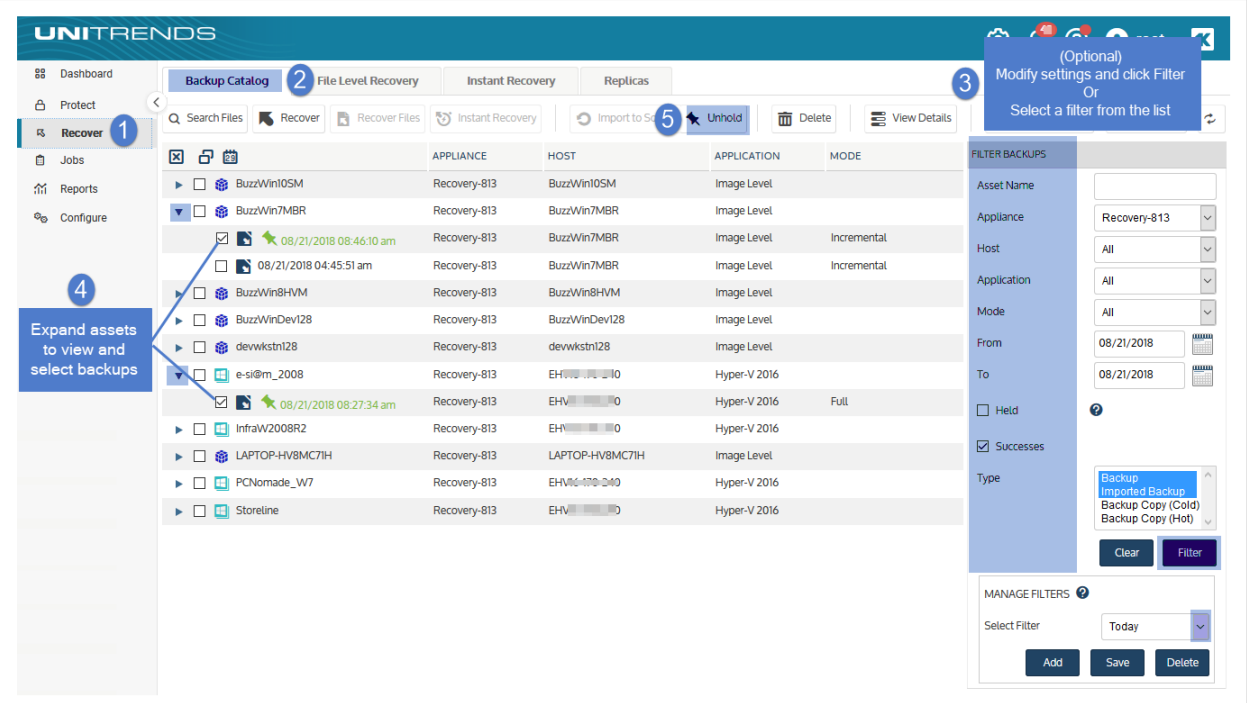

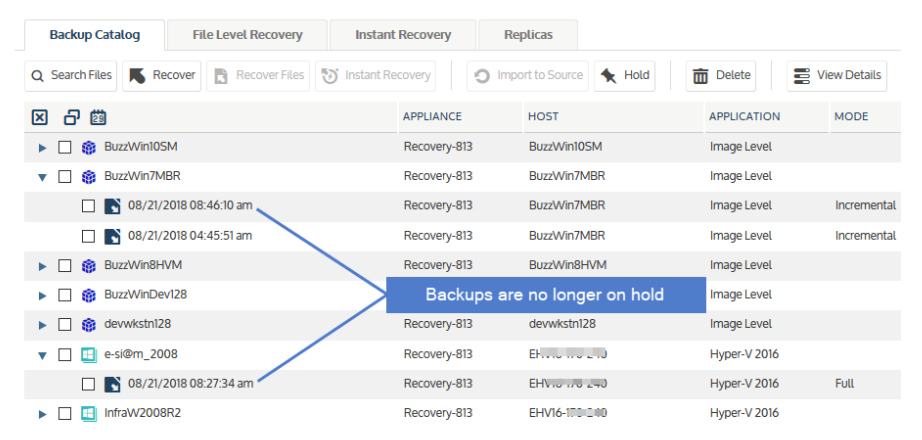

To place backups on hold by using the Backup Browser

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

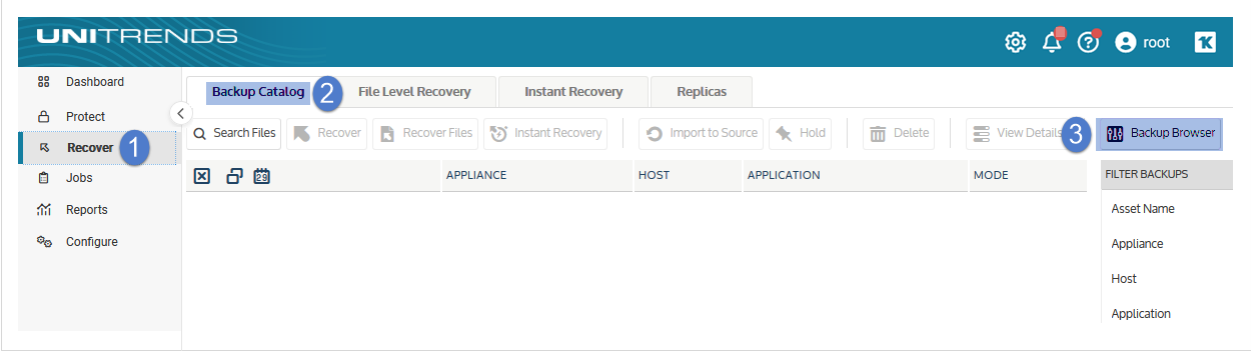

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display:

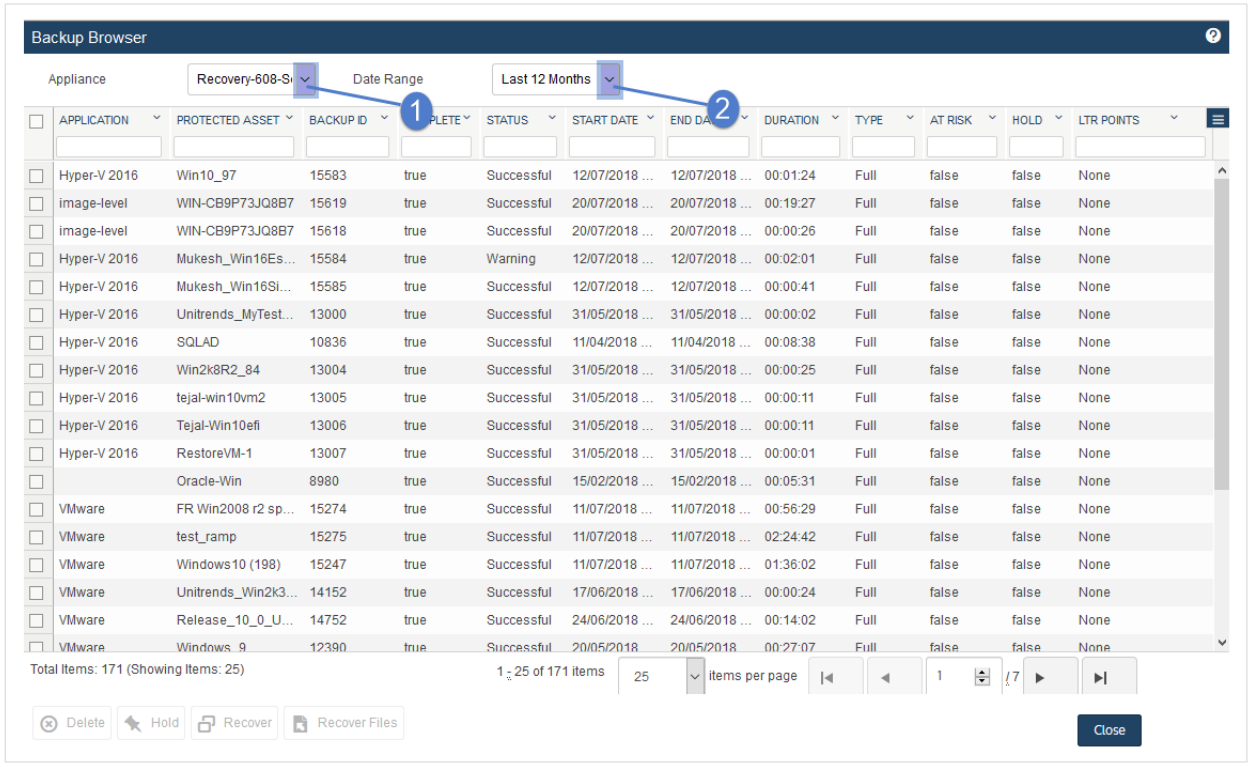

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

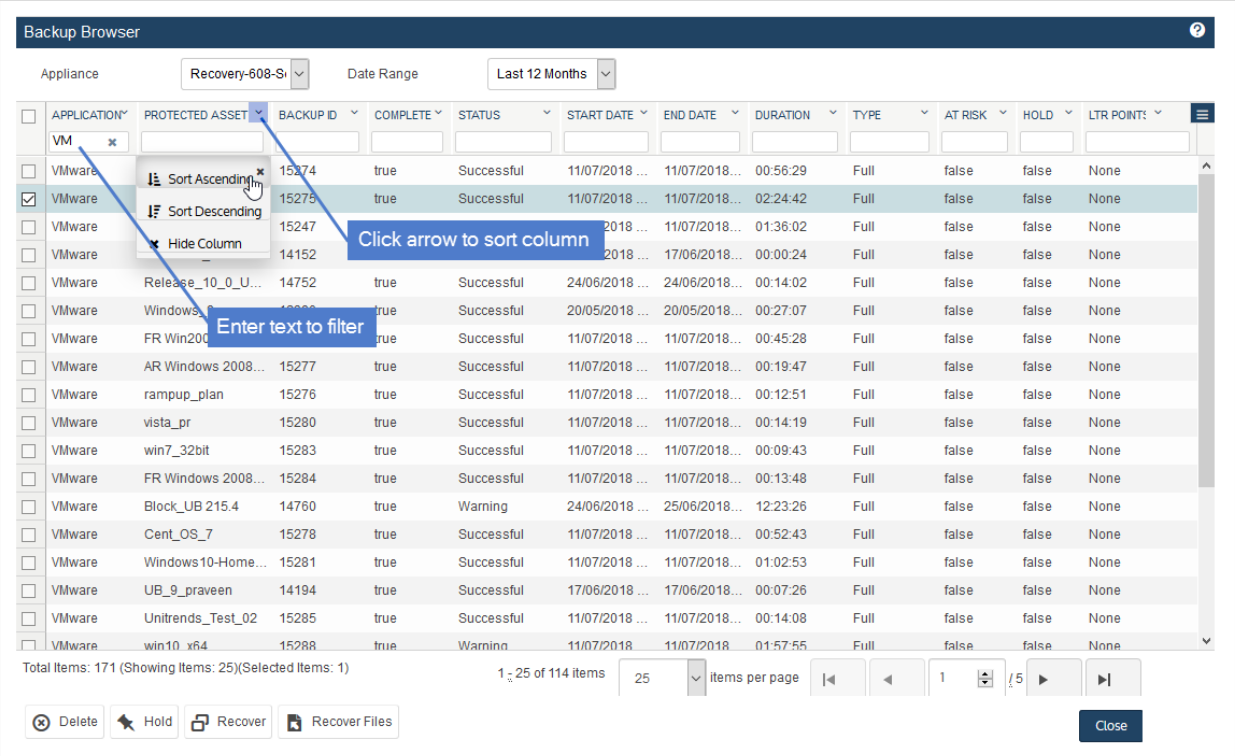

5 Check boxes to select the backups. (To select all, check the box at the top of the column.)

6 Click **Hold**.

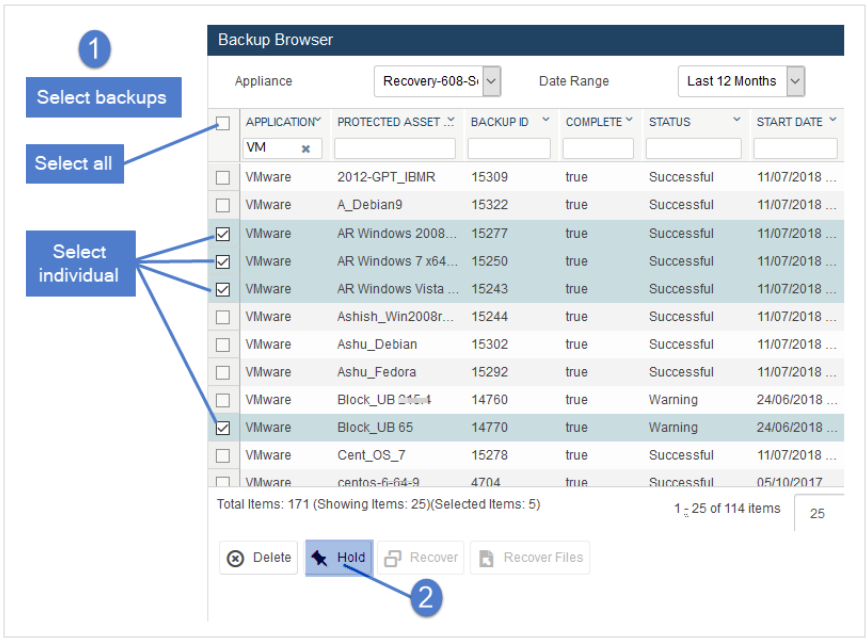

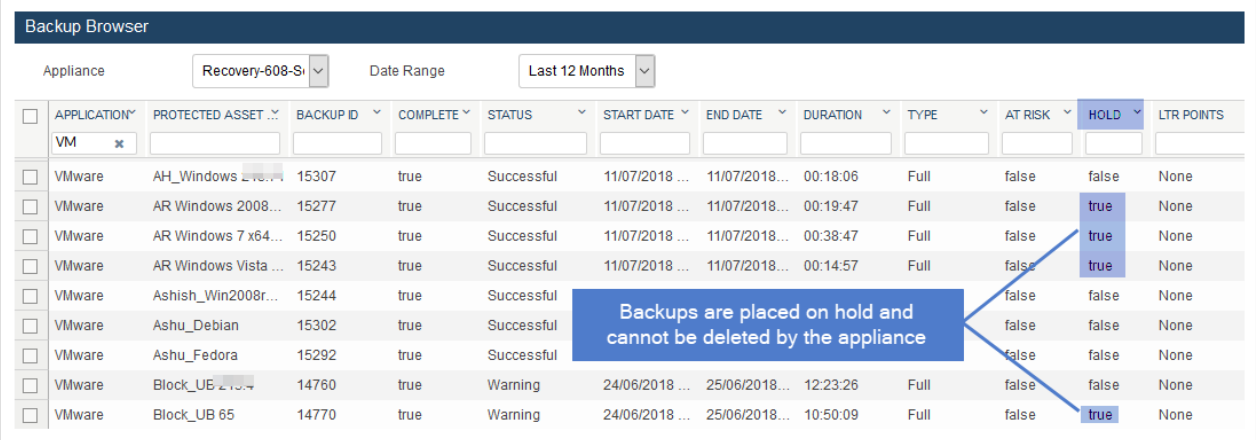

To unhold backups by using the Backup Browser

Note: The asset's retention policy is applied to the backup upon removing the hold.

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog** and click **Backup Browser**.

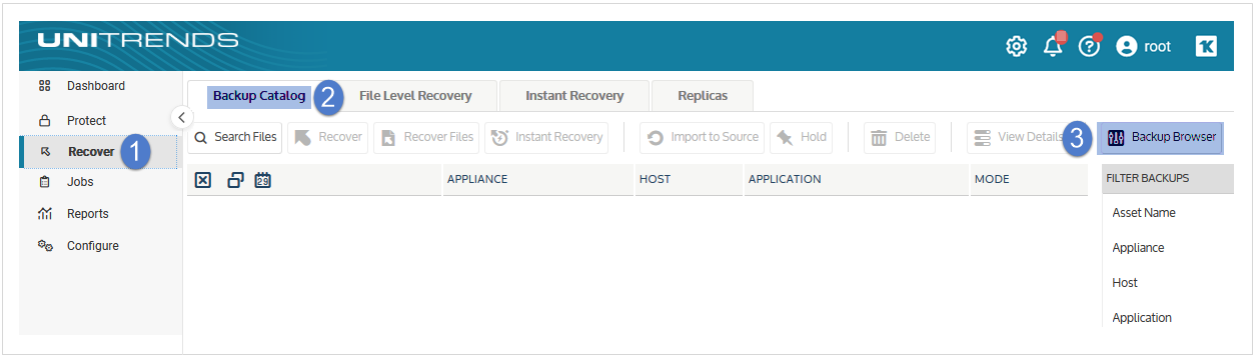

3. Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display:

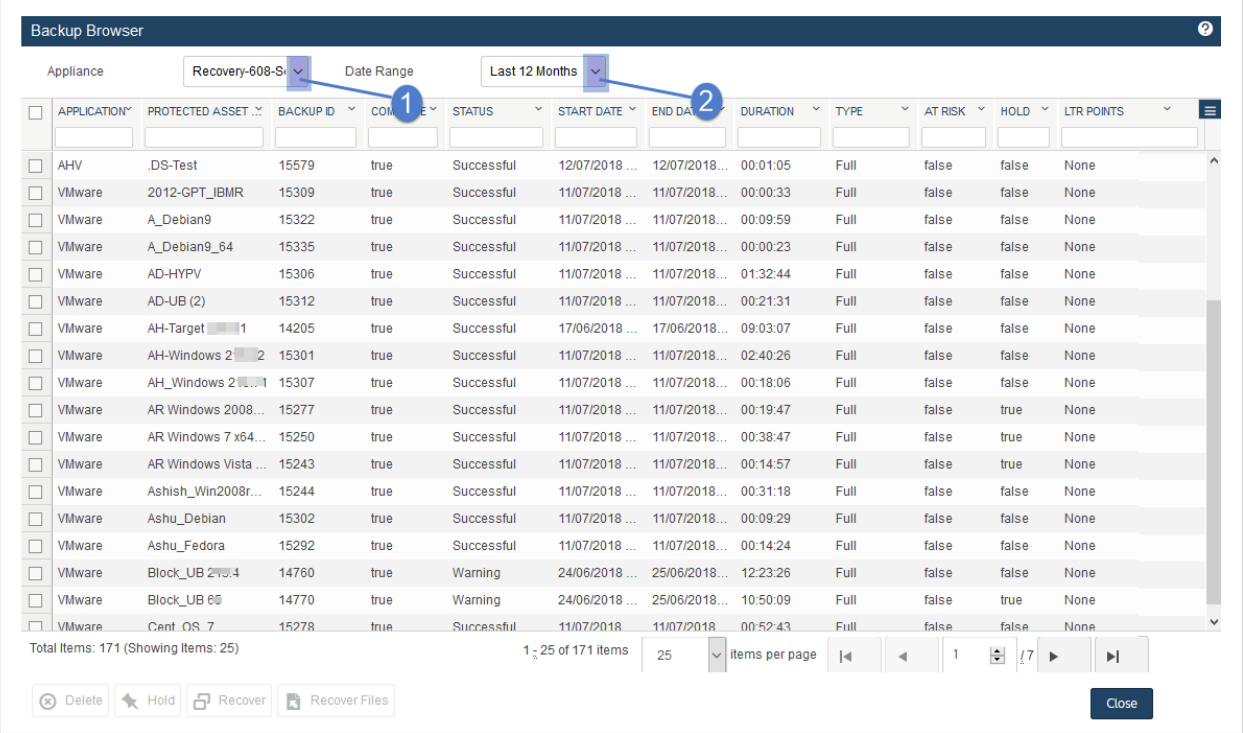

4 Enter **true** in the **Hold** column field to display only held backups.

5 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

6 Check boxes to select held backups. (To select all, check the box at the top of the column.)

7 Click **Unold**.

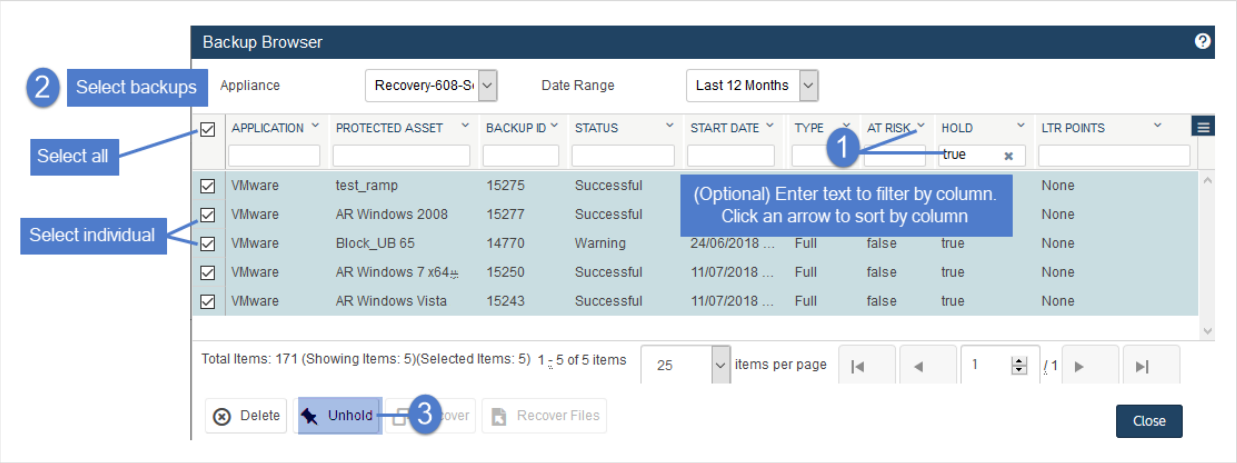

Working with cold backup copy sets

A group of one or more backups that are copied to a target in a given cold backup copy job is called a *set*. When a set is created, reference information about the set is added to a catalog on the backup appliance. This reference information is needed so the appliance can discover and display the set's cold backup copies in the Backup Catalog. A cold copy must display in the Backup Catalog before you can import it to the backup appliance.

To access a backup copy that was run by another Unitrends appliance, you must import its set's reference information. To do this, connect the media where the set resides to a cold copy target on the appliance. (If needed, you can add a cold copy target to the appliance as described in "Backup copy targets" on page 214.) You can also reimport a set's reference information to the original appliance, if needed.

Use the following procedures to manage cold backup copy sets:

- "To view or import reference information for sets on connected media"
- "To view all sets in the cold copy catalog" on page 645
- "To remove a set's reference information from the cold copy catalog" on page 648

See these topics for additional cold backup copy procedures:

- "Managing backup copy targets" on page 260
- "Deleting backups and backup copies" on page 627

To view or import reference information for sets on connected media

Use this procedure to view the cold copy sets that reside on media that is currently connected to the backup copy target and to import a set's reference information to the cold copy catalog on the appliance.

Notes:

- Backup copy jobs mount and unmount the target automatically. The target remains offline unless a copy job is running. To view sets, you must enable the target, which also imports reference information for any backup copies that were not found on the appliance.
- A new cold copy format was introduced beginning in Unitrends release 10.1. If you attempt to import a cold copy and receive the message *Failed to read archive file*, verify that the appliance is running version 10.1 or higher. If not, install appliance updates, then try the import again.
- To view additional cold copy job details, use the "To view all sets in the cold copy catalog" on page 645 procedure instead.

1 Connect the media containing the cold copies to the target appliance that will import the reference information.

Notes:

If you need to remove a drive from an eSATA or USB dock, be sure to:

- Power down the dock.
- Swap the drive.
- Power on the dock.

2 Log in to the target appliance.

3 On the **Configure > Appliances** page, select the appliance.

4 Click the **Backup Copy Targets** tab below.

5 Click **Scan for Media**.

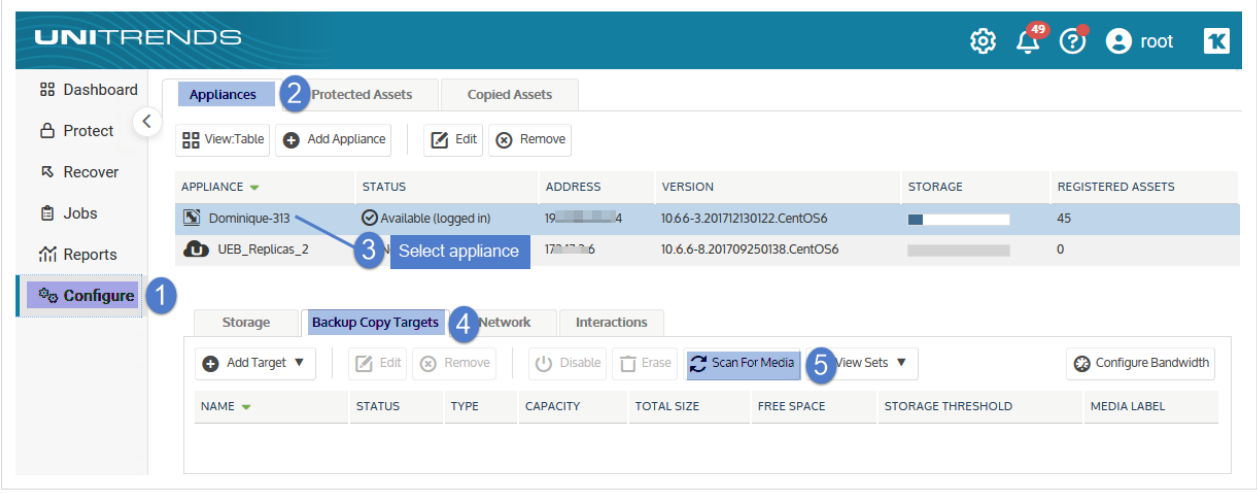

6 Select the cold backup copy target.

7 (If needed) Click **Enable** to bring the target online. (If the target is already enabled, you see a Disable button instead and you can skip this step.)

Notes:

If you receive a message that the media has not been initialized, you must erase or prepare the media before you can enable the target. Erasing or preparing the media removes all backup copies stored on the media. For details, see one of the following:

- "To prepare tapes for use with an autochanger device" on page 273 for tape autochanger targets.
- "To initialize and erase cold backup copy media" on page 271 for all other cold backup copy targets, including single tape drive devices.

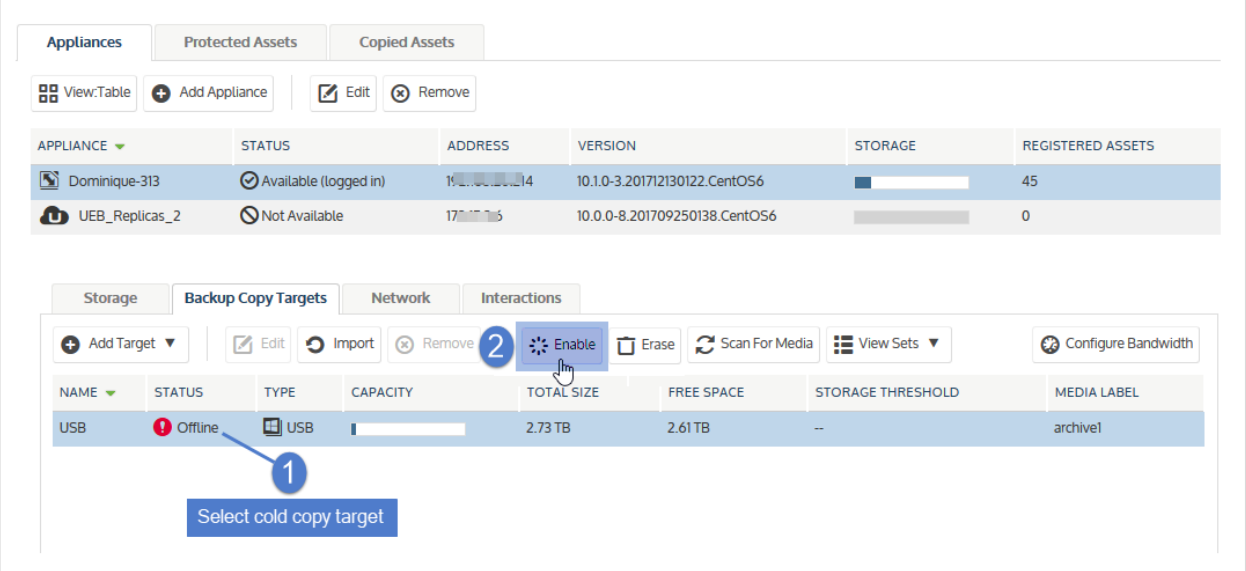

8 Select **On Media** from the **View Sets** list:

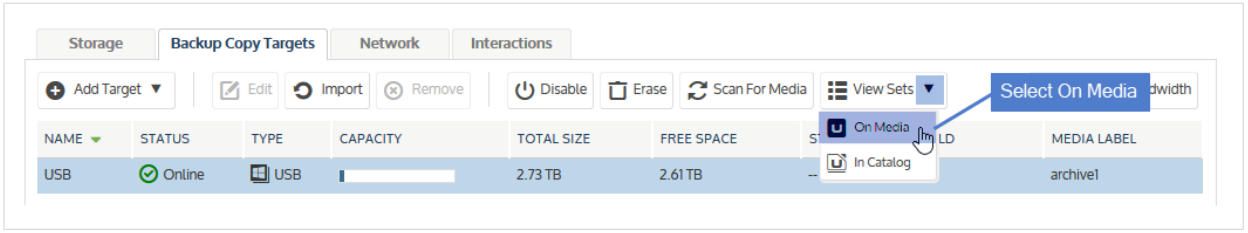

9 Sets that are stored on the media display in a list. The following is given for each set:

- Description – Name of the backup copy job that created the set.
- Date – Date and time that the set was copied to the media.
- Imported – Indicates whether the set's reference information has been imported to the cold target's catalog on this backup appliance. (Reference information must be imported so the set can be discovered in the Backup Catalog, where you can then opt to import the set's backup copies to the appliance. For details, see "To import a cold backup copy" on page 786.)

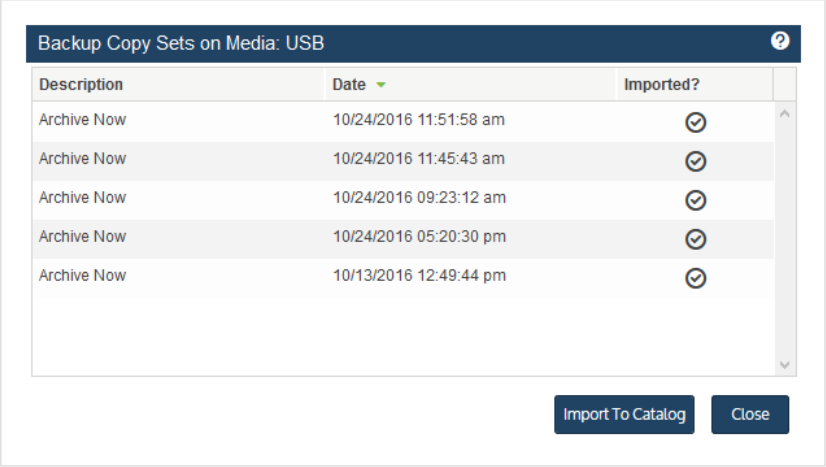

10 Do one of the following:

- To import a set's reference information, select it in the list and click **Import to Catalog**.
- To exit, click **Close**.

To view all sets in the cold copy catalog

Use this procedure to view all sets in the appliance's cold copy catalog.

Note: Backup copy jobs mount and unmount the target automatically. The target remains offline unless a copy job is running. To view sets, you must enable the target, which also imports reference information for any backup copies that were not found on the appliance.

1 On the **Configure > Appliances** page, select the appliance.

2 Click the **Backup Copy Targets** tab below.

3 Click **Scan for Media**.

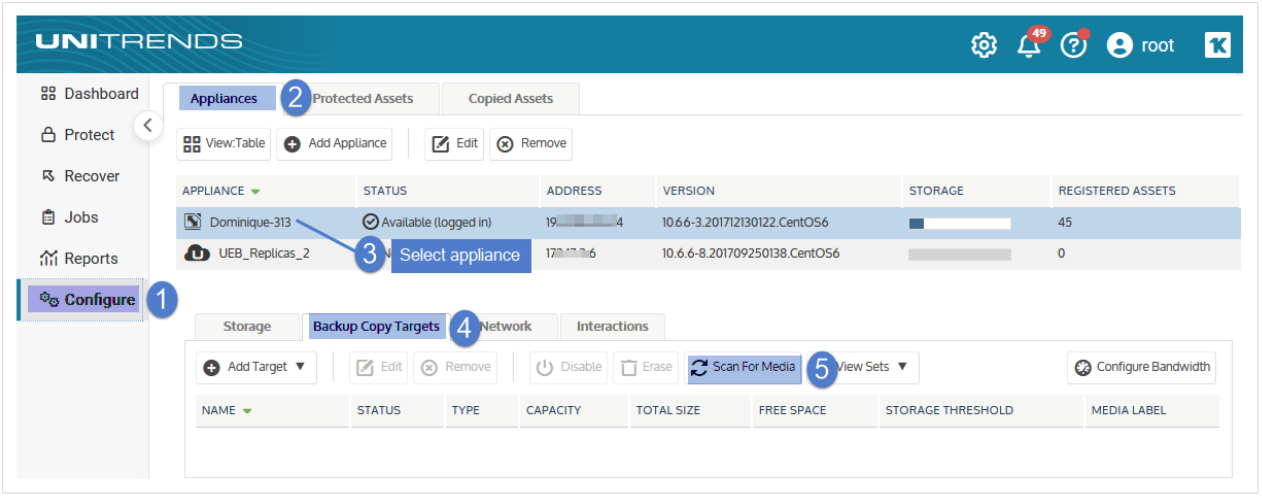

4 Select the cold backup copy target.

5 (If needed) Click **Enable** to bring the target online. (If the target is already enabled, you see a Disable button instead and you can skip this step.)

Notes:

If you receive a message that the media has not been initialized, you must erase or prepare the media before you can enable the target. Erasing or preparing the media removes all backup copies stored on the media. For details, see one of the following:

- "To prepare tapes for use with an autochanger device" on page 273 for tape autochanger targets.
- "To initialize and erase cold backup copy media" on page 271 for all other cold backup copy targets, including single tape drive devices.

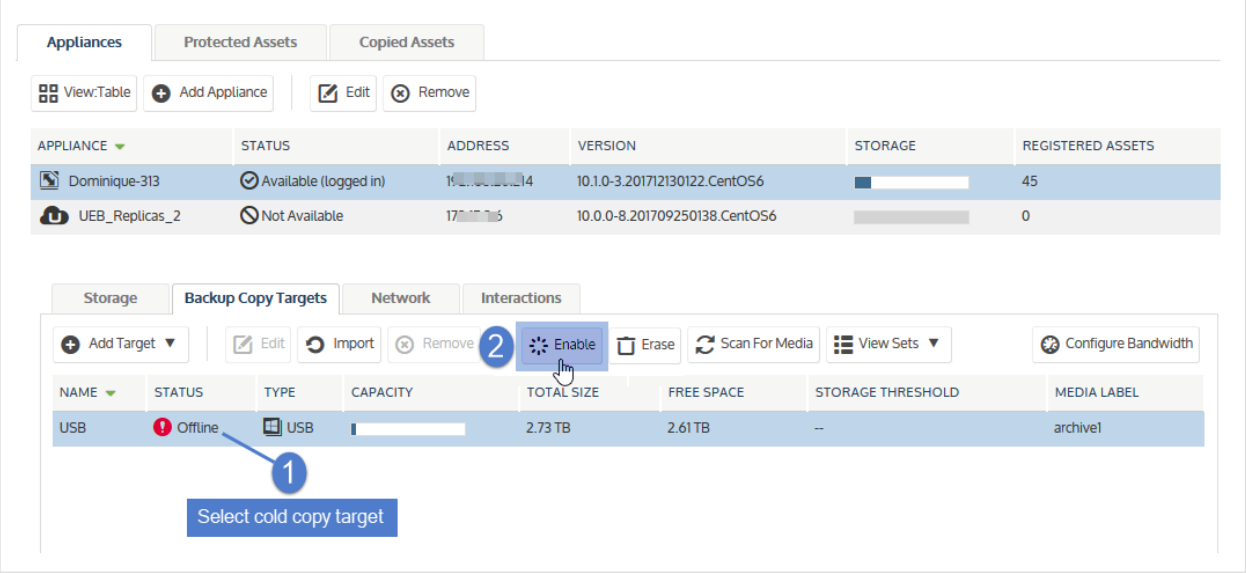

6 Select **In Catalog** from the **View Sets** list:

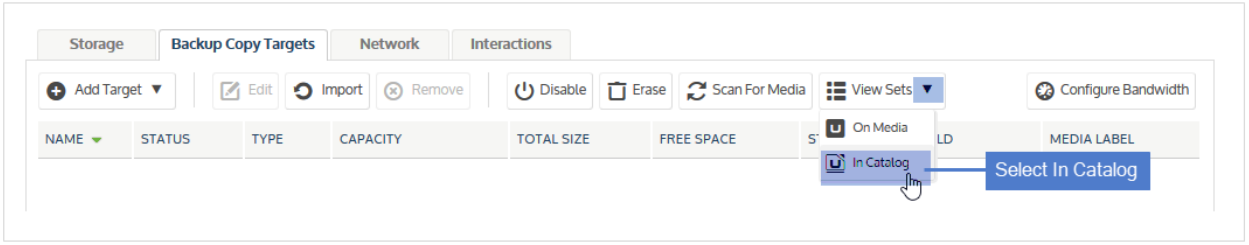

7 Sets in the catalog display in a list. The following is given for each set:

- Job – Name of the backup copy job that created the set.
- Date – Date and time that the set was copied to the target.

- Status – Status of the cold backup copy job that created the set. If needed, hover over the Status column to view the entire message.

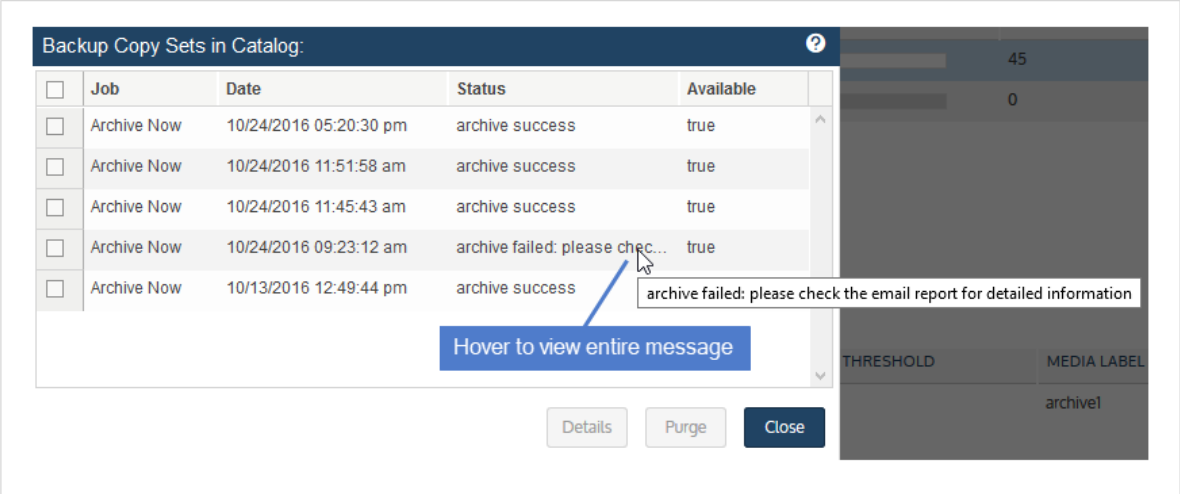

- Available – Indicates whether the set is available on the media.
- To view job details, select a set in the list and click **Details**. Click **Close** to exit the Set Details dialog.

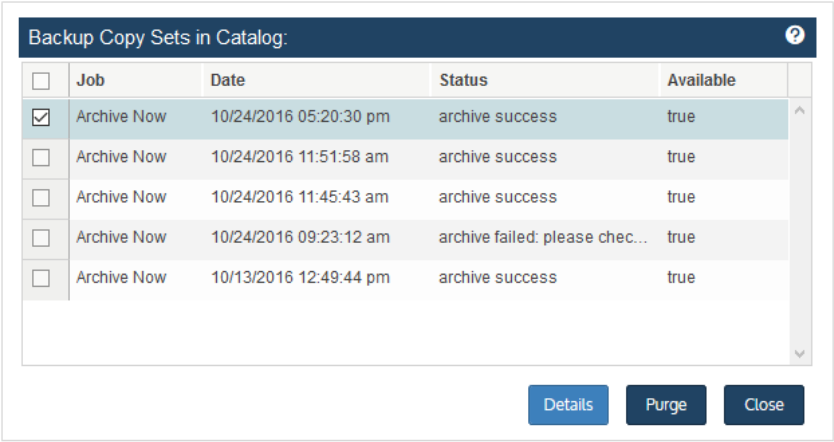

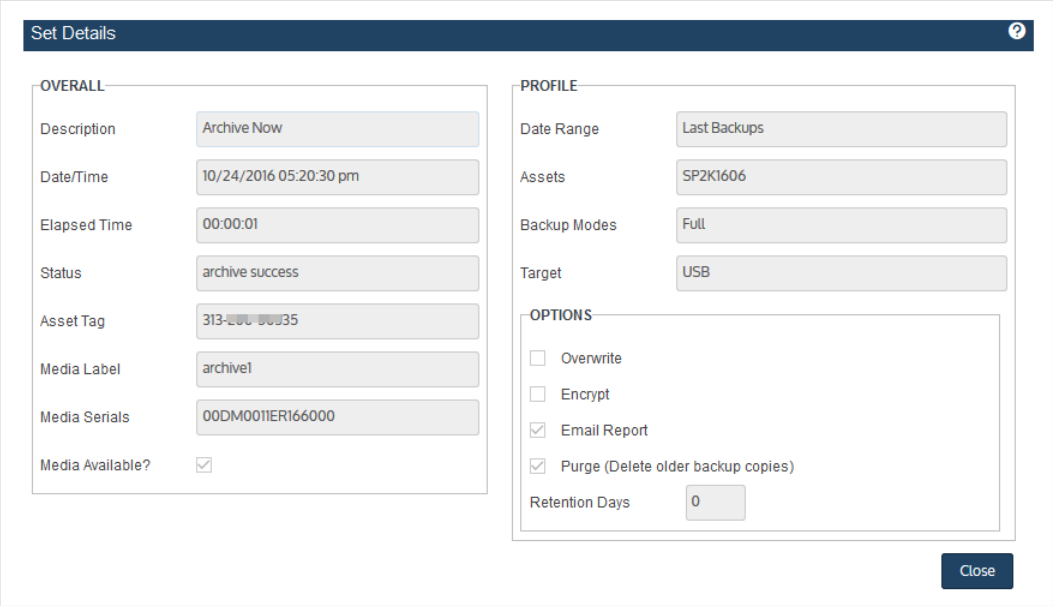

8 Click **Close** to exit:

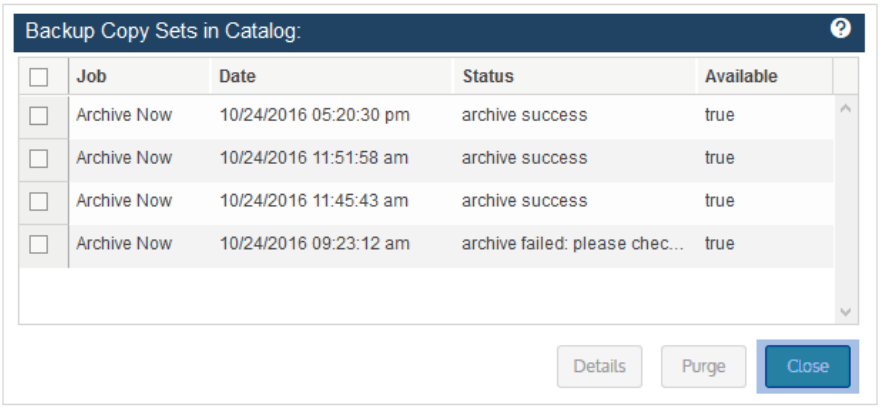

To remove a set's reference information from the cold copy catalog

Use this procedure to remove a set's reference information from the appliance.

Notes:

- Backup copy jobs mount and unmount the target automatically. The target remains offline unless a copy job is running. To view sets, you must enable the target, which also imports reference information for any backup copies that were not found on the appliance.
- In most cases, this procedure removes reference information only and does not remove the set from the cold copy target. The only exceptions are third-party cloud and NAS backup copy targets that are configured with the

setting *[CloudHook] RemoveFiles=1* or *[Archive] NASRemoveFiles=1*(located in **Edit Appliance > Advanced > General Configuration**). If this setting is configured for a third-party cloud or NAS backup copy target, the set's reference information is removed from the cold copy catalog on the appliance and the set is removed from the cold copy target.

1. On the **Configure > Appliances** page, select the appliance.
2. Click the **Backup Copy Targets** tab below.
3. Click **Scan for Media**.

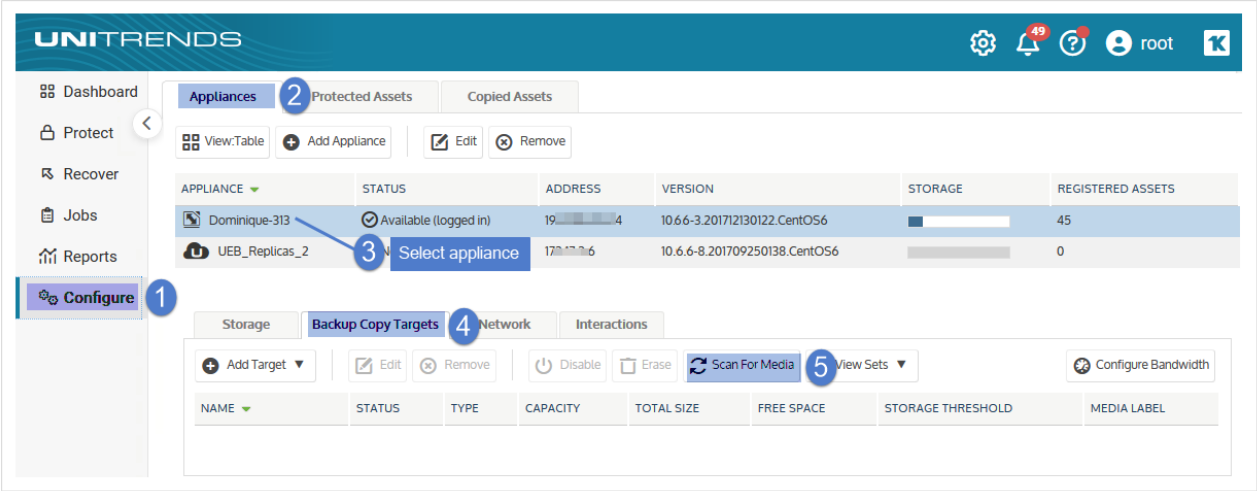

4. Select the cold backup copy target.
5. (If needed) Click **Enable** to bring the target online. (If the target is already enabled, you see a Disable button instead and you can skip this step.)

Notes:

If you receive a message that the media has not been initialized, you must erase or prepare the media before you can enable the target. Erasing or preparing the media removes all backup copies stored on the media. For details, see one of the following:

- "To prepare tapes for use with an autochanger device" on page 273 for tape autochanger targets.
- "To initialize and erase cold backup copy media" on page 271 for all other cold backup copy targets, including single tape drive devices.

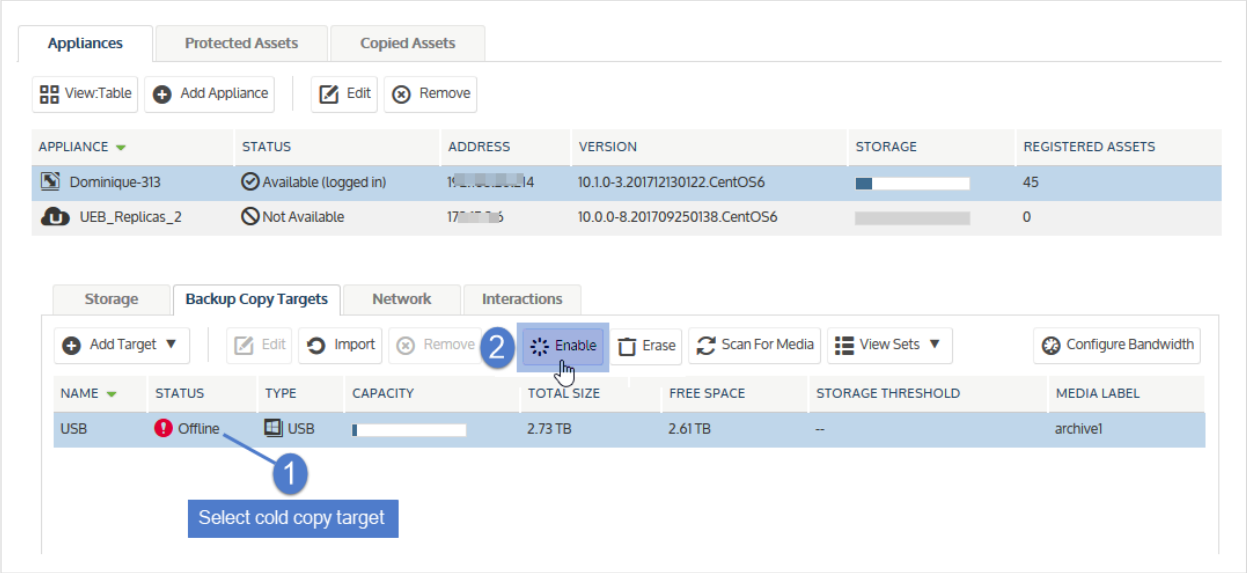

6 Select **In Catalog** from the **View Sets** list:

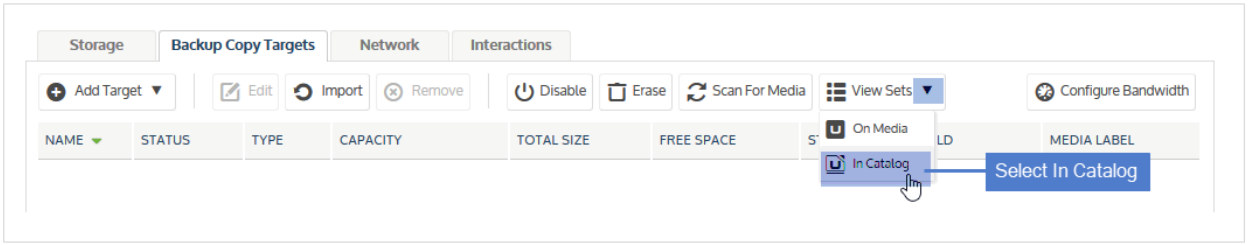

7 Select the set in the list and click **Purge**:

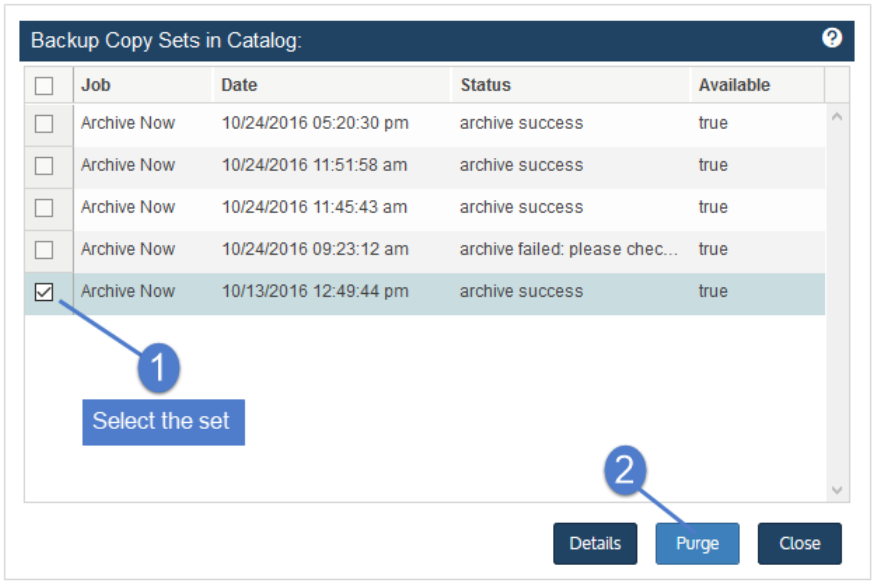

8 Check the box to confirm, then click **Remove Set**:

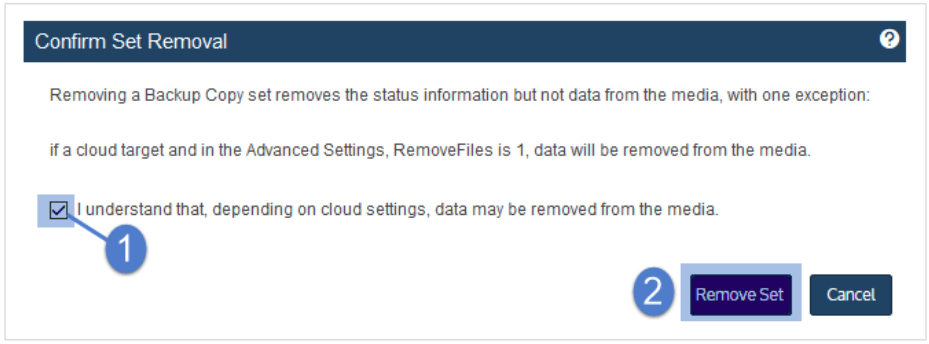

9 Click **Close** to exit:

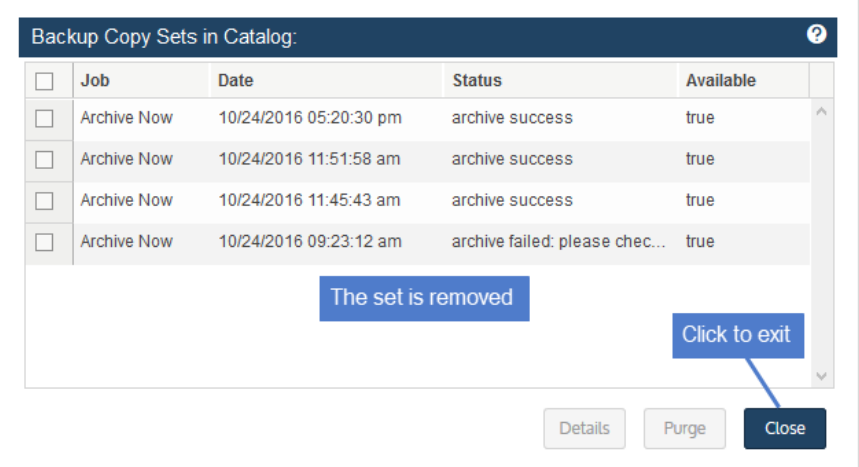
Backup Copy Sets in Catalog:
Job
Date
Status
Available
Archive Now
10/24/2016 05:20:30 pm
archive success
true
Archive Now
10/24/2016 11:51:58 am
archive success
true
Archive Now
10/24/2016 11:45:43 am
archive success
true
Archive Now
10/24/2016 09:23:12 am
archive failed: please chec...
true
The set is removed
Click to exit
Details
Purge
Close

Chapter 6: Host-level Backups Overview

This section provides information for implementing host-level protection of VMware, Hyper-V, Nutanix AHV, and Citrix XenServer environments. Host-level backups protect hosted VMs by leveraging host snapshots. To run host-level backups, you do not install an agent on the guest VMs. Simply add the virtual host to the Unitrends appliance and select VMs to protect. For details on recovery options, see "Recovering Host-level Backups" on page 793. For a general overview of Unitrends protection, see "Protection Overview" on page 91.

Notes:

- If you install a Unitrends agent on your VM and use file-level protection, the Unitrends appliance treats the VM as a physical asset. See "File-level Backups Overview" on page 703 for information on protecting VMs with the agent.
- Host-level protection is not supported for these environments: Azure and Amazon Web Services (AWS). Instead, use agent-based file-level and application backups.

Review the topics in this section to determine which features you want to use. Ensure also that all requirements have been met before you begin protecting your virtual machines. After you have verified that all applicable requirements have been met, see these topics to set up host-level backups:

- "Protected assets" on page 279 to add your virtual host to the appliance
- "Backup Administration and Procedures" on page 425 to run host-level backups

Hyper-V virtual machines

This section provides considerations and requirements for protecting Hyper-V environments.

Preparing for Hyper-V backups

Following is a summary of the high-level steps for backing up Hyper-V virtual machines. The information includes links to detailed instructions for each procedure.

Step 1: Review "Best practices and requirements for Hyper-V protection".

Step 2: Install the Unitrends Windows agent on your Hyper-V host. See "Installing the Windows agent" on page 362.

Note: For most Windows servers, the appliance can push-install the agent when you add the asset. If you will be push-installing the agent, skip to Step 3:. For push-install requirements, see "Windows agent requirements" on page 362.

Step 3: Add the Hyper-V host to your Unitrends appliance. See "Adding a virtual host" on page 308.

Step 4: Create backup jobs for your VMs:

- To create a job manually, see "To create a Hyper-V backup job" on page 462.
- To create a job by using an SLA policy, see "To create an SLA policy for Hyper-V assets" on page 556.

- For a comparison of the manual and SLA policy job creation methods, see "About creating backup and backup copy jobs" on page 426.

Best practices and requirements for Hyper-V protection

Review the information in these topics before implementing Hyper-V host-level protection:

- "Hyper-V best practices" on page 654
- "General Hyper-V requirements" on page 655
- "Additional Hyper-V requirements" on page 658

Hyper-V best practices

Follow these recommendations:

- Adhere to Microsoft's best practices for virtualization. For a list of Microsoft documents on virtualization, see Microsoft Virtualization: Hyper-V best practices.
- Install the latest Windows agent on your Hyper-V host for best performance. (If needed, upgrade the appliance first to ensure it is running an equal or later version.)
- To protect the file system and operating system of the Hyper-V host, you must run file-level backups. For details, see "File-level Backups Overview" on page 703. Any files belonging to the Hyper-V application are automatically excluded from file-level backups of the Hyper-V host.
- In some cases, you may want or need to protect VMs using file-level backups. For recommendations, see "Protecting Hyper-V virtual machines with file-level backups" on page 661.
- To protect a VM with both host-level and file-level (agent-based) backups, ensure that the VM's host-level and file-level jobs do not overlap. Running both simultaneously may lead to undesirable results.
- If recovery time objectives are important, set up "Virtual machine instant recovery" to quickly to spin up a failed VM from host-level backups.
- Full and incremental backups are supported for Hyper-V VMs.
- A new full backup is required in these cases:
  - The VM configuration has changed since the last backup. This includes any configuration changes made to a VM in the Hyper-V manager, such as creating or deleting a snapshot, adding a new disk, or converting a disk from VHD to VHDX format.
  - The VM's configuration version is not present in the last backup.
  - The VM's configuration version has changed since the last backup.

  If the VM configuration has changed since the last backup, the next incremental fails. After this failure, the appliance promotes the next scheduled backup to a full (or displays a message indicating a full is required if an on-demand incremental is attempted). Once a full backup succeeds, subsequent incrementals run as scheduled.

- Beginning in release 10.0, additional VM configuration version checking is used to support Hyper-V's VMCX format (introduced in Windows 2016). If you are using hosts running Windows 2016 or later, or protecting VM configuration version 6.2 or higher, see "Virtual machine configuration" on page 656 for additional requirements.

General Hyper-V requirements

The following requirements must be met for host-level protection of Hyper-V virtual machines.

Item	Description
Hyper-V host	The following are required for the Hyper-V host: • The Hyper-V host must be a supported version listed in the Unitrends Compatibility and Interoperability Matrix. • The Unitrends Windows agent must be installed on the host as described in "Installing the Windows agent" on page 362. (It is not necessary to install agents on your virtual machines.) • The Hyper-V PowerShell module must be installed on the host. In most cases, this module is installed by default. In some Windows Core OS versions, this module is not included in the default installation. • To protect VMs running on Hyper-V 2022, both the Unitrends appliance and the Windows agent on the Hyper-V host must be running version 10.6.2 or higher. • To protect VMs running on Hyper-V 2019, both the Unitrends appliance and the Windows agent on the Hyper-V host must be running version 10.3.8 or higher. • To protect VM configuration versions 6.2 or higher, both the Unitrends appliance and the Windows agent on the Hyper-V host must be running version 10.0 or higher. (If using a host that is running Windows 2016 or later, note that VMs are created with configuration version 8.0 by default.) • For cluster configurations: – Each node in the cluster must be running agent version 10.6.9 or higher. – Be sure to install the same agent version on all nodes in the cluster.
Microsoft VSS	Microsoft's Volume Shadow Copy Service (VSS) and the Hyper-V VSS writer must be installed and running on the Hyper-V host.
Integration Services	To avoid VM downtime, Unitrends recommends online backups. To perform online backups, you must install Integration Services in the guest operating system to enable the VM to create a child state snapshot. The host then uses this snapshot to perform an online backup of the virtual machine. For online backups, the following conditions must be met on the protected VMs:

Item	Description
	• The latest version of backup Integration Services must be installed and enabled. For a list of guest operating systems for which Integration Services is supported, see the Microsoft document Hyper-V Overview. • The VM's VHD(X) files and snapshot file location must be set to the same volume in the host operating system. • All volumes must reside on basic disks and the VMs cannot have dynamic disks. • All disks must use a file system that supports snapshots (for example, NTFS). If an online backup cannot be performed, the VM is temporarily put in a saved state. In saved state there is a brief downtime during the backup.
Virtual machine configuration	VMs must meet the requirements listed below. For VMs that do not meet these requirements, install the Windows agent on the VM and run file-level backups. • If the VM configuration is 6.2 or higher, both the Unitrends appliance and the Windows agent on the Hyper-V host must be running version 10.0 or higher. If the host is running agent version 10.0 or higher but the appliance is running an older version, backup and recovery jobs fail for VM configuration version 6.2 or higher. **Notes:** • Starting with VM configuration 6.2, Microsoft implemented a binary-based configuration format (VMCX). On Windows Server 2016 and later, VMs are created by default with this VMCX configuration. VMCX requires new backup and recovery methods that are available in Unitrends release 10.0 and later. • Additional VM configuration version checking is used to support the VMCX format. A new full backup is required in these cases: the VM's configuration version is not present in the last backup, or the VM's configuration version has changed since the last backup. In these cases, the appliance cannot run an incremental until a full backup runs and: – If you attempt an on-demand incremental, the appliance promotes the backup to a full. – If the next job is a scheduled incremental, the job fails due to the configuration change. After this failure, the appliance promotes the next scheduled run to a full. Once this full succeeds, subsequent incrementals run as scheduled. (If you upgrade a VM's configuration version, you can opt to run an on-demand full, rather than letting the appliance fail then

Item	Description
	promote the next incrementals.) • The VM cannot be configured with shared VHDX disks. • The VM cannot be configured with pass-through disks. • The VM cannot be configured in a VHDS cluster or in a VHD Set. • The VM cannot be configured as a shielded VM. • The VM cannot be configured in a Hyper-V container.
VM name and location	The VM name and path to the virtual machine on the Hyper-V host must not contain escape or special characters. Backups fail if escape or special characters are present in the pathname.

Additional Hyper-V requirements

These additional requirements may apply to your environment.

Item	Description
Faster incrementals for VM configuration version 5.0 and higher (2012, 2012 R2, 2016, 2019, and 2022 hosts)	For VM configuration version 5.0 and higher, Unitrends leverages a Hyper-V changed-block-tracking (CBT) driver that greatly increases incremental backup performance. To use this driver: 1 If using a Windows server to host your VMs, enable the Hyper-V role. **Note:** For Windows servers, you must enable the Hyper-V role before you install the Windows agent. If the Hyper-V role is not enabled, the CBT driver is not installed with the Windows agent. 2 Install the Windows agent on your Hyper-V host. The CBT driver is automatically installed with the Windows agent. To verify that the CBT driver was used, view backup details and look for the following in the Output: *CBT DRIVER ACTION IS ENABLED*. If the driver has been uninstalled or corrupted, backups complete with a warning to indicate that the CBT driver was not used.
Windows agent push	To use the push feature to install the Windows agent and agent updates on the Hyper-V host, see "Installing the Windows agent" on page 362 for additional requirements.
Microsoft Hyper-V Replicas feature	Microsoft's Hyper-V Replicas feature is not compatible with the Hyper-V CBT driver. If you are using this feature, you must manually uninstall the Hyper-V CBT driver. For details, see "Manually installing and uninstalling the Hyper-V CBT driver" on page 384. Once the driver has been uninstalled, Hyper-V incrementals are supported but do not use the driver. The Hyper-V CBT driver is installed each time the Windows agent is installed or updated. Manually uninstall the driver as needed after updating or reinstalling the Windows agent.
Virtualized Active Directory servers	To ensure database consistency, you must set up the virtualized Active Directory (AD) server in accordance with Microsoft best practices. If all Microsoft considerations are not addressed, backup and recovery of the virtual machine may yield undesired results. If you prefer not to research these best practices, it is recommended to install the agent on the VM and protect it as you would a physical server (leveraging Microsoft's VSS writers).
Distributed File System environments	Distributed File System (DFS) Namespaces and DFS Replication offer high-available access to geographically dispersed files. Because of the replication and syncing operations in DFS environments, you must set up the virtual machine in accordance

Item	Description
	with Microsoft best practices to ensure database consistency. If all Microsoft considerations are not addressed, backup and recovery of the virtual machine may yield undesired results. If you prefer not to research these best practices, it is recommended to install the agent on the VM and protect it as you would a physical server (leveraging Microsoft's VSS writers).
Multi-point Services Role and RDS User Profile Disks	Windows VMs with the Multi-point Services Role and RDS User Profile Disks are not supported.
Cluster shared volumes (CSVs)	Adhere to the following requirements when using CSVs: • A cluster with a single cluster shared volume does not follow Microsoft's best practices and may be unreliable. If you have VMs in a cluster with a single CSV, protect them as you would physical machines, by using file-level backups. • A VM's data cannot reside on both CSV and SMB shares. Using a combination of CSV and SMB shares for a single VM is not supported. • Avoid backing up a set of VMs in the same schedule (i.e., at the same time) where some of the VMs are based on SMB storage and others are based on CSV storage. This can lead to backup failures.
Storage on SMB 2.0 shares for Hyper-V 2022/2019/2016	For Hyper-V 2022, 2019, and some later Hyper-V 2016 versions, additional configuration is required to protect VMs with data that resides on SMB 2.0 shares. For details, see Protecting Hyper-V 2019/2016 VMs with disks that reside on SMB 2.0 shares.
Storage on SMB 3.0 shares	Unitrends can protect virtual machines with disk storage located on SMB 3.0 shares. When these VMs are backed up, the Hyper-V agent creates a VSS snapshot on the remote server and exposes it to the Hyper-V host through the SMB share pathing. The agent then backs up the VM's files from the remote snapshot location. When the backup completes, all VSS snapshots created for the backup are removed from the server hosting the SMB share. The following are required to protect Hyper-V VMs with SMB 3.0 file storage: • The File Server and the File Server VSS Agent Service roles must be installed on the server hosting the SMB 3.0 shares. For instructions on installing these roles, see How do I install the File Server and File Server VSS Agent Service roles on a server hosting SMB shares?. • The Unitrends Hyper-V agent installed on the Hyper-V host must be granted read/write access to remote SMB 3.0 shares. The most secure option to provide all necessary access is to change the login account for the Unitrends Hyper-V agent service from *bpagent* to the domain

Item	Description

	administrator account.
	If permissions for the domain administrator do not allow access to all files for file-level backups of the Hyper-V host, run the agent as a local system account on the Hyper-V host and grant it read/write permission for the SMB shares. For instructions, see Running the Windows agent as local system account on Hyper-V server and granting account read/write permissions for SMB shares.
	• The servers hosting the VMs and SMB shares must belong to the same Windows domain.
	• The VM can contain one or more disks on SMB 3.0 shares. Disks can reside on the same share or different shares hosted by one or more servers in the same domain. All servers participating in the VM backup must belong to the same domain.
	• If multiple Hyper-V hosts are using the same SMB server, avoid backing up VMs on different hosts at the same time. A single SMB server can only service one VSS snapshot request at a time from remote hosts.
	• A VM's data cannot reside on both CSV and SMB shares. Using a combination of CSV and SMB shares for a single VM is not supported.
	• Avoid backing up a set of VMs in the same schedule (i.e., at the same time) where some of the VMs are based on SMB storage and others are based on CSV storage. This can lead to backup failures.
	• If the VM contains multiple disks on SMB 3.0 shares, the storage paths to all disks must use the same naming pattern (all must include the NETBIOS or FQDN or IP). Using a mix of naming patterns is not supported.
	Example of mixed naming that causes the backup job to fail:
	`\\HVSMBServer\Share1\WinVM\Disk1.vhd`
	`\\HVSMBServer.qatest.local\Share1\WinVM\Disk2.vhd`
	`\\192.168.199.215\share1\WinVM\Disk3.vhd`
	Examples of consistent naming where the backup job succeeds:
	`\\HVSMBServer\Share1\WinVM\Disk1.vhd`
	`\\HVSMBServer\Share1\WinVM\Disk2.vhd`
	`\\HVSMBServer\share1\WinVM\Disk3.vhd`
	or
	`\\HVSMBServer.qatest.local\Share1\WinVM\Disk1.vhd`
	`\\HVSMBServer.qatest.local\Share1\WinVM\Disk2.vhd`

Item	Description

	`\\HVSMBServer.qatest.local\share1\WinVM\Disk3.vhd`
	or
	`\\192.168.199.215\Share1\WinVM\Disk1.vhd`
	`\\192.168.199.215\Share1\WinVM\Disk2.vhd`
	`\\192.168.199.215\share1\WinVM\Disk3.vhd`

Protecting Hyper-V virtual machines with file-level backups

In most cases, Unitrends recommends that you use host-level backups to protect your Hyper-V virtual machines. However, in some instances, you might wish to protect your VMs at the guest level in the same way you would protect physical machines, using file-level backups. Host- and file-level backups provide you with different options.

Use the following topics to determine whether to run host- or file-level backups of Hyper-V virtual machines:

- "General features of Hyper-V host-level and file-level protection" on page 662
- "File-level protection examples" on page 662

General features of Hyper-V host-level and file-level protection

General features of Hyper-V host-level and file-level protection are given here:

Hyper-V protection strategy	Considerations
Host-level backups	• Quickest setup, do not need to add VMs individually or install a Unitrends agent on each VM. • Automatically include new VMs in backup schedules. (Not supported for SLA policy schedules.) • Recover individual files from backups for VMs running Windows or Linux. • Rapid disaster recovery of a failed VM using "Virtual machine instant recovery" on page 904.
File-level backups	• Backup appliance treats the VM like a physical asset. • All backup options are supported, including options to exclude files, directories, or volumes from backup, and run pre- and post-backup commands. Recommended for VMs where more granular exclusion of data is required. • Provide application and operating system consistent backup and recovery. • Protect VM configurations that cannot be protected by host-level backups (such as shared VHDX disks, pass-through disks, and VHDS clusters). • Support Windows replicas to quickly spin up a virtual replica of a failed Windows asset. File-level application backups provide these benefits: **Note:** Application backups schedules cannot be created through SLA policies. • SQL, Exchange, Oracle, and SharePoint backups perform application-level post backup processing, such as log truncation. • Support all SQL database recovery models. Must run file-level application backups for all recovery models other than *simple*. • Support backup of multi-node SharePoint farms.

File-level protection examples

Specific instances when you might want to protect VMs at the asset level are described below. For instructions on setting up file-level protection, see "File-level Backups Overview" on page 703.

Note: To protect a VM with both host-level and file-level (agent-based) backups, ensure that the VM's host-level and file-level jobs do not overlap. Running both simultaneously may lead to undesirable results.

VM type	Protection considerations
Hosted applications	
Hosted applications for which you need more granular control.	Use file-level application backups to select individual databases to back up and recover.
Application versions that are not supported by Integration Services	Use file-level application backups to protect the databases. Use file-level protection for the VMs' file and operating systems.
Exchange	Use either host-level or file-level application backups. (Use file-level application backups if more granular control is needed.)
SQL	For *simple* recovery model databases, use either host-level or file-level application protection. For *full* or *bulk-logged* recovery model databases, use file-level application protection. (Host-level protection is not supported.)
SharePoint	Use file-level application protection. (Host-level protection is not supported.)
Oracle	Use file-level application protection. (Host-level protection is not supported.)
Other VM considerations	
VMs running operating systems that are not supported by Integration Services	With host-level backups, these VMs are temporarily put in a saved state for a brief time during the backup. If you cannot permit a brief VM downtime during the backup, use file-level protection instead.
VMs in a cluster configuration with only one cluster shared volume	Unitrends recommends that you use file-level protection.
VM configurations that cannot be protected with host-level	For these VM configurations, use file-level protection. (See "Virtual machine configuration" on page 656 to determine whether host-level

VM type	Protection considerations
backups.	backups are supported for your VM.)
VMs for which you would like to exclude volumes or large numbers of files when running backups.	Use file-level protection and exclude files from backups.
VMs functioning as large file servers for which you may need to frequently recover files.	Use file-level protection so you can search for files to recover by name.
Windows VMs that you would like to protect with Windows replicas.	Use file-level protection.
Virtualized Active Directory (AD) servers for which you are not following Microsoft's best practices	Use file-level protection.
VMs in Distributed File System environments for which you are not following Microsoft's best practices.	Use file-level protection.

Working with Hyper-V servers

To begin protecting your Hyper-V virtual machines, add to your Unitrends appliance the servers that host them. You must install the Windows agent on the Hyper-V server. For most versions of Windows, this agent is automatically installed when you add the server to the appliance. For details, see "Installing the Windows agent" on page 362. If the VMs you want to protect are clustered (configured as highly available), you must add the cluster and each individual node (server).

When the Hyper-V host is added to the appliance, all hosted VMs are discovered and available for protection. The Windows asset displays on the **Configure > Protected Assets** page. Expand this asset to see the Hyper-V application and hosted VMs. When a cluster is added, only the Hyper-V application displays.

Special considerations for adding Hyper-V clusters

The Unitrends appliance must be able to resolve the name and IP address of every node in a Hyper-V cluster. When adding a cluster node to the appliance, you must enter the correct IP address and the exact name of the node. If you enter an incorrect IP address or a name that does not exactly match the name of the node, backups will fail because the appliance will be unable to determine the owner of the VMs in the cluster configuration. Be sure to enter the correct hostname and IP address for every node in the cluster.

Selecting Hyper-V VMs to protect

Review these guidelines and tips before running Hyper-V backups.

- A separate backup is created for each VM you select.
- A VM may be included in only one schedule. If you attempt to add a VM to a second schedule, you cannot save that schedule. Remove the VM from the first schedule before adding it to another. If the schedule was created by an SLA policy, remove the VM from the policy instead of editing the schedule directly.
- A scheduled or on-demand job can contain one or more VMs that reside on a single host. You cannot run a job that contains VMs from multiple hosts. Create a separate job for each host instead.
- For VMs hosted on Windows Server versions later than 2008 R2, backups are executed concurrently. The number of jobs that run concurrently varies by the resource load of the system, and Monitor the resource utilization on the Hyper-V server to determine whether its backups should be staggered.

The following apply to Hyper-V clusters only:

- Non-clustered VMs hosted on a cluster node do not display when you select the cluster. To protect these VMs, you must select the host node.
- VMs that are hosted on a cluster node, but are not configured as highly available (i.e. not part of the cluster) cannot be included in the same job as the clustered VMs. Instead, you must create a separate job for the node that hosts these non-clustered VMs.
- If multiple virtual machines in a clustered environment are running on Windows Server 2008 R2, the system serializes the backups. Jobs are queued but run one at a time. This is a Windows limitation.

VMware virtual machines

This section provides considerations and requirements for protecting VMware environments.

Preparing for VMware backups

When you add a VMware virtual host to the appliance, all VMs are discovered and available for host-level protection. Unitrends uses VMware's vStorage API for Data Protection (VADP) to communicate with ESXi hosts directly or through a vCenter server. You can add ESXi hosts, vCenter servers, or both, to the Unitrends appliance to protect your VMs. Some features require a vCenter (for details, see "Additional VMware requirements" on page 671).

The following information summarizes the high-level steps that protect VMware virtual machines. The information includes links to detailed instructions for each procedure.

Step 1: Review the "Best practices and requirements for VMware protection".

Step 2: Add the VMware host to your Unitrends appliance. See "Adding a virtual host" on page 308.

Step 3: Create backup jobs for your VMs:

- To create a job manually, see "To create a VMware backup job" on page 455 or "To create a VMware backup schedule by using regular expression filters" on page 458.
- To create a job by using an SLA policy, see "To create an SLA policy for VMware assets" on page 551.

- For a comparison of the manual and SLA policy job creation methods, see "About creating backup and backup copy jobs" on page 426.

Best practices and requirements for VMware protection

Review the information in these topics before implementing VMware host-level protection:

VMware host best practices and considerations

To protect hosted virtual machines, use the "To add a virtual host asset" on page 311 procedure to add the following VMware servers to the Unitrends appliance:

Server	Description
vCenter and managed ESXi servers	If ESXi servers belong to a vCenter and both are accessible on the network, Unitrends recommends that you add the vCenter and its ESXi servers to the appliance. This enables the appliance to contact the vCenter for management operations (including vMotion support) and to directly contact the ESXi servers for backup and recovery, potentially improving performance by reducing network traffic around the vCenter server.
vCenter only	If the ESXi servers are accessible through a vCenter, adding the vCenter to the Unitrends appliance automatically detects all of the associated ESXi servers and their hosted virtual machines. This also enables the Unitrends appliance to be compatible with vMotion, a process through which VMs can migrate among the vCenter's ESXi servers. In this case, the appliance detects when VMs move between ESXi servers in a cluster and contacts the appropriate server to perform backups.
ESXi server only	If ESXi servers are not accessible through a vCenter, or if only a subset of the VMs hosted on the vCenter's ESXi servers are to be protected, you can add individual ESXi servers. In this case, the appliance contacts the ESXi servers directly for backup and recovery.

VMware virtual machine best practices and considerations

Follow these best practices to protect your VMware virtual machines:

- Adhere to VMware's best practices.
- If you are adding an ESXi or vCenter server to multiple Unitrends appliances, be sure to back up each VM on only one appliance. Backing up the same VM on multiple appliances causes problems with the Change Block Tracking (CBT) used for incremental and differential backups.
- If you add a vCenter, Unitrends recommends also adding the individual ESXi hosts managed by the vCenter.

- Full, differential, and incremental backups are supported for VMs configured with hardware version 7 or higher. CBT must be enabled for differentials and incrementals. See "Change Block Tracking (CBT)" on page 669 for details.
- For VMware Windows machines, you can opt to index backups so you can search the VM's backups by filename to quickly recover individual files or folders. See "Indexed VMware Windows backups" on page 669 for details.
- A new full backup is required if the VM configuration has changed since the last backup. This includes any configuration changes made to a VM through the hypervisor, such as creating or deleting a snapshot, or adding a new disk.

  If the VM configuration has changed since the last backup, the next incremental or differential fails. After this failure, the appliance promotes the next scheduled backup to a full (or displays a message indicating a full is required if an on-demand incremental is attempted). Once a full backup succeeds, subsequent incrementals and differentials run as scheduled.
- In some cases, you may want or need to protect VMs by using file-level backups. For recommendations, see "Protecting VMware virtual machines with file-level backups" on page 674.
- To protect a VM with both host-level and file-level (agent-based) backups, be sure to adhere to the following:
  - Ensure that the VM's host-level and file-level jobs do not overlap. Running both simultaneously may lead to undesirable results.
  - If protecting hosted SQL or Exchange databases with agent-based application backups, do not use application-aware protection for host-level backups.
- To protect hosted Exchange or SQL simple recovery model applications, use the application-aware feature for host-level backups. See "VMware application-aware protection" on page 678 for details.
- If recovery time objectives are important, set up "Virtual machine instant recovery" to quickly to spin up a failed VM from host-level backups.
- For Recovery Series and Recovery MAX appliances protecting ESXi hosts whose datastores are located on an external SAN, use the SAN-direct backup feature for host-level backups. These backups run more quickly since network bandwidth is not a hit to performance. See "VMware SAN-direct backups" on page 683 for details.
- For Unitrends Backup on VMware appliances protecting ESXi hosts whose datastores are located on an external SAN, use the HotAdd backup feature for host-level backups. These backups run more quickly since network bandwidth is not a hit to performance. See "VMware HotAdd backups" on page 679 for details.
- For virtual disks hosted on a NAS datastore, running a full backup captures the complete disk (entire virtual disk size).
- Backup failures can occur after a VM’s disks are converted from VHD to VMDK using a third-party tool. For details and solutions for resolving this issue, see Best Practices for Converting Virtual Machine Disks from VHD (Hyper-V) to VMDK (VMware).
- Host-level protection is not supported for the following (use file-level backups instead):
  - VMs in a cluster configuration with a fault tolerant disk.
  - VMs with dynamic MAC addresses.

- Independent and pass-through disks. These disks are automatically excluded from host-level backups.
- Physical Raw Disk Mapping (RDM) disks. These disks are automatically excluded from host-level backups. (Virtual-mode raw device mapped disks are supported with host-level protection.)
- Sparse disks.
- VMs that use VMware NSX storage.
- VMs hosted on free ESXi versions.

General VMware requirements

The following requirements must be met for host-level protection of VMware virtual machines.

Item	Description
ESXi host	Must be running a paid, licensed version listed in the Unitrends Compatibility and Interoperability Matrix. **Notes:** • Host-level protection is not supported for free ESXi versions. To protect VMs hosted on free ESXi, you must install the agent on each VM and run file-level backups. • Additional requirements and limitations apply to ESXi versions 6, 6.5, 6.7, 7.0, and 8.0. See "Additional VMware requirements" on page 671 for details.
vCenter	Must be running a licensed version listed in the Unitrends Compatibility and Interoperability Matrix. **Note:** Additional requirements and limitations apply to vCenter 6, 6.5, 6.7, 7.0, and 8.0. See "Additional VMware requirements" on page 671 for details.
vCenter or ESXi account privileges	An account with full administrative privileges is required. The user or group must have the role *administrator*. You supply these credentials when adding the vCenter or ESXi server to the backup appliance.
Virtual machine configuration	Verify the following VM configuration settings: • VM hardware version must be 4 or higher, and the version must be supported by the VM's ESXi host server. • VMware tools must be installed and running in the guest operating system to ensure file system and application consistency. • The VM must not be configured to use VM encryption.

Item	Description
	Note: VMware introduced new SATA Virtual Hardware Controllers with vSphere 5.5 and VM Hardware Version 10. See Cannot backup VMs with unsupported devices on ESXi 5.5 host for details on selecting the correct controller when creating new VMs in version 5.5.
Change Block Tracking (CBT)	CBT is required to run incremental and differential backups. Running a full backup enables CBT on the VM disks as long as: • VMware tools are installed and running. • No snapshots are present on the VM prior to running the full backup. **Notes:** Only full backups are supported for the following: • VMs configured with hardware version 4. • VMs that have more than 16 disks.
Indexed VMware Windows backups	You can opt to index a VMware Windows VM's backups so you can quickly search for and recover individual files or folders. With indexed backups, you can use wild card searches to find files/folders across all backups of the virtual machine, rather than mounting and browsing each backup individually. Simply select the files/folders and recover them to any Windows agent-based asset that has been added to your backup appliance. To start indexing a VMware VM's backups, select the **Index Backups** option in the Edit Assets dialog (see "To edit a virtual machine asset" on page 317 for details). Subsequent backups of the VM are indexed.

Item	Description
	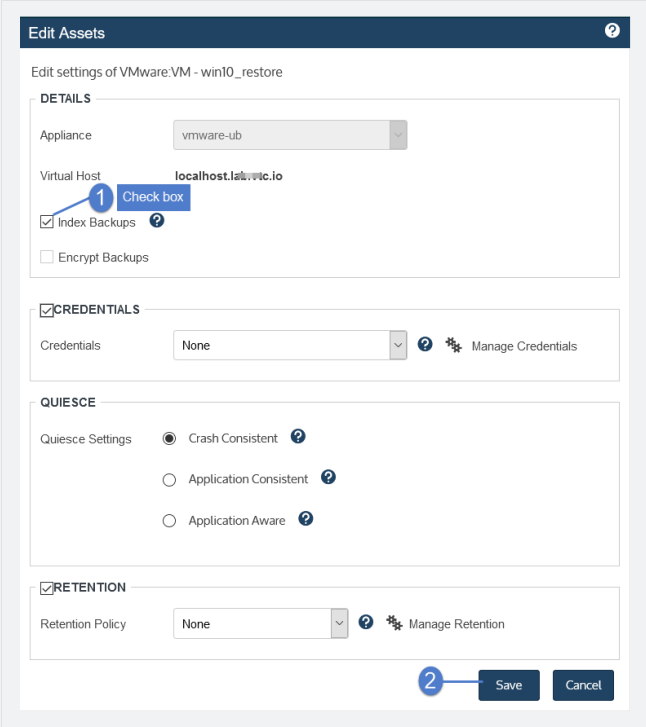 The following requirements and considerations apply: • The backup appliance must be running release 10.5 or higher. • Indexing must be enabled on the appliance. Indexing is enabled by default (Configure > Appliances > Edit Appliance > Advanced > General Configuration > *indexImageLevelBackups* = *Yes*): 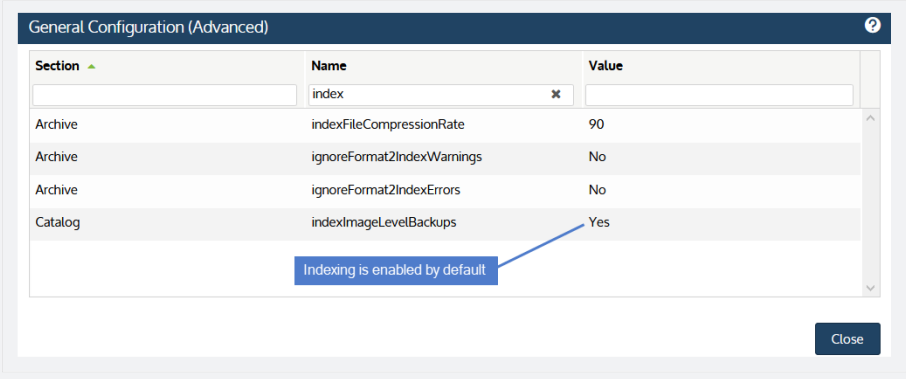 • The target asset where files are recovered must be a Windows asset that has been

Item	Description
	added to the appliance as an agent-based asset. To recover files to the original location or to another Windows VM, install the Unitrends agent on the VM (see "Installing the Windows agent" on page 362), then add the VM to the backup appliance as an agent-based asset (see "To add an agent-based asset" on page 289). • Assets with high-frequency backups or with very large file counts can add considerable load to the appliance. Consider appliance load when enabling the index option for these types of assets. • Filename search of indexed VMware backups is not supported for recovery from imported backup copies. Recover by mounting the imported copy instead (see "Windows file-level recovery" on page 817). • The index feature is not supported for recovery of ReFS filesystems. Recover by mounting and browsing the backup instead (see "Windows file-level recovery" on page 817). • To index the backup, the job creates and mounts an object. If a file recovery object is already mounted for the asset, the backup runs but no index is created (as only one object per asset can be mounted at any given time). The resulting backup completes in warning status, with a message indicating that no index was created.

Additional VMware requirements

These additional requirements may apply to your environment.

Item	Description
vSphere 8.0	To protect VMs hosted in vCenter 8.0 or ESXi 8.0, the following requirements and limitations apply: • The Unitrends appliance must be running release 10.7.5 or higher. • The vCenter must not be configured to use the Server High Availability feature. High Availability is not supported. • Hosted VMs must not be configured to use VM encryption. The VM encryption feature and features that require VM encryption (such as virtual TPM) are not supported. • Hosted VMs must not be configured to use NVDIMM. NVDIMM devices and controllers are not supported. • Hosted VMs must not be configured to use Microsoft Virtualization Based

Item	Description
	Security (VBS). VBS is not supported. • If a vCenter 8.0 VM migrates to a different vCenter, that VM is no longer protected on the original Unitrends schedule. You must manually add it to a new schedule to resume protection.
vSphere 7.0	To protect VMs hosted in vCenter 7.0 or ESXi 7.0, the following requirements and limitations apply: • The Unitrends appliance must be running release 10.4.6 or higher. • The vCenter must not be configured to use the Server High Availability feature. High Availability is not supported. • Hosted VMs must not be configured to use VM encryption. The VM encryption feature and features that require VM encryption (such as virtual TPM) are not supported. • Hosted VMs must not be configured to use NVDIMM. NVDIMM devices and controllers are not supported. • Hosted VMs must not be configured to use Microsoft Virtualization Based Security (VBS). VBS is not supported. • If a vCenter 7.0 VM migrates to a different vCenter, that VM is no longer protected on the original Unitrends schedule. You must manually add it to a new schedule to resume protection.
vSphere 6.7	To protect VMs hosted in vCenter 6.7 or ESXi 6.7, the following requirements and limitations apply: • The Unitrends appliance must be running release 10.3 or higher. • The vCenter must not be configured to use the Server High Availability feature. High Availability is not supported. • Hosted VMs must not be configured to use VM encryption. The VM encryption feature and features that require VM encryption (such as virtual TPM) are not supported. • Hosted VMs must not be configured to use NVDIMM. NVDIMM devices and controllers are not supported. • Hosted VMs must not be configured to use Microsoft Virtualization Based Security (VBS). VBS is not supported.

Item	Description
	• If a vCenter 6.7 VM migrates to a different vCenter, that VM is no longer protected on the original Unitrends schedule. You must manually add it to a new schedule to resume protection.
vSphere 6.5	To protect VMs hosted in vCenter 6.5 or ESXi 6.5, the following requirements and limitations apply: • The Unitrends appliance must be running release 9.1.1 or higher. • The vCenter must not be configured to use the Server High Availability feature. High Availability is not supported. • Hosted VMs must not be configured to use VM encryption. The VM encryption feature is not supported. • If a vCenter 6.5 VM migrates to a different vCenter, that VM is no longer protected on the original Unitrends schedule. You must manually add it to a new schedule to resume protection.
vSphere 6	For VMs hosted in vCenter 6 or ESXi 6, the following limitation applies: If a vCenter 6 VM migrates to a different vCenter, that VM is no longer protected on the original Unitrends schedule. You must manually add it to a new schedule to resume protection.
VMware clusters	To protect VMware clustered environments, you must add the vCenter to your Unitrends appliance.
VMware templates	To protect VMware templates, you must add the vCenter to your Unitrends appliance.
Virtual-mode raw device mapped disks	Raw device mapping (RDM) is a feature of ESXi that allows a virtual disk in a VM to be created on a remote iSCSI LUN rather than on a datastore local to the ESX server. VMs with virtual-mode raw device mapped disks are supported with the following limitations: • The size of the full backup is equal to the entire allocated VM disk size, rather than the used size, since change tracking is not used for RDM backups. • Any RDM disks are recovered as standard virtual disks.
SAN-direct backup for Recovery Series and Recovery	For ESXi hosts whose datastores are located on an external SAN, configure SAN-direct backups. This configuration enables the job to move data directly from the external SAN to the backup appliance during the backup. This direct connection increases backup performance and decreases network bandwidth utilization, affording greater

Item	Description
MAX appliances	scheduling flexibility as the production network is not used during the backup. See "VMware SAN-direct backups" on page 683 for requirements and setup procedures.
HotAdd backup for Unitrends Backup on VMware appliances	For ESXi hosts whose datastores are located on an external SAN, configure backups to use the HotAdd transport mode. This configuration enables the job to move data directly from the external SAN to the appliance during the backup. This direct connection increases backup performance and decreases network bandwidth utilization, affording greater scheduling flexibility as the production network is not used during the backup. See "VMware HotAdd backups" on page 679 for requirements and setup procedures.
Application-aware backups	For Exchange and SQL simple recovery model databases that are hosted on VMware virtual machines, apply local administrator credentials to the VM to run application-aware backups. For details, see "Protecting VMware virtual machines with file-level backups" on page 674.
Virtualized Active Directory servers	To ensure database consistency, you must set up the virtualized Active Directory (AD) server in accordance with Microsoft best practices. If all Microsoft considerations are not addressed, backup and restore of the virtual machine may yield undesired results. If you prefer not to research these best practices, install the agent on the VM and protect it as you would a physical server (leveraging Microsoft's VSS writers).
Distributed File System environments	Distributed File System (DFS) Namespaces and DFS Replication offer high-available access to geographically dispersed files. Because of the replication and syncing operations in DFS environments, you must set up the virtual machine in accordance with Microsoft best practices to ensure database consistency. If all Microsoft considerations are not addressed, backup and restore of the virtual machine may yield undesired results. If you prefer not to research these best practices, install the agent on the VM and protect it as you would a physical server (leveraging Microsoft's VSS writers).

Protecting VMware virtual machines with file-level backups

In most cases, Unitrends recommends that you use host-level backups to protect your VMware virtual machines. However, in some instances, you might wish to protect your VMs at the guest level in the same way you would protect physical machines, using file-level backups. Host- and file-level backups provide you with different options.

Use the following topics to determine whether to run host- or file-level backups of VMware virtual machines:

- "General features of VMware host-level and file-level protection" on page 675
- "File-level protection examples" on page 677
- "Protecting VMware virtual machines with file-level backups" on page 674

General features of VMware host-level and file-level protection

General features of VMware host-level and file-level protection are given here:

Notes:

To protect a VM with both host-level and file-level (agent-based) backups, be sure to adhere to the following:

- Ensure that the VM's host-level and file-level jobs do not overlap. Running both simultaneously may lead to undesirable results.
- If protecting hosted SQL or Exchange databases with agent-based application backups, do not use application-aware protection for host-level backups. Doing so may compromise log truncation changes and lead to other undesirable results.

VMware protection strategy	Considerations
Host-level backups	• Quickest setup, do not need to add VMs individually or install a Unitrends agent on each VM. • Automatically include new VMs in backup schedules. (Not supported for SLA policy schedules.) • Leverages VMware's VADP framework to perform application and operating system consistent backup and recovery. • Application-aware protection of Exchange or SQL simple recovery model applications. • For Recovery Series and Recovery MAX appliances, supports SAN-direct backup. • For Unitrends Backup on VMware appliances, supports HotAdd backup. • Supports backup of VMware templates. • Supports excluding disks from a backup. If you have a requirement to exclude data at the directory- or file-level, or if you don't have space in your VMFS datastores for snapshots of your VMs, consider using file-level backups. • Supports recovering individual files from backups for VMs running Windows or Linux. • Supports VMware instant recovery to quickly spin up a failed VM.
File-level backups	• Backup appliance treats the VM like a physical asset.

VMware protection strategy	Considerations
	• All backup options are supported, including options to exclude files, directories, or volumes from backup, and run pre- and post-backup commands. Recommended for VMs where more granular exclusion of data is required. • Provide application and operating system consistent backup and recovery. • Support Windows replicas to quickly spin up a virtual replica of a failed Windows asset. File-level application backups provide these benefits: **Note:** Application backups schedules cannot be created through SLA policies. • SQL, Exchange, Oracle, and SharePoint backups perform application-level post backup processing, such as log truncation. • Support all SQL database recovery models. Must run file-level application backups for all recovery models other than *simple*. • Support backup of multi-node SharePoint farms.

File-level protection examples

Specific instances when you might want to protect VMs at the asset level are described below. For instructions on setting up file-level protection, see "Protected assets" on page 279.

VM type	Protection considerations
Hosted applications	
Hosted applications for which you need more granular control.	Use file-level application backups to select individual databases to back up and recover.
Exchange	Do one of the following: • Use host-level protection with the application-aware feature (see "VMware application-aware protection" on page 678 for details). • Use file-level application protection if more granular control is needed.
SQL	For simple recovery model databases, do one of the following: • Use host-level protection with the application-aware feature (see "VMware application-aware protection" on page 678 for details). • Use file-level application protection for more granular control. For full or bulk-logged recovery model databases, use file-level application protection. (Host-level protection is not supported)
SharePoint	Use file-level application protection. (Host-level protection is not supported.)
Oracle	Use file-level application protection. (Host-level protection is not supported.)
Disk configuration	
Cluster with fault tolerant disks	Use file-level protection. (Host-level protection is not supported.)
Physical RDM disks	Use file-level protection. (These disks are automatically excluded from host-level backups.)

VM type	Protection considerations
Independent or pass-through disks	Use file-level protection. (These disks are automatically excluded from host-level backups.)
Sparse disks	Use file-level protection. (Host-level protection is not supported.)
Other VM considerations	
Dynamic MAC address	Use file-level protection. (Host-level protection is not supported.)
VMs hosted on free ESXi versions	Use file-level protection. (Host-level protection is not supported.)
Virtualized Active Directory (AD) servers for which you are not following Microsoft's best practices	Use file-level protection.
VMs in Distributed File System environments for which you are not following Microsoft's best practices	Use file-level protection.
VMs for which you would like to exclude volumes or large numbers of files when running backups	Use file-level protection and exclude files from backups. (With host-level you can exclude virtual disks only. File-level provides more granular control.)
VMs functioning as large file servers for which you may need to frequently recover files	Use file-level protection so you can search for files to recover by name.
Windows VMs that you would like to protect with Windows replicas	Use file-level protection.

VMware application-aware protection

To provide application-aware protection of Windows VMs, the appliance requires local administrator credentials to interface with the VM's application-specific VSS writers. For Windows servers that are running UAC, you must either disable UAC or use the server's default local administrator account when applying credentials. (UAC does not apply to the default local administrator account unless specified by a Group Policy, but does apply to other accounts that belong to the Local Administrator group).

See these topics for details on creating and applying credentials to the Windows asset:

- "To add a credential" on page 322
- "To apply a credential to an asset" on page 326

Once credentials have been applied to the Windows VM, the appliance discovers any hosted SQL or Exchange applications, and leverages VSS writers to quiesce data and perform any necessary post-backup processing.

To protect Windows VMs hosting Exchange or SQL simple recovery model applications, Unitrends recommends that you set credentials to ensure an application consistent backup. Log file truncation is handled by VMware application-aware backups as described here:

Application	Log file truncation with VMware application-aware backup
Exchange	Logs are truncated with VMware full and incremental backups.
SQL	Logs are NOT truncated with VMware application-aware backups. Do the following: • Simple recovery model - No SQL logs are created. Use VMware application-aware backups. • Full recovery model - Use agent backups or use VMware application-aware backups with separate transaction log backups to truncate logs. (Schedule periodic transaction log backups using a SQL Maintenance Plan. Do not use a SQL Maintenance Plan with agent backups.) • Bulk-logged recovery model - Use agent backups. See "Recommendations for bulk-logged recovery model" on page 748 for details.

Note: Application-aware backups cannot be used to protect VMware templates or VMs on non-Windows operating systems.

Once you have configured and enabled credentials for a Windows VM, application-aware backups are run. If the appliance cannot gain access using these credentials, the backup fails.

If credentials have not been enabled for the Windows VM, the appliance does not attempt application-aware backup. Application data is included in the host-level backup.

VMware HotAdd backups

For Unitrends Backup appliances that are deployed on VMware, Unitrends recommends that you configure backups to use the HotAdd transport mode for ESX hosts whose datastores are located on an external SAN. This configuration enables the appliance to move data directly from the external SAN to the appliance during the backup. This direct connection increases backup performance and decreases network bandwidth utilization, affording greater scheduling flexibility as the production network is not used during the backup.

See the following topics to set up HotAdd backups:

- "VMware HotAdd requirements"
- "To configure a Unitrends Backup on VMware appliance for HotAdd backups" on page 680

VMware HotAdd requirements

In addition to the "Best practices and requirements for VMware protection" on page 666, these requirements must be met to run VMware backups using the HotAdd transport mode:

- The Unitrends appliance must be a Unitrends Backup on VMware appliance. (For Recovery Series and Recovery MAX appliances, use the SAN-direct feature instead. See "VMware SAN-direct backups" on page 683 for details.)
- Both the Unitrends Backup VM and the VMs to protect must be running in an ESXi environment with the following storage configurations:
  - SAN storage mounted on the ESXi server as a VMFS datastore.
  - SAN storage mounted directly by the ESXi guest as a virtual RDM disk.

  Other storage configurations, such as VSANs and VVoL datastores, are not supported.
- Both the Unitrends Backup VM and the VMs to protect must be added to the ESXi host through a vCenter. Directly adding a VM to the ESXi host is not supported.
- The VMware vSphere version must either be version 5.1 or later, or 5.0 with an advanced license that supports the HotAdd feature. (See VMware documentation for details.)
- The HotAdd transport mode can be used for all guest operating system versions.
- The HotAdd transport mode cannot be used to protect VMware templates.
- Do not use the HotAdd transport mode to protect Linux VMs configured with Logical Volume Manager (LVM) disks.

  IMPORTANT! If your datastores are located on SAN LUNs, work with Unitrends Support to disable the HotAdd feature on the appliance itself to ensure the HotAdd transport mode is not used for backups of your LVM disks.
- For optimal performance, it is recommended to use VAAI-compatible storage arrays with the HotAdd feature.

  For non-VAAI storage arrays, the ESXi host may perform certain operations on the datastore while the HotAdd backup is running, causing a SCSI reservation conflict. In this case, the backup continues over the network instead of the SAN, which increases both backup time and network traffic on the LAN. To reduce the probability of reservation conflicts, follow VMware's recommendations in the article Frequently Asked Questions for vStorage APIs for Array Integration (1021976).

To configure a Unitrends Backup on VMware appliance for HotAdd backups

Note: No configuration of the Unitrends storage subsystem is required. All configuration is done in the VMware and SAN environments.

1. Configure the ESXi server(s) to access the datastore(s) on the SAN LUNs.

   In clustered environments, all ESXi servers in the cluster must be configured to access the datastores.
2. Physically connect the ESXi host machine to the SAN using an iSCSI or Fibre Channel cable.
3. Run backups for the guest VMs as described in "To create a VMware backup job" on page 455.

   The appliance detects the SAN storage configuration and uses the HotAdd transport method during backup.

Notes:

- If the appliance cannot directly access the SAN, the backup runs using the network connection, moving the guest's data from the external SAN through the ESXi file system to the backup appliance. If using network bandwidth is a problem during certain hours, schedule your backups accordingly.
- Note that when recovering HotAdd backups, VMware determines the most efficient transport method. Data is usually restored directly to the ESXi server as the VMFS datastore must be utilized.

4 After running the first HotAdd backup, view backup details to verify that the HotAdd transport method was used. View details by doing one of the following:

- Click **Jobs > Recent Jobs**, select the job and click **View Log**. Look for the transport method in the Output area.

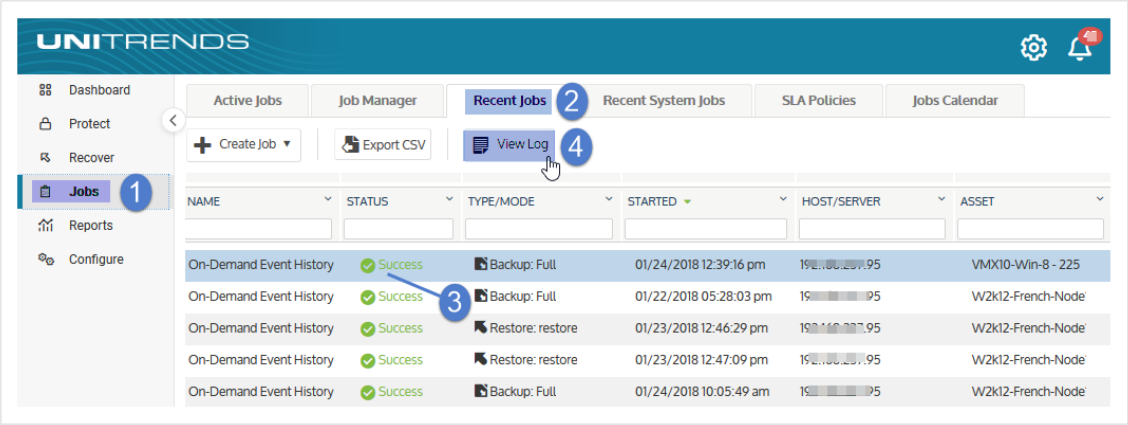

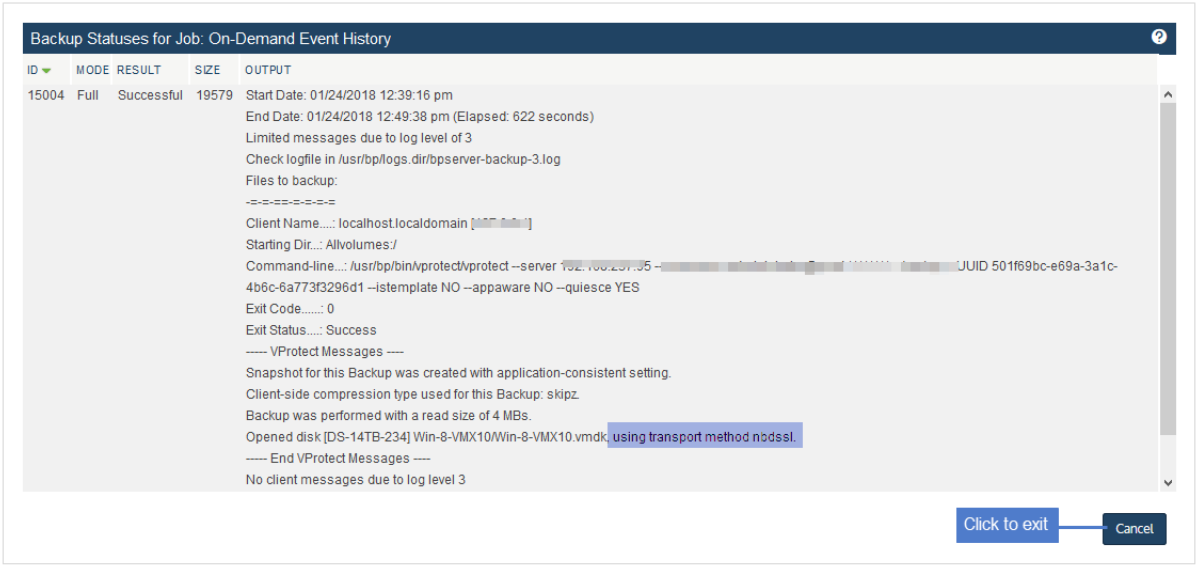

- Click **Reports > Backup > Backup History**. Run the report by entering a date range and clicking **Generate Report**. Select the job and look for the transport method in the Output area in the Backup Status Report Entry dialog.

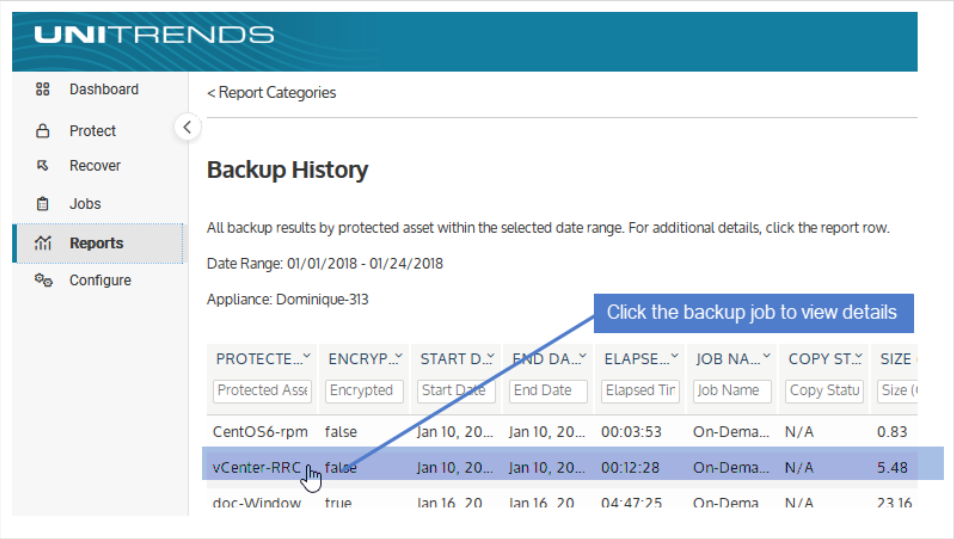

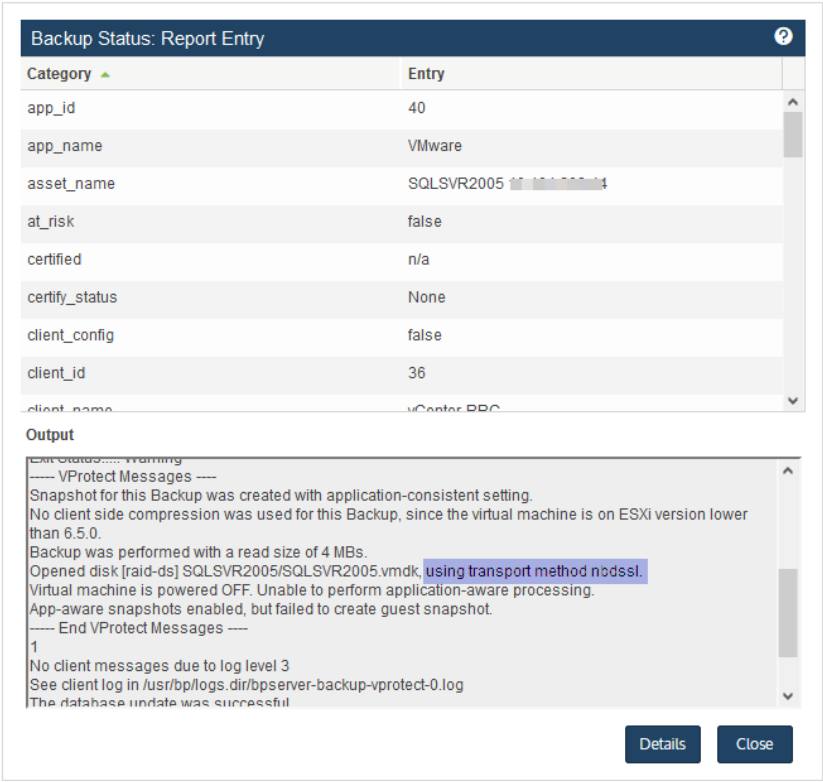

The transport method message indicates whether the HotAdd method was used:

- *using transport method hotadd* indicates the HotAdd transport mode was used for the backup.
- *using transport method nbd* or *using transport method nbdssl* indicates the regular network connection was used for the backup. Check your physical cable connection and the iSCSI or FC configuration for issues.

VMware SAN-direct backups

For Unitrends Recovery Series, Recovery MAX, and ION+ physical appliances, Unitrends recommends that you configure SAN-direct backups for ESXi hosts whose datastores are located on an external SAN. This configuration enables backups to move data directly from the external SAN to the backup appliance. This direct connection increases backup performance and decreases network bandwidth utilization, affording greater scheduling flexibility as the production network is not used during the backup.

See the following topics to set up SAN-direct backups:

- "VMware SAN-direct requirements"
- "To configure a Recovery Series, Recovery MAX, or ION+ appliance for SAN-direct backups" on page 683

VMware SAN-direct requirements

In addition to the "Best practices and requirements for VMware protection" on page 666, these requirements must be met to run VMware backups by using the SAN transport mode:

- The Unitrends appliance must be a Recovery Series, Recovery MAX, or ION+ appliance. (For Unitrends Backup on VMware appliances, use the HotAdd transport mode instead. See "VMware HotAdd backups" on page 679 for details.)

  Note: SAN-direct backups are not supported on ION appliances.

- VMs to protect must be running in an ESXi environment with the following storage configurations:
  - SAN storage mounted on the ESXi server as a VMFS data store.
  - SAN storage mounted directly by the ESXi guest as a virtual RDM disk.

  Other storage configurations, such as VSANs and VVoL datastores, are not supported.
- SAN-direct backup can be used for all guest operating system versions.
- SAN-direct backup cannot be used to protect VMware templates.
- For optimal performance, it is recommended to use VAAI-compatible storage arrays with the SAN-direct feature.

  For non-VAAI storage arrays, the ESXi host may perform certain operations on the datastore while the SAN-direct backup is running, causing a SCSI reservation conflict. In this case, the backup continues over the network instead of the SAN, which increases both backup time and network traffic on the LAN. To reduce the probability of reservation conflicts, follow VMware's recommendations in the article Frequently Asked Questions for vStorage APIs for Array Integration (1021976).

To configure a Recovery Series, Recovery MAX, or ION+ appliance for SAN-direct backups

1. Physically connect the Unitrends appliance to the SAN using an iSCSI or Fibre Channel cable.

2 Log in to the backup appliance.

3 On the **Configure > Appliances** page, select the appliance and click **Edit**.

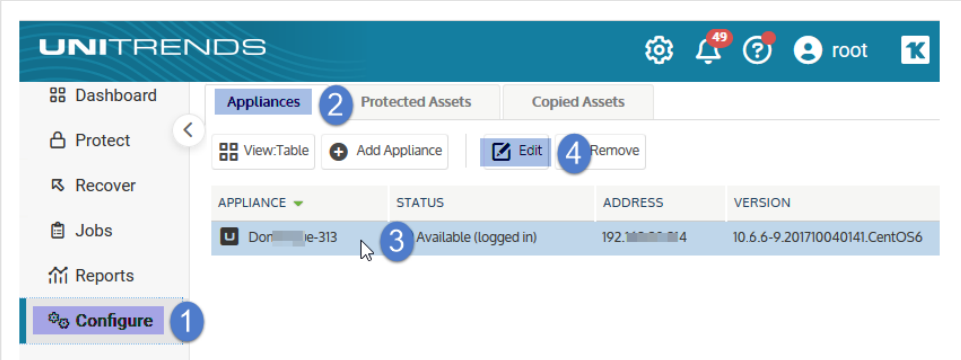

4 Click **Advanced** and select **Add** in the San Direct Details area.

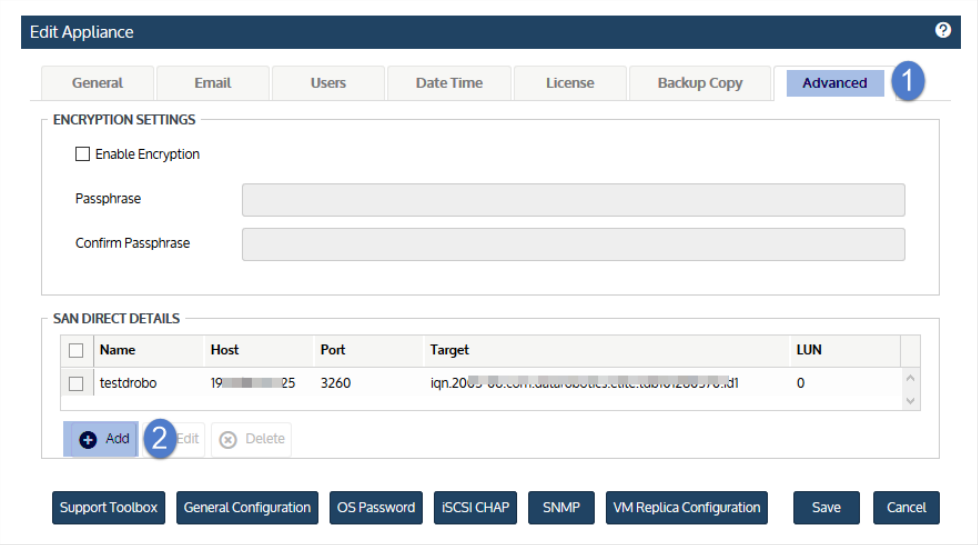

5 Enter a unique **Name** for the storage device. This name cannot contain spaces.

6 Enter the IP address of the SAN storage array in the **Host** field.

7 The default port used for iSCSI communication is 3260. If the LUN is configured to use a different port, enter it in the **Port** field.

8 Click **Scan for targets** to retrieve a list of targets on the remote storage array, then choose one from the list.

Notes: If you do not see the LUN in the list, go to your SAN manager and check the following:

- Verify that you can see the Unitrends appliance in your SAN manager.
- Verify that you have a LUN assigned to the Unitrends appliance with the correct permissions.

- Check with your Storage Administrator for more information.

9 Click **Scan for LUNs** and select one from the list.

Note: If you receive an error indicating CHAP authentication has failed, CHAP has been configured on the target and either CHAP has not been enabled on the Unitrends appliance, or the Unitrends CHAP credentials do not match those of the target. To configure the appliance to use CHAP, see "To configure iSCSI CHAP authentication" on page 182.

10 Click **Save**.

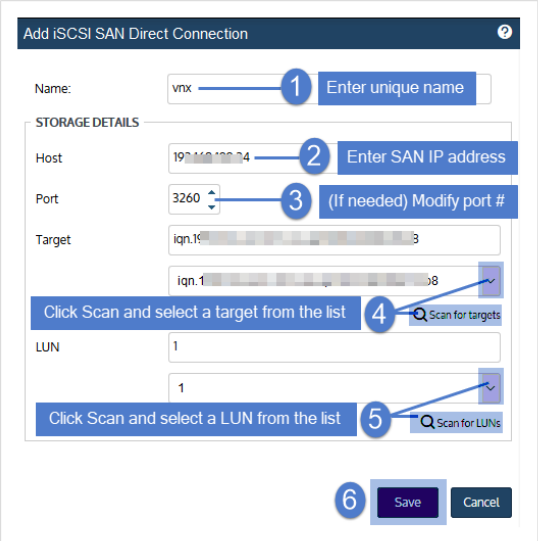

11 Click **Save** to exit.

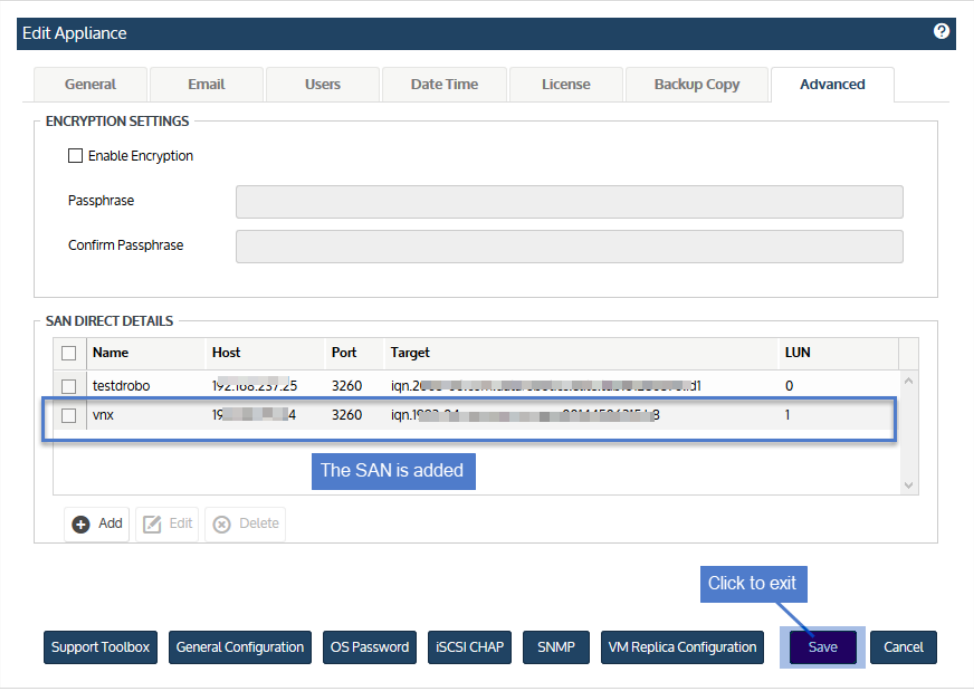

12 Run backups for the guest VMs as described in "To create a VMware backup job" on page 455.

The appliance detects the SAN storage configuration and uses the SAN-direct transport method during backup.

Notes:

- If the appliance cannot directly access the SAN, the backup runs using the network connection, moving the guest's data from the external SAN through the ESXi file system to the backup appliance. If using network bandwidth is a problem during certain hours, schedule your backups accordingly.
- Note that when recovering SAN-direct backups, VMware determines the most efficient transport method. Data is usually restored directly to the ESXi server as the VMFS datastore must be utilized.

13 After running the first SAN-direct backup, view backup details to verify that the SAN transport method was used. View details by doing one of the following:

- Click **Jobs > Recent Jobs**, select the job and click **View Log**. Look for the transport method in the Output area.

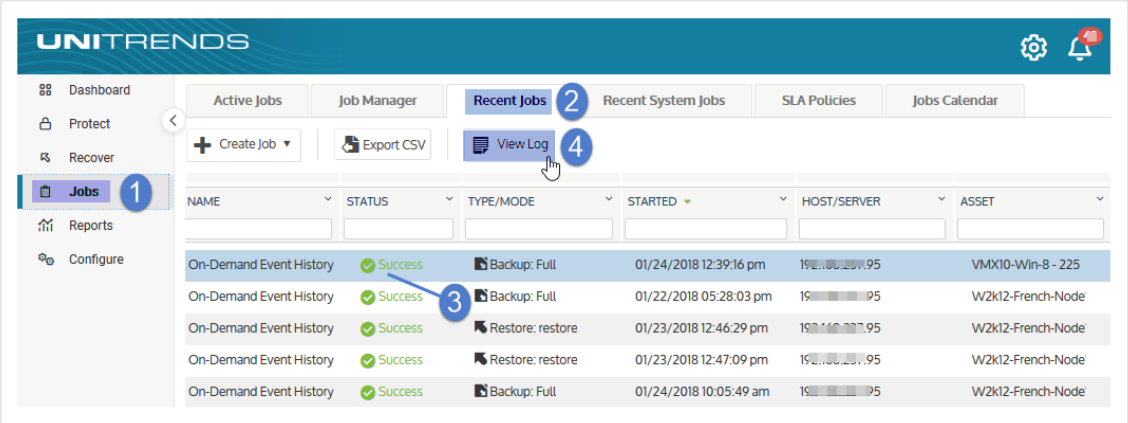

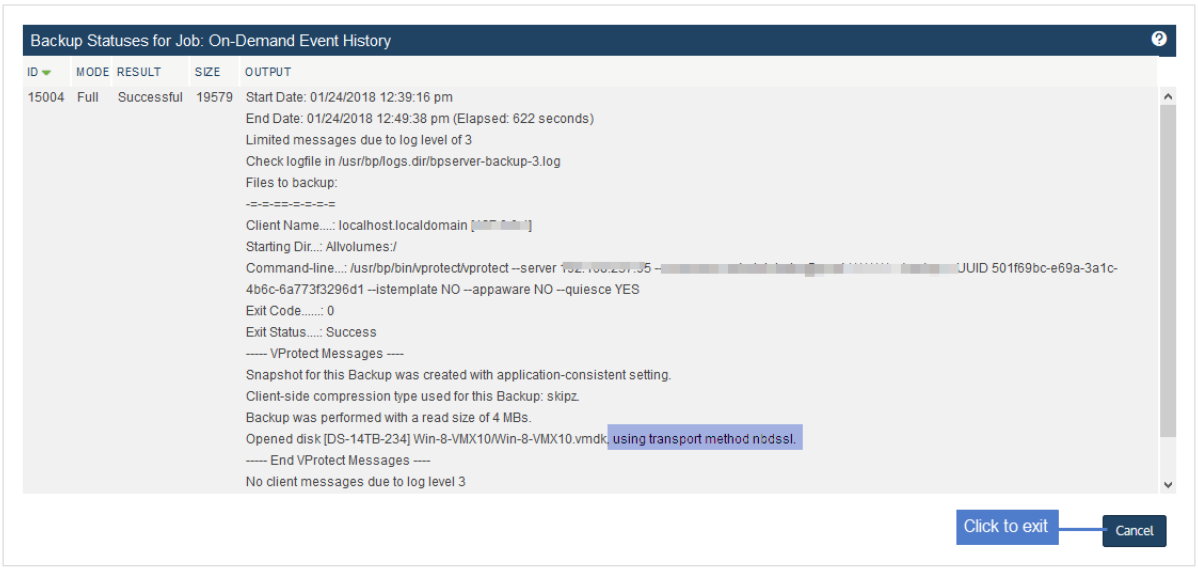

- Click **Reports > Backup > Backup History**. Run the report by entering a date range and clicking **Generate Report**. Select the job and look for the transport method in the Output area in the Backup Status Report Entry dialog.

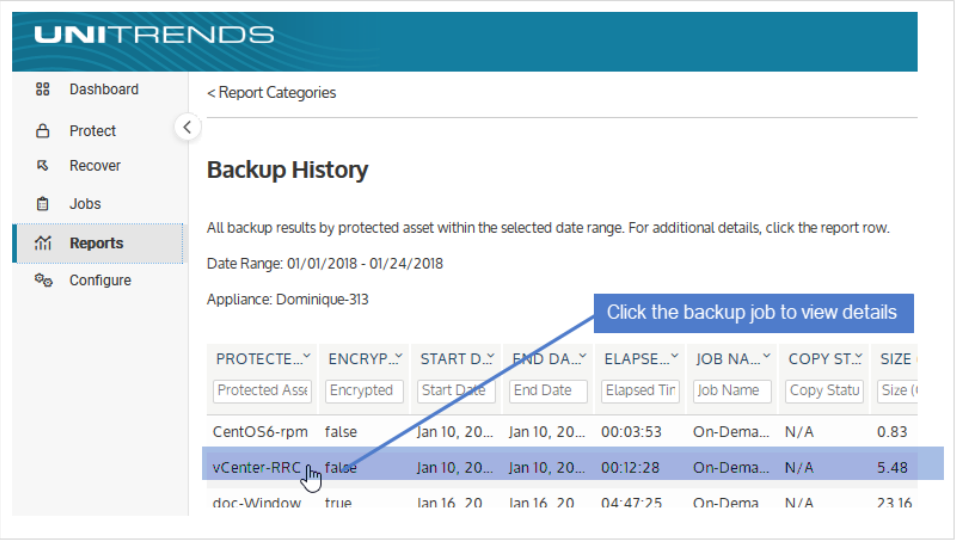

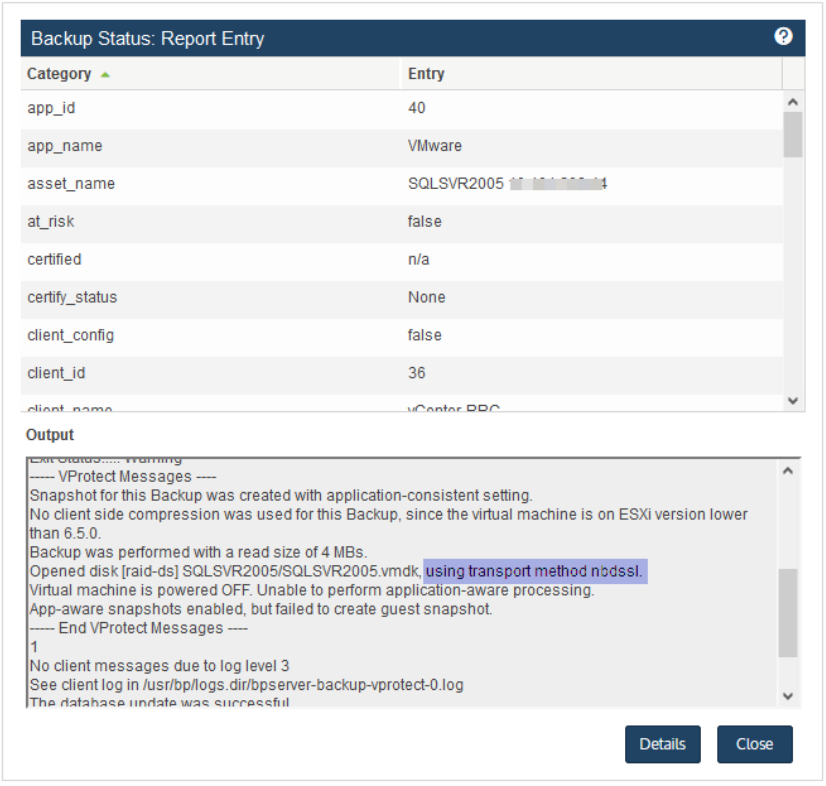

The transport method message indicates whether the SAN-direct connection was used:

- *using transport method san* indicates the SAN-direct connection was used for the backup.

- *using transport method nbd* or *using transport method nbdssl* indicates the regular network connection was used for the backup. Check your physical cable connection and the iSCSI or FC configuration for issues.

Citrix XenServer virtual machines

This section provides details and requirements for protecting Citrix XenServer environments with host-level backups.

Preparing for XenServer backups

When you add a XenServer host to the Unitrends appliance, its VMs are discovered and available for host-level protection. Following is a summary of the high-level steps that protect XenServer virtual machines. Included are links to detailed instructions for each procedure.

1. Review the "Best practices and requirements for XenServer protection" on page 689.
2. Add the XenServer host to your Unitrends appliance. See "Adding a virtual host" on page 308.
3. Create backup jobs for your VMs. See "Creating backup jobs" on page 433.

Best practices and requirements for XenServer protection

Review the information in these topics before implementing XenServer host-level protection:

- "XenServer host best practices and considerations" on page 689
- "XenServer virtual machine best practices and considerations" on page 690

XenServer host best practices and considerations

Host-level protection of XenServer VMs is supported on Unitrends Backup on Citrix XenServer appliances only.

> **Note:** To protect XenServer VMs using other Unitrends appliance types, use agent-based backups instead. For details, see "File-level Backups Overview" on page 703.

The XenServer host must be running a licensed version listed in the Unitrends Compatibility and Interoperability Matrix. The host must be one of the following:

- A XenServer pool master host meeting both of these criteria:
  - The Unitrends Backup VM resides either on the pool master host itself or on one of the pool master's slave hosts.
  - The Unitrends Backup VM has been granted access to the shared storage used by the pool master host.
- A stand-alone XenServer host where the Unitrends Backup VM resides.

Only one XenServer host can be added to the Unitrends Backup appliance. (See the procedure "To add a virtual host asset" on page 311.) See the following for considerations by host type:

Host type	Description
XenServer pool master host	• By adding the pool master host to the Unitrends Backup appliance, you can protect the following: – VMs on the host where the Unitrends Backup VM resides. This can be the pool master host itself or one of the pool master's slave hosts. – VMs that reside on shared storage in the resource pool. • While adding the pool master host to the Unitrends Backup appliance, enter its username and password credentials in the Add Virtual Host dialog. These credentials are needed because the XenServer APIs have to communicate via the pool master. The appliance can then discover hosted VMs and VM slaves, as well as track any VM changes.
Stand-alone XenServer host	• By adding the XenServer host to the appliance, you can protect its hosted VMs. • While adding the XenServer host to the Unitrends Backup appliance, enter its username and password credentials in the Add Virtual Host dialog.

XenServer virtual machine best practices and considerations

Follow these best practices to protect your XenServer virtual machines:

- Adhere to all Citrix XenServer best practices. This information can be found in your Citrix XenServer Administrator's guide.

  Note: Failure to comply with Citrix recommendations for quiesced snapshots can result in backup failures.

- If you are deploying multiple Unitrends Backup appliances in your XenServer environment, be sure to back up each VM from one appliance only. This ensures that a given VM will not be backed up simultaneously by multiple appliances, which can cause undesirable results.
- Backups can protect VMs, VMs on slave XenServer hosts, and user-configured templates. Only full backups are supported.
- To back up VM disks, the Unitrends Backup appliance must have access to the storage repositories where the disks reside. If the appliance cannot access any VM disk, the backup fails.
- In some cases, you may want or need to protect VMs using file-level backups. If you choose to protect a VM with both host-level and file-level (agent-based) backups, ensure that the VM's host-level and file-level jobs do not overlap. Running both simultaneously may lead to undesirable results.

Protecting XenServer VMs with file-level backups

In most cases, Unitrends recommends that you use host-level backups to protect your XenServer virtual machines. However, in some instances, you might wish to protect your VMs at the guest level in the same way you would protect physical machines, by using file-level backups. Host- and file-level backups provide you with different options.

Use the following topics to determine whether to run host- or file-level backups of XenServer virtual machines:

- "General features of XenServer host-level and file-level protection"
- "File-level protection examples" on page 693

General features of XenServer host-level and file-level protection

General features of XenServer host-level and file-level protection are given here:

Note: To protect a VM with both host-level and file-level (agent-based) backups, ensure that the VM's host-level and file-level jobs do not overlap. Running both simultaneously may lead to undesirable results.

XenServer protection strategy	Considerations
Host-level backups	• Quickest setup, do not need to add VMs individually or install a Unitrends agent on each VM. • Automatically include new VMs in backup schedules. (Not supported for SLA policy schedules.) • Supports excluding disks from a backup. If you have a requirement to exclude data at the directory- or file-level, or if you don't have space in your containers for snapshots of your VMs, consider using file-level backups. • Supports recovering individual files from backups for VMs running Windows or Linux.
File-level backups	• Backup appliance treats the VM like a physical asset. • All backup options are supported, including options to exclude files, directories, or volumes from backup, and run pre- and post-backup commands. Recommended for VMs where more granular exclusion of data is required. • Provide application and operating system consistent backup and recovery. • Support Windows replicas to quickly spin up a virtual replica of a failed Windows asset. File-level application backups provide these benefits: **Note:** Application backups schedules cannot be created through SLA policies.

XenServer protection strategy	Considerations
	• SQL, Exchange, Oracle, and SharePoint backups perform application-level post backup processing, such as log truncation. • Support all SQL database recovery models. Must run file-level application backups for all recovery models other than *simple*. • Support backup of multi-node SharePoint farms.

File-level protection examples

Specific instances when you might want to protect VMs at the asset level are described below. For instructions on setting up file-level protection, see "Protected assets" on page 279.

VM server type	Protection considerations
Application server hosting applications for which you need more granular control.	Use file-level application backups to select individual databases to back up and recover.
Exchange server	Use file-level application protection if more granular control is needed.
SQL server	• For simple recovery model databases, use file-level application protection for more granular control. • For full or bulk-logged recovery model databases, use file-level application protection. (Host-level protection is not supported)
SharePoint server	Use file-level application protection. (Host-level protection is not supported.)
Oracle server	Use file-level application protection. (Host-level protection is not supported.)
Virtualized Active Directory (AD) servers for which you are not following Microsoft's best practices	Use file-level protection.
VMs in Distributed File System environments for which you are not following Microsoft's best practices.	Use file-level protection.
VMs for which you would like to exclude volumes or large numbers of files when running backups.	Use file-level protection and exclude files from backups. (With host-level you can exclude virtual disks only. File-level provides more granular control.)
VMs functioning as large file servers for which you may need to frequently recover files.	Use file-level protection so you can search for files to recover by name.
Windows VMs that you would like to protect with Windows replicas.	Use file-level protection.

AHV virtual machines

This section provides considerations and requirements for protecting virtual machines hosted in Nutanix Acropolis Hypervisor (AHV) environments.

Note: AHV host-level backups include the VM file data but not the Acropolis File Services data that is internal to the Nutanix share. To protect this internal data, you can add the Nutanix share to the Unitrends appliance and run Unitrends NAS backups over the NFS protocol. Start by reviewing the requirements in "NAS protection using CIFS/NFS" on page 726. Next, add the share to the Unitrends appliance as described in "To add a NAS CIFS or NFS asset" on page 296. Then create a backup job as described in "To create a NAS CIFS or NFS backup job" on page 470.

Preparing for AHV backups

When you add an AHV host cluster to the appliance, all VMs are discovered and available for host-level protection. Unitrends leverages Nutanix REST APIs to communicate with AHV hosts.

The following information summarizes the high-level steps that protect AHV virtual machines. The steps include links to detailed instructions for each procedure.

Step 1: Review the "Best practices and requirements for AHV protection".

Step 2: Add the AHV host cluster to your Unitrends appliance. See "Adding a virtual host" on page 308.

Step 3: Create backup jobs for your VMs:

- To create a job manually, see "To create a Nutanix AHV backup job" on page 465.
- To create a job by using an SLA policy, see "To create an SLA policy for AHV assets" on page 559.
- For a comparison of the manual and SLA policy job creation methods, see "About creating backup and backup copy jobs" on page 426.

Best practices and requirements for AHV protection

Review the information in these topics before implementing AHV host-level protection:

AHV best practices and considerations

Follow these best practices to protect your AHV virtual machines:

- Adhere to Nutanix best practices.
- Full and incremental backups are supported for AHV VMs.
- A new full backup is required if the VM configuration has changed since the last backup. This includes any configuration changes made to a VM through the hypervisor, such as creating or deleting a snapshot, or adding a new disk.

If the VM configuration has changed since the last backup, the next incremental fails. After this failure, the appliance promotes the next scheduled backup to a full (or displays a message indicating a full is required if an on-demand incremental is attempted). Once a full backup succeeds, subsequent incrementals run as scheduled.

- Due to a Nutanix limitation, Unitrends AHV snapshots do not display in the Nutanix AHV hypervisor. Note the following:
  - The first time a VM is backed up, the job creates a new snapshot of the AHV VM that remains with the VM after the job completes. During subsequent backups, the job creates a new snapshot of the AHV VM, performs the backup, then removes either the previous snapshot (if the job was successful) or the current snapshot (if the job failed). If a job ends ungracefully (such as due to a power outage) the unneeded snapshot may remain on the hypervisor. A Unitrends cleanup process runs hourly to check for and remove any unneeded snapshots.
  - If you are no longer protecting a VM on this Unitrends appliance, any leftover snapshot that has not been removed will remain on the hypervisor. This applies even if you begin protecting the VM with another Unitrends appliance. If you are no longer protecting a VM with the original Unitrends appliance, contact Support for assistance removing any unneeded snapshots.
- In some cases, you may want or need to protect VMs by using agent based, file-level or image-level backups. For recommendations, see "Protecting AHV virtual machines with file-level backups" on page 698.
- To protect a VM with both host-level and file-level (agent-based) backups, ensure that the VM's host-level and file-level jobs do not overlap. Running both simultaneously may lead to undesirable results.

General AHV requirements

The following requirements must be met for host-level protection of AHV virtual machines.

Item	Description
Unitrends appliance version	The Unitrends appliance must be running version 10.2 or higher. The Unitrends appliance must be running an Enterprise or Enterprise Plus license.
Nutanix AHV host cluster version	The AHV host must be running Acropolis Operating System (AOS) version 5.1.4 or a higher version supported version listed in the Unitrends Compatibility and Interoperability Matrix.
AHV host account privileges	While adding the AHV cluster to the Unitrends appliance (described in "To add a virtual host asset" on page 311), you must enter the username and password credentials of one of the following AHV cluster accounts: • The Nutanix **cluster admin** account – You must use this account if the AHV cluster is not configured to use directory services authentication and the cluster is running a pre-5.5 AOS release. Other user accounts with full administrative privileges are not supported. • Any local Nutanix cluster account that has been assigned the **user admin** or

Item	Description
	cluster admin role – Use a local account with either of these roles if the AHV cluster is not configured to use directory services authentication and the cluster is running AOS release 5.5. • An LDAP user that has the **cluster admin** role – Use this account if the AHV cluster is configured to use directory services authentication. See these topics for additional requirements: "Requirements for directory services authentication in AOS 5.1" or "Requirements for directory services authentication in AOS 5.5". **Requirements for directory services authentication in AOS 5.1** These additional requirements apply to Nutanix AHV clusters running in AOS 5.1 that are configured to use directory services authentication: • Set up authentication (as described in this Nutanix document: Configuring Authentication) to use these settings: – In the Directory List, add a new directory of type **Active Directory** and connection **LDAP**. For the Directory URL, specify **ldap://<ip-address>:<port>** – Create a role mapping for the LDAP user and assign the **cluster admin** role. – In the self service portal (SSP), set or update the SSP administrators to the user@domain. Use fully qualified domain names. – SSP will need to query the active directory for details of users. Ideally a service account with no time limit should be used. This account must have privileges for listing the users in the Directory server. • While adding the AHV cluster to the Unitrends appliance (described in "To add a virtual host asset" on page 311), you must specify a domain in addition to the username. The username and domain are case sensitive. Be sure to match the case that was entered in the self service portal (SSP). In the Username field, enter the username and domain in this format: *user@domain*. For example, **jalvarez@unitrends.com** **Requirements for directory services authentication in AOS 5.5** These additional requirements apply to Nutanix AHV clusters running in AOS 5.5 that are configured to use directory services authentication: • Set up authentication (as described in this Nutanix document: Configuring Authentication) to use these settings: – In the Directory List, add a new directory of type **Active Directory** and connection **LDAP**. For the Directory URL, specify **ldap://<ip-address>:<port>** – Create a role mapping for the LDAP user and assign the **cluster admin** role.

Item	Description
	• While adding the AHV cluster to the Unitrends appliance (described in "To add a virtual host asset" on page 311), you must specify a domain in addition to the username. In the Username field, enter the username and domain in this format: *user@domain*. For example, **jalvarez@unitrends.com**
iSCSI target access	AHV backup and recovery jobs access the AHV host over the iSCSI protocol. Ensure the following: • The Unitrends appliance is able to connect to the iSCSI targets on the Nutanix storage LAN. • iSCSI Data Services are configured for the Nutanix AHV cluster. To configure this setting: 1 In the Nutanix Prism interface, select **Cluster Details** from the Options menu. 2 Enter the iSCSI Data Services IP address. 3 Click **Save**. 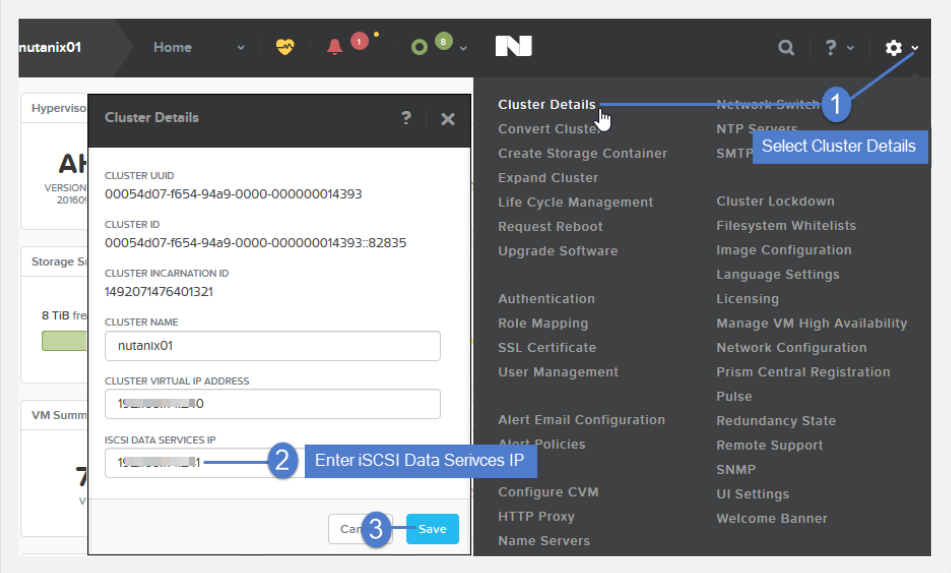
Virtual machine storage	The following requirements apply to virtual machine storage: • Virtual machine storage must be disk storage allocated on a storage container. VM disks that are attached to a Volume Group are not included in the backup. • Host-level protection is not supported for independent and pass-through disks. To

Item	Description
	protect these disks, you must install a Unitrends agent and use file-level backups instead.
Virtual machine configuration	The following VM configuration requirements must be met for Unitrends host-level protection: • Nutanix recommends installing Nutanix Guest Tools (NGT) in the guest operating system to ensure file system and application consistency. • NGT tools must be installed and running to enable application consistent quiesce. If NGT is not running, crash consistent quiesce is used. For details, see this Nutanix document: Nutanix Guest Tools. • For Windows guests, Nutanix recommends installing VirtIO drivers for enhanced performance and stability. For details, see this Nutanix document: Nutanix Virtio for Windows.
Virtualized Active Directory servers	To ensure database consistency, you must set up the virtualized Active Directory (AD) server in accordance with Microsoft best practices. If all Microsoft considerations are not addressed, backup and restore of the virtual machine may yield undesired results. If you prefer not to research these best practices, install the agent on the VM and protect it as you would a physical server (leveraging Microsoft's VSS writers).
Distributed File System environments	Distributed File System (DFS) Namespaces and DFS Replication offer high-available access to geographically dispersed files. Because of the replication and syncing operations in DFS environments, you must set up the virtual machine in accordance with Microsoft best practices to ensure database consistency. If all Microsoft considerations are not addressed, backup and restore of the virtual machine may yield undesired results. If you prefer not to research these best practices, install the agent on the VM and protect it as you would a physical server (leveraging Microsoft's VSS writers).

Protecting AHV virtual machines with file-level backups

In most cases, Unitrends recommends that you use host-level backups to protect your AHV virtual machines. However, in some instances, you might wish to protect your VMs at the guest level in the same way you would protect physical machines, by using agent-based file-level or Windows image-level backups. Host-level and agent-based backups provide you with different options.

Use the following topics to determine whether to run host-level or agent-based backups of AHV virtual machines:

- "General features of AHV host-level and agent-based protection"
- "File-level protection examples" on page 701

General features of AHV host-level and agent-based protection

General features of AHV host-level and agent-based protection are given here:

Note: To protect a VM with both host-level and agent-based backups, ensure that the VM's host-level and agent-based jobs do not overlap. Running both simultaneously may lead to undesirable results.

AHV protection strategy	Considerations
Host-level backups	The following apply to host-level backups: • Quickest setup, do not need to add VMs individually or install a Unitrends agent on each VM. • Automatically include new VMs in backup schedules. (Not supported for SLA policy schedules.) • Leverages Nutanix REST APIs to perform application and operating system consistent backup and recovery. Application consistent protection requires that Nutanix Guest Tools (NGT) are installed and running on the virtual machine. • Supports excluding disks from a backup. If you have a requirement to exclude data at the directory- or file-level, or if you don't have space in your containers for snapshots of your VMs, consider using file-level backups. • Supports recovering individual files from backups for VMs running Windows or Linux.
Agent based file-level backups	The following apply to file-level backups: • Backup appliance treats the VM like a physical asset. • All backup options are supported, including options to exclude files, directories, or volumes from backup, and run pre- and post-backup commands. Recommended for VMs where more granular exclusion of data is required. • Provide application and operating system consistent backup and recovery. • Support Windows replicas to quickly spin up a virtual replica of a failed Windows asset. File-level application backups provide these benefits: **Note:** Application backups schedules cannot be created through SLA policies. • SQL, Exchange, Oracle, and SharePoint backups perform application-level post backup processing, such as log truncation.

AHV protection strategy	Considerations
	• Support all SQL database recovery models. Must run file-level application backups for all recovery models other than *simple*. • Support backup of multi-node SharePoint farms.
Agent based Windows image-level backups	The following apply to Windows image-level backups: • Assets are protected and the disk and volume level. • Backups include the 'in use' regions of the Windows disk or volume only. Deleted regions are not included. • Provide application consistent backup and recovery for NTFS and ReFS filesystems. • Provide crash consistent backup and recovery for FAT, FAT32, and exFAT filesystems. • Opt to include or exclude volumes from backup. • Opt to run commands on the asset before or after the backup job. • Opt to index the backup so you can search an asset's backups by filename to recover individual files. This option is set by asset (see "To edit an agent-based asset" on page 293 for details).

File-level protection examples

Specific instances when you might want to protect VMs at the asset level are described below. For instructions on setting up file-level protection, see "Protected assets" on page 279.

VM type	Protection considerations
Application server hosting applications for which you need more granular control.	Use file-level application backups to select individual databases to back up and recover.
Exchange server	Use file-level application protection if more granular control is needed.
SQL server	• For simple recovery model databases, use file-level application protection for more granular control. • For full or bulk-logged recovery model databases, use file-level application protection. (Host-level protection is not supported)
SharePoint server	Use file-level application protection. (Host-level protection is not supported.)
Oracle server	Use file-level application protection. (Host-level protection is not supported.)
Virtualized Active Directory (AD) servers for which you are not following Microsoft's best practices	Use file-level protection.
VMs in Distributed File System environments for which you are not following Microsoft's best practices.	Use file-level protection.
VMs for which you would like to exclude volumes or large numbers of files when running backups.	Use file-level protection and exclude files from backups. (With host-level you can exclude virtual disks only. File-level provides more granular control.)
VMs functioning as large file servers for which you may need to frequently recover files.	Use file-level protection so you can search for files to recover by name.
Windows VMs that you would like to protect with Windows replicas.	Use file-level protection.

This page is intentionally left blank.

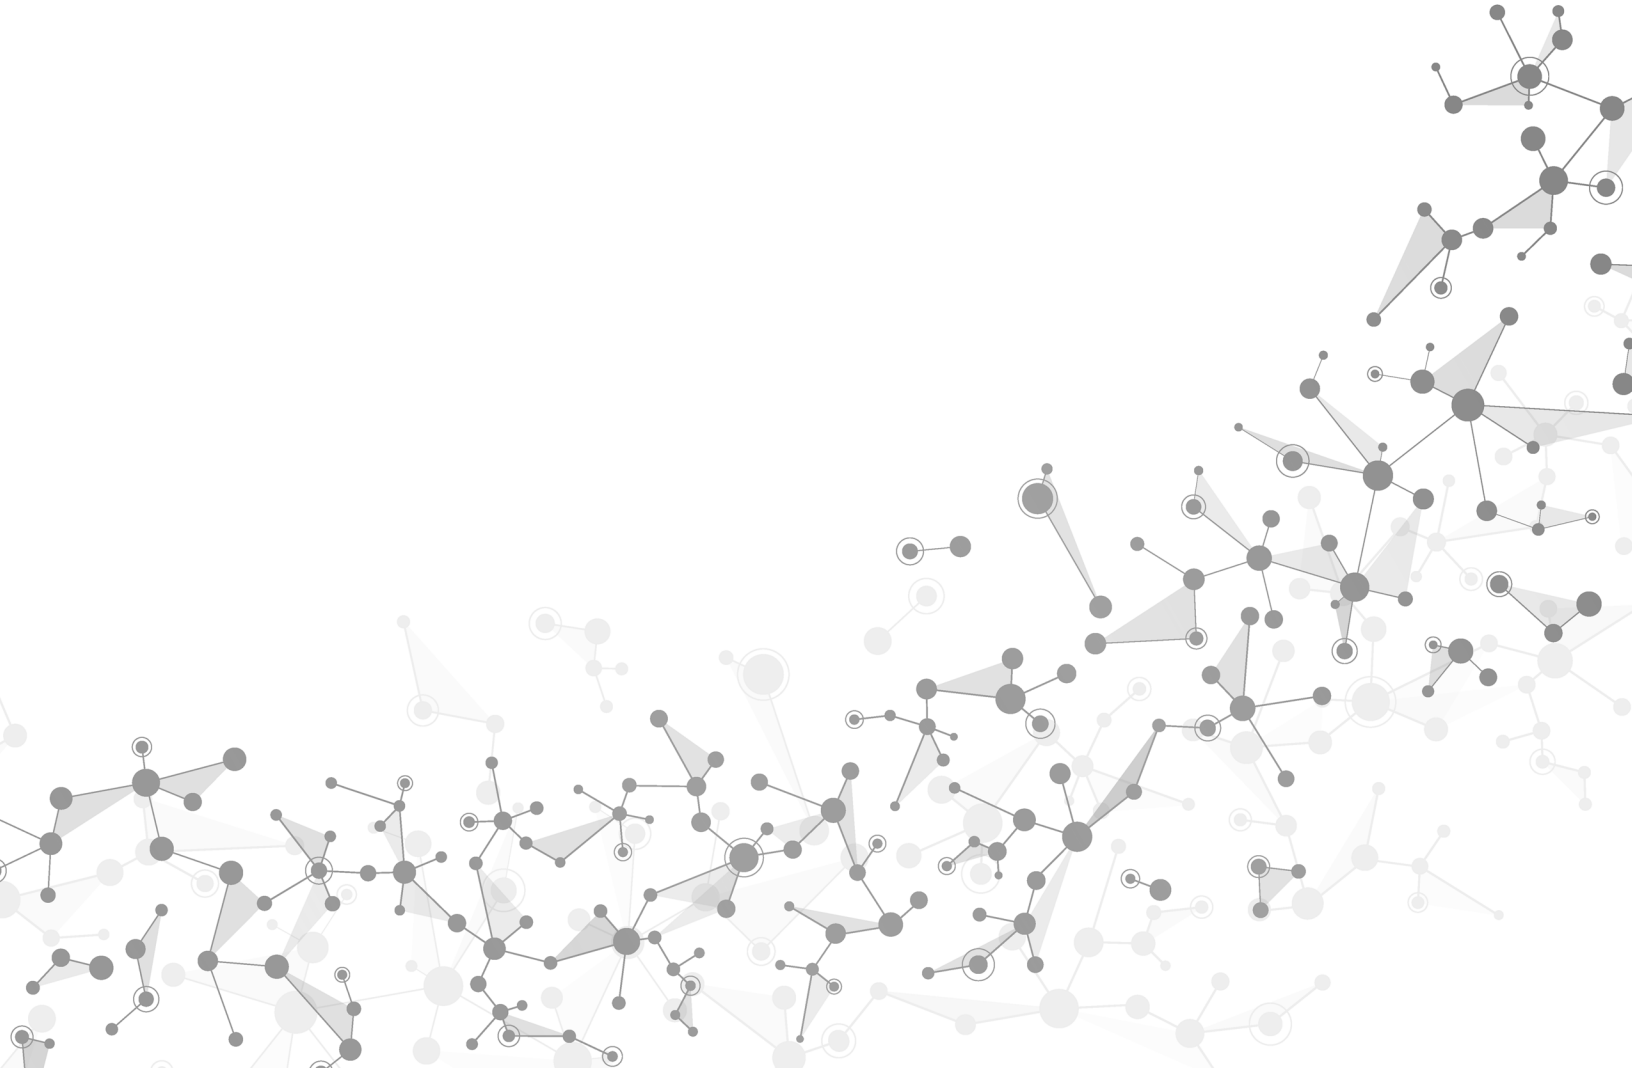

Chapter 7: File-level Backups Overview

This section provides details and requirements for file-level backups. A file-level backup protects an asset's file system and operating system. You can select to include or exclude files from file-level backups.

File-level backups protect physical assets. For virtual machine assets, you can choose host-level or file-level protection. Host-level backups capture files, application data, and virtual hardware. With file-level protection, the appliance treats the VM as a physical asset to run file-level and application backups. For more information on determining which backup type to use for a VM, see these topics:

- "Protecting Hyper-V virtual machines with file-level backups" on page 661
- "Protecting VMware virtual machines with file-level backups" on page 674
- "Protecting AHV virtual machines with file-level backups" on page 698
- "Protecting XenServer VMs with file-level backups" on page 691

File-level protection requires installing a Unitrends agent on the asset. Agent installation procedures vary by operating system. For details on installing the agent on various operating systems, see "Installing the Unitrends agent" on page 280. After installing the required agent, add the asset to the appliance as described in "To add an agent-based asset" on page 289, then proceed to "Backup Administration and Procedures" on page 425 to set up backup jobs.

Note: iSeries assets are not protected using an agent. For details on protecting iSeries, see "iSeries Backups Overview and Procedures" on page 767 instead.

Requirements and considerations for file-level backups

The following are required to run file-level backups:

- You must install a Unitrends agent on the asset. Agent installation procedures vary by operating system. For instructions, see "Installing the Unitrends agent" on page 280.
- You must add the asset to the backup appliance. After installing the required agent, add the asset as described in "To add an agent-based asset" on page 289.

Additional requirements may apply. Review the following topics for details:

- "Suspend, hibernate, and sleep modes" on page 704
- "Maximum file pathname lengths" on page 704
- "Default exclusions from file-level backups of Windows servers" on page 704
- "Default exclusions from file-level backups of Linux servers " on page 705
- "Additional Windows limitations" on page 705
- "Linux maximum UID and GID for 64-bit assets" on page 707

- "Exclude active databases from file-level backups" on page 707
- "Mac OS X SIP directories" on page 708

Once you've installed the agent, added the asset to the appliance, and reviewed the applicable additional requirements, proceed to "Backup Administration and Procedures" on page 425 to set up backup jobs.

Suspend, hibernate, and sleep modes

Protected assets must be awake for the full duration of the backup job.

Maximum file pathname lengths

Some Unitrends agents have a maximum file pathname size limitation. The backup does not include file pathnames that exceed this limit. The following table lists the agents affected by this restriction and the supported maximum file pathname lengths.

Unitrends agent	Maximum file pathname length
Windows	32 KB
Linux	4 KB
Solaris	1 KB
Mac OS X	1 KB

Default exclusions from file-level backups of Windows servers

By default, file-level backups of Windows severs exclude certain directories and files. These exclusions are in addition to any exclusions you have applied to the Windows server's backups.

File-level backups of Windows servers exclude:

- Any mapped network drives
- */RECYCLER*
- */$Recycle.Bin*
- *%TMP%*
- *%TEMP%*
- **.tmp*
- **.temp*
- *%AllUsersProfile%\Microsoft\Network\Downloader\Cache*
- *%WINDIR%\System32\Config*

- *%WINDIR%\System32\Catroot2*
- *%WINDIR%\win386.swp*
- Contents of the server's *DataDirectory* as specified by the registry key *HKLM\Software\Microsoft\Windows*
- Contents of the server's *DefaultDataDirectory* as specified by the registry key *HKLM\Software\Microsoft\Windows*
- Files specified by the registry key *HKLM\System\CurrentControlSet\Control\BackupRestore\FilesNotToBackup*
- Additionally, the following profile directories specified by the registry key *HKLM\Software\Microsoft\WindowsNT\CurrentVersion\ProfileList*\ProfileImagePath* are also excluded:
  - *\AppData\Local\Temp*
  - *\Local Settings\Temp*
  - *\Local Settings\Temporary Internet Files*

Default exclusions from file-level backups of Linux servers

By default,file-level backups of Linux exclude certain directories and files. These exclusions are in addition to any exclusions you have applied to the Linux client's backups.

> **Note:** If you need to include any of the system-excluded directories in your environment, see How to adjust the Linux default exclusion list.

Default Linux directories excluded from backup:

- Any network mounts
- */proc*
- */sys*
- */var/tmp*
- */home/*/.gvfs*
- */var/lib/nfs*
- */rpc_pipefs*
- */lib/modules/*/volatile/**
- */usr/bp/incremental_forever*

Additional Windows limitations

The following table describes additional limitations that may apply to your Windows environment.

OS or feature	Limitation
Windows Server 2022, 2019, and 2016	The following Windows Server 2022, 2019, and 2016 features are not yet supported: • Nano Servers (this deployment option does not install VSS components, which are required for agent backups) • Windows Containers, such as Kubernetes clusters • Shielded VMs • Storage Spaces Direct • Hyper-Converged Infrastructure (HCI) deployments
File server clusters	For clustered configurations: • Each node in the cluster must be running agent version 10.6.9 or later. • Be sure to install the same agent version on all nodes in the cluster.
Multi-point Services Role and RDS User Profile Disks	The Multi-point Services Role and RDS User Profile Disks are not supported. (You can disable the Multi-point Services Role to run agent backups.)
Windows Storage Server 2008	Single Instance Storage (SIS) is not supported and must be disabled to run agent backups.
BitLocker encrypted volumes	Backup and recovery of BitLocker encrypted volumes is supported only if the Windows LocalSystem account can access the volume. After each reboot of the Windows server, the Windows Administrator may need to manually re-enable access. If the Unitrends Windows agent cannot access the volume during backup or recovery, the job skips the volume and either fails or runs with a warning. Check the backup details for the Windows API error and refer to Microsoft's documentation for the steps needed to correct the issue.
Continuous Availability File Shares (CAFS) and Cluster Shared Volumes (CSVs)	In some environments, CSVs must be excluded from backups of Windows 2012 R2 and 2016 servers that use the CAFS feature. If CSVs are included and the backup fails with the following error, use the steps in " To protect a Windows server and its CSVs if backups fail with VSS error": *(VSS_E_VOLUME_NOT_SUPPORTED_BY_PROVIDER)* *The given shadow copy provider does not support shadow copying the specified volume.* **To protect a Windows server and its CSVs if backups fail with VSS error** 1 Add the Windows server to the appliance as described in "To add an agent-based

OS or feature	Limitation
	asset" on page 289. 2 Create a backup schedule for the Windows server as described in "To create a file-level backup job" on page 437. In the schedule: • Exclude all CSVs. • Include the system state. 3 Create an alias for each CSV as described in "Creating aliases for agent-based assets" on page 291. 4 Add each CSV to the appliance as described in "To add an agent-based asset" on page 289. 5 Create a separate backup schedule for each CSV as described in "To create a file-level backup job" on page 437. In each schedule: • Include only the CSV volume. • Exclude the system state.

Linux maximum UID and GID for 64-bit assets

On Linux assets, each file has a user ID (UID) and group ID (GID) that indicates the user and group that owns the file. On 64-bit Linux assets only, the Unitrends agent enforces a maximum value of 2097151 for these IDs. While backing up a 64-bit asset, the agent resets any UID or GID that is greater than 2097151 back to 0. Resetting these IDs does not prevent backup or recovery of the files.

After recovering a file whose UID or GID exceeded the maximum limit, an administrator can modify the user and/or group ID to its original value.

To prevent this condition on a 64-bit Linux asset, you can configure the maximum UID and GID values by setting the UID_MAX and GID_MAX values to 2097151 in this file: */etc/login.defs*.

Exclude active databases from file-level backups

Unitrends recommends excluding active databases from file-level backups. Run application-level backups to protect active databases. For information on application backups, see "Application Backups Overview" on page 733.

Note: SQL files for system databases (such as master, model, and msdb) are always included to support Windows replicas. Do not exclude these if you will be using Windows replicas for hosted SQL databases.

Some active databases are automatically excluded with file-level backups, as described here:

- Exchange — All transaction log files (.LOG files), the Exchange database (.EDB files), and streaming content files (.STM files) are excluded. See "Automatic exclusion of application data during file-level backups" on page 735.

- SQL — The following extensions are excluded from SQL user databases if the SQL VSS component is running on the Windows asset: *.mdf*, *.ldf*, and *.ndf*. Files in SQL *database/log* directories are excluded. See "Automatic exclusion of SQL data during file-level backups" on page 748.

Mac OS X SIP directories

Starting with Mac 10.11, Apple introduced a new security feature called System Integrity Protection (SIP). With SIP, the following directories are write-protected: */System*, */bin*, */sbin*, */usr* (except for */usr/local*), and any default applications in */Applications*. Only Apple-signed packages can write to these protected directories.

Unitrends handles these protected directories as follows:

- Backups include the contents of these directories (unless you have explicitly excluded them).
- File-level recovery operations recover all data other than the contents of these SIP directories. If you need to recover SIP-protected files, perform a bare metal recovery instead (as described in the Bare Metal Protection and Recovery Guide).

Chapter 8: Windows Image-level Backups Overview

Windows image-level protection employs block acceleration to deliver lightning-fast backup performance for Windows servers, particularly those with millions of files. Because Windows assets are backed up at the disk and volume level, image-level protection yields faster backup performance than file-level protection. With this increased backup speed, you can snapshot more frequently to reduce data loss in the event of an outage.

Note: Image-level backups are not supported for virtual machines in these environments: Azure and Amazon Web Services (AWS). Instead, use agent-based file-level and application backups.

See these topics to get started with image-level protection:

Features of Windows image-level and file-level protection

Features of image-level and file-level protection are given here:

Feature	Image-level backup	File-level backup
Data protected	The following apply to image-level protection: • Assets are protected and the disk and volume level. • Backups include the 'in use' regions of the Windows disk or volume only. Deleted regions are not included. • Provide application consistent backup and recovery for NTFS and ReFS filesystems. • Provide crash consistent backup and recovery for FAT, FAT32, and exFAT filesystems.	The following apply to file-level protection: • Assets are protected at the file system and operating system level. (Backups must include system state to protect the OS and for disaster recovery of the entire asset.) • Provide application and operating system consistent backup and recovery. • Protect configurations that cannot be protected by

Feature	Image-level backup	File-level backup
		image-level backups, such as cluster shared volumes (CSVs) and VHDS clusters. • Protect Azure and AWS virtual machines. (Image-level backups are not supported for these VMs.)
Backup options	Image-level backups support these options: • Opt to include or exclude volumes from backup. • Opt to run commands on the asset before or after the backup job. • Opt to index the backup so you can search an asset's backups by filename to recover individual files. This feature is available in release 10.4.8 and higher. This option is set by asset (see "To edit an agent-based asset" on page 293 for details).	File-level backups support these options: • Opt to include or exclude files, folders, or volumes from backup. • Opt to exclude system state. • Opt to run commands on the asset before or after the backup job. • Recommended where more granular exclusion of data is required.
Backup and recovery job performance	Image-level jobs yield faster backup and recovery: • Backup and recovery jobs use multiple parallel streams for fast performance. (Resource availability on the Windows asset impacts the number of streams used.) • Backup performance is not impacted by the size or number of individual protected files.	File-level jobs yield slower backup and recovery: • Backup and recovery jobs use a single stream. • Backup performance is impacted by the size or number of individual protected files. • Switch to image-level backups to avoid creating aliases for large assets or assets with many small files.

Feature	Image-level backup	File-level backup
Backup size	Image-level full backups may be larger than file-level fulls of the same asset. Smaller deduplication ratio than with file-level backups. (Fewer duplicate blocks found, more unique blocks stored.) This may result in decreased on-appliance retention.	File-level full backups may be smaller than image-level fulls of the same asset. Greater deduplication ratio than with image-level backups. (More duplicate blocks found, fewer unique blocks stored.) This may result in greater on-appliance retention.
Recovery	These recovery procedures are supported for image-level backups: • Recover individual files by browsing the contents of a backup. • Opt to enable **Index Image-Level Backups** on the Edit Asset page so you can search the asset's backups by filename to recover individual files. This feature is available in release 10.4.8 and higher. For details, see "To edit an agent-based asset" on page 293. **Notes:** • The index feature is not supported for recovery of ReFS filesystems. Recover by browsing the backup instead. • Filename search of indexed Windows image-level backups is not supported for recovery from imported backup copies. Recover by browsing the imported copy instead. • Assets with high-frequency backups or with very large file counts can add considerable load to the appliance. Consider appliance load when enabling the index option for these types of assets. • To index the backup, the job creates and mounts an object. If a file recovery object is already mounted for the asset,	These recovery procedures are supported for file-level backups: • Recover an entire backup. This restores the asset to the point-in-time at which the backup was taken. • Recover individual files by searching by filename or browsing the contents of a backup. • Use Windows file-level replicas to quickly spin up a virtual replica of a failed Windows asset. **Notes:** • Only one Windows replica can exist per Windows asset. You cannot run both an image-level replica and a file-level replica of the same asset at the same time. If a replica exists, you must tear it down before creating another for the asset. • You can opt to run

Feature	Image-level backup	File-level backup
	the backup runs but no index is created (as only one object per asset can be mounted at any given time). The resulting backup completes in warning status, with a message indicating that no index was created. • Use Windows image-level replicas to quickly spin up a virtual replica of a failed Windows asset on your VMware or Hyper-V host. **Notes:** • Only one Windows replica can exist per Windows asset. You cannot run both an image-level replica and a file-level replica of the same asset at the same time. If a replica exists, you must tear it down before creating another for the asset. • You can opt to run image-level replicas in the Unitrends Cloud. Contact your Account Manager for assistance. • Use instant recovery for rapid disaster recovery of a failed asset. • Use bare metal recovery to recover a failed asset to identical physical hardware or to a virtual machine. For details, see "Recovering Windows Image-level Backups" on page 1031.	file-level replicas in the Unitrends Cloud. Contact your Account Manager for assistance. • Use bare metal recovery to recover a failed asset to identical physical hardware, to dissimilar physical hardware, or to a virtual machine. For details, see "Recovering File-level Backups" on page 925.
Hosted applications	Exchange and SQL applications are included in image-level backups. Image-level backups can be taken with VSS copy-only snapshots (which do NOT truncate application logs) or with VSS full snapshots (which do truncate application logs by default). Image-level backups provide application protection but fewer recovery options than Unitrends application backups.	Active Exchange databases and active SQL user databases are automatically excluded from file-level backups. (SQL system databases are always included to support Windows replicas.) File-level backups are taken with VSS full snapshots. Full file-level backups do NOT truncate application logs.

Feature	Image-level backup	File-level backup
	Note: Oracle on Windows – You must use Windows file-level backups to protect the Oracle server and Oracle application backups to protect hosted applications. Windows image-level backups cannot be used for Oracle. Use one of these strategies to protect an Exchange or SQL server with image-level backups: **Note:** Additional recovery considerations apply for SQL clustered instances, SQL availability groups, and Exchange DAG nodes. For details, see "Considerations for recovering SQL clusters, SQL availability groups, and Exchange DAGs" on page 1032. • Run image-level backups using VSS copy-only snapshots, along with Unitrends application backups. Note the following: – This strategy provides additional options for application recovery. For example, you can recover a database, transaction log, storage group, or Exchange item from an application backup. (See "Application backups" below for details.) – Image-level backups use VSS copy-only snapshots by default. You do not need to configure any special setting to use copy-only snapshots. – Image-level backups that are run with VSS copy-only snapshots do not affect any application backup chain and do not truncate any application log files. – Logs are automatically truncated by these Unitrends application backups: Exchange fulls, Exchange incrementals, and SQL transaction log backups.	Run file-level backups along with Unitrends application backups to protect hosted applications. (See "Application backups" below for details.)

Feature	Image-level backup	File-level backup
	• Run image-level backups using VSS full snapshots to protect both the system files and hosted applications. Note the following: – Each image-level backup truncates any Exchange and SQL transaction logs for any online user databases on the server. **Note:** Image-level backups run using VSS full snapshots do NOT truncate logs for the following: offline databases, databases whose parent application instance is offline, and SQL system databases (master, model, and tempdb). – From an image-level backup, you can recover individual files, use replicas or instant recovery to quickly recover a failed or corrupted Windows machine, or use bare metal recovery to recover a failed asset to identical physical hardware or to a virtual machine. – VSS full snapshots are not used for image-level backups by default. You must configure this setting by checking these boxes in the Edit Asset dialog: **Show Image Level Backup Settings** and **Allow application aware**. For details, see "To edit an agent-based asset" on page 293. **Notes:** • For hosted Exchange and SQL applications, you cannot run both application backups and image-level backups that are configured with the **Allow application aware** setting. This would result in backup failures (because VSS full	

Feature	Image-level backup	File-level backup
	snapshots truncate all application logs). • There is no way to recover SQL databases or Exchange instances separately. You can recover Exchange and SQL files from image backups, but you cannot easily recover a database or Exchange instance after data loss or corruption. • Once an asset is configured with the **Allow application aware** setting: – Any existing SQL and Exchange schedules are disabled. – Windows *file-level* replicas of the asset cannot include SQL or Exchange applications. An error displays if you attempt to configure an application when adding or editing the replica. – Any existing *file-level* replicas that include SQL or Exchange should be manually deleted as their application data will become stale over time (since there will be no future application backups being restored to them). You can then recreate the replica without SQL and Exchange. • Run image-level backups using VSS copy-only snapshots. Note the following: – Application logs are NOT truncated. The database administrator must manually truncate the application logs. – From an image-level backup, you can	

Feature	Image-level backup	File-level backup
	recover individual files or use instant recovery to quickly recover a failed or corrupted Windows machine, or use bare metal recovery to recover a failed asset to identical physical hardware or to a virtual machine. – There is no way to recover SQL databases or Exchange instances separately. You can recover Exchange and SQL files from image backups, but you cannot easily recover a database or Exchange instance after data loss or corruption.	
Application backups	Unitrends application backups provide these benefits: • SQL, Exchange, Oracle, and SharePoint backups perform application-level post backup processing, such as log truncation. • Support all SQL database recovery models. Must run asset-level application backups for all recovery models other than *simple*. • Support backup of multi-node SharePoint farms. For more on application protection, see "Application Backups Overview" on page 733.	
SLA policies	SLA policies supported for image-level backups.	SLA policies are supported for file-level backups.
Copy Data Management	Image-level backups and image-level replicas that reside on a VMware or Hyper-V host can be used with the Copy Data Management feature.	File-level backups can be used to create Windows replicas. Windows replicas that reside on a VMware or Hyper-V host can be used with the Copy Data Management feature.

Best practices and considerations for Windows image-level protection

Follow these best practices to protect your Windows assets with image-level backups:

- Adhere to Microsoft best practices.
- Full and incremental backups are supported for image-level backups.
- Opt to index backups so you can search the asset's backups by filename to quickly recover individual files. See "Recovering from an indexed image-level backup by using Search Files" on page 1033 for details.

- A new full backup is required if the disk or volume configuration has changed since the last backup. This includes any change to the number, size, or properties of the disks or to the number, size, or properties of the volumes on a disk.

  If the disk or volume configuration has changed since the last backup, the appliance promotes the next scheduled backup to a full (or displays a message indicating a full is required if an on-demand incremental is attempted). Once a full backup succeeds, subsequent incrementals run as scheduled.

  Note: Protecting a given image-level asset by using more than one appliance is not recommended. Each time a backup is run on an one appliance, the integrity of the incremental backup chain on any other appliance is compromised. To correct a compromised backup chain, the appliance automatically promotes the next incremental to a full. Because of this, running incrementals on multiple machines is likely to result in many full backups.

- In-place Windows operating system upgrades disable the Volume CBT driver that is used to run incremental backups. To run incremental backups after performing an in-place OS upgrade, you must uninstall then re-install the Volume CBT driver:
  - Use Windows **Control Panel > Programs > Programs and Features** to uninstall Unitrends Volume CBT Driver:

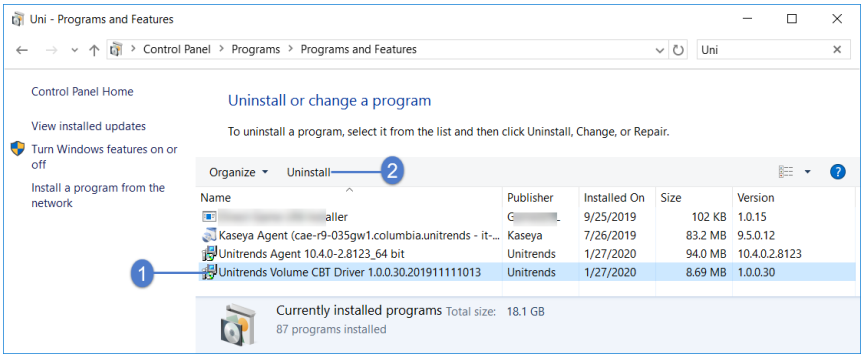

  - To re-install, run the MSI installer (*C:\PCBP\Installers\uvcbt.msi*). After installing the driver, you must enable it by rebooting the Windows asset.
  - After re-installing the Volume CBT driver, the next incremental is promoted to a full backup.

Considerations for protecting hosted applications with Windows image-level backups

Follow these best practices to protect your hosted applications with image-level backups:

- Adhere to Microsoft best practices.
- Full and incremental backups are supported for image-level backups.
- Do not use Windows image-level backups for Oracle. You must use Windows file-level backups to protect the Oracle server and Oracle application backups to protect hosted applications.
- Additional recovery considerations apply for SQL clusters, SQL availability groups, and Exchange DAGs. For details, see "Considerations for recovering SQL clusters, SQL availability groups, and Exchange DAGs" on page 1032.

- A new full backup is required if the disk or volume configuration has changed since the last backup. This includes any change to the number, size, or properties of the disks or to the number, size, or properties of the volumes on a disk.

  If the disk or volume configuration has changed since the last backup, the appliance promotes the next scheduled backup to a full (or displays a message indicating a full is required if an on-demand incremental is attempted). Once a full backup succeeds, subsequent incrementals run as scheduled.

- A new full backup is required if the disk or volume configuration has changed since the last backup. This includes any change to the number, size, or properties of the disks or to the number, size, or properties of the volumes on a disk.

  If the disk or volume configuration has changed since the last backup, the appliance promotes the next scheduled backup to a full (or displays a message indicating a full is required if an on-demand incremental is attempted). Once a full backup succeeds, subsequent incrementals run as scheduled.

  Note: Protecting a given image-level asset by using more than one appliance is not recommended. Each time a backup is run on an one appliance, the integrity of the incremental backup chain on any other appliance is compromised. To correct a compromised backup chain, the appliance automatically promotes the next incremental to a full. Because of this, running incrementals on multiple machines is likely to result in many full backups.

Requirements for Windows image-level protection

The following requirements must be met for image-level protection of Windows assets:

Item	Description
Unitrends appliance	These requirements apply to the Unitrends backup appliance: • The appliance must be running release 10.3 or higher (10.5.1 or higher to use the Windows image-level replicas feature). • Port 443 must be open inbound to the appliance from the protected Windows asset for the TCP protocol.
Windows agent	The Windows asset must be running Unitrends agent version 10.3 or higher with the Volume CBT driver. (Agent version 10.5.1 or higher is required to use the Windows image-level replicas feature). During agent installation, you have the option to install the Volume CBT driver. The Volume CBT driver is needed to run image-level incremental backups. After you install the Windows agent and Volume CBT driver, you must reboot the Windows asset to enable the driver. For details, see "Installing the Windows agent" on page 362. **Notes:** • If the Volume CBT driver has not been installed or has not been enabled, image-level incrementals are not supported. Any scheduled incremental is automatically promoted to a full backup. If you attempt to run an on-demand

Item	Description
	incremental, you receive a message indicating that only full backups are supported. • There is a known issue where the Hyper-V CBT driver cannot be installed on Hyper-V servers that are running a pre-10.1.0-3 agent. In this case, you must manually uninstall the older Windows agent before installing the latest agent. For details, see "To uninstall the Windows agent" on page 382.
Windows asset	See these rows below for Windows requirements: **Note:** Additional Windows requirements apply for instant recovery. For details, see "Windows asset requirements for IR" on page 1065. • "Operating systems" • "Firmware interface type" • "Disk configuration" • "Disk partition type" • "File system configuration" • "Volume configuration" • "Multi-point Services Role and RDS User Profile Disks" • "Domain controllers" • "Unsupported Windows features"
Operating systems	The operating systems listed below are supported. (Additional version limitations apply. See the Compatibility and Interoperability Matrix for details.) Supported client operating systems: • Windows 11, 64-bit only • Windows 10, 64-bit only • Windows 8.1, 64-bit only • Windows 8, 64-bit only • Windows 7 with SP1, 64-bit only Supported server operating systems: • Windows 2022, 64-bit only

Item	Description
	• Windows 2019, 64-bit only • Windows 2016, 64-bit only • Windows 2012 R2, 64-bit only • Windows 2012, 64-bit only • Windows 2008 R2 with SP1, 64-bit only **Notes:** These additional requirements apply to Windows 2008 R2 SP1: • These Windows security updates must be installed: Update for Windows Server 2008 R2 x64 Edition (KB2533623) and Security Update for Windows 7 for x64-based Systems (KB3033929). (If these updates have not been installed, you are prompted to install them during agent installation.) • The Unitrends Volume CBT driver (used to run image-level incremental backups) cannot be installed along with the Unitrends Windows agent. You must install it manually. During agent installation the Volume CBT installer is placed here: *C:\PCBP\Installers\uvcbt.msi*. To install the driver, simply run *uvcbt.msi*. After installing the driver, you must enable it by rebooting the Windows asset.
Firmware interface type	Image-level protection is supported for BIOS- and UEFI-based assets.
Disk configuration	Image-level protection is supported for Windows machines configured with basic disks only. Dynamic disks are not supported. Offline disks are included in image-level backups.
Disk partition type	Image-level protection is supported for GUID Partition Table (GPT) partitions and Master Boot Record (MBR) partitions.
File system configuration	Image-level protection is supported for these file systems: NTFS, FAT, FAT32, exFAT, and ReFS.

Item	Description
	Note: Due to a Microsoft limitation, VSS snapshots cannot be taken of these volumes: FAT, FAT32, and exFAT. Backups of these volumes may contain data that is not in a consistent state if data changes during the backup job.
Volume configuration	Image-level protection is not supported for read-only disks. You must exclude all volumes on read-only disks from the backup job or run file-level backups. Image-level backups fail if read-only volumes have not been excluded. (For details on excluding volumes, see step 4 in "To create an image-level backup job" on page 449.) **Note:** Removable media is automatically excluded from image-level backups. You do not need to exclude volumes on a read-only disk that resides on removable media. Image-level protection is not supported for VHD or VHDX files that are mounted as local volumes. For details, see "VHD or VHDX files that are mounted as local volumes".
Multi-point Services Role and RDS User Profile Disks	The Multi-point Services Role and RDS User Profile Disks are not supported.
Domain controllers	Image backups are supported for Active Directory (AD) when running a standalone, single domain controller instance only. Microsoft system state backups are also required for bare metal recovery into other environments, including Unitrends DRaaS. For all other configurations, use Unitrends file-level backups for proper AD application protection. **Notes:** • For more on protecting domain controllers with image-level backups, see Detailed Options for Protecting Domain Controller (DC) and and Restoring Active Directory (AD). • For more on running Windows system state backups, see this Microsoft article: AD Forest Recovery - Backing up the System State data.
Unsupported Windows features	Image-level protection is not supported for the following features. Use file-level protection instead: • Windows Storage Spaces • Cluster shared volumes (CSVs)

Item	Description
	• Windows Server Failover Clusters (WSFCs) • Distributed File System environments – Distributed File System (DFS) Namespaces and DFS Replication offer high-available access to geographically dispersed files. Because of the replication and syncing operations in DFS environments, you must set up the Windows machine in accordance with Microsoft best practices to ensure database consistency. If all Microsoft considerations are not addressed, backup and restore of the Windows machine may yield undesired results. If you prefer not to research these best practices, protect the machine by running file-level backups instead.

VHD or VHDX files that are mounted as local volumes

Image-level protection is not supported for VHD or VHDX files that are mounted as local volumes.

If you include one of these locally mounted volumes in an image-level backup, recovery from the backup yields undesired results.

Protecting unmounted VHD or VHDX files is supported. For example:

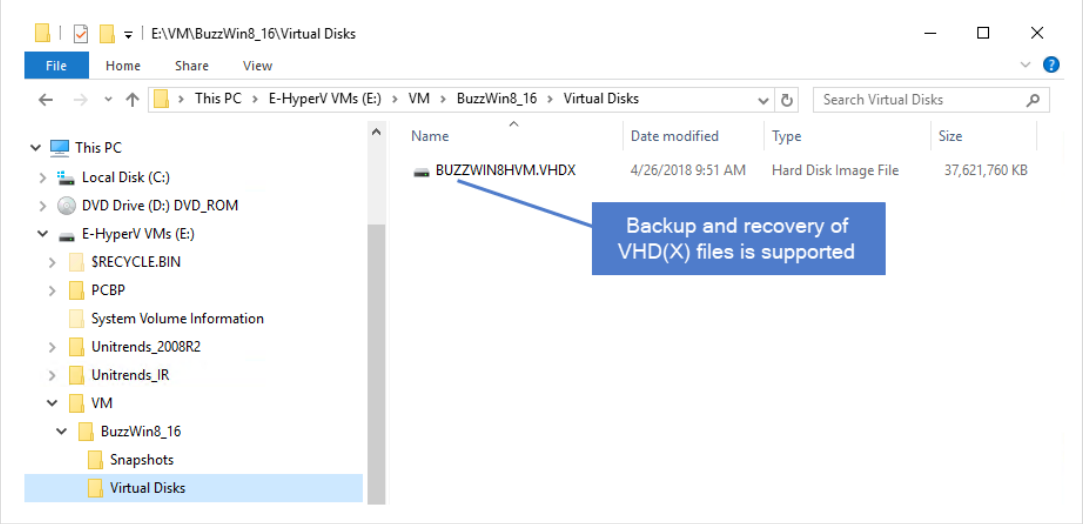

Protecting a volume that was created by mounting a local VHD(X) file is not supported. For example:

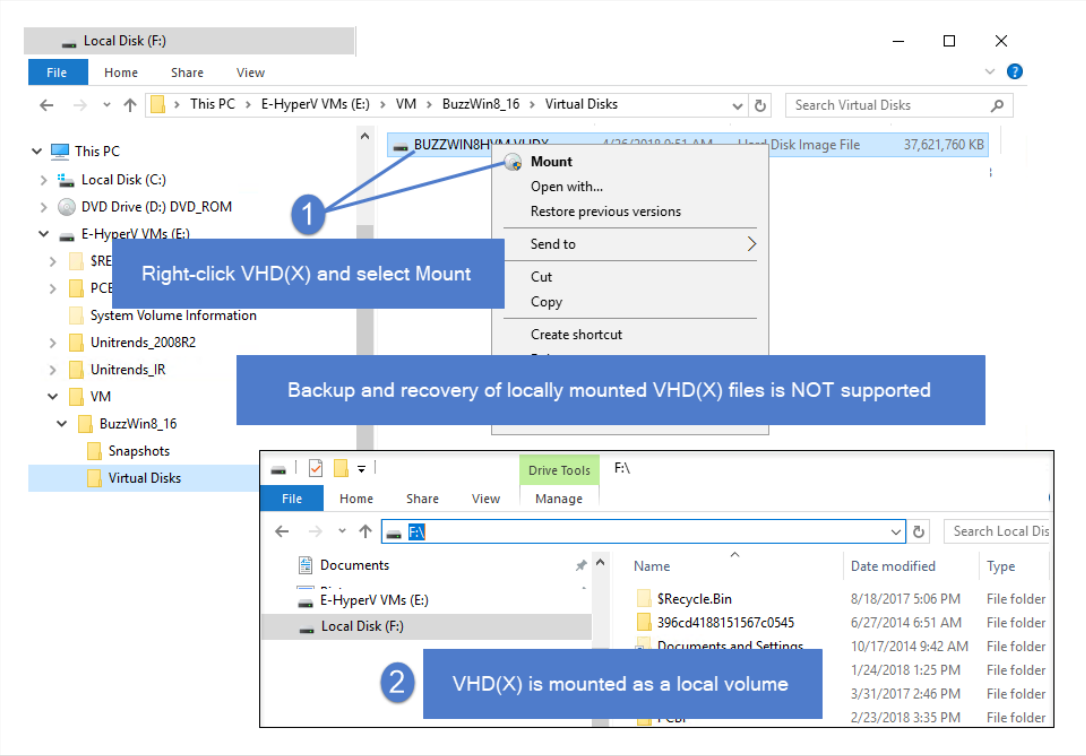
Local Disk (F:)
File
Home
Share
View
This PC
E-HyperV VMs (E:)
VM
BuzzWin8_16
Virtual Disks
Search Virtual Disks
Local Disk (C:)
DVD Drive (D:) DVD_ROM
System Volume Information
Unitrends_2008R2
Unitrends_IR
Snapshots
37,621,760 KB
Mount
Open with...
Restore previous versions
Send to
Cut
Copy
Create shortcut
1
Right-click VHD(X) and select Mount
Backup and recovery of locally mounted VHD(X) files is NOT supported
Drive Tools
F:\
Manage
Search Local Dis
Documents
Name
Date modified
Type
$Recycle.Bin
8/18/2017 5:06 PM
File folder
396cd4188151567c0545
6/27/2014 6:51 AM
File folder
10/17/2014 9:42 AM
File folder
1/24/2018 1:25 PM
File folder
3/31/2017 2:46 PM
File folder
2/23/2018 3:35 PM
File folder
2
VHD(X) is mounted as a local volume

Chapter 9: NAS Backups Overview

This section provides details and requirements for protecting the data stored on Network Attached Storage (NAS). Review this information to determine the best strategy for your environment and to ensure requirements have been met before you start protecting your NAS data.

Unitrends uses the following protocols to protect data stored on NAS devices:

- Common Internet File System (CIFS)
- Network File System (NFS)
- Network Data Management Protocol (NDMP)

To protect this data, add the NAS share or NDMP device to the backup appliance as a protected asset. The NAS is then backed up through the network connection. For NAS shares, data is backed up just like any other internal directory or volume. Data transfers more quickly than if you simply mount the share on another protected asset.

Requirements and considerations for NAS protection vary depending on the protocol you are using. See the following for details:

- "Determining which NAS protocol to use" on page 724 to determine which approach best suits your needs.
- "NAS protection using CIFS/NFS" on page 726 for additional CIFS and NFS requirements and considerations.
- "NAS protection using NDMP" on page 727 for additional NDMP requirements and considerations.
- "Start protecting the NAS asset" on page 732 for next steps.

Determining which NAS protocol to use

There are benefits to both approaches Unitrends offers for protecting a NAS. The recommended approach for you depends on your business requirements. Use the following comparison to determine how to protect your NAS:

Function	NDMP	CIFS and NFS
Backup	Features of NDMP backups: • Application backups. Protected at the volume level. Each volume is protected in a separate backup job. Captures Access Control Lists (ACL) and other file attributes. • Full, Differential, and Incremental backup modes. Automatically promotes every 10th incremental to a differential. • Shorter backup windows, especially if protecting many small files.	Features of CIFS and NFS backups: • File-level backups. Protected at the NAS share level. • Full, incremental, differential, and selective backup modes.

Function	NDMP	CIFS and NFS
Recover	Features of NDMP recovery: • Recover to NDMP devices of the same vendor. See vendor documentation for additional compatibility limitations. • Point-in-time recovery of the entire backup group is supported. • Recovering individual files from a backup is supported for some filers.	Features of CIFS and NFS recovery: • Recover to the same CIFS or NFS device or to an alternate CIFS or NFS device. • Point-in-time recovery of the entire backup group is supported. • Recovery of selected files is supported.
Hot backup copy	Configure volumes on the NDMP device for backup copy to the Unitrends Cloud or to another Unitrends appliance.	Configure the CIFS or NFS asset for backup copy to the Unitrends Cloud or to another Unitrends appliance. Better deduplication and backup copy performance. Longer retention possible because of smaller backup copy footprints.
Cold backup copy	Backup copy at the file-level or by volume.	Backup copy at the file-level.

NAS protection using CIFS/NFS

The following table describes features and limitations to consider when planning your NAS CIFS/NFS protection strategy.

Feature	Description
Backup	The following apply to NAS CIFS/NFS backups: • The NAS is protected at the share level. Backups start at the NAS mount point and do not include files in other system directories. You specify the desired mount point when adding the NAS asset to the backup appliance. If you want more granular control: – For a given NAS share, add separate mount points to the appliance, each as a separate asset. Create jobs for each asset you add. – When creating jobs, select folders and/or files to include or exclude from the backup. (Wildcards are not supported for inclusion lists.) • Open files are not included in the backup. Be sure to schedule jobs to run when file activity is at its lowest level. • Permissions of the files as seen when mapped to the backup appliance are not exactly the same as those on the NAS share. • If the NAS share is configured for authentication, you must supply credentials to access the specified mount point. If in your environment you only have credentials to access a parent directory, enter the full path to the parent directory and specify desired folders and files to include in the backup. • The NFSv4 protocol is not supported.
Recovery	The following apply to NAS CIFS/NFS recovery: • Point-in-time recovery of the entire backup group is supported. • Recovery of select files is supported. • You can recover to the original location, to another location on the original NAS, or to another NAS CIFS/NFS asset.
Nutanix AHV shares	You can protect Nutanix AHV shares by running Unitrends NAS backups over the NFS protocol . (Due to a Nutanix limitation, Unitrends NAS backups are not supported for Nutanix SMB shares.) To protect a Nutanix AHV share over the NFS protocol, you must first enable the *multi-protocol management access for NFS* option. To configure this setting: 1 In the Nutanix Prism interface, select the share.

Feature	Description
	2 Click **Update** and check the **multi-protocol management access for NFS** box. **Note:** If the box is grayed out, go to the **protocol management** option and select **user mapping**. Complete the mapping steps to enable the checkbox.
Backup Copy	Backup copy to an off-site target is supported.

NAS protection using NDMP

The following table describes features and limitations to consider when planning your NAS NDMP protection strategy.

Feature	Description
Appliance requirements	The Unitrends appliance must meet the following requirements to protect NAS devices using the NDMP protocol: • Must be running Unitrends release 9.0.0-13 or higher. • Must be licensed for the NDMP feature. Check the appliance license string for NDMP=X, where *X* equals the number of NDMP licenses purchased. To view the license string, select **Configure > Appliances > Edit > License**.
NDMP requirements	The following NDMP requirements apply: • Unitrends protects NDMP version 4.0. • Unitrends currently certifies devices from NetApp and EMC (Celerra, VNX, and VNXe). Devices from other vendors can be added as "Generic" NDMP NAS assets. Consider vendor specific limitations when protecting generic assets. • It is important to be familiar with your vendor's documentation and limitations because they can affect Unitrends protection of your NDMP device. • NDMP password: – Your NDMP NAS device must be configured with an MD5 password. Clear text passwords are not supported. – For VNXe devices only, the password cannot contain the following characters: **&** and *****.
Network	The following network requirements apply:

Feature	Description
requirements	• Unitrends uses a single, customer-specified IP address when protecting an NDMP asset. NDMP operations to and from multiple isolated IP networks are not supported. • These ports must be open for bi-directional traffic: – Port range 32768 - 61000 - Unitrends dynamically assigns ports in this range when protecting NDMP devices. If your environment is configured with a firewall, make sure the ports in this range are open. – Port 10000 - Unitrends appliances use this control port when protecting NDMP. Port 10000 is open for the following security levels: None, Low, and Medium. You cannot protect NDMP devices if you set your Unitrends appliance to High security. For details, see "To view or edit port security settings" on page 113.
Backup and recovery	See these topics for backup and recovery requirements and considerations: • "All jobs" on page 728 • "Backup jobs" on page 728 • "NetApp cluster protection" on page 729 • "Advanced configuration settings" on page 729 • "Recovery jobs" on page 730
All jobs	The following apply to all NDMP jobs: • Because NDMP NAS devices normally have a limited number of NDMP connections, backup and recovery jobs for NDMP assets are queued and run as NDMP connections become available. • Non-UTF-8 compatible characters cause backups to run more slowly. If your NAS share contains non-UTF-8 compatible characters, it is recommended to convert the NAS share to support UTF-8.
Backup jobs	The following apply to NAS NDMP backups: • NAS NDMP assets are protected at the volume level. A separate backup runs for each volume. • A recurring full backup must be in the schedule. • NDMP only supports nine consecutive incremental backups between

Feature	Description
	successful fulls and differentials. Schedules with more than nine consecutive incremental backups result in automatically promoted differential backups. For details, see "Automatic promotions of NDMP Incremental backups" on page 731. • To protect NetApp high availability C-mode clusters, additional requirements apply. See "NetApp cluster protection" below for details. • Additional configuration may be required in your environment. See "Advanced configuration settings" below for details.
NetApp cluster protection	Additional configuration is needed to protect NetApp high availability C-mode clusters. Once you have configured your clusters for Unitrends protection, they can be backed up and recovered using the standard Unitrends NDMP procedures. The following requirements must be met to protect NetApp clusters: • NetApp ONTAP must be version 7.*x* or 8.*x*. • NDMP must be enabled for both the cluster and the Vserver. See the NetApp configuration documentation and Configuring NDMP NetApp clusters for backup for details. • Volumes to protect must be exported through an LIF. We recommend assigning a unique IP address to each volume you wish to backup. In NetApp cluster environments, volumes may migrate over to a different node. If your NDMP schedule has the **Auto-include new NDMP Volumes** box checked, migrated volumes are automatically included in the backup schedule. Generally, adding a migrated volume to the schedule causes the next backup to run as a full backup, even though a full of this of this migrated volume may exist on the appliance. If desired, you can prevent a new full backup by migrating the volume back to the original node or by unchecking the **Auto-include new NDMP Volumes** box in the backup schedule.
Advanced configuration settings	Because each NDMP vendor has different limitations, there are some advanced configuration settings that might be required to protect your NDMP device. To access the advanced configuration options, go to the **Configure > Appliances > Edit > Advanced > General Configuration** page, and scroll down to the NDMP section. The following advanced settings are available: • DAR - Unitrends uses Direct Access Recovery (DAR) to recover NDMP backups. DAR is on (*DAR=1*) by default. For NetApp devices, DAR only works with ONTAP version 8.0 and later. If using an earlier version of ONTAP, disable DAR by setting *DAR=0*. • IPv4 Address - Blank by default. Unitrends automatically attempts to use the

Feature	Description
	eth0 or *seth0* IPv4Address. If your environment is configured with either of these IP addresses, it is retrieved and used. If you do not have *eth0* or *seth0* configured in your environment, you must enter an IP address in this field and restart NDMP services as described in "To restart NDMP services" on page 731. (This is most common in the case of bonded NICs.) Entering an IP address in this field will override *eth0* or *seth0*. • Username - The NDMP daemon username defaults to *ndmp*. If you change this username, NDMP services must be restarted as described in "To restart NDMP services" on page 731. • Password - The NDMP daemon password defaults to *unitrendsndmp*. If you change this password, NDMP services must be restarted as described in "To restart NDMP services" on page 731. • Maximum Running NDMP Jobs (Only accessible from the terminal — Configuration Options section) - The maximum number of running NDMP sessions per NAS NDMP asset defaults to 2. For more information, see Maximum Running NDMP Jobs.
Recovery jobs	The following limitations apply to NAS NDMP recovery: • NDMP backups can only be recovered to NDMP devices of the same vendor. • Supported recovery targets vary by vendor. Recovery to the original location is supported for all vendors. See the vendor documentation to determine whether you can recover to another location on the original NDMP device, or to another NDMP device that has been added to the appliance as an asset. • Point-in-time recovery of the entire backup group is supported. • Recovery of selected files is supported for some NDMP devices from the certified vendors. See the vendor documentation for compatibility limitations. • When performing point-in-time recovery of an NDMP volume, you cannot specify files to include or exclude. The volume is recovered exactly as it was at the selected recovery point. • Recovery of selected files that contain non-UTF-8 compatible characters is not supported. Instead you must recover the entire backup.
Backup Copy	Backup copy to an off-site target is supported.

To restart NDMP services

NDMP services must be restarted on the Unitrends appliances if any of the following are changed: the IPv4 address, the NDMP daemon username, or the NDMP daemon password.

1. Using a terminal emulator, such as PuTTY, connect to the appliance using the following:
   - Appliance IP address
   - Port 22
   - SSH connection type
2. Log in as user *root*.
3. Enter the following command:

```
# service unitrends-ndmp restart
```

Automatic promotions of NDMP Incremental backups

NDMP limits the number of incrementals that can occur between fulls to 9. This limitation is enforced by assigning and tracking levels of each backup mode. It does so in the following way:

- Fulls are always counted as level *0*.
- Differentials are always counted as level *1*.
- Incrementals are counted by increasing the previous backup's level by 1. These can be counted as levels 1-9, with 9 being the maximum level allowed by the protocol.

Automatic promotion for schedules

For schedules, the NDMP level assignments described above result in incremental backups being automatically promoted to differentials if there is already a level-9 backup in that volume's current backup group. The promotion to a differential resets the level to 1. After the automatic promotion, the schedule resumes running the jobs as expected.

Note: Only 8 incrementals run between automatically promoted differentials because the count starts from 1 rather than 0 (as it does with full backups).

One-time incremental backups and automatic promotion

On-demand incremental backups also affect the backup level for the volume. If you attempt to run a one-time incremental backup and the backup level is less than 9, the job is queued and the backup level of the group is increased by one. However, if the volume's backup level is already 9, the job is not queued. Instead you are notified that you have reached the maximum limit of consecutive incremental backups for this volume and a full must be run.

Note: Only 8 incrementals run between automatically promoted differentials because the count starts from 1 rather than 0 (as it does with full backups).

Differential backups and automatic promotion

Because differential backups are always counted as level 1, they do not have the same limitations as incremental backups. Any number of differentials can be run between successful full backups of an NDMP volume.

Start protecting the NAS asset

After ensuring all requirements have been met, do the following to start protecting your NAS device:

Step 1: Add the NAS to the Unitrends appliance as described in "To add a NAS CIFS or NFS asset" on page 296 or "To add a NAS NDMP asset" on page 298.

Step 2: Run backup jobs as described in "To create a NAS CIFS or NFS backup job" on page 470 or "To create a NAS NDMP backup job" on page 475.

Chapter 10: Application Backups Overview

This section provides detailed information for protecting applications. Application backups only protect applications. To protect the application's host server, you must create separate file-level or image-level backups.

Once you have reviewed this information, install the required agent, add the application server to the Unitrends appliance, then proceed to "Backup Administration and Procedures" on page 425 to set up backup jobs. Considerations and requirements vary by application. See the following for details on the desired application:

Exchange backup requirements and considerations

Consider the following before implementing your Exchange protection strategy:

Exchange agent requirements

Exchange application backups are run using the Unitrends Windows agent. Before you install the agent, enure that the Exchange server is running the latest service packs and that these services are installed and running on the Exchange server:

- Microsoft Exchange VSS Writer. If the Exchange VSS Writer is not installed or is not running, an error message displays. The Exchange VSS Writer must be running to continue the backup operation.
- Microsoft VSS Service.

It is best practice to run the latest Unitrends appliance and agent software versions to protect your Exchange environment. Older versions do not support all current Unitrends features:

- To protect Exchange 2016, the appliance and Windows agent must be running release 9.0.0-13 or later.
- To protect Exchange 2013, the appliance and Windows agent must be running release 8.0.0-4 or later.

Supported Exchange environments

Unitrends protects the Exchange environments listed in the Unitrends Compatibility and Interoperability Matrix. For Microsoft requirements, see these Microsoft articles:

- Exchange 2003 system requirements
- Exchange 2007 system requirements
- Exchange 2010 system requirements
- Exchange 2013 system requirements
- Exchange 2016 planning and deployment

Recommended Exchange configurations

The following configurations are recommended for optimal protection and recovery:

Recommendation	Description
Disable circular logging	This enables you to run differential or incremental backups of Exchange. If you do not disable circular logging, only full backups are supported. See "Circular logging setting" on page 735 for more information.
Do not allow the physical or virtual machine hosting the Exchange server to be a domain controller	This enables much simpler and faster Exchange restores since you will not first have to restore Active Directory on the same server.
Make sure that the physical or virtual machine hosting the Exchange server is a member of a domain that has at least two domain controllers	This enables faster recovery. Active Directory information is replicated if there is more than one domain controller, which means that if one domain controller fails the other can be used to recover missing transactions after the failed domain controller is restored.
Separate transaction log files from the Exchange server database	Exchange performs much more efficiently if the Exchange database and transaction logs are placed on different physical storage devices. In addition, by separating these two important components, recovery of failed storage is eased.
Disable the write cache on any hard drive or RAID adapters being used in the system that is hosting the Exchange server	This prevents data corruption by ensuring that any Exchange write operation is committed to secondary storage (i.e., disk) correctly.

Exchange backup considerations and requirements

Consider the following when planning your Exchange protection strategy:

Automatic exclusion of application data during file-level backups

When you run file-level backups of the Windows server hosting Exchange, certain Exchange-related files are automatically excluded. For example, all transaction log files (i.e., *.LOG* files), the Exchange database (i.e., *.EDB* files), and streaming content files (i.e., *.STM* files) are excluded.

Using incremental backups

Exchange versions 2007 or higher can use incremental backups. Exchange incrementals offer the following benefits:

- Incrementals can run more quickly and frequently than differentials since they include only the changes since the last successful full or incremental backup. This enables you to meet more aggressive RPOs than with differentials, which contain all changes since the last full backup.
- Upon completion of a successful incremental, unneeded transaction log files are automatically truncated, freeing space on the Exchange server. Automatic log truncation does not occur with Exchange differentials.

When creating an Exchange job that includes incremental backups:

- The same schedule cannot contain differentials and incrementals.
- The schedule must contain a full backup. The Exchange job does not support the incremental forever strategy.

Circular logging setting

Circular logging is an Exchange feature that enables overwriting transaction log files. Unitrends recommends disabling circular logging. You must disable circular logging to run differentials or incrementals. If you enable circular logging, you can only run full backups. If you disable circular logging, the transaction logs are used to create differential or incremental backups.

- With differentials, these transaction logs accumulate until a successful full backup runs.
- With incrementals, unneeded logs are removed after each successful backup.

The removal of unneeded truncation logs is typically termed *transaction log truncation*. Transaction log truncation removes unneeded logs but does not reclaim space. Reclaiming space is a separate operation that must be performed periodically by the Exchange system administrator.

Microsoft snapshots

Unitrends leverages the Microsoft snapshot feature to protect Exchange. Our protection of Exchange with these snapshots is not supported for the following:

- Any type of NAS configuration (SAN configurations are supported).

- The Exchange 2003 Recovery Storage Group feature.

Protecting databases and storage groups

Unitrends protects databases and storage groups as follows:

- Databases - For Exchange 2016, 2013, and 2010, you can back up multiple databases or a single database. Backups protect locally deployed databases only. Remote databases, such as Office 365 or Hybrid deployments, cannot be protected by Unitrends backups.
- Storage groups- For Exchange 2007 and 2003, you can back up multiple storage groups or an individual storage group. You cannot back up individual databases within a storage group. The reason for this is the transaction logs for the entire storage group are backed up for each database selected. Thus a full backup must be run on every database in a storage group in order for the transaction logs to be properly handled for full/differential backups.

Protecting clustered Exchange environments

Unitrends supports protection of Exchange 2007 clusters in either CCR or SCR configurations and Exchange 2016, 2013, and 2010 in the DAG configuration.

Note: The Microsoft Exchange Replication Service must be running in order to protect the above cluster configurations.

Unitrends recommends adherence to the following best practices when protecting clustered Exchange environments:

- Add each Exchange server node you wish to protect to the Unitrends appliance using its native server IP address. Do not use the cluster hostname/IP address used to access the active Exchange server.
- Do not add the cluster hostname/IP as a separate asset.
- Do not backup multiple copies of the same database simultaneously. Each full backup of an active or passive database results in database log truncation. The truncation of logs is replicated to the other members of the cluster where that database exists. If the same database is undergoing a full backup on two nodes, the log truncation of each could interfere with the other. To avoid this, schedule backup of replicated databases at staggered times across the cluster.
- Backup only passive copies to reduce workload on the active server(s).

Start protecting Exchange

After ensuring all requirements have been met, do the following to start protecting your Exchange environment:

Step 1: Install the Windows agent on the Exchange server as described in "Installing the Windows agent" on page 362.

Step 2: Add the Exchange server to the Unitrends appliance as described in "To add an agent-based asset" on page 289.

Step 3: Run backup jobs as described in "To create an Exchange backup job" on page 478.

SQL backup requirements and considerations

Review the following before implementing your SQL protection strategy:

Supported SQL features

The following table describes the SQL features that Unitrends protects.

Feature	Description
SQL system and user databases	Unitrends protects SQL system databases and user databases. Unitrends best practices for backup and recovery vary by SQL database type and recovery model. See these topics for details: • "SQL recovery models" on page 746 • "SQL system databases" on page 747 • "Example SQL backup strategies" on page 747
Always Encrypted databases	The Always Encrypted feature encrypts SQL data at rest and in motion. Unitrends supports protection of SQL Always Encrypted databases.
Stretch databases	A Stretch database consists of a database that resides on the local SQL server and a paired database that resides in Microsoft Azure. Unitrends supports protection of Stretch databases with some limitations. For example, Unitrends backups capture the data on the local SQL server only (and do not include any data in the paired Azure database).
Cluster shared volumes	Unitrends supports protection of databases that reside on clustered shared volumes. Nodes in the Windows Server Failover Cluster must be running Windows 2012 or later.
SMB 3.0 shares	SQL Server 2012 and higher can host SQL instances with disk storage located on SMB 3.0 shares. Unitrends supports protection of databases that reside on SMB 3.0 shares.

Feature	Description
SQL clusters	In a SQL cluster, one or more SQL Server failover cluster instances are installed into a Windows Server Failover Cluster (WSFC). Unitrends protects a variety of cluster configurations, such as: • Always On Failover Cluster Instances • Always On availability groups • Resource group storage on cluster shared volumes (CSV) and SMB 3.0 shares

Requirements for SQL protection

The requirements for protecting your SQL databases vary based on the configuration of your SQL servers and the SQL features used in your environment. The "Agent requirements for Microsoft SQL" on page 738 and "SQL system requirements" on page 741 apply to all SQL protection. If you are protecting SQL clusters, availability groups, data on SMB 3.0 shares, Always Encrypted databases, or Stretch databases, additional requirements apply. See the following topics for details:

Agent requirements for Microsoft SQL

The Unitrends Windows agent is needed to protect hosted SQL databases. Before you install the agent, the following must be installed on the SQL server:

- The SQL Server VSS Writer, SQL Server Browser, and BP Agent services must be installed and running to perform backup and restore operations. If the SQL Server VSS Writer or SQL Server Browser services are not started when you install the Windows agent, the agent cannot detect the SQL instance.
  - The SQL Server VSS Writer must be started and set to automatic startup.
  - The SQL Server Browser must be started and set to automatic startup.
  - The BP Agent service is installed when the Windows agent is installed on the SQL server.
- The Volume Shadow Copy service must be installed and can be set to manual or automatic startup.

- The NT AUTHORITY\SYSTEM account must be configured as sysadmin. This account is used to perform SQL backup and recovery jobs.

  Note: Beginning in SQL Server 2012, SQL does not grant NT AUTHORITY\SYSTEM sysadmin privileges by default. For SQL Server 2012 and later versions, you must manually add NT AUTHORITY\SYSTEM as a system administrator. For details, see the Microsoft Knowledge Base.

It is best practice to run the latest Unitrends appliance and agent software versions to protect your SQL environment. Older versions do not support all current Unitrends features:

- To protect SQL Server 2022 on Windows, the SQL server must be running Windows 2022 and agent version 10.7.8 or later. The appliance must be running release 10.7.8 or later. (SQL Server 2022 on Linux is not supported.)
- To protect SQL Server 2019 on Windows, the appliance and Windows agent must be running release 10.4.4 or later. (SQL Server 2019 on Linux is not supported.)
- To protect SQL Server 2017 on Windows, the appliance and Windows agent must be running release 10.1 or later. (SQL Server 2017 on Linux is not supported.)
- To protect Always On availability groups, the appliance and Windows agent must be running release 10.1 or later.
- To protect SQL Server 2016, the appliance and Windows agent must be running release 9.0.0-13 or later.
- To protect SQL Server 2014, the appliance and Windows agent must be running release 8.0.0-4 or later.

Detecting newer SQL versions upon upgrading the agent

After upgrading the Windows agent, you must re-save the SQL server asset and sync inventory to enable the appliance to detect any newly supported SQL versions. Follow these steps to detect new SQL versions:

1 Select **Configure > Protected Assets**. Select the SQL host asset and click **Edit**:

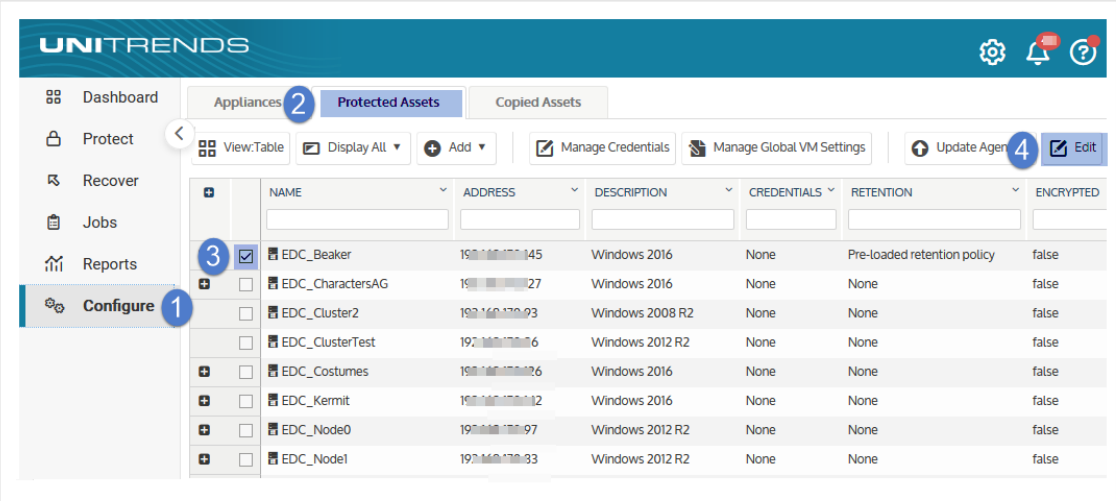

2 Click **Save**:

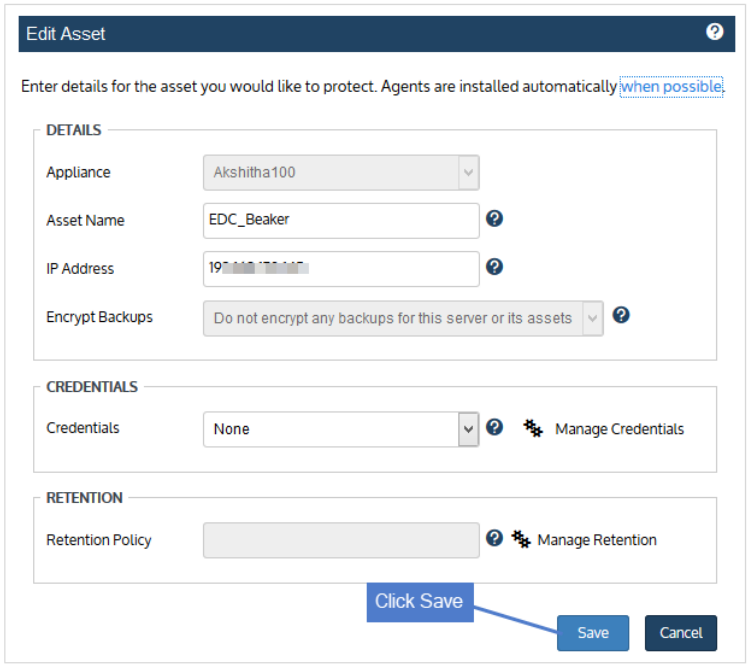

3 Click the Gear icon (Options menu) and select **Inventory Sync**:

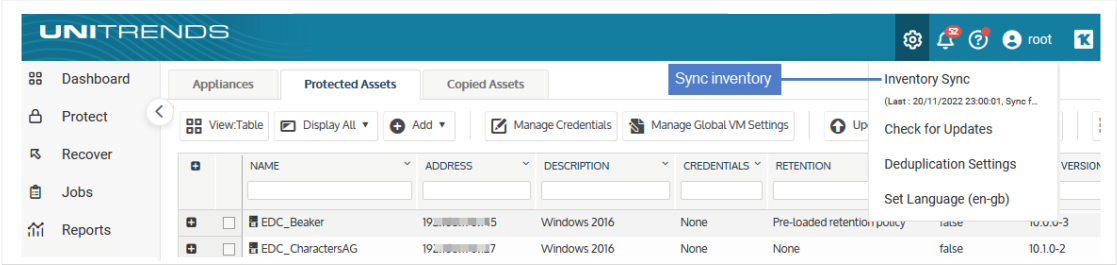

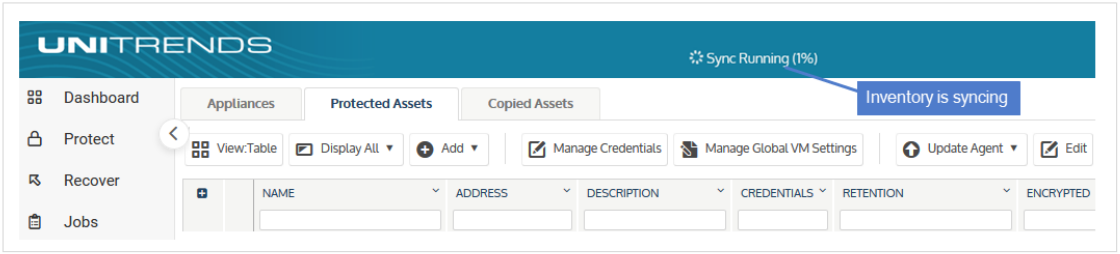

After the inventory sync completes, any newly supported SQL versions hosted on the selected asset display in the UI, and you can begin protecting the new applications.

SQL system requirements

In addition to the agent requirements above, the following requirements must be met for Unitrends protection of all SQL environments:

- The SQL application must be a supported version listed in the Unitrends Compatibility and Interoperability Matrix.
- The SQL server must be running a supported Windows Server operating system listed in the Unitrends Compatibility and Interoperability Matrix. Windows Server 2012 or later is required if the configuration uses cluster shared volumes (CSV).

  Note: SQL application backups are supported for SQL on Windows Server OS only. SQL on Windows Workstation or Linux OS is not supported.

- The SQL application and Windows Server must be set up in a supported Microsoft deployment configuration.
- TLS version 1.0 or 1.2 must be installed and enabled on the SQL server.
- SQL database names must not contain single or double quotes (' or ").

  Note: Databases that contain quotes in the name cannot be protected due to a Microsoft VSS limitation. These databases do not display in the Unitrends UI.

Requirements for SQL clusters and availability groups

In a SQL cluster, one or more SQL Server failover cluster instances are installed into a Windows Server Failover Cluster (WSFC). Unitrends protects SQL cluster configurations that meet the requirements in the following table.

Requirement	Description
Cluster configuration	The SQL cluster must be set up in a supported Microsoft configuration.
SQL application	The SQL applications must be supported versions listed in the Unitrends Compatibility and Interoperability Matrix.
SQL server	The SQL servers must be running supported Windows operating systems listed in the Unitrends Compatibility and Interoperability Matrix. Windows 2012 or later is required if the configuration uses cluster shared volumes (CSV). **Note:** On Windows 2008 R2, there is a Microsoft VSS limitation that allows only one snapshot at a time for CSV volumes. Because SQL jobs typically run in parallel, Unitrends requires support for multiple VSS snapshots.
Windows agent	The Unitrends Windows agent must be installed on every SQL server node in the WSFC. (This applies even if your SQL servers are virtual machines.) These agent version requirements apply: • Each node in the WSFC must be running agent version 10.6.5 or earlier. (The

Requirement	Description
	secure pairing feature introduced in agent release 10.6.6 does not yet support WSFC configurations.) • Every node in the WSFC must be running the same agent version. It is best practice to upgrade agents to release 10.6.5 to take advantage of performance enhancements and fixes. • Agent version 10.0 or later is required to protect Always On availability groups. (The appliance must also be running Unitrends version 10.0 or later. Upgrade the appliance before installing the agent.) **Note:** It is important that the appliance version is the same or higher than the agent versions of its protected assets. For SQL, API changes introduced in release 10.0 require that the appliance is upgraded first.
Clustered SQL instance	Because a clustered SQL instance can move between SQL server nodes in the cluster, you must add each cluster node and the SQL cluster itself to the backup appliance, each as a separate asset. You then create backup jobs by selecting the SQL cluster asset. On the backup appliance, there must be only one SQL cluster asset for each clustered SQL instance. To add a SQL cluster asset to the appliance, you enter its SQL cluster IP address. (For details, see "Protecting SQL clusters and availability groups" on page 748.) If a clustered SQL instance is configured with multiple IP addresses, add a single SQL cluster asset to the appliance by using one static IP address. Do not add multiple assets for a given SQL cluster by using its other IP addresses. If backups start failing after the SQL clustered instance has failed over, follow Microsoft's failover cluster troubleshooting steps. Examples of common post- failover issues for Always On Failover Cluster Instances (FCI) are given below, along with links to applicable Microsoft SQL documents. (If your cluster is not an Always On FCI, see the applicable Microsoft SQL documents for your environment.) • For general troubleshooting information, see these Microsoft documents: Failover Cluster Troubleshooting and Always On Failover Cluster Instances (FCI). • There are some circumstances where a manual restart of the new primary database is required. See the following Microsoft document for details: Failover Policy for Failover Cluster Instances. • When utilizing SQL failover clusters, the databases are protected at the SQL instance level, where one set of database files is saved on a shared storage device. The failover process takes as long as necessary to write all dirty pages in the cache to disk. For information on cutting down your SQL failover time, see the following SQL documentation: Indirect Checkpoints.
Always On	With this feature, one or more user databases are configured in an Always On

Requirement	Description
availability groups	availability group, where mirrored copies of each database reside on the secondary server nodes in the cluster. Availability group databases on the primary (active) cluster node reside in a *primary replica*. Mirrored copies reside in *secondary replicas*. Unitrends backs up databases in the primary replica only. Because any secondary replica can become the primary in a failover, you must add each cluster node and the availability group itself to the backup appliance, each as a separate asset. You then create backup jobs by selecting the availability group asset. The following are required to protect Always On availability groups: • The availability group must be set up in a supported Microsoft configuration. The following configurations are *not* supported: – Clusterless Availability Groups (introduced in SQL 2017) – Always On Basic Availability Groups – Availability groups with secondary replicas in Microsoft Azure • An availability group listener is required. • The availability group listener must be configured to use a static IP address and cannot be configured to use DHCP. • The cluster nodes and availability groups must be added to the Unitrends appliance as described in "Protecting SQL clusters and availability groups". To add an availability group, you must supply its listener IP address. • On the backup appliance, there must be only one asset for each availability group. To add an availability group asset to the appliance, you enter its listener IP address. (For details, see "Protecting SQL clusters and availability groups" on page 748.) If an availability group is configured with multiple listener IP addresses, add a single asset to the appliance by using one static listener IP address. Do not add multiple assets for a given availability group by using its other listener IP addresses. • After a failover, a new full backup is required. If the next job is a differential or transaction log backup, it is automatically promoted to a full. Once this full completes, subsequent differential and transaction log backups run as scheduled. **Note:** If you cannot allow a full backup to run at certain times in your environment (such as during peak business hours), you can opt to have the appliance fail any scheduled differential and transaction log backups until a full runs. To implement this option, set the SQL_AutoPromote flag to FALSE in the *C:\PCBP\MASTER.INI* file on each SQL server node. (If you installed the Windows agent in a custom location, check for the file under that location instead.)

Requirements for SQL databases located on SMB 3.0 shares

SQL Server 2012 and higher can host SQL instances with disk storage located on SMB 3.0 shares.

Prerequisites and considerations for protecting SQL databases located on SMB 3.0 shares

The following prerequisites must be met to protect SQL databases located on SMB 3.0 shares:

- The File Server and the File Server VSS Agent Service roles must be installed on the server hosting the shares. For instructions on installing these roles, see How do I install the File Server and File Server VSS Agent Service roles on a server hosting SMB shares?.
- The Windows agent installed on the SQL server must be granted read/write access to remote SMB 3.0 shares. For instructions on granting this access, see "Granting the Windows agent read/write access to remote SMB 3.0 shares" on page 744.
- The SQL server hosting the databases and the server hosting the SMB shares must belong to the same Windows domain.
- The database can contain one or more files located on SMB 3.0 shares. All files can reside on the same SMB 3.0 share or on different shares hosted by one or more servers in the same domain. All servers participating in the database backup must belong to the same domain.
- For files located on remote SMB 3.0 shares, the Windows agent creates a VSS snapshot on the remote server and then exposes it to the SQL server through the SMB share pathing. The agent then backs up the database files from the remote snapshot location. When the backup completes, all VSS snapshots created for the backup are removed from the server hosting the SMB share.

Granting the Windows agent read/write access to remote SMB 3.0 shares

The Windows agent installed on the SQL server must be granted read/write access to remote SMB 3.0 shares. Grant this access using one of the following methods:

- On the SQL server, change the login account for the Unitrends Windows agent service "bpagent" to the domain administrator account. Using these credentials provides all necessary access to the SMB shares. This is the most secure option for SMB access. Note, however, that backups of the SQL server may encounter files whose permissions do not allow domain administrator access. If this is the case for your SQL server and SMB share security is less of an issue, then the method below is recommended.
- Run the agent as local system account on the SQL server and grant it read/write permission for the SMB shares. For instructions, see Running the Windows agent as local system account on Hyper-V server and granting account read/write permissions for SMB shares.

Once you have satisfied the SMB 3.0 prerequisites and have granted the Windows agent access to the SMB 3.0 shares, run backups as described in "Backup Administration and Procedures" on page 425.

Requirements for SQL Always Encrypted databases

With the SQL Always Encrypted feature, data is encrypted at rest and in motion. Plaintext is exposed only within the SQL application itself. In addition to the agent and system requirements, the following apply to protecting SQL Always Encrypted databases:

Item	Always Encrypted database requirement or consideration
Unitrends appliance version	Must be running release 9.0.0-13 or higher.
Unitrends agent version	The SQL server must be running Windows agent release 9.0.0-13 or higher.
Encrypted data	Data is encrypted at the client level and not at the database level. Encrypted databases cannot be viewed in SQL Management Studio.
SQL Column Encryption Keys	These keys are included in SQL backups and are restored when a SQL backup is recovered.
SQL Column Master Keys (CMKs)	Each Always Encrypted database has CMKs that are stored in a trusted key store located on the local SQL server. The CMKs are not included in SQL backups. After recovering backups of Always Encrypted databases, the CMKs must be available on the recovery target so you can access the recovered data. If these keys are not available, you must install them after you recover the backup. In most environments: • You will not need to install CMKs if recovering to the original database or to another database on the original instance. • You will need to install CMKs if recovering to a different SQL server. • You may need to install CMKs if recovering to a different instance on the original server. • See Microsoft's documentation for instructions on installing the CMKs.

Requirements for SQL Stretch databases

A Stretch database consists of a database that resides on the local SQL server and a paired database that resides in Microsoft Azure. For each table being stretched, an identical table exists in both the Azure and SQL databases. SQL Server moves data from the local tables to the Azure tables based on a user-defined function that acts as a filter.

In addition to the agent and system requirements, the following apply to protecting SQL Stretch databases:

Item	Stretch database requirement or consideration
Unitrends appliance version	Must be running release 9.0.0-13 or higher.
Unitrends	The SQL server must be running Windows agent release 9.0.0-13 or higher.

Item	Stretch database requirement or consideration
agent version	
Data protected	Unitrends backups capture the data on the local SQL server only. Data that was migrated to Azure before the backup runs is not included in the SQL backup.
Data recovered	Recovering a Stretch database backup recovers the part of the database that was backed up on the local SQL server only. You must recover the Unitrends backup to the original database on the original instance. Recovering to an alternate database, instance, or SQL server is not supported. After you recover to the original database, you must reconcile the local data with data that has been migrated to Azure. For instructions, follow Microsoft's Stretch database recovery recommendations in the article Backup and restore Stretch-enabled databases. This requires reconnecting the local recovered database to the remote Azure database using the SQL Master Key and the original credentials that were created when the database was stretched.
SQL Master Key	Each Stretch database has a SQL Master Key that is stored in a certificate located on the local SQL server. This key is not included in SQL backups. After recovering a Stretch database backup, you need to use this key to connect to the Azure database and reconcile the local recovered data with data that has been migrated to Azure.

SQL recovery models

The recovery model of your SQL database determines what type of Unitrends backups are supported. See the table below for descriptions of the SQL recovery models that are supported by Unitrends. See the Microsoft article Recovery Models (SQL) for additional information on recovery models and how to choose the best recovery model for your environment.

Recovery Model	Backups Supported	Considerations
Simple	• Full • Differential	No SQL logs created.
Full	• Full • Differential • Transaction log	Schedule weekly transaction log backups to truncate logs. See "Recommendations for full recovery model" on page 748 for details.

Recovery Model	Backups Supported	Considerations
Bulk-Logged	• Full • Differential	Run a transaction log backup before switching from the full recovery model to the bulk-logged recovery model. See "Recommendations for bulk-logged recovery model" on page 748 for details.

SQL system databases

The following table provides descriptions of the SQL system databases and how they can be protected with Unitrends.

Database	Description	Compatible recovery model and strategy
master	Stores all system-level information, such as logon accounts, configuration settings, and metadata.	Only uses the simple recovery model and must be protected with full backups. Before recovering this database, all other databases must be stopped.
msdb	Used to schedule alerts, jobs, and broker services for database mail. Records backup and restore history.	Uses the simple recovery model by default, but can be configured to use the full recovery model. (Recommended only if msdb history is used when recovering backups.)
model	Acts as a template for any new databases that are created. Content of the model is copied to each new database.	By default it is configured to use the full recovery model, and new databases inherit this setting. It is only backed up when settings are changed.
resource	Contains internal system objects. (Read-only)	This database cannot be backed up or recovered.
tempdb	A temporary workspace used by any session connected to the SQL Server instance. Used to hold intermediate or temporary data, such as temporary tables, cursors, and data for sorting.	Every time SQL Server starts, this database is re-created. Backups are not needed since there is no reason to preserve this database.
distribution	Stores metadata and history data in support of SQL Server replication.	Present only if replication is configured.

Example SQL backup strategies

This section provides example strategies for protecting your SQL databases with Unitrends software.

Database	Backup Strategy
System databases	Weekly full backups
User databases using the full recovery model	Weekly full, daily differential, and hourly transaction logs
User databases using the simple recovery model	Bi-weekly full backups with daily differentials

Recommendations for full recovery model

When using the SQL full recovery model, transaction log backups must be performed to truncate log files. If not truncated, log files continue to grow until the space on your disk is full, resulting in system failure. To prevent runaway transaction log files, make sure that you create a schedule with frequent transaction log backups.

Recommendations for bulk-logged recovery model

The SQL bulk-logged recovery model is used as a temporary recovery model to enhance performance when running bulk jobs. Unitrends does not support log backups while a database is in the bulk-logged recovery model because they are unnecessarily large. For compliance with Unitrends best practices, perform the following steps:

1. Run a log backup while the database is still in full recovery model.
2. Switch to the bulk-logged model.
3. Perform the bulk operation. (For example, importing new labels, copying data from one table to another, or creating an index.)
4. Switch back to the full recovery model.

Automatic exclusion of SQL data during file-level backups

When you run file-level backups of the Windows server hosting SQL, certain SQL-related files are automatically excluded:

- The following extensions are excluded from SQL user databases if the SQL VSS component is running on the Windows asset: *.mdf*, *.ldf*, and *.ndf*.

  Note: If the VSS component is not running, these files are included. SQL files for system databases (such as master, model, and msdb) are always included to support Windows replicas.

- Files in SQL *database/log* directories are excluded.

Protecting SQL clusters and availability groups

SQL Failover Clustering is a high availability and disaster recovery solution where one or more SQL Server failover cluster instances are installed into a Windows Server Failover Cluster (WSFC). Optionally, SQL failover clusters may include Always On failover cluster instances and Always On availability groups.

The WSFC has one active Windows server node. Its other servers are inactive members of the server cluster. If the active node fails, a secondary node is automatically promoted to the primary (active) node.

Unitrends backs up clustered instances and availability groups on the active server node. To provide seamless protection, job schedules must be attached to the clustered instance or availability group, and not to the individual nodes. This enables backups to continue after a failover, providing uninterrupted protection of your clustered SQL environment.

Note: After a failover, a new full backup is required. If the next job is a differential or transaction log backup, it is automatically promoted to a full. Once this full completes, subsequent differential and transaction log backups run as scheduled.

In the Unitrends UI, each SQL server, clustered SQL instance, and availability group is treated as a separate asset. After you add each SQL server node to the appliance, you use the same Add Asset dialog to add each clustered SQL instance and availability group. You add a clustered SQL instance by entering its virtual IP address (note that this is the SQL cluster IP address, not the general cluster IP address). You add an availability group by entering its listener IP address.

In the Unitrends UI, each SQL instance or availability group displays under its host asset, and each database displays under its instance or availability group:

- For SQL server assets, hosted non-clustered instances display.
- For SQL cluster assets, the clustered database instance displays.
- For availability group assets, the availability group displays.

Note: Filtering the display of instances and availability groups in the UI by SQL asset type was introduced in Unitrends release 10.0.0-6. If your schedules were backing up clustered instances or availability groups from a SQL server node asset, those instances and availability groups are no longer present in existing schedules upon upgrading to release 10.0.0-6 or higher. Follow the instructions in "Start protecting SQL clusters and availability groups" on page 749 to add your clustered instance and availability group assets and modify your job schedules.

Start protecting SQL clusters and availability groups

Do the following to start protecting your clustered SQL environment:

Step 1: Ensure all applicable requirements have been met. For details, see "Requirements for SQL protection" on page 738.

Step 2: Install the Windows agent on each server node in the WSFC. For details, see "Installing the Windows agent" on page 362.

Note: For most Windows servers, the appliance can push-install the agent when you add the asset. If you will be push-installing the agent, skip to Step 3:. For push-install requirements, see "Windows agent requirements" on page 362.

Step 3: Add each WSFC server node to the Unitrends appliance. For details, see "To add a SQL server asset".

Step 4: Add each clustered SQL instance to the Unitrends appliance. For details, see "To add a SQL cluster asset" on page 751.

Step 5: Add each availability group to the Unitrends appliance. For details, see "To add a SQL availability group asset" on page 753.

Step 6: Create backup jobs as described in "To create a SQL backup job" on page 483.

To add a SQL server asset

You must add each server node to the appliance before adding any clustered instance or availability group assets.

1 Select **Configure** > **Protected Assets**, then **Add** > **Asset**:

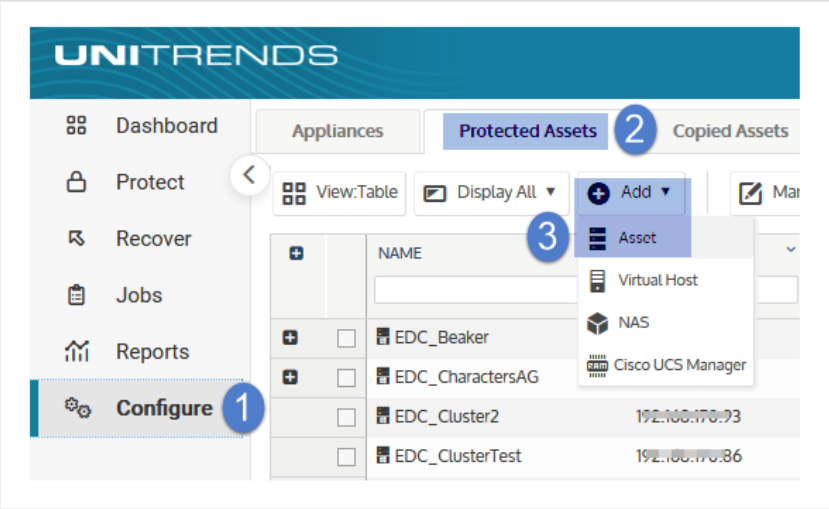

2 Enter the asset's hostname.

3 Enter the asset's IP address. This is optional in some cases, as described here:

- For Windows assets, you can use DNS rather than entering a static IP address.
- DNS registration should be used for assets that obtain their network settings through DHCP. It is optional for assets with static IP addresses.
- If you do not enter a static IP address, make sure that both the asset and the appliance have DNS entries and that reverse lookup is configured.
- If you enter a static IP address, the appliance attempts to connect using this address. If the attempt fails, it tries again using DNS.

4 Enter or select optional settings.

- To push-install the Windows agent, you must supply administrative credentials and check the **Install Agent** box.
- If you have already installed the Windows agent, uncheck the **Install Agent** box.

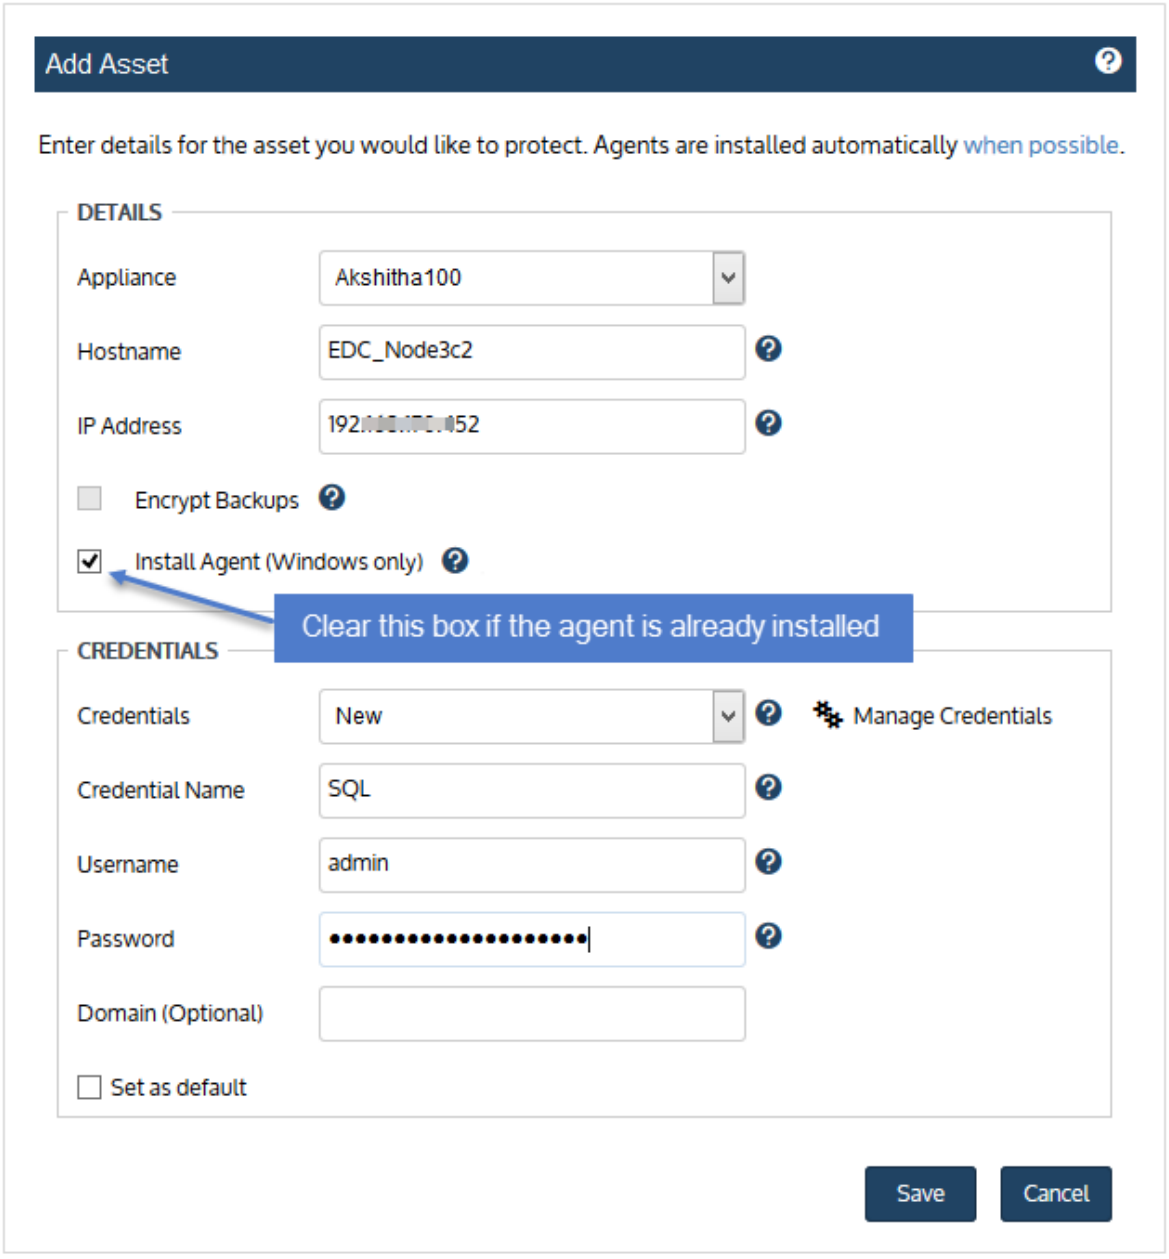

5 Click **Save**.

Repeat this procedure until you have added every server node in the WSFC.

To add a SQL cluster asset

1 Select **Configure > Protected Assets**, then **Add > Asset**:

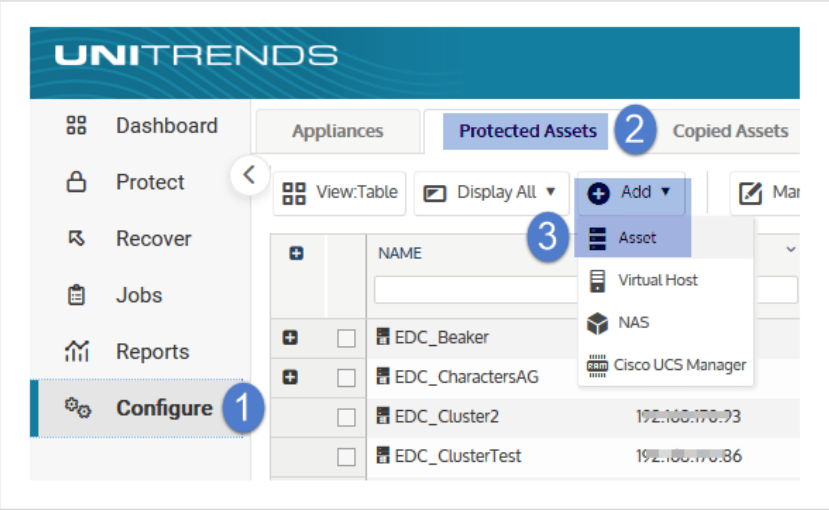

2 Enter a hostname for the clustered instance.

Because assets display by hostname in the Unitrends UI, it is recommended to use a naming convention that identifies the instance and its cluster.

3 Enter the IP address of the clustered SQL instance. (This is the virtual IP address used to connect to the SQL server.)

4 Uncheck the **Install Agent** box. (The agent is installed on SQL servers only).

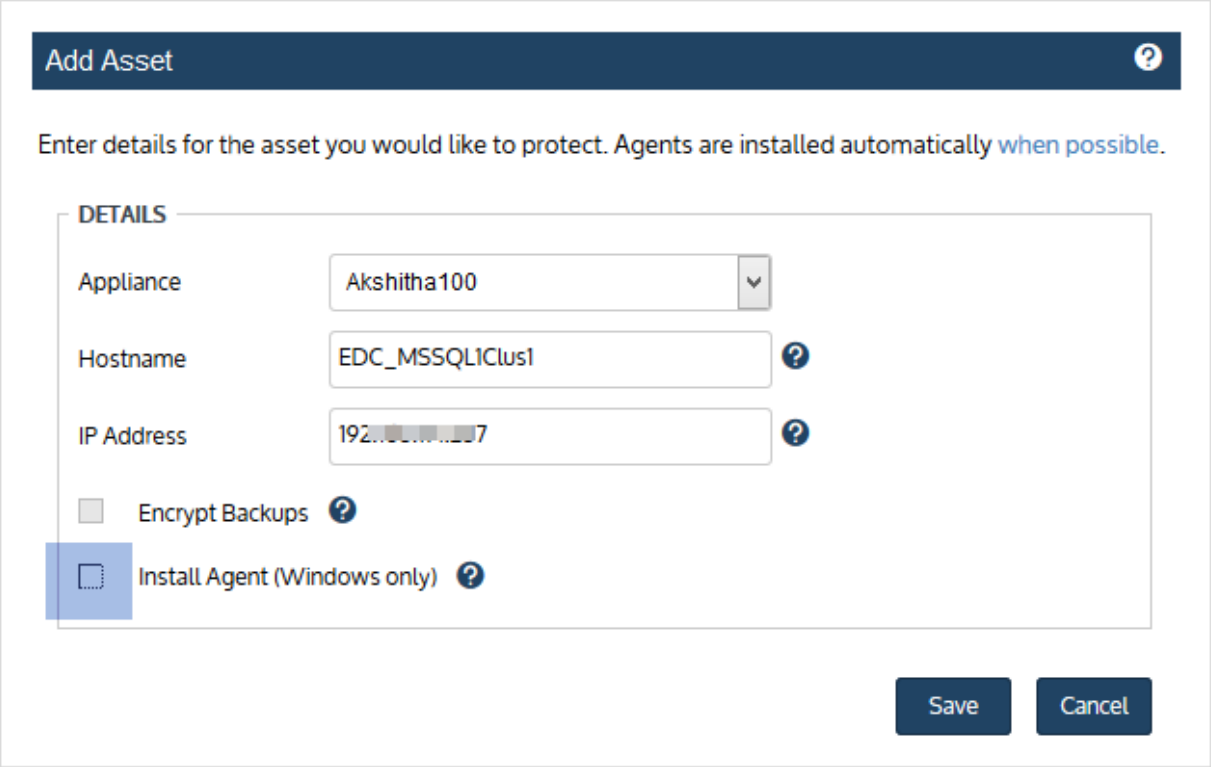

5 Click **Save**.

To add a SQL availability group asset

1 Select **Configure** > **Protected Assets**, then **Add** > **Asset**:

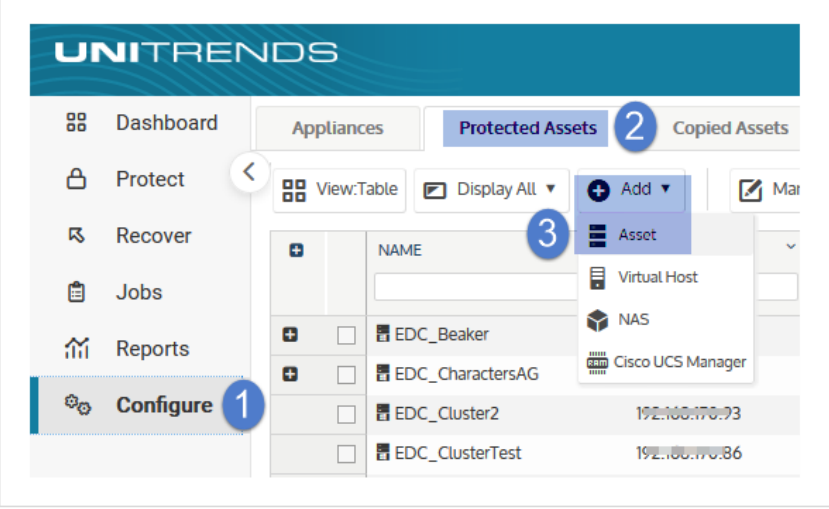

2 Enter a hostname for the availability group asset.

Because assets display by hostname in the Unitrends UI, it is recommended to use a naming convention that identifies the availability group and its cluster.

3 Enter the IP address of availability group listener.

4 Uncheck the **Install Agent** box. (The agent is installed on SQL servers only).

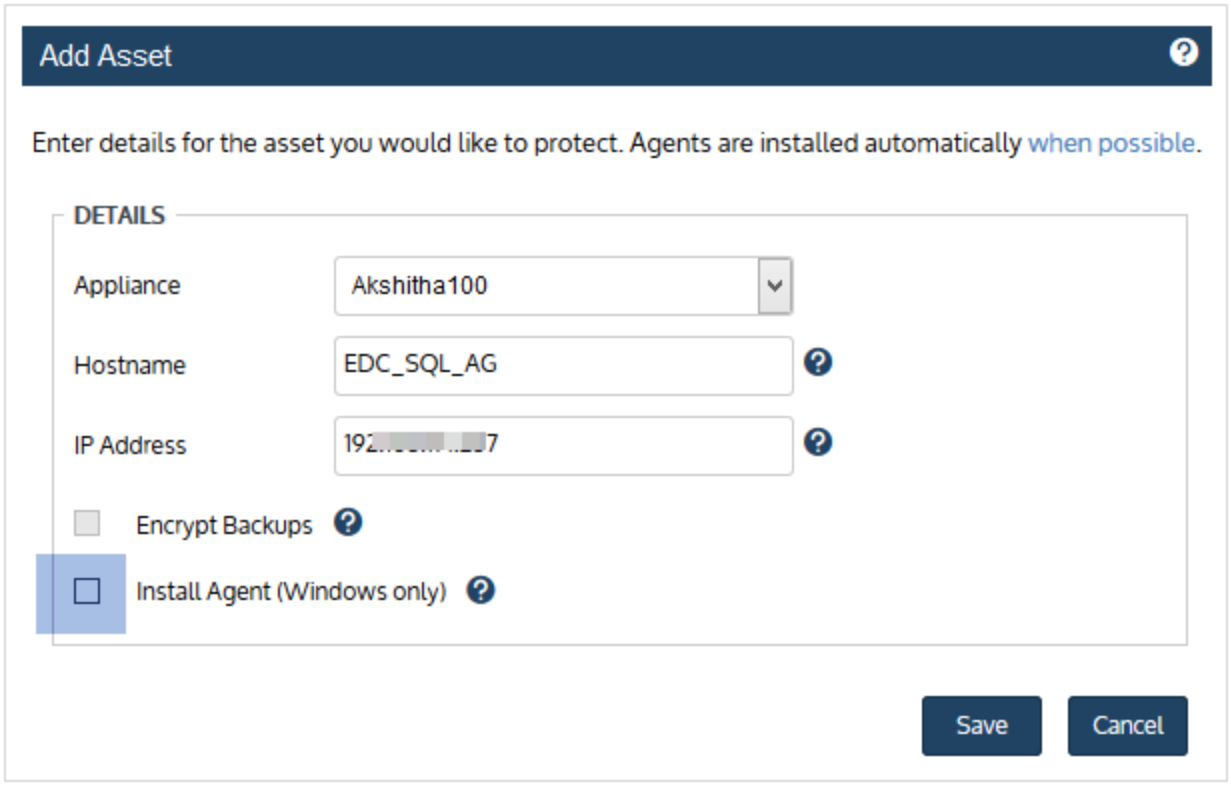

5 Click **Save**.

Start protecting non-clustered SQL environments

Do the following to start protecting your SQL environment:

Step 1: Ensure all applicable requirements have been met. For details, see "Requirements for SQL protection" on page 738.

Step 2: Install the Windows agent on the SQL server as described in "Installing the Windows agent" on page 362.

> **Note:** For most Windows servers, the appliance can push-install the agent when you add the asset. If you will be push-installing the agent, skip to Step 3:. For push-install requirements, see "Windows agent requirements" on page 362.

Step 3: Add the SQL server to the Unitrends appliance as described in "To add a SQL server asset" on page 750.

Step 4: Run backup jobs as described in "To create a SQL backup job" on page 483.

SharePoint backup requirements and considerations

Unitrends protects the following SharePoint environments:

- Farm deployments where the SharePoint installation type is *full farm* (all SharePoint releases) or *single server farm* (SharePoint 2016). In farm deployments, the SharePoint data and components may reside on one server or on multiple servers. For details, see these Microsoft articles: Install SharePoint Server 2016 across multiple servers, Install SharePoint 2013 across multiple servers for a three-tier farm, or Install SharePoint 2016 on a single server with SQL Server.
- Single server deployments where the SharePoint installation type is *single server* (SharePoint 2010 and 2013). In single server deployments, all SharePoint data and components reside on one server. For details, see this Microsoft article: Install SharePoint 2013 on a single server with SQL Server.

For all installations, the Primary SharePoint server runs the Central Administration website service, which can be accessed using **`http://<machine name>:<admin port>`**

The Unitrends Windows agent provides protection of services and resources in a Microsoft standalone or multi-server SharePoint farm.

In a SharePoint deployment, the primary node installs SharePoint services on other member servers and initiates administrative commands to manage the farm. The Central Administration service runs on the primary node to perform farm management. All nodes directly access the SharePoint central configuration database for configuration of services, features, database connections, and the like. The central configuration database resides either on the primary node or on a stand-alone SQL server. Unitrends protects the farm from the primary node, where administrative commands are run to coordinate the backup of data across other nodes in the farm.

To ensure application consistency, the agent leverages SharePoint's STSADM and PowerShell (SharePoint 2013 and higher) tools to run backup and recovery jobs . The agent invokes commands on the SharePoint primary node and supplies STSADM or PowerShell with a local share target (*/backups/rae/<client_name>/<instance>*) so that jobs run on the backup appliance itself.

The agent works with STSADM or PowerShell to back up the SharePoint-specific data and files on each node in the farm. STSADM or PowerShell discovers the online nodes and performs backup operations to the local backup appliance share. If a node is not available, the backup continues without error. The resulting backup does not include any nodes that were unavailable when the backup ran.

Notes:

- SharePoint protection includes SharePoint data only. To protect an entire node in the farm, add the node to the backup appliance and run file-level backups.
- Full catastrophic farm recovery can only be performed for SharePoint 2013 and 2010 deployments where the installation type is *single server*. For *full farm* (all SharePoint releases) or *single server farm* (SharePoint 2016) installations, you must recover items instead. To check your installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 757.

Consider the following before implementing your SharePoint protection strategy:

SharePoint agent requirements

The following requirements must be met for SharePoint protection:

- SharePoint must be running a supported version listed in the Unitrends Compatibility and Interoperability Matrix.
- Unitrends supports on-premise farm deployments only. Hybrid deployments, such as integration with Office 365, cannot be protected by Unitrends backups.
- The SharePoint farm configuration must adhere to Microsoft best practice standards. An SPFarmBackup domain account that is a member of the *local administrators* group must be configured on each node in the farm.

  Note: Farms containing a single server may have been set up as a full farm or as a single server during installation. For SharePoint 2013 and 2010, protection procedures vary by installation type. To check the installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 757.

- SharePoint administration and timer services must be running on the primary node.
- The SharePoint administration and timer services must have local administrator privileges. Be sure the service is a member of the necessary Windows security groups or SharePoint groups.
- SharePoint must be able to access the appliance's Samba share:
  - SMB 2.0 – The SMB 2.0 security option is enabled by default on Unitrends appliances that were originally imaged or deployed with version 10.4.8 or higher. To perform backup and recovery operations, SMB 2.0 must be enabled on each node in the farm.

    Notes:

    - SharePoint 2007 on Windows 2003 and prior is not supported on SMB 2.0 appliances. (To configure your appliance to use SMB 1.0, see How Unitrends supports SMBv2.)
    - SharePoint may require custom client configuration for use with SMB 2.0. If SharePoint backups do not run successfully, see this Microsoft article for client configuration details: SharePoint Ports, Proxies and Protocols...An overview of farm communications.

  - SMB 1.0 – The SMB 1.0 security option is enabled by default on appliances that were originally imaged or deployed with a pre-10.4.8 version. SMB 1.0 must be enabled on each node in the farm.

    Note: Upgrading from a pre-10.4.8 version does not change the SMB 1.0 setting. (To configure your appliance to use SMB 2.0, see How Unitrends supports SMBv2.)

- Prerequisite configuration steps must be performed on the primary node, as described in "SharePoint configuration prerequisites" on page 757.

- Trust credentials are needed to back up the SharePoint database. Adhere to the following requirements when creating SharePoint credentials:

  Note: Beginning in Unitrends release 9.1, trust credentials are required for both single server and full farm installations. In 9.0, credentials were not supported for single server installations.

  - Credentials must be applied to the database instance.
  - To ensure sufficient privilege, the credential user must be a member of the administrators group on the local computer for each member of the farm, and a member of the farm administrator's SharePoint group.
  - The SharePoint user must have permission to log on as *batch job* and log on as *service*.
  - Create the credentials and apply them using the procedures in "Managing asset credentials" on page 322. When applying credentials, be sure to expand the SharePoint server and select the **Full Farm** or **Single Server** application instance.
  - If you experience backup errors using new credentials, see the following Knowledge Base articles for more information: How Unitrends supports SMBv2, SharePoint backup failed when user account name or password is incorrect, SharePoint backup failed when user account password has been expired, SharePoint backup failed when user account has not been granted the requested logon type, SharePoint backup failed when user has not been granted the requested logon as service, and SharePoint backup failed with error: 'Object reference not set to an instance of an object'.

SharePoint configuration prerequisites

You must perform one of these procedures before you can begin protecting your SharePoint environment:

- For all SharePoint 2016 deployments, see "To configure a farm for Unitrends protection" on page 758.

  Note: Unitrends protects the SharePoint 2016 *single server farm* installation type just like any *full farm* installation. Use the standard farm procedures for your SharePoint 2016 environment, for both *single server farm* and *full farm* deployments.

- For SharePoint 2013 and 2010 deployments where the SharePoint installation type is *full farm* and the SharePoint data and components reside on one or more servers, see "To configure a farm for Unitrends protection" on page 758.
- For SharePoint 2013 and 2010 deployments where the SharePoint installation type is *single server* and all SharePoint data and components reside on one server, see "To configure services on a standalone SharePoint 2013 or 2010 server" on page 758.
- If you are unsure of the installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 757.

To determine the installation type for SharePoint 2013 and 2010 deployments

1. On the **Configure > Appliances** page, select the appliance that is protecting the farm.
2. Click **Protected Assets** and expand the SharePoint asset to view the farm instance:

- If the instance is *Full Farm*, it was configured as a multi-server installation. Note that you must use the Untirends multi-farm procedures to protect this farm, even if there is only one physical server in the SharePoint installation.
- If the instance is *Single Server*, it was configured as a single server installation.

To configure services on a standalone SharePoint 2013 or 2010 server

Use this procedure for standalone SharePoint 2013 or 2010 servers where the installation type is *single server*. For the *full farm* installation type, see "To configure a farm for Unitrends protection" on page 758. To check your SharePoint installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 757.

1. Install the Unitrends Windows agent on the SharePoint server as described in "Manually installing the Windows agent" on page 366.
2. Add the SharePoint server to the Unitrends appliance as described in "To add an agent-based asset" on page 289.
3. Log in to the SharePoint server.
4. Open **Services** and verify that the following services are running. If not, start them.
   - SharePoint 2010/2013 Timer or Windows SharePoint Services Timer
   - SharePoint 2010/2013 Administrator or Windows SharePoint Services Administration
5. For each of the above services, set the startup type to *automatic*.
6. Proceed to "To create a SharePoint backup job" on page 487 to start protecting your SharePoint environment.

To configure a farm for Unitrends protection

Use this procedure to configure a *full farm* (all SharePoint releases) or *single server farm* (SharePoint 2016) deployment containing one to many servers. For SharePoint 2013 or 2010 *single server* installations, see "To configure services on a standalone SharePoint 2013 or 2010 server" on page 758. To check your SharePoint installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 757.

This procedure assumes your SharePoint environment has been setup with a SPFarmBackup domain account that is a member of the local administrators group, in accordance with Microsoft best practices.

1. Install the Unitrends Windows agent on the primary node as described in "Manually installing the Windows agent" on page 366.

   The primary node is the one running the Central Administration service. To see services on each node, log in to any node in the farm, and select **All Programs > Microsoft SharePoint Products > SharePoint Central Administration**. On the Central Administration page, select **System Settings > Manage servers in the farm**.
2. Log in to the primary node, and open **Services**.
3. Verify that the following services are running. If not, start them.
   - SharePoint 2010/2013/2016 Timer or Windows SharePoint Services Timer
   - SharePoint 2010/2013/2016 Administrator or Windows SharePoint Services Administration

4 For each of the above services, set the startup type to *automatic*.

5 Add the primary node to the Unitrends appliance (as described in "To add an agent-based asset" on page 289) and apply administrative trust credentials to the *Full Farm* database instance.

6 Proceed to "To create a SharePoint backup job" on page 487 to start protecting your SharePoint environment.

Oracle backup requirements and considerations

Use application backups to protect Oracle Database on Windows, Linux, and Solaris platforms. Application backups ensure database consistency, whereas file-level backups of the Oracle server are likely to contain database inconsistencies since only data that has been flushed to disk is included.

With Oracle protection, the Unitrends agent leverages Oracle's Recovery Manager (RMAN) utility for backup and recovery jobs to:

- Ensure a consistent database snapshot is captured.
- Perform standard Oracle database backup operations, such as saving redo logs and quiescing buffers.

The agent invokes commands on the Oracle server and supplies RMAN a Samba share target (*/backups/rae/<client_name>/<instance>*) so that jobs save directly to the backup appliance.

Oracle protection requirements and considerations vary by platform and Oracle Database version. See the following for details:

- "Oracle server, instance, and job requirements" on page 759
- "Guidelines for creating Oracle credentials" on page 762
- "Start protecting Oracle" on page 763
- "Upgrading to newer Oracle versions" on page 763

Oracle server, instance, and job requirements

These requirements must be met for Oracle protection:

Oracle Requirement	Description
Oracle server	Oracle platform, agent, server, and credential requirements are described below.
Platform	Verify the server is running a supported Windows, Linux, or Solaris version listed in the Unitrends Compatibility and Interoperability Matrix.
Agent	Install the applicable Unitrends agent as described in "Installing the Unitrends agent" on page 280. The agent must have access to the Samba share on the Unitrends appliance to perform backup and recovery operations:

Oracle Requirement	Description
	• SMB 2.0 – The SMB 2.0 security option is enabled by default on Unitrends appliances that were originally imaged or deployed with version 10.4.8 or higher. For Windows servers, SMB 2.0 must be enabled. For Oracle on Solaris, Samba must be enabled and a Samba key must be added for the backup appliance (see "Oracle on Solaris 11" for details). • SMB 1.0 – The SMB 1.0 security option is enabled by default on appliances that were originally imaged or deployed with a pre-10.4.8 version. For Windows servers, SMB 1.0 must be enabled. For Oracle on Solaris, Samba must be enabled (see "Oracle on Solaris 11" for details). **Note:** Upgrading from a pre-10.4.8 version does not change the SMB 1.0 setting. (To configure your appliance to use SMB 2.0, see How Unitrends supports SMBv2.)
Oracle server	Add the Oracle server to the Unitrends appliance as described in "To add an agent-based asset" on page 289.
Credentials	Configure trust credentials for each application instance you wish to protect as described in "Managing asset credentials" on page 322. • For Windows, the credential user must be a member of the *ora_dba* group. • For Linux and Solaris, the user must be a member of the group that owns the Oracle database instance. • For additional considerations, see "Guidelines for creating Oracle credentials" on page 762.
Instances	The following requirements apply to all Oracle instances. (See additional requirements for Oracle on Windows, Linux, and Solaris below.) • Must be online and in *OPEN* status. Modes such as *MOUNTED*, *NOT MOUNTED*, and *SHUTDOWN* are not supported. • Must be running and configured in *ARCHIVELOG* mode. This enables archiving (backup) of the Oracle redo log which guarantees you can recover all committed transactions, and also enables Unitrends to back up the database while it is open and in normal system use. Archived redo log files are deleted from Oracle each time a full backup completes successfully. This keeps the logs from overrunning tablespace.

Oracle Requirement	Description
	• Each Oracle SID on an Oracle server must be unique. • Each instance must have a defined Oracle Home and only one Oracle Home per instance. • Instances containing non-UTF-8 compatible characters are not supported. • Oracle database instances must be deployed using the File System storage type. Other configurations are not supported. • Oracle databases must be configured as single instances. Clustered configurations, such as Oracle single-server Real Application Clusters (RAC) and Oracle multi-server RACs, are not supported.
Oracle on Windows	Version must be 19c, 18, 12c, or 11g.
Oracle on Linux	• Version must be 19c, 12cM, or 11g. • Must install the Oracle Dependency as described in "Installing and updating the Linux agent" on page 387.
Oracle on Solaris 11	• Must be version 12c or 11g. • A Samba client for Solaris must be enabled. See How Unitrends supports SMBv2 for details. • A Samba key must be added for the backup appliance. To add the key, issue this command (the default password is *samba*): `smbadm add-key -u samba@<applianceIP>` • Ensure the Solaris client has sufficient memory available. See Oracle Database on Solaris: Backup Fails with Out of Memory for details. • Full pathname to each Solaris object cannot exceed 1024 characters. For details, see Solaris: File name too long.
Oracle jobs	The following apply to Oracle backup and recovery jobs: • A given Oracle database can be protected by one Unitrends appliance only and cannot be included in an Oracle Enterprise Manager schedule.

Oracle Requirement	Description
	• Free space equivalent to twice the size of the backup is required on the remote share. If adequate space is not available, the backup fails • Only one backup or restore job per Oracle instance can run at any given time. • For a given database, any job initiated while another job is in progress will fail. Once the job completes, another can be run for the given database. • For Oracle on Windows and Oracle on Solaris, Unitrends supports full backups and level 1 incremental backups. The incremental forever backup strategy is not supported. • For Oracle on Linux, Unitrends supports full backups, level 1 incremental backups, and the incremental forever backup strategy. Additional setup is required to use the incremental forever backup strategy. For details, see Oracle Database: Incremental Forever Schedules on Linux Platforms.

Guidelines for creating Oracle credentials

Credentials are required to perform Oracle backup and restore operations. If no credentials are available, or if credentials are incorrect, the job fails with a *TNS permission denied* error.

Follow the guidelines below when applying Oracle credentials. After reviewing the guidelines, proceed to "Managing asset credentials" on page 322 to create and apply credentials.

Oracle platform	Guidelines and requirements
Oracle on Linux or Oracle on Solaris	Apply credentials to each application instance you wish to protect. The credential user must be a member of the group that owns the Oracle database instance.
Oracle on Windows	Choose one of the following strategies: • If the Windows *NT AUTHORITY\SYSTEM* user is a member of the *ora_dba* group, you do not need to use Oracle credentials. Oracle backups and restores are performed using the *NT AUTHORITY\SYSTEM* account. • If you are using the push feature to install and update the Windows agent on the Oracle server, administrative credentials have been applied to the Windows server asset. If this Windows credential user is a member of the *ora_dba* group, these

Oracle platform	Guidelines and requirements
	credentials can be used for Oracle protection as well. If not, you must also apply credentials to each application instance you wish to protect. **Note:** If credentials have been applied to the Oracle server and its application instances, the appliance uses instance-level credentials for Oracle backups and restores. If instance-level credentials are incorrect, the job fails without attempting to use the server-level credential. • If you are not using the Windows agent push feature, apply credentials to each application instance you wish to protect. The credential user must be a member of the *ora_dba* group.

Start protecting Oracle

After ensuring all requirements have been met, do the following to start protecting your Oracle environment:

Step 1: Install the applicable agent on the Oracle server as described in "Installing the Windows agent" on page 362, "Installing and updating the Linux agent" on page 387, or "Installing and updating the Solaris agent" on page 401.

Step 2: Add the Oracle server to the Unitrends appliance as described in "To add an agent-based asset" on page 289.

Step 3: Run file-level backups to protect the Oracle server, as described in "To create a file-level backup job" on page 437.

Note: Oracle on Windows – You must use Windows file-level backups to protect the Oracle server. Windows image-level backups cannot be used for Oracle.

Step 4: Run Oracle backups to protect the application, as described in "To create an Oracle backup job" on page 480.

Upgrading to newer Oracle versions

If you have upgraded an existing protected Oracle database instance to a newer Oracle database version, follow this procedure to begin protecting your new instance:

1. Ensure all requirements are met for your new version of Oracle database. See "Oracle backup requirements and considerations" on page 759.
2. Select **Options > Inventory Sync** to discover the new instance.
3. Schedule and begin running backups of your Oracle new databases as described in "Backup Administration and Procedures" on page 425.

4 (Optional) If you no longer need to backup the older databases, disable or delete backup schedules for the older instance.

This is necessary because Oracle creates a new database instance when you upgrade, and does not remove or overwrite any older instances.

> **Note:** The appliance does not purge the last successful backup group for the older databases (see "Backup groups" on page 98 for details). If you no longer need any backups of the older databases, you can delete them manually.

5 (Optional) Once you have gained the desired retention on your new instance, you can manually delete backups of the older instance.

Cisco UCS service profile backup requirements and considerations

Review the following before implementing your service profile protection strategy:

- "About protecting Cisco UCS service profiles" on page 764
- "Service profile protection requirements" on page 765
- "Start protecting Cisco UCS service profiles" on page 766

About protecting Cisco UCS service profiles

You can use your Unitrends appliance to back up and recover Cisco UCS service profiles and related configuration objects. In the event of a disaster, you can use this feature to quickly recover your service profiles, greatly reducing the recovery time objective (RTO) of reconfiguring your network and servers.

The Cisco UCS environment provides a “virtual chassis” that enables you to create and assign hardware profiles to individual logical servers. You can then bring up the logical server on dedicated hardware that you can easily migrate to another server in the case of hardware failure, or migrate between servers that do not require 24/7 up-time for efficient hardware reuse.

For UCS B-Series blade servers and C-Series rack-mount servers, allocation of UCS resources and hardware is managed at the domain level by the Cisco UCS manager. Each server in the UCS is a “logical server” that utilizes various resources as defined in the server’s service profile, and there is a one-to-one relationship between a service profile and a physical server. The service profile references hardware requirements, such as hardware identifiers, firmware, state, configuration, connectivity and behavior, but is completely separate from the physical UCS environment. Once a service profile is instantiated and associated with a given blade, rack-mount server, or server in a server pool, you configure a PXE server or map a bootable ISO image to the virtual-media CDROM drive to install the desired hypervisor or operating system (OS). See the Cisco document Cisco UCS Manager Configuration Common Practices and Quick Start Guide for details.

A service profile may be associated with a template and various policies. A service profile template can be used to quickly create additional service profiles. Policies can be used to enforce rules to help ensure consistency. For example, a boot policy defines how a server boots, including boot devices, methods, and boot order.

Because service profiles are essential to managing the servers in your Cisco UCS environment, it is important that you protect these configurations. Unitrends leverages native Cisco UCS data protection for profile backup and recovery, utilizing the Cisco XML API. Unitrends UCS profile backups capture all supported profiles, templates, pools, and policies in your UCS environment. For a description of each supported object that may be included in the UCS profile backup, see "Identifying files in UCS service profile backups" on page 1202. Once you have a UCS profile backup, you can easily recover these items to quickly spin up your Cisco UCS environment in the event of a disaster, greatly reducing RTO.

Notes:

- The following objects are not included in Unitrends UCS profile backups: BIOS defaults, IPMI access policies, management firmware policies (deprecated, replaced by host firmware packages), and iSCSI authentication profiles.
- UCS profile backups capture only service profiles, templates, pools, and policies. To protect UCS servers themselves, add them as assets to the Unitrends appliance and schedule file-level backups.

Service profile protection requirements

The following requirements must be met to protect Cisco UCS service profiles, templates, pools, and policies:

- The Unitrends appliance must be running version 9.0.0-13 or higher.
- The UCS environment must utilize the Cisco UCS manager for resource and hardware allocation.
- The Cisco UCSM firmware must be version 2.0 or higher.

  Note: UCS manager is used for B-Series and C-Series UCS servers. Profile backups of E-Series UCS servers is not supported, but you can protect the servers in your E-Series environment.

- The Cisco UCS manager must be turned on.
- The Cisco UCS manager must be added to the Unitrends appliance with administrative trust credentials that support native backup and recovery of UCS service profiles.
- Only application-level backups of your service profiles and other configuration objects are supported, since the UCS manager is not a server.
- The Cisco UCS may be configured as a stand-alone system, or as a cluster to support failover in the event of an outage. Which IP and name you supply when adding the UCS manager to the Unitrends appliance varies depending on this configuration:
  - The stand-alone configuration consists of one physical UCS fabric interconnect that runs a single UCS manager. To add the UCS manager asset to the appliance, you must either supply the IP address of this node or, if DNS is setup in your environment, you can add the asset by node name only.
  - The cluster configuration is comprised of two physical Cisco UCS fabric interconnects, one active and one standby. A UCS manager runs on each. To add the UCS manager asset to the Unitrends appliance, you must either supply the cluster IP address or, if DNS is setup in your environment, you can add the asset by cluster node name only. Be sure to add the asset by cluster name or cluster IP. Do not use the IP or name of either

fabric interconnect. With this approach, Unitrends can connect to the UCS manager regardless of which fabric interconnect is currently active.

Start protecting Cisco UCS service profiles

Unitrends recommends running weekly or daily full backups of your UCS profiles, templates, pools, and policies. If your profile data changes frequently, you can schedule fulls to run throughout each day at any desired frequency. If you schedule the backup every few minutes, be aware that if the last backup is still running, the next backup is added to the queue and will be started once the last run completes.

After ensuring all requirements have been met, do the following to start protecting your service profiles:

Step 1: Add the Cisco UCS manager to the Unitrends appliance as described in "To add a UCS manager asset" on page 302.

Step 2: Run backup jobs as described in "To create a UCS service profile backup job" on page 490.

Chapter 11: iSeries Backups Overview and Procedures

This section provides requirements and procedures for protecting iSeries. The software for the iSeries is designed to aid in the recovery of lost or corrupted files on this platform. It can be used to back up many types of libraries and objects, including security and configuration files, user programs, and the Integrated File System (IFS). Protection for the iSeries platform is agentless. The iSeries software uses the FTP protocol to back up files from the iSeries to the appliance.

See these topics for details:

Start protecting iSeries

Following is a summary of the high-level steps for setting up iSeries protection. Each step includes a link to detailed instructions.

Step 1: Review "Requirements and considerations for iSeries protection" on page 767 to ensure all prerequisites have been met.

Step 2: Add the iSeries server to the Unitrends appliance as described in "To add an iSeries asset" on page 771.

Step 3: Run backup jobs as described in "Creating iSeries backup jobs" on page 773.

Requirements and considerations for iSeries protection

The following requirements and considerations apply to iSeries protection:

Item	Requirement or consideration
Appliance version	The Unitrends appliance must be running release 9.0.0-13 or higher.
iSeries version	The iSeries must be a supported version listed in the Unitrends Compatibility and Interoperability Matrix.
Backup and	The following apply to backup and recovery jobs:

Item	Requirement or consideration
recovery jobs	• There can be no other active jobs running on the iSeries. Only one backup or recovery job can be running at a time. Use the **WRKACTJOB** command to monitor all active jobs on the iSeries. • Backup and recovery jobs must run without conflict or interruption. • Performance of iSeries backup and recovery jobs is influenced heavily by the following: – Commercial Processing Workload (CPW) of the iSeries server(s) – Amount of library data – Amount of Integrated File System (IFS) data – Available network bandwidth – For recommendations on performance enhancements for the iSeries FTP server, see the IBM article Improving FTP server performance with configurable subsystem support.
Disk space	There must be adequate disk space available on your iSeries asset for a backup to complete successfully. When the backup runs, it backs up the library file system using one thread and the IFS using a second thread. Normally, these threads run in parallel for increased performance. For each thread a SAVF file is created in QTEMP, which consumes disk space. You must have enough available disk space to create these SAVF files or the backup fails. If parallel processing requires too much space in your environment, you can opt to use serial processing. To determine the minimum space required: • Parallel processing - Use the size of the largest library + size of the largest IFS file. • Serial processing - Use the size of the largest library or the size of the largest IFS file, whichever is greater. To modify the processing mode: 1 In the Unitrends UI, select **Configure > Appliances > Edit > Advanced > General Configuration**. 2 Scroll down to the **iSeriesAgent** section. 3 Click **Threading**, change this setting to *1* for parallel or *0* for serial, then click **Save**. 4 Click **Close** to exit.
FTP	The FTP server must be configured and running on the iSeries. The FTP protocol is used

Item	Requirement or consideration
	by backup and recovery jobs.
Maximum file size	If any IFS file exceeds 500MB, that file is backed up individually and is not included in the backup. To prevent this, you can increase the MaxBlockSize setting to accommodate the largest IFS directory. To increase the MaxBlockSize: 1 In the Unitrends UI, select **Configure > Appliances > Edit > Advanced > General Configuration**. 2 Scroll down to the **iSeriesAgent** section. 3 Click **MaxBlockSize**, enter the desired size in bytes, and click **Save**. • To avoid backing up files individually, set this value to accommodate the largest IFS directory. • To use an unlimited MaxBlockSize, set this value to *-1*. • To back up all files individually, set this value to *0*. 4 Click **Close** to exit.
Data protected and disaster recovery	Unitrends software cannot backup all iSeries data, such as licensed internal code and certain system libraries. The iSeries backup cannot be used to recover an iSeries system to its original state in the event of hardware or software failure. To enable disaster recovery, you must perform periodic GO SAVE option 21 or option 22 system backups to use for disaster recovery. GO SAVE backups contain critical files needed to recover the system. For more information on GO SAVE, see the IBM article GO SAVE command menu options.
Locked objects	The iSeries software invokes the save-while-active option when performing backup operations. These operations require a brief lock in order to reach a stable checkpoint. An object with a prolonged conflicting lock may not be able to reach a valid checkpoint. When a library contains an object that fails to reach a checkpoint the default behavior is to skip the entire library. In this case, it is logged that *N* files were not saved, but the names of specific files skipped cannot be determined. To change this behavior, in the Unitrends UI select **Configure > Appliances > Edit > Advanced > General Configuration** and change the iSeriesAgent PreCheck setting from *1* (default) to *0*. With this change, only objects that failed to reach a checkpoint are skipped and the remainder of the library is backed up. If an object is consistently skipped in this manner it may be a protected system object. In this case it can only be backed up in a restricted state.
Pseudo objects	Be sure to carefully configure your iSeries backups to exclude active system files. • /Security Data – Contains the save file from the **SAVSECDTA** command. This object is included in backups. If you do not want to back up or recover this object, you must exclude it when creating the iSeries profile. Unless excluded, it will

Item	Requirement or consideration
	always be the first object in the backup file. During a full recovery, it is the first object recovered (via the **RSTUSRPRF** command), then a **RSTAUT** command is executed after everything else has been recovered. • /System Configuration – Contains the save file from the SAVCFG command. This object is included in backups. If you do not want to back up or recover this object, you must exclude it when creating the iSeries profile. Unless excluded, it is recovered before any other objects, except /Security Data, during a full recovery. It is recovered using the **RSTCFG** command.
Wildcard support	Supported wildcards include: • *: Zero or more characters • ?: Exactly one character • [abc]: Exactly one character from list • [a-c]: Exactly one character from range • [!abc]: Exactly one character not from list • [!a-c]: Exactly one character not in the range Wildcards cannot be used in these cases: • Backup Include List: Object Name • Backup Exclude List: Path Name • Backup Exclude List: Object Name • Backup Include List: Path Name • Any recovery
User privileges	The user performing the backup or recovery job must, at a minimum, have *SECADM privileges added to their profile.
File attributes	Files to be recovered must have read-write attributes. This is accomplished on the OS400 operating system by granting object authority to the user performing the restore command. Following is an example of modifying security privileges in the QGPL and QUSRSYS libraries for user QSECOFR: **# GRTOBJAUT OBJ(QGPL/*ALL) OBJTYPE(*ALL) USER(QSECOFR) AUT(*ALL)** **# GRTOBJAUT OBJ(QUSRSYS/*ALL) OBJTYPE(*ALL) USER(QSECOFR) AUT(*ALL)**

Item	Requirement or consideration
Encryption and compression	iSeries backups are not encrypted on the appliance and backups are compressed post-transmission.

Managing iSeries assets

Use these procedures to manage iSeries assets:

- "To add an iSeries asset"
- "To edit an iSeries asset"
- "To edit retention settings for an iSeries asset" on page 772
- "Removing an iSeries asset" on page 772

To add an iSeries asset

To add the iSeries asset, create an iSeries profile using this procedure:

1. If the iSeries is not accessible via DNS, add the iSeries to the hosts file of the appliance as described in "To view or edit the hosts file" on page 110.
2. Using a terminal emulator, such as PuTTY, connect to the appliance using the following:
   - Appliance IP address
   - Port 22
   - SSH connection type
3. Log in as user *root*.
4. Enter the following command to access the console menu:

   ```
   # dpuconfig
   ```
5. Select option **4** for Advanced Options.
6. Select option **2** for IBM iSeries Backup and Recovery.
7. Select option **1** for Create iSeries Profile.
8. Follow the prompts on the screen to create your profile.

   To specify files to included or exclude, enter the full pathname of the file. The syntax you enter is case sensitive. Example: */QSYS.LIB/RNISSIMOV.LIB*
9. The iSeries is added to the appliance and can be viewed in the appliance UI. To start protecting the iSeries, proceed to "Creating iSeries backup jobs" on page 773.

To edit an iSeries asset

Use this procedure to edit these settings: asset name, backup strategy, IP, or credentials. You will be required to create a new profile or overwrite an existing profile.

1. Using a terminal emulator, such as PuTTY, connect to the appliance using the following:
   - Appliance IP address
   - Port 22
   - SSH connection type
2. Log in as user *root*.
3. Enter the following command to access the console menu:

   ```
   # dpuconfig
   ```
4. Select option **4** for Advanced Options.
5. Select option **2** for IBM iSeries Backup and Recovery.
6. Select option **1** for Create iSeries Profile.

   To specify files to included or exclude, enter the full pathname of the file. The syntax you enter is case sensitive. Example: */QSYS.LIB/RNISSIMOV.LIB*
7. Follow the prompts on the screen to update the profile.

To edit retention settings for an iSeries asset

1. In the appliance UI, select **Configure > Protected Assets**.
2. Select the desired iSeries asset and click **Edit**.
3. Click **Manage Retention** and edit the settings.
4. Click **Save**.

Removing an iSeries asset

CAUTION! When an asset is removed, all associated backups of that asset are also deleted. Please use caution when removing an asset.

Preparing to remove an asset

Before removing an asset, you must

- Remove the iSeries asset from any backup copy job schedules by using the appliance UI. See "To view or edit a backup copy job" on page 579 for details.
- Remove the iSeries asset from any backup schedules by using the dpuconfig menu-based console.

To remove an asset

1. In the appliance UI, select **Configure > Protected Assets**.

2 Select the asset you want to remove.

3 Click **Remove > Confirm**.

Creating iSeries backup jobs

Before running jobs, be sure to add the iSeries asset as described in "To add an iSeries asset" on page 771. Start protecting your iSeries by creating a backup schedule or running an on-demand backup job, as described in these topics:

- "To create an iSeries backup schedule" on page 773
- "To run an iSeries backup on-demand" on page 774

To create an iSeries backup schedule

iSeries backup jobs are created through the menu-based console. There can only be one backup or recovery job running at a time. Be sure to schedule the job to run at a time when no other jobs will be running.

1 Using a terminal emulator, such as PuTTY, to connect to the appliance using the following:
   - Appliance IP address
   - Port 22
   - SSH connection type

2 Log in as user *root*.

3 Enter the following command to access the console menu:

```
# dpuconfig
```

4 Select option **4** for Advanced Options.

5 Select option **2** for IBM iSeries Backup and Recovery.

6 Select option **4** to Schedule Backup.

7 Select option **1** to select the profile to use for your scheduled backup job. Follow the prompts to select a profile or to create a new profile.

   To specify files to included or exclude, enter the full pathname of the file. The syntax you enter is case sensitive. Example: */QSYS.LIB/RNISSIMOV.LIB*

8 Follow the prompts to set additional schedule options, such as days and times the job will run, and to save the schedule.

9 (Optional) Set up a backup copy job to copy iSeries backups to an off-appliance target. For details see:
   - "Backup copy targets" on page 214 to add a backup copy target to your backup appliance.
   - "Creating backup copy jobs" on page 491 to create a job to copy iSeries backups to the target.

To run an iSeries backup on-demand

Notes:

- Only one backup or recovery job can be running at any given time. Ensure that there are no jobs running before creating an on-demand backup job.
- To run an on-demand job using the default iSeries profile, you can use the procedure below or run the job from the appliance UI by selecting **Configure > Appliances > Edit > Advanced > Support Toolbox > On-Demand iSeries Backup**.

1. Using a terminal emulator, such as PuTTY, to connect to the appliance using the following:
   - Appliance IP address
   - Port 22
   - SSH connection type
2. Log in as user *root*.
3. Enter the following command to access the console menu:

   ```
   # dpuconfig
   ```
4. Select option **4** for Advanced Options.
5. Select option **2** for IBM iSeries Backup and Recovery.
6. Select option **2** for Backup iSeries.
7. Follow the prompts to create and run the job.

   To specify files to included or exclude, enter the full pathname of the file. The syntax you enter is case sensitive. Example: */QSYS.LIB/RNISSIMOV.LIB*

Chapter 12: Recovery Overview

Unitrends' recovery features ensure that you can recover data quickly and easily. You can recover files, databases, entire assets, or perform an instant recovery.

Note: Your Unitrends user account determines which recovery procedures you can run. Procedures that you cannot run do not display or are disabled in the UI.

To meet low RTOs, recover from local backups on the Unitrends appliance. For details on recovering from a given local backup, see the Recovery section for the applicable backup type:

- "Recovering Host-level Backups" on page 793
- "Recovering File-level Backups" on page 925
- "Recovering Windows Image-level Backups" on page 1031
- "Recovering Application Backups" on page 1147
- "Recovering NAS Backups" on page 1121
- "Recovering iSeries Backups" on page 1205

If a local backup is not available, you can recover from a backup copy as described in "Recovering Backup Copies" on page 777.

This page is intentionally left blank.

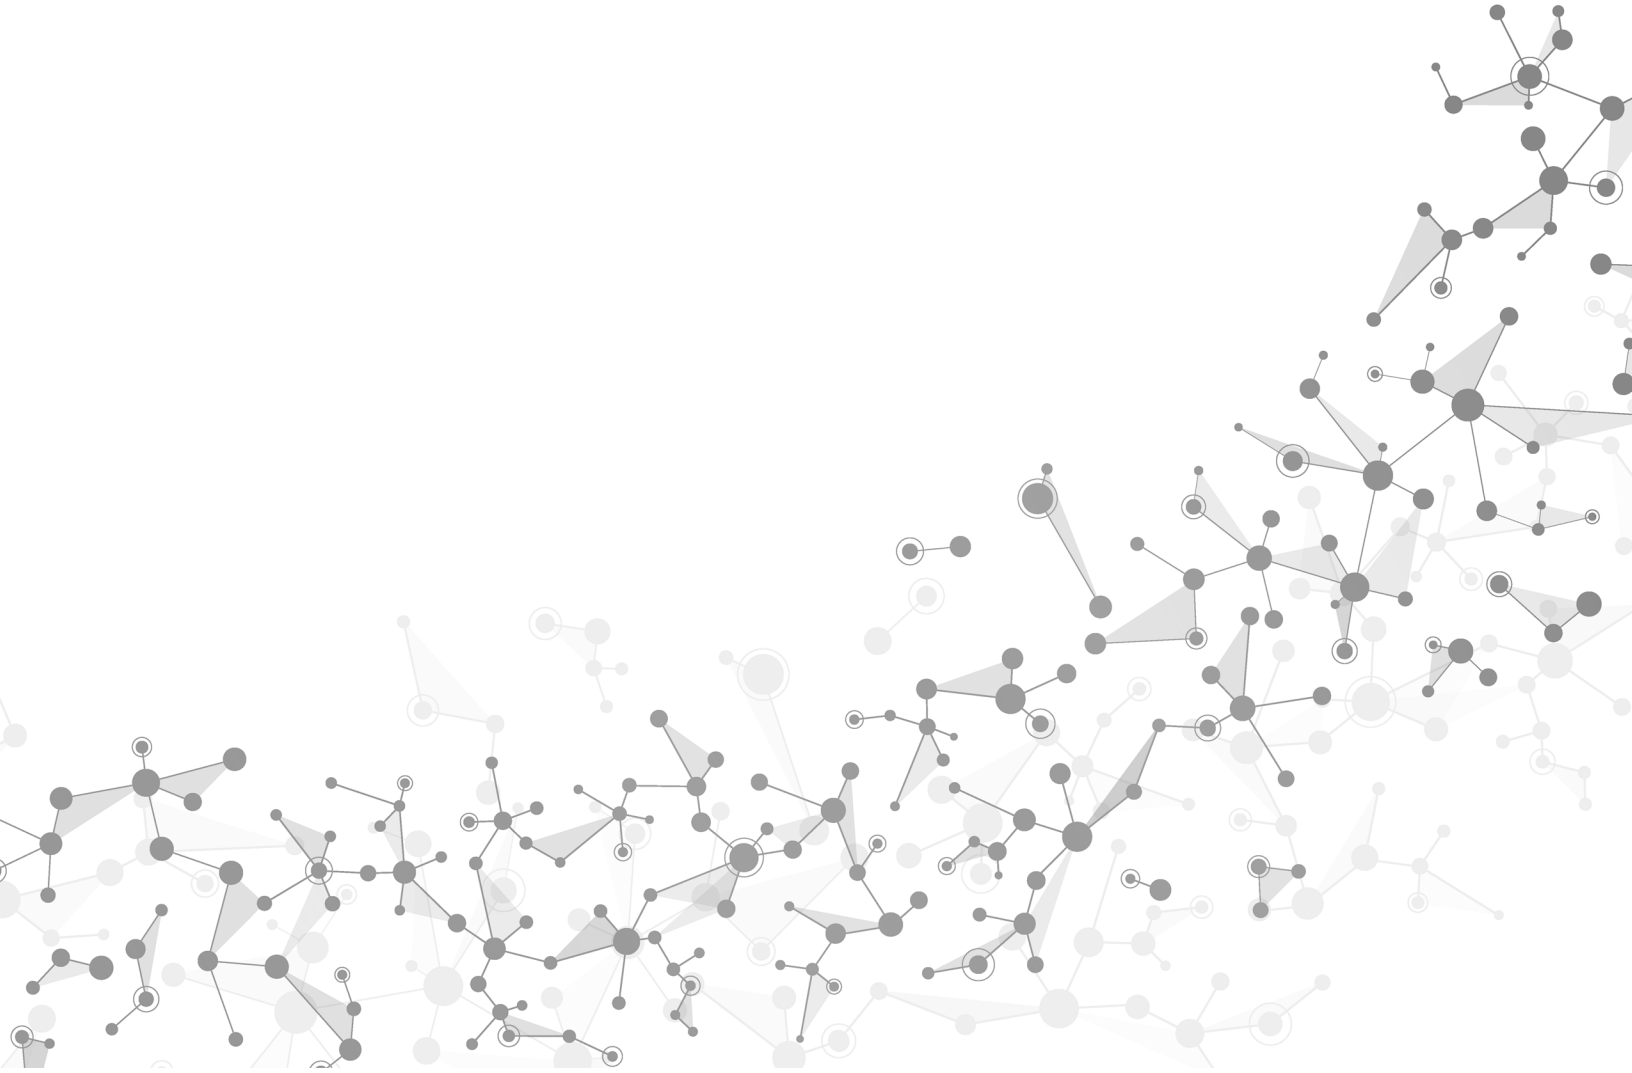

Chapter 13: Recovering Backup Copies

If a local backup is not available, you can recover from a backup copy. In most cases, the same recovery operations that are used for local backups are supported for backup copies, but recovering from a backup copy requires additional steps. These steps vary depending on the backup copy target. See the following table for procedures by backup copy target:

Backup copy target	Recovery procedure
Unitrends Cloud	See "Recovering hot copies by using the source backup appliance".
Managed service provider	Contact the service provider.
Unitrends appliance	See the following: • "Recovering hot copies by using the source backup appliance". • "Recovering hot copies by using the target appliance" on page 784.
Third-party cloud (Amazon, AWS, Google, or Rackspace)	See "Recovering cold backup copies".
NAS	See "Recovering cold backup copies".
FC	See "Recovering cold backup copies".
iSCSI	See "Recovering cold backup copies".
Attached disk	See "Recovering cold backup copies".
Tape	See "Recovering cold backup copies".

Recovering hot copies by using the source backup appliance

You can run procedures from your source backup appliance to recover files or entire backup copies that reside in the Unitrends Cloud or on your target appliance. After reviewing the considerations and requirements, proceed to one of the recovery procedures to recover the hot backup copy.

Considerations and requirements:

- "How do I use my backup appliance to recover from hot copies that reside on a target appliance or in the Unitrends Cloud?" on page 778
- "Requirements and limitations for recovering hot copies by using the source backup appliance" on page 779

Recovery procedures:

- "To import a hot backup copy" on page 780
- "To recover files from a file-level backup copy by using the File Browser" on page 985
- "To recover files from a file-level backup copy by using Search Files" on page 988
- " Recovering files from virtual machine backups" on page 808
- "Recovering files by browsing a Windows image-level backup" on page 1040

How do I use my backup appliance to recover from hot copies that reside on a target appliance or in the Unitrends Cloud?

To recover hot copies, you either import the backup copy to the source appliance or recover files directly from the hot copy in the Unitrends Cloud or on the appliance target. Recovery procedures vary by backup type and whether the copy has been imported to the source backup appliance. See the following for details:

Hot backup copy type and location	Recovery options
All types, imported to source backup appliance	Import the backup copy from the Cloud or appliance target to the source appliance, as described in "To import a hot backup copy" on page 780. In most cases, once the backup copy has been imported, you can recover from it just like you would from any other local backup. See the applicable recovery chapter in this guide for recovery procedures.
File-level hot backup copy that resides on an appliance target or in the Unitrends Cloud	Use the following options to recover files directly from the backup copy. Before you start, review the "Requirements and limitations for recovering hot copies by using the source backup appliance". • Browse a backup copy and download selected file(s) in a *.zip* file. For details, see "To recover files from a file-level backup copy by using the File Browser" on page 985. • Search the backup copy for files and download selected file(s) in a *.zip* file. For details, see "To recover files from a file-level backup copy by using Search Files" on page 988.
Image-level hot backup copy that resides on an appliance target or	Browse a backup copy and download selected file(s) in a *.zip* file. Before you start, review the "Requirements and limitations for recovering hot copies by using the source backup appliance". For details, see "Recovering files by browsing a Windows image-level backup" on page 1040.

Hot backup copy type and location	Recovery options
in the Unitrends Cloud	
Host-level hot backup copy that resides on an appliance target or in the Unitrends Cloud	For host-level backups of Windows or Linux VMs, recover files directly from the backup copy by downloading selected file(s) in a *.zip* file. Before you start, review the "Requirements and limitations for recovering hot copies by using the source backup appliance". For details, see " Recovering files from virtual machine backups" on page 808.

Requirements and limitations for recovering hot copies by using the source backup appliance

The following requirements and limitations apply:

Requirement or limitation	Description
Appliance	The source appliance must meet the following requirements: • Must be running release 9.1 or higher (10.3 or higher for Windows image-level backup copies). • Must be the source appliance that copied the backup to the Unitrends Cloud or target appliance. You cannot recover backup copies that were created by another source appliance. • Must have the Unitrends Cloud or target appliance configured as a backup copy target. The target appliance must be running release 9.1 or higher (10.3 or higher for Windows image-level backup copies).
Recovering files from the Unitrends Cloud or appliance target	You can recover files directly from backup copies that reside in the Unitrends Cloud or on a target appliance. The selected file(s) are downloaded in a *.zip* file. See these rows below for requirements: • Supported backup types – The backup that was copied must be one of the following: – A file-level backup (run using a Unitrends agent). File recovery is supported for most agent-based assets. While running the recovery procedure, eligible assets are presented in the UI. – An image-level backup of a Windows asset (run using the Unitrends Windows

Requirement or limitation	Description
	agent). – A VMware, Hyper-V, AHV, or XenServer host-level backup of a Windows or Linux VM. • Other requirements – A *.zip* file is created in the default download location of the browser where you are running the appliance UI. Once the browser presents the *.zip* file, you can extract the downloaded files. Persistent browser and UI sessions are required during file recovery. While the recovery is in progress, closing the browser or logging out of the appliance UI prevents *.zip* file creation in the browser's default download location. If you close the browser or UI session during the recovery and the *.zip* file is not created, you must run a new recovery job.
Recovering entire backup copies	To recover entire backup copies, you must first import the backup copy to the source appliance. Once imported, you recover from the imported backup copy as you would from any other local backup. **Notes:** • Importing a backup copy from the Cloud or appliance target can take quite some time. The duration of the import is impacted by various factors, such as the size of the backup copy, bandwidth, and download speed. • Persistent browser and UI sessions are NOT required while importing a backup copy from the Cloud or appliance target. While an import is running, closing the browser or UI does not impact the job.
Maximum number of recovery jobs in the Unitrends Cloud	The number of active recovery jobs that can run simultaneously on a Unitrends Cloud target is limited to prevent the target from becoming overloaded. If you attempt a recovery and this limit has been met, a message displays indicating that the recovery task will not start. If you see this message, try again later.

To import a hot backup copy

1. Log in to the source appliance.
2. Click **Recover > Backup Catalog**.
3. In the Filter Backups area to the right, select **Backup Copy (Hot)** in the Type list.

   You must select this Type to view backup copies in the Unitrends Cloud or on a target appliance.
4. (Optional) Enter other filter options. For details, see "Working with custom filters" on page 67.

5 Click **Filter**.

- Assets with backup copies meeting the filter options you specified display in the Backup Catalog list.
- Expand an asset to view its backup copies.
- Backup copies that reside in the Unitrends Cloud or on a target appliance are purple and the description *Backup Copy (Hot) on Target* displays when you hover over the backup copy icon.
- If your source appliance is also being used as a backup copy target appliance, the catalog lists both the copies of local backups that are stored on the remote Unitrends Cloud or appliance target and the hot backup copies that are stored on this appliance (that were received from another appliance). Hover over the backup copy icon to determine whether this backup copy resides on the remote target or on this appliance. *Backup Copy (Hot)* indicates that the backup copy is stored on this appliance. *Backup Copy (Hot) on Target* indicates that the backup copy resides in the Cloud or on the remote appliance target.

6 Select the *Backup Copy (Hot) on Target* copy to import.

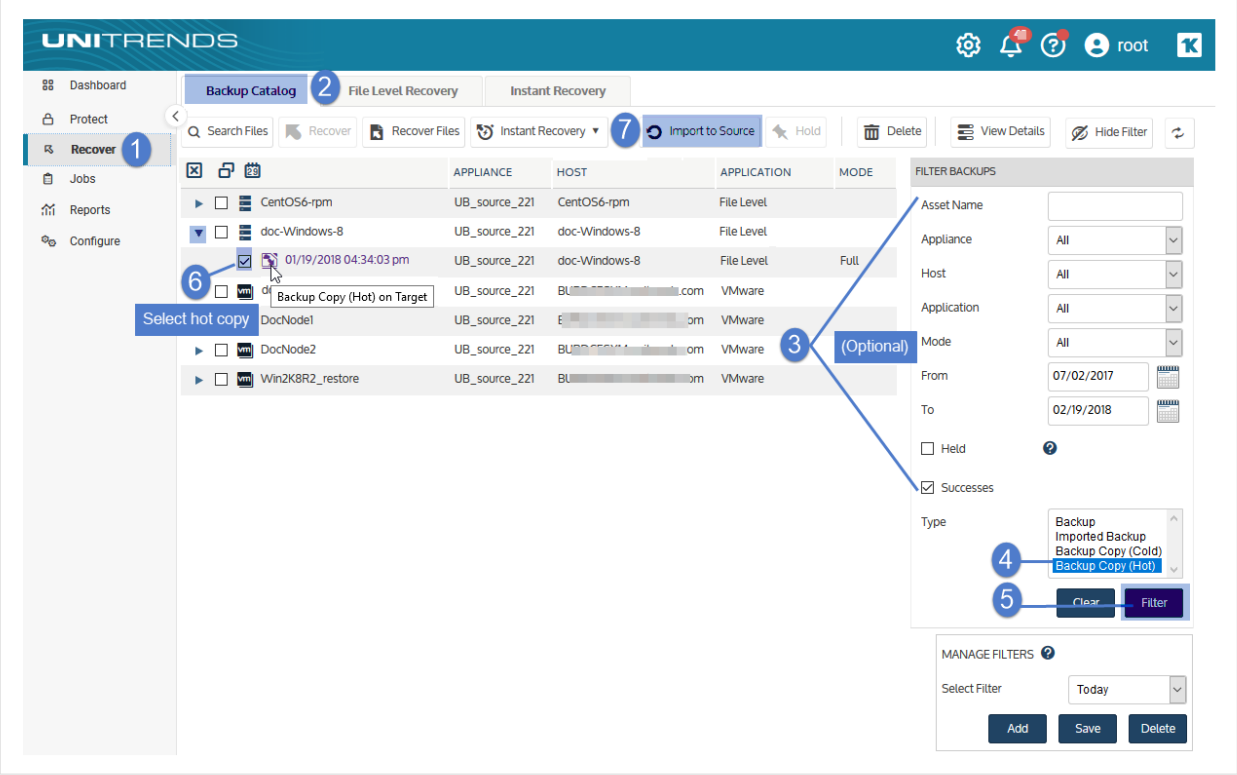

7 Click **Import to Source**.

The Import Backup Copies dialog lists the backup copies to import.

Notes:

- Selecting a copy of a full backup imports only the full backup.

- Selecting a copy of an incremental backup imports the full backup and the incrementals up to and including the selected backup.
- Selecting a copy of a differential backup imports the differential and the associated full backup.
- For more information about backup groups, see "Backup groups" on page 98.

8 Click **Import**.

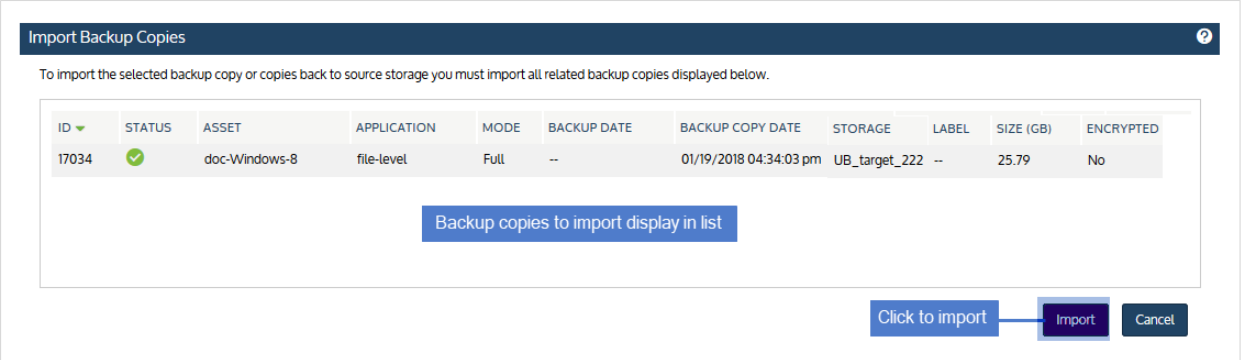

9 Selected backup copies are imported to the appliance. Click **View Jobs** to monitor the status of the import.

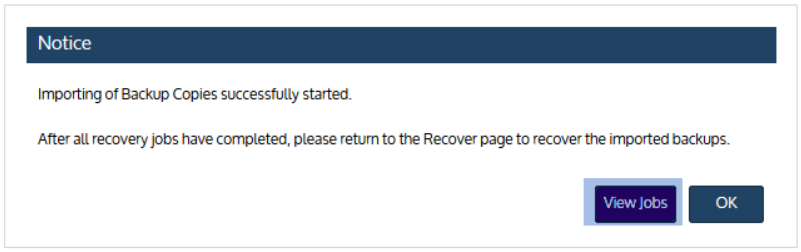

10 Select the job and click **View Details**.

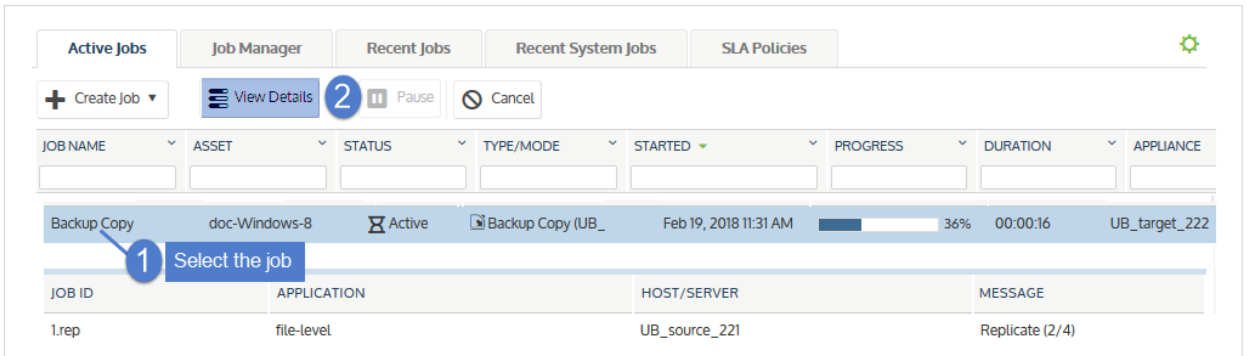

The import is complete when the job status changes to Success:

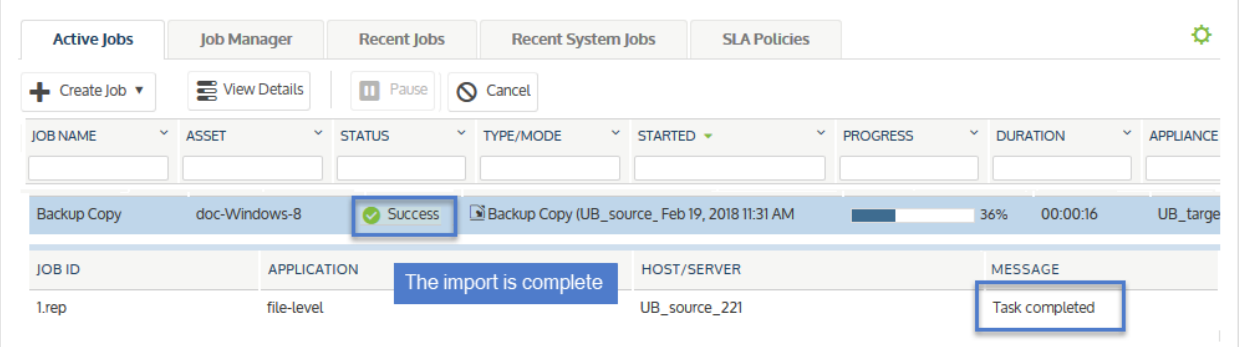

11 (Optional) Recover from the imported backup copy as you would from any regular backup.

Notes:

- The appliance purges the oldest backups when space is needed. Because imported backup copies are often older than others on the appliance, they are retained for 72 hours before becoming eligible for purging. Be sure to recover from imported backups within the first 72 hours.
- Imported SQL transaction log and differential backups must be recovered as the original name and to the original location.

- Filter the display to view imported backups (Type = Imported Backup).
- Imported hot backup copies are purple and the description *Imported Backup from Target* displays when you hover over the backup copy icon.

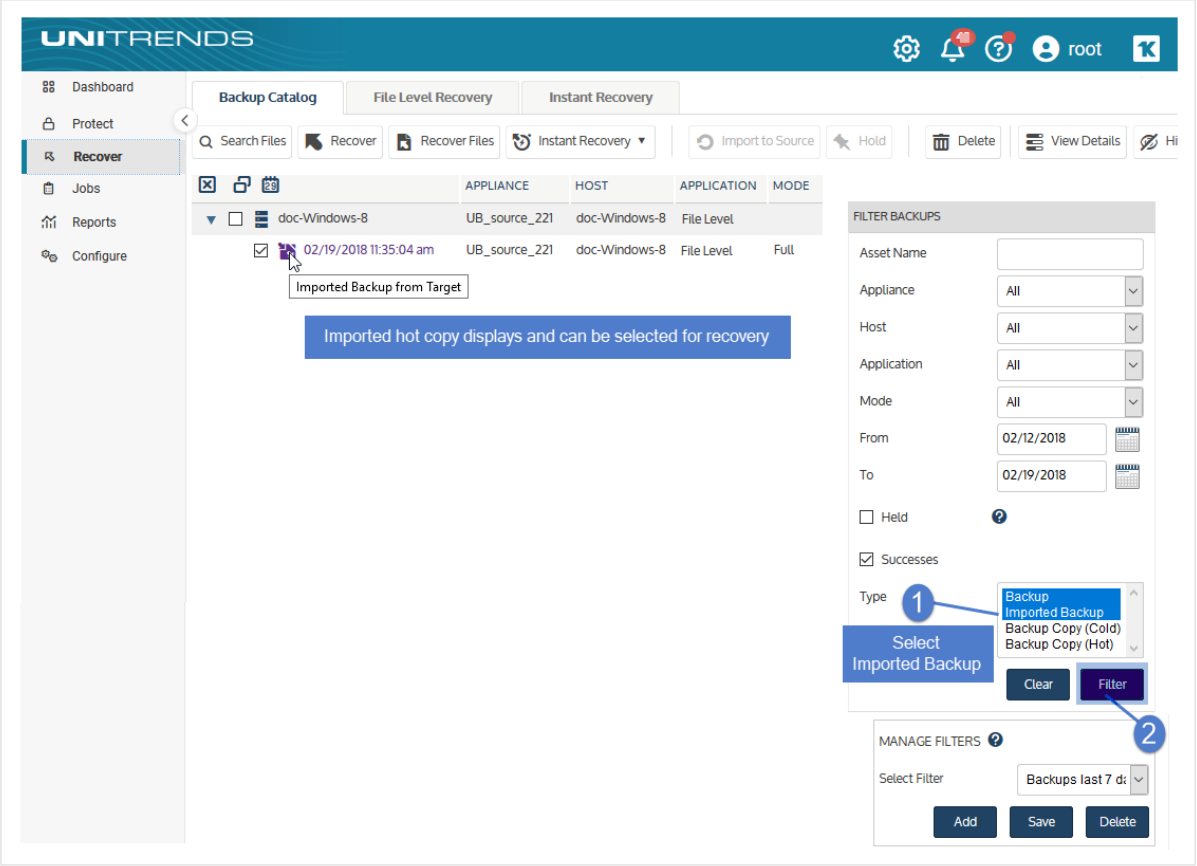

Recovering hot copies by using the target appliance

You can recover from hot backup copies that are stored on a Unitrends target appliance by running the standard recovery procedures on the target appliance. Use the applicable recovery procedures, but run them from the target appliance and select a hot backup copy (instead of a regular backup or imported backup). For details on viewing hot backup copies, see "To view the hot backup copies stored on the target appliance" below.

Recovery procedures require that you select a target asset where the backup copy will be recovered. Only assets that have been added to the target appliance can be used as recovery targets. Be sure to add the desired target asset before running the recovery procedure. (See the adding assets procedures in "Managing protected assets" on page 286 for details.)

Once you've added the target asset, proceed to one of these topics for detailed recovery procedures:

- "Recovering Host-level Backups" on page 793
- "Recovering File-level Backups" on page 925
- "Recovering Windows Image-level Backups" on page 1031
- "Recovering Application Backups" on page 1147

- "Recovering NAS Backups" on page 1121

Note: For iSeries, you cannot recover directly from a hot backup copy. Instead, you must import the backup copy to the source appliance. See "To import a hot backup copy" on page 780 for details.

To view the hot backup copies stored on the target appliance

1. Log in to the target appliance and select **Recover > Backup Catalog**.
2. In the Filter Backups area to the right, select **Backup Copy (Hot)** in the Type list.
3. (Optional) Enter other filter options. For details, see "Working with custom filters" on page 67.
4. Click **Filter**.
   - Assets with backup copies meeting the filter options you specified display in the Backup Catalog list. The source appliance where the backup originated displays in the Appliance column.
   - Expand an asset to view its backup copies.
   - Hot backup copies are purple and the description *Backup Copy (Hot)* displays when you hover over the backup copy icon,
   - If your target appliance is also being used as a backup appliance and its local backups are being copied to a hot backup copy target, the catalog lists both the hot backup copies stored on this appliance and any backups that were copied from this appliance to the hot backup copy target. (The hot backup copy target could be another appliance or the Unitrends Cloud).
     - You can recover the backup copies that are stored on this appliance as you would any other local backup.
     - To determine whether the backup copy is stored on this appliance, hover over the backup copy icon to display more information. If the backup copy is labeled *Backup Copy (Hot)*, it can be recovered using the standard backup recovery procedures. If the backup copy is labeled *Backup Copy (Hot) on Target*, you must recover it using the procedures in "Recovering hot copies by using the source backup appliance" on page 777.

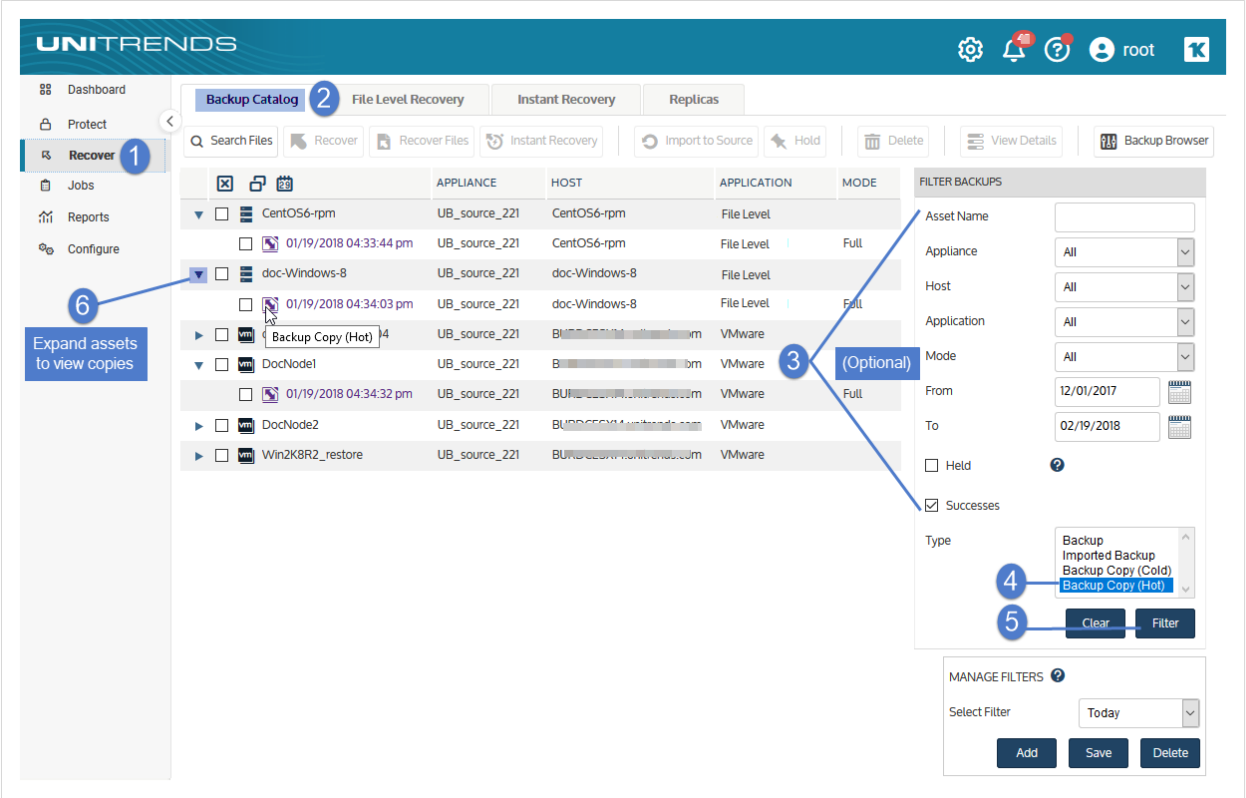

Recovering cold backup copies

Backup copies stored on external media are known as *cold backup copies*. Cold backup copies reside on cloud storage managed by third-party vendors or on other off-site targets, such as eSATA, tape, and NAS devices.

Before you can recover from a cold backup copy, you must import the data from the backup copy target to the source backup appliance. You can either import the entire backup copy or import selected files (supported for file-level backup copies only):

- To import the entire backup copy, see "To import a cold backup copy". Once you have imported the backup copy, it displays in the Backup Catalog with the label *Imported Backup*. You can then select this backup and follow the same recovery steps you would use for recovering from a local backup.
- To import selected files from a file-level backup copy, see the following procedures. With these procedures, the appliance creates and imports a selective backup containing the files you picked:
  - "Recover files from a cold backup copy by using Search Files" on page 957
  - "Recover files from one cold backup copy by using the File Browser" on page 966

To import a cold backup copy

Use this procedure to import a backup copy that currently resides on the cold backup copy target.

1. Verify the following:
   - The backup copy target is connected and accessible. To check this:
     - On the **Configure > Appliances** page, select the source backup appliance.
     - Click the **Backup Copy Targets** tab below.
     - Click **Scan For Media**. The target displays on the Backup Copy Targets tab.
       - For NAS devices, the target must be in Online status to import the copy. If necessary, select the target and click **Enable** to bring the target online.
       - For removable media, such as USB, eSTATA and tape devices, the target can be in Offline or Online status to import the copy.
   - If you are using removable media, the tape(s) or disk(s) where the backup copy is stored must be loaded in the target.
     - If the job copied to multiple drives or tapes, be sure to load all drives or tapes that were loaded when the backup copy job ran. The appliance writes across all drives or tapes and all must be present to perform the import.
     - If you are using tapes that are not labeled with barcodes, each tape must be inserted into the slot where it resided during the backup copy job.
2. Log in to the backup appliance.
3. Click **Recover > Backup Catalog**.
4. In the Filter Backups area to the right, select **Backup Copy (Cold)** in the Type list.
5. (Optional) Enter other filter options. For details, see "Working with custom filters" on page 67.
6. Click **Filter**.
   - Assets with backup copies meeting the filter options you specified display in the Backup Catalog list.
   - Expand an asset to view its backup copies.
7. Click to select the backup copy.

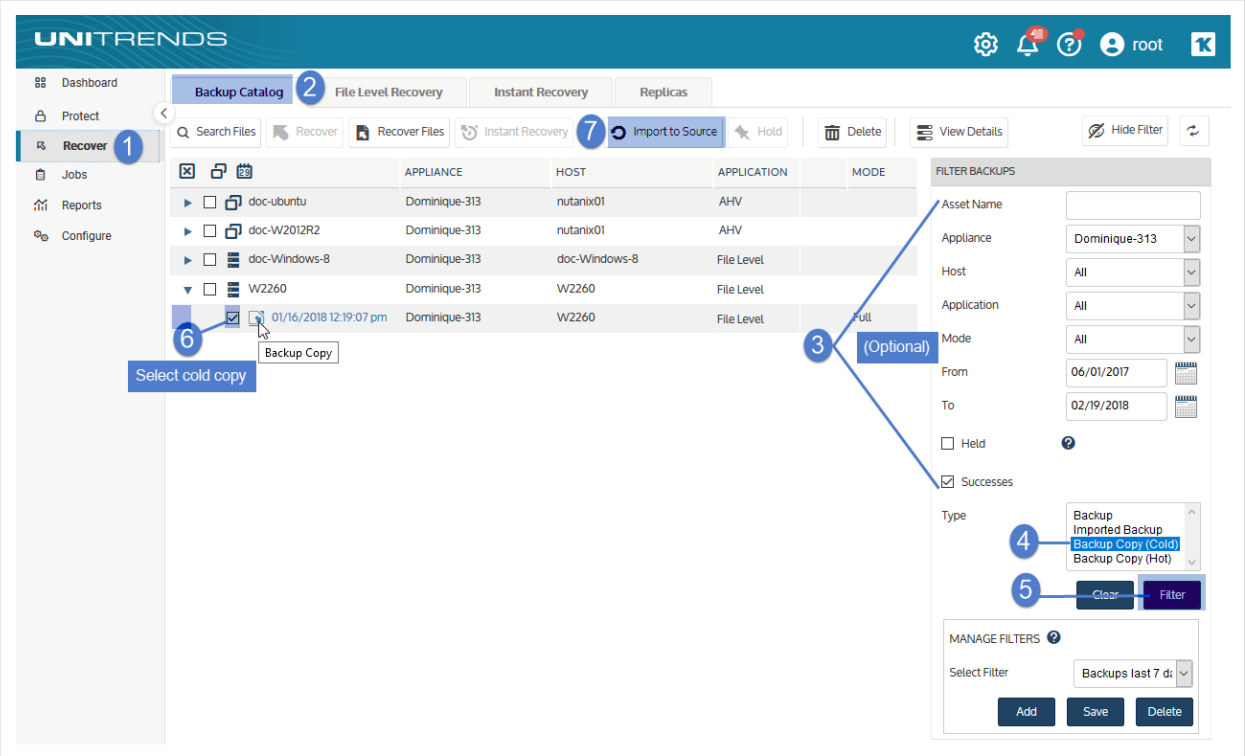

8 Click **Import to Source**.

The Import Backup Copies dialog lists the backup copies to import.

Notes:

- Selecting a copy of a full backup imports only the full backup.
- Selecting a copy of an incremental backup imports the full backup and the incrementals up to and including the selected backup.
- Selecting a copy of a differential backup imports the differential and the associated full backup.
- For more information about backup groups, see "Backup groups" on page 98.

9 Click **Import**.

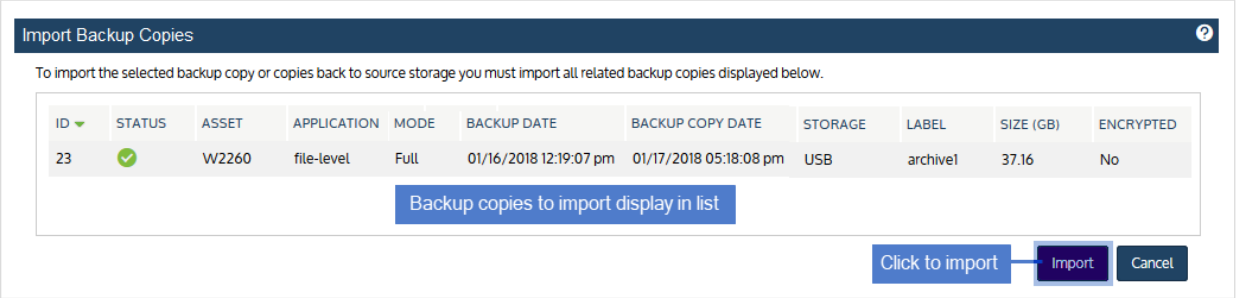

10 Selected backup copies are imported to the appliance. Click **View Jobs** to monitor the status of the import.

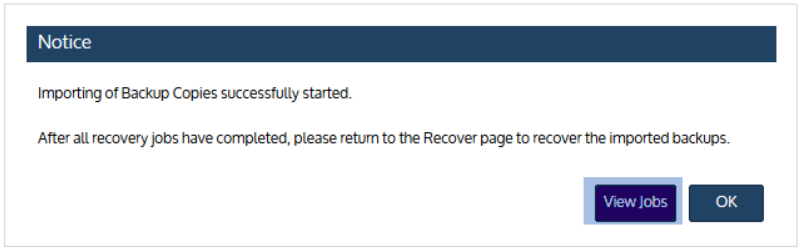

11 Select the job and click **View Details**.

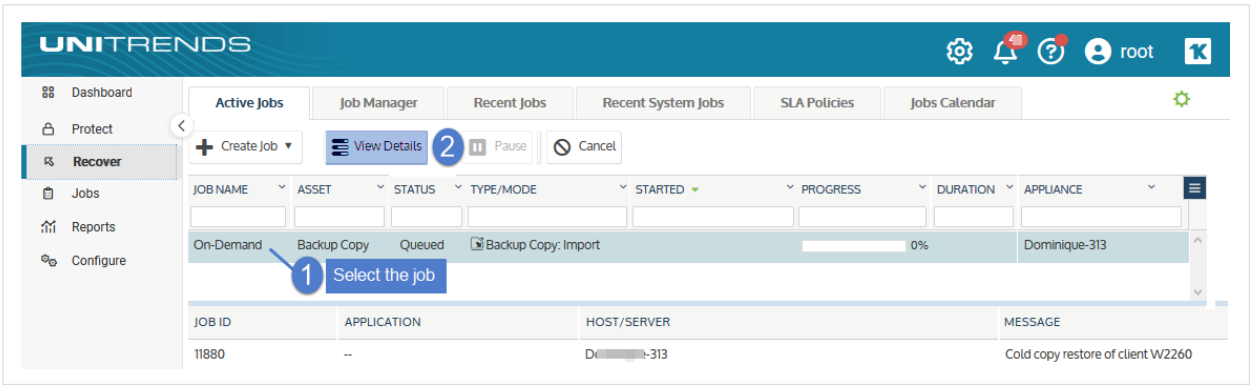

The import is complete when the job status changes to Success:

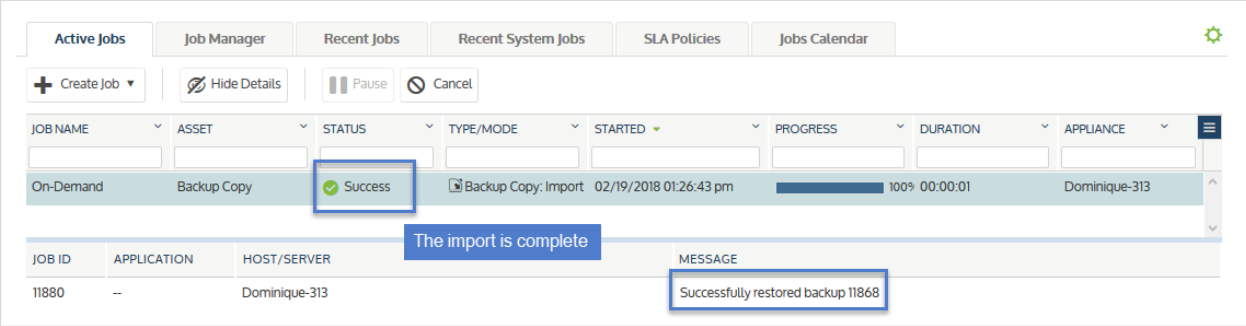

12 (Optional) Recover from the imported backup copy as you would from any regular backup.

Notes:

- The appliance purges the oldest backups when space is needed. Because imported backup copies are often older than others on the appliance, they are retained for 72 hours before becoming eligible for purging. Be sure to recover from imported backups within the first 72 hours.
- Imported SQL transaction log and differential backups must be recovered as the original name and to the original location.

- Filter the display to view imported backups (Type = Imported Backup).
- Imported cold backup copies are light blue and the description *Imported Backup* displays when you hover over the backup copy icon, as shown here:

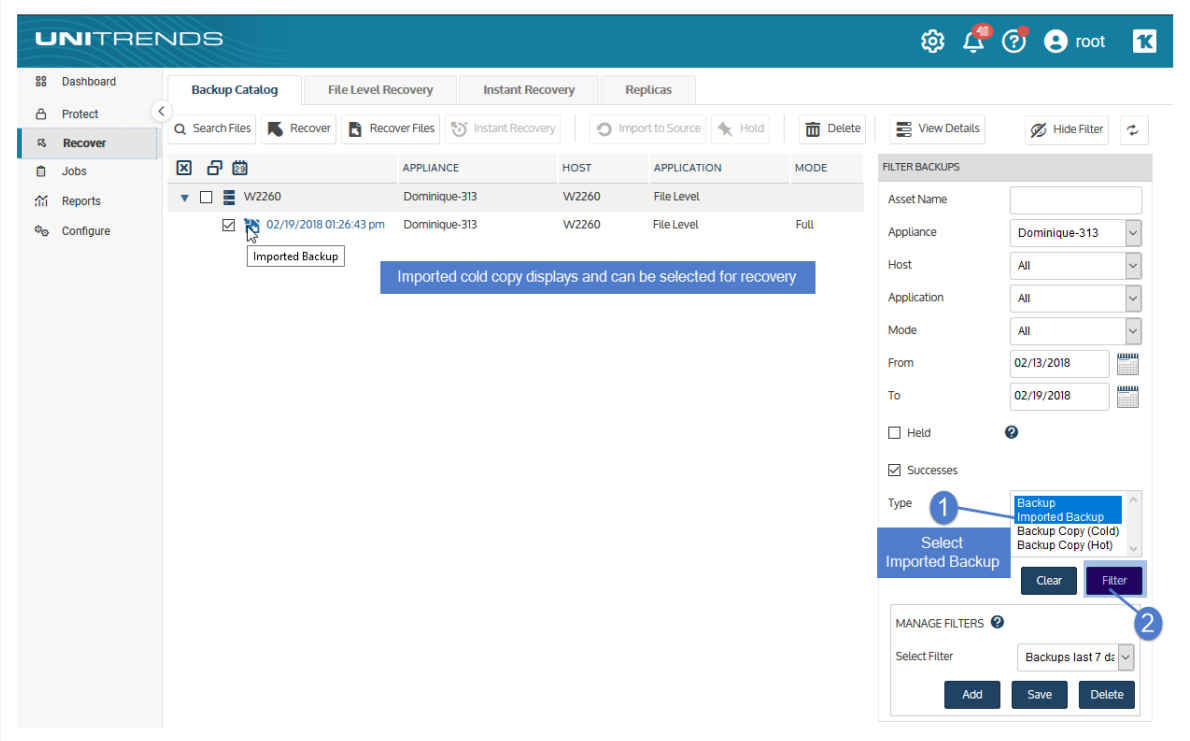
UNITRENDS
root
Dashboard
Protect
Recover
Jobs
Reports
Configure
Backup Catalog
File Level Recovery
Instant Recovery
Replicas
Search Files
Recover
Recover Files
Instant Recovery
Import to Source
Hold
Delete
View Details
Hide Filter
APPLIANCE
HOST
APPLICATION
MODE
W2260
Dominique-313
W2260
File Level
02/19/2018 01:26:43 pm
Dominique-313
W2260
File Level
Full
Imported Backup
Imported cold copy displays and can be selected for recovery
FILTER BACKUPS
Asset Name
Appliance
Dominique-313
Host
All
Application
All
Mode
All
From
02/13/2018
To
02/19/2018
Held
Successes
Type
1
Backup
Imported Backup
Backup Copy (Cold)
Backup Copy (Hot)
Select
Imported Backup
Clear
Filter
2
MANAGE FILTERS
Select Filter
Backups last 7 da
Add
Save
Delete

This page is intentionally left blank.

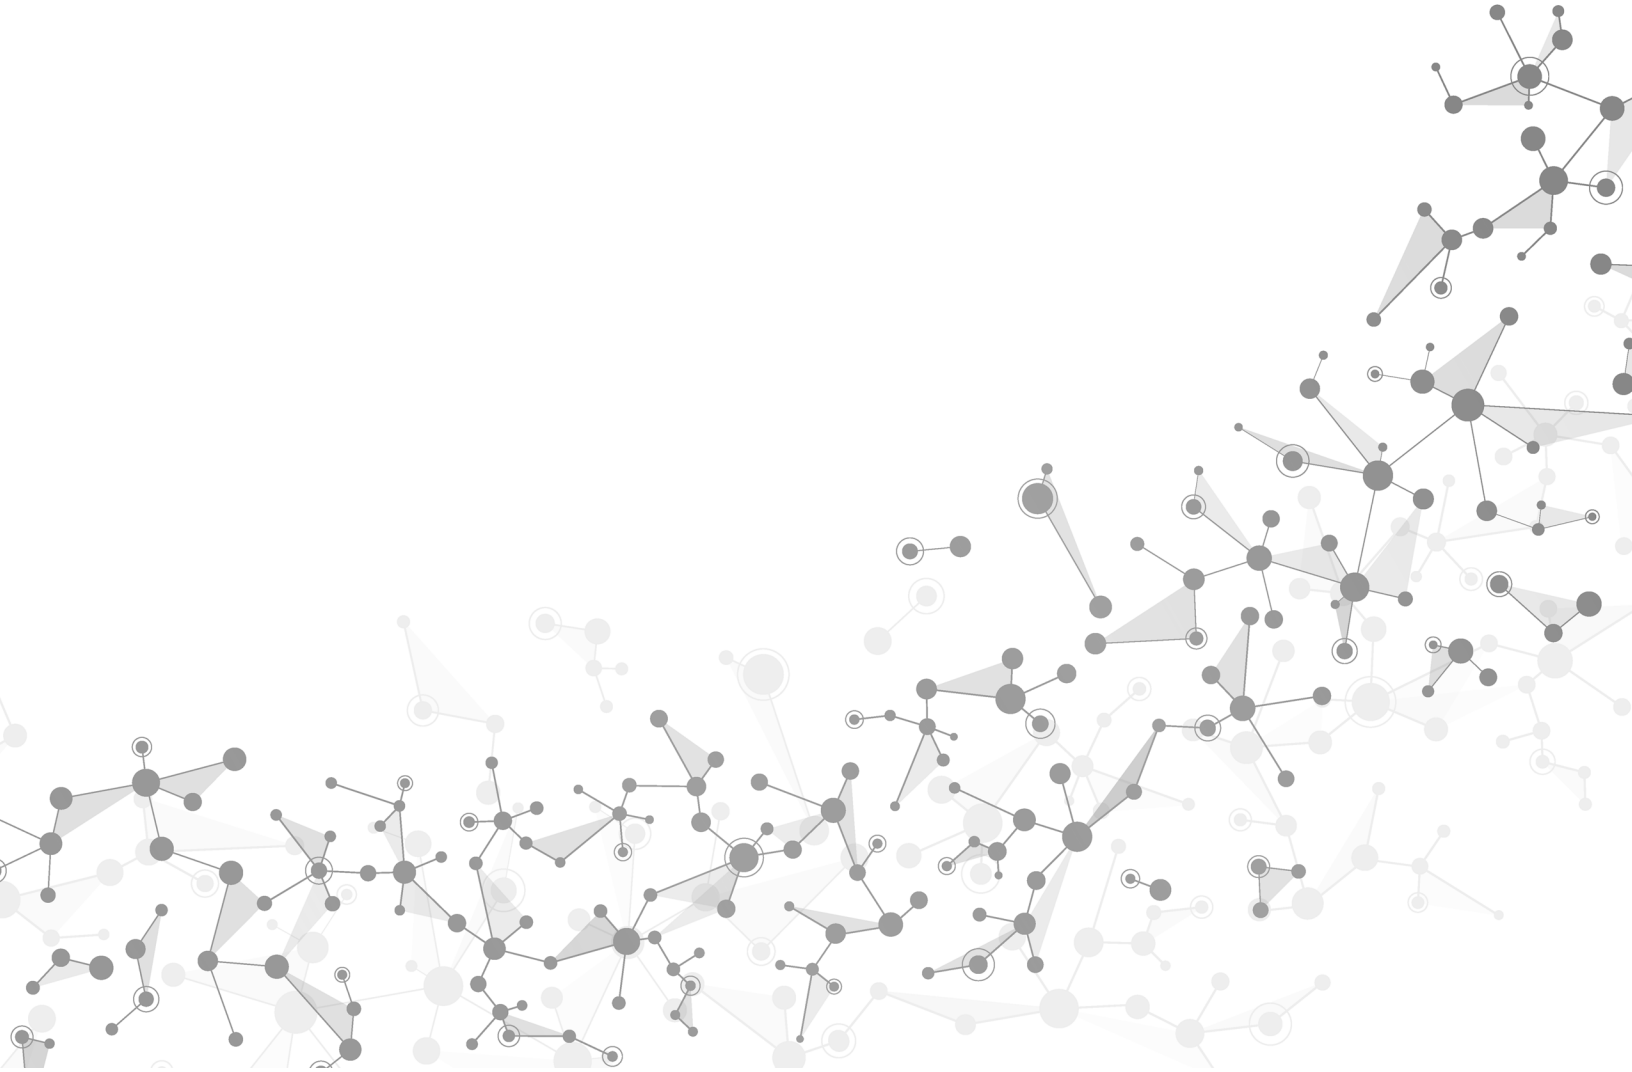

Chapter 14: Recovering Host-level Backups

Unitrends provides a variety of methods for recovering host-level backups of virtual machines. You can recover entire virtual machines or selected files from backup. For quick recovery of critical VMs, you can create virtual machine replicas or use the instant recovery feature.

See the following table for descriptions of each recovery method:

Recovery method	Description
Virtual machine recovery	Recovers the entire virtual machine from its host-level backup to a target virtual host, which can be the original host or an alternate host. Supported for host-level backups of all VM guest operating systems. For detailed requirements and recovery procedures, see "Recovering a virtual machine".
File recovery	Recovers selected files from a backup of a Windows or Linux VM. For detailed requirements and recovery procedures, see " Recovering files from virtual machine backups".
Virtual machine replicas	Creates a stand-by replica of a VMware VM that you can bring online in minutes. As backups of the original VM run, they are applied to the replica to keep it up-to-date. Take the replica 'live' to assume the role of a failed VM. To meet near-zero RTOs, set up the replica before the VM fails. For detailed requirements and procedures, see "VM replicas" on page 876. For a comparison of the virtual machine replica and instant recovery features, see "Replicas or instant recovery?"
Virtual machine instant recovery	Recovers a failed VMware or Hyper-V VM in minutes. The recovered VM can immediately assume the role of the original, failed machine. Unitrends recommends that you plan for instant recovery (IR) before a VM fails, by reviewing requirements, allocating IR storage, and using audit mode to test IR of your critical VMs. For detailed requirements and procedures, see "Virtual machine instant recovery". Because instant recovery uses appliance resources that can impact the performance of other jobs, Unitrends recommends using virtual machine replicas for VMware VMs and standard VM recovery for non-critical VMs. For a comparison of the virtual machine replica and instant recovery features, see "Replicas or instant recovery?"

Replicas or instant recovery?

Replicas and instant recovery provide ways to recover your critical VMs in minutes. Instant recovery is supported for Hyper-V and VMware. Replicas are supported for VMware only.

For VMware, you can use either feature, or a combination of both, to reduce production downtime. The following table provides a high-level comparison of the virtual machine replicas and instant recovery features. Consider these differences when deciding which feature to use for your critical VMware machines.

Item	VMware replicas	VMware instant recovery (IR)
Scale	Because creating and updating replicas consumes few appliance resources, this feature can be used for 1000s of VMs.	Appliance resources are used to create each IR session. The number of VMs that can be recovered simultaneously is limited by appliance load and available IR space.
Automated and orchestrated recovery	Integrates seamlessly with Copy Data Management to perform automated and orchestrated recovery of a single VM or a subset of the entire virtual infrastructure. For details, see "Recovery Assurance" on page 1263.	Not supported. VM IR must be run manually for each VM.
Recovery Time Objective (RTO)	Fastest, near-zero recovery time. Just click **Go Live** to boot into live mode, then configure network settings to bring the replica online in production to assume the role of the failed VM. **Note:** After creating a replica, you can set up a Copy Data Management job to fully automate testing and failover. This way, you can failover by using pre-configured network settings. An Enterprise Plus license is required to use this feature. For details, see "Recovery Assurance" on page 1263.	Slower recovery time than replica VMs. You select a backup and configure the IR job. The appliance then creates a local disk image and a VM on the virtual host. This can take some time, especially for large incremental backup chains. The appliance powers on the VM and begins migrating data from the local disk image. At this point the VM is fully operational and can assume the role of the failed VM.
Recovery Point Objective (RPO)	Recover from any recovery point stored with the replica VM on the hypervisor. By default, only the latest recovery point is retained. You can modify this setting to retain up to 31 recovery points.	Recover from any host-level backup or imported backup copy on the Unitrends appliance.
On-appliance retention	Does not impact on-appliance retention; all backup storage is used for backups.	Decreases on-appliance retention; a portion of backup storage is allocated to IR.

Item	VMware replicas	VMware instant recovery (IR)
Appliance performance	Little impact to appliance performance. Replicas are created and updated by running regular recovery operations. A replica's compute and storage are allocated by the ESXi host.	VM IR has a greater impact to appliance performance than VM replicas. Each IR session uses appliance storage and compute resources to create and run a disk image from the selected recovery point. This disk image must remain on the appliance until all data has been copied to the VM on the hypervisor. Once data migration is complete, you can tear down the IR session to free up appliance resources.
ESXi host performance	Impact to the ESXi host for a given VM is about the same for a live replica as for a live VM IR. A replica is created on the specified ESXi host as a cold stand-by and remains powered off, even as backups are applied. A portion of the ESXi host's storage and compute resources are allocated to each replica VM, but no compute resources are used until you boot the replica into live or audit mode. Because an appliance can manage 1000s of replicas, it is important that you monitor ESXi host resources carefully and make adjustments as needed.	Impact to the ESXi host for a given VM is about the same for an IR VM as for a live replica. Note that Storage vMotion runs on the hypervisor to copy data from the appliance to the IR VM, which may temporarily impact host performance.
vCenter and Storage vMotion	A replica can be hosted on a stand-alone ESXi server or on a server that is managed by a vCenter. Replicas do not use Storage vMotion and do not require a vCenter. **Note:** To run Copy Data Management jobs for the replica, the ESXi host must be managed by a vCenter.	The IR target must be an ESXi server that is managed by a vCenter and must have a license that supports Storage vMotion.
Setup for multiple VMs	You can select multiple VMs in the Create Replica VMs dialog, to quickly set up replicas for a group	Not supported. VM IR must be run manually for a single VM.

Item	VMware replicas	VMware instant recovery (IR)
	of virtual machines. You can even include VMs that have not yet been backed up. In this case, when a successful backup runs, it is automatically applied to complete the replica creation process for that VM.	

Recovering from a backup copy

If a local backup is not available, you can recover from a backup copy. Procedures vary by backup copy type, as described here:

- To recover from a cold backup copy, you must first import it to the source backup appliance as described in "To import a cold backup copy" on page 786. Once the backup copy has been imported, use the standard host-level recovery procedures to recover files or the entire VM.
- To recover from a Unitrends appliance backup copy, use the standard host-level recovery procedures. If recovering from the source appliance, you must first import the backup as described in "To import a hot backup copy" on page 780. If recovering from the target appliance, you can recover from the hot backup copy directly.
- To recover from a Unitrends Cloud backup copy, you either import the backup copy and use standard host-level recovery procedures or recover files directly from the Unitrends Cloud as described in " Recovering files from virtual machine backups" on page 808.

Recovering a virtual machine

Virtual machine recovery enables you to recover VMs running any operating system. This method restores the entire VM and associated metadata with the configured peripherals, from any given Unitrends recovery point. The appliance uses the backup or backup copy to recreate the VM on the recovery target. The recovery target can be the original VM's host or an alternate host.

You select the backup or backup copy to use for the recovery. With an incremental backup, the appliance uses all dependent backups in the group to recreate the VM.

Recovery time depends on various factors, such as the amount of data on the VM and other tasks running on the appliance. For example, recovering an incremental that has a long chain of dependent incrementals can take extra time.

Preparing to recover a virtual machine

Unitrends supports recovery of VMware, Hyper-V, AHV, and XenServer virtual machines. You can recover the VM to the original host or to another host that has been added to the backup appliance. If necessary, add the target host as described in "Adding a virtual host" on page 308 before recovering a VM.

About recovering VMware VMs

Review the following information on recovering VMware VMs:

- The target ESXi host must be added to the appliance as an asset. See "Adding a virtual host" on page 308.
- The VM must be recovered to an ESXi host running the same version as the original host, or to a higher supported version listed in the Compatibility and Interoperability Matrix.
- The target ESXi host must support the operating system (OS) of the VM you are recovering. (See the VMware documentation for details.) For example, you cannot recover a Windows 2016 VM to ESXi 5.1.
- The target ESXi host must have sufficient space and compute resources for the new VM.
- A recovered VM is configured with the latest hardware version supported by the target ESXi host.
- A recovered VMware VM is created with the following default name: *<original_VM_name>_restore*. You can edit this name when you create the recover job.
- Any virtual-mode raw device mapping (RDM) disks recover as standard virtual disks.

About recovering Hyper-V VMs

Review the following information on recovering Hyper-V VMs:

- The target Hyper-V host must be added to the appliance as an asset. See "Adding a virtual host" on page 308.
- The VM must be recovered to a Hyper-V host running the same version as the original host, or to a higher supported version listed in the Compatibility and Interoperability Matrix.
- The target Hyper-V host must support the operating system (OS) of the VM you are recovering. (See the Hyper-V documentation for details.) For example, you cannot recover a Windows 2016 VM to Hyper-V 2008 R2.
- The target Hyper-V host must have sufficient space and compute resources for the new VM.
- A recovered VM is configured with the latest hardware generation version supported by the target Hyper-V host.
- A recovered VM's configuration version matches that of the source Hyper-V backup used for the recovery.
- A recovered Hyper-V VM is created with the same name as the original VM and no suffix. Due to Hyper-V limitations, it is not possible to rename the VM during the recovery.
- VMs are uniquely identified by GUID (not by VM name), and VMs are recovered with the same GUID as the VM that was backed up. (A recovered VMs GUID matches that of the source Hyper-V backup used for the recovery.)
- If the recovery target is hosting a VM with the same GUID as the one that is being recovered, the original VM is overwritten by the recovery operation.
- If the recovery target is hosting a VM with the same name but a different GUID as the one that is being recovered, the original VM is not overwritten by the recovery operation. Once the recovery operation completes, the host has two VMs with the same name (but different GUIDs).
- VM configured with the Hyper-V native replication feature – Native replication settings may not be recovered. After recovery, use Hyper-V Manager to check whether replication settings were recovered. If not, delete any existing replica VM from the replica server and re-enable replication for the recovered VM on the primary server.

About recovering AHV VMs

Review the following information on recovering AHV VMs:

- The target AHV host cluster must be added to the appliance as an asset. See "Adding a virtual host" on page 308.
- Unitrends recommends recovering to an AHV cluster that is running the same version as the original cluster (but you can recover to a higher supported version listed in the Compatibility and Interoperability Matrix if necessary).
- The target AHV host cluster must support the operating system (OS) of the VM you are recovering. (See this Nutanix document for details: Supported Guest VM Types for AHV.)
- The target AHV host cluster must have sufficient space and compute resources for the new VM.
- AHV recovery jobs access the target AHV host over the iSCSI protocol. The Unitrends appliance must be able to connect to the iSCSI targets on the Nutanix storage LAN.
- A recovered AHV VM is created with the following default name: <*original_VM_name*>_*restore*. You can edit this name when you create the recover job.
- VMs are uniquely identified by UUID (not by VM name). VMs are recovered with a new UUID. A recovered VMs UUID does not match that of the source AHV backup used for the recovery.
- If the recovery target is hosting a VM with the same name as the one that is being recovered, the original VM is not overwritten by the recovery operation. Once the recovery operation completes, the host has two VMs with the same name (but different UUIDs).

About recovering XenServer VMs

Review the following information on recovering XenServer VMs:

- The target XenServer host must be added to the appliance as an asset. See "Adding a virtual host" on page 308.
- Unitrends recommends recovering to a XenServer host that is running the same version as the original host (but you can recover to a higher supported version listed in the Compatibility and Interoperability Matrix if necessary).
- The target XenServer host must support the operating system (OS) of the VM you are recovering. (See the Citrix XenServer documentation for details.)
- The target XenServer host must have sufficient space and compute resources for the new VM.
- A recovered XenServer VM is created with the following default name: <*original_VM_name*>_*restore*. You can edit this name when you create the recover job.

Recovering a VM

To perform the recovery, you will start by selecting a backup or backup copy. For backups, you can do this in the Backup Catalog or in the Backup Browser. For backup copies and imported backups, you must use the Backup Catalog.

In the Backup Catalog, backups and backup copies display under the protected asset. You can modify the display by entering filter criteria. Expand an asset to view its backups and backup copies:

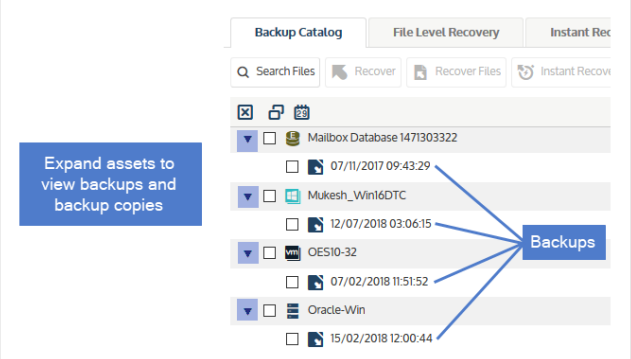

The Backup Browser provides advanced search and filter options. Backups are not grouped under the protected asset. Search for backups by selecting an appliance and date range. Filter the display by entering text in the column fields. Click an arrow to sort by column:

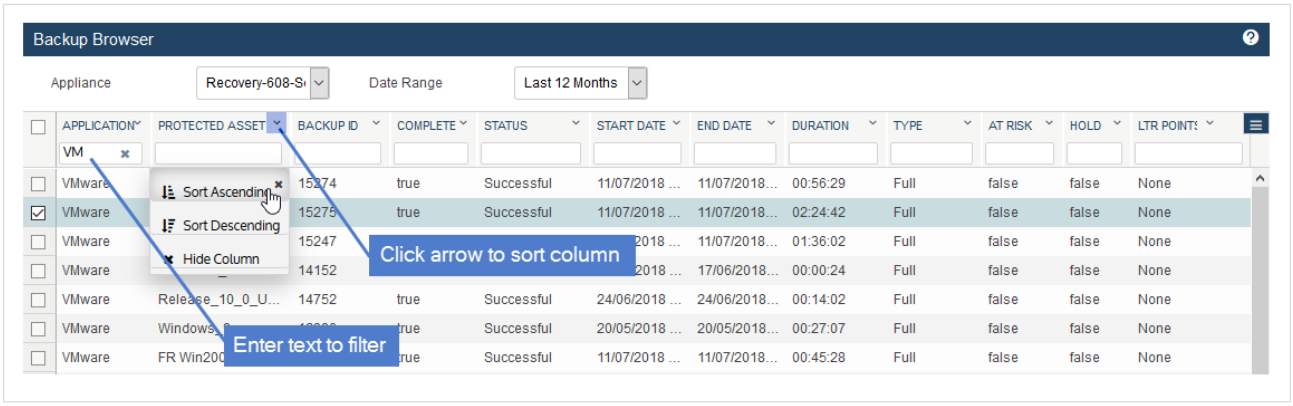

For more on working with these features, see "Backup Catalog tab" on page 60.

To recover a VM, use one of these procedures:

- "To recover a VM by using the Backup Catalog"
- "To recover a VM by using the Backup Browser" on page 803

To recover a VM by using the Backup Catalog

Use this procedure to recover an entire virtual machine by using the Backup Catalog.

1. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
2. Select **Recover** and click the **Backup Catalog** tab.
3. Locate the backup to use for the recovery

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

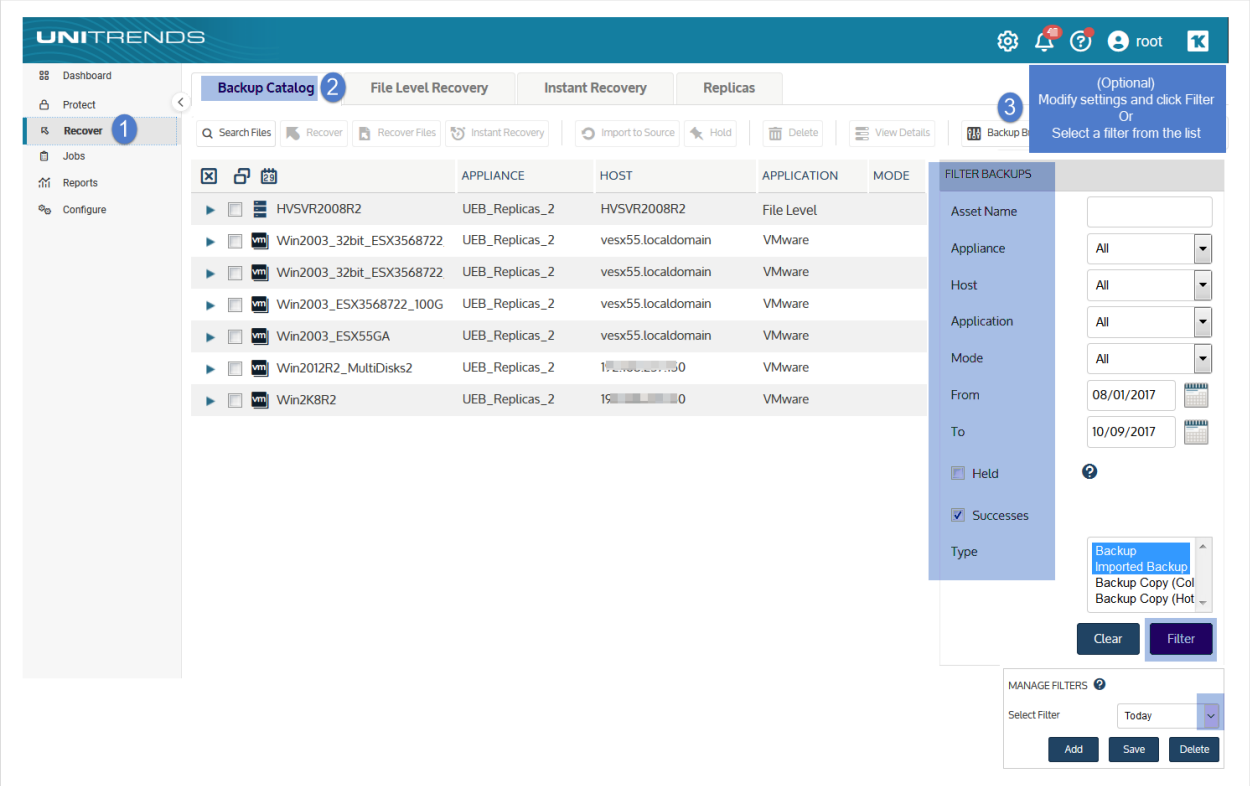

4 Expand the VM asset and select one of the following to use for the recovery:

- A host-level backup.
- An imported host-level backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
- A hot backup copy (supported if performing the recovery on the target appliance where the hot copy resides).

5 Click **Recover**.

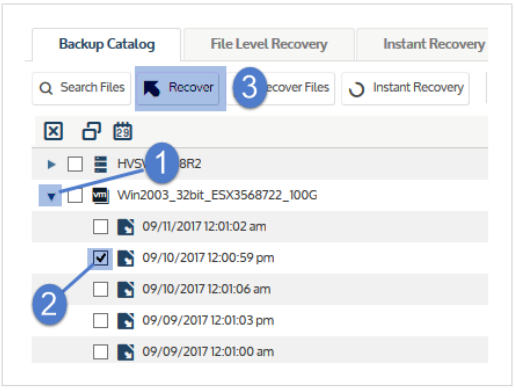

6 Select these Recovery Options:

Recovery Options	Description
Target Location	Select the host where the VM will be recovered. The list contains hosts that have been added to the appliance and are compatible with the VM being recovered. Incompatible hosts do not display in the list. To add a host, see "Adding a virtual host" on page 308. For Hyper-V clusters – To create a clustered VM, you must select a cluster as the target Hyper-V host.
Target Resource	(Optional) Select a resource pool. This field displays only if the Target Location is an ESXi host that has resource pools.
Target Storage	Select a datastore (ESXi host), a volume (Hyper-V host), a storage container (AHV host), or a Storage Repository (XenServer host).
Recover excluded disk metadata	(Optional) Check this box to recover the metadata for disks that were excluded from the backup.

7 Click **Next**.

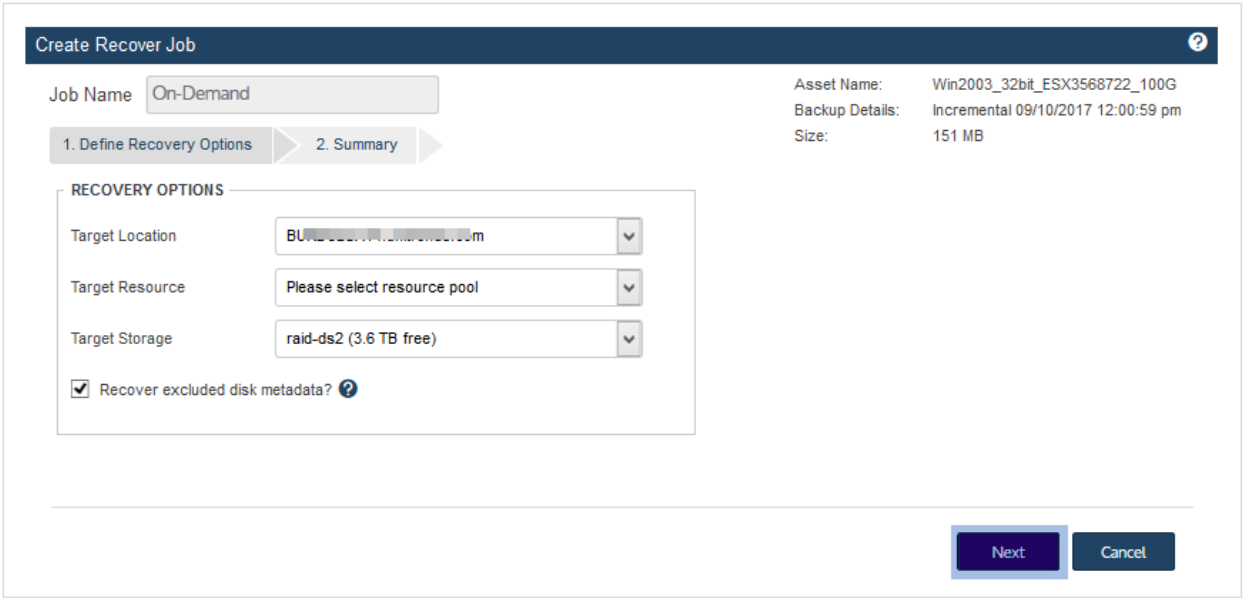

A summary of the selected recovery options display.

8 (Optional) Modify the VM Name by clicking it in the Assets to Recover list and entering a new name. (Supported for VMware, AHV, and XenServer only.)

9 Click **Save**. The job is queued immediately.

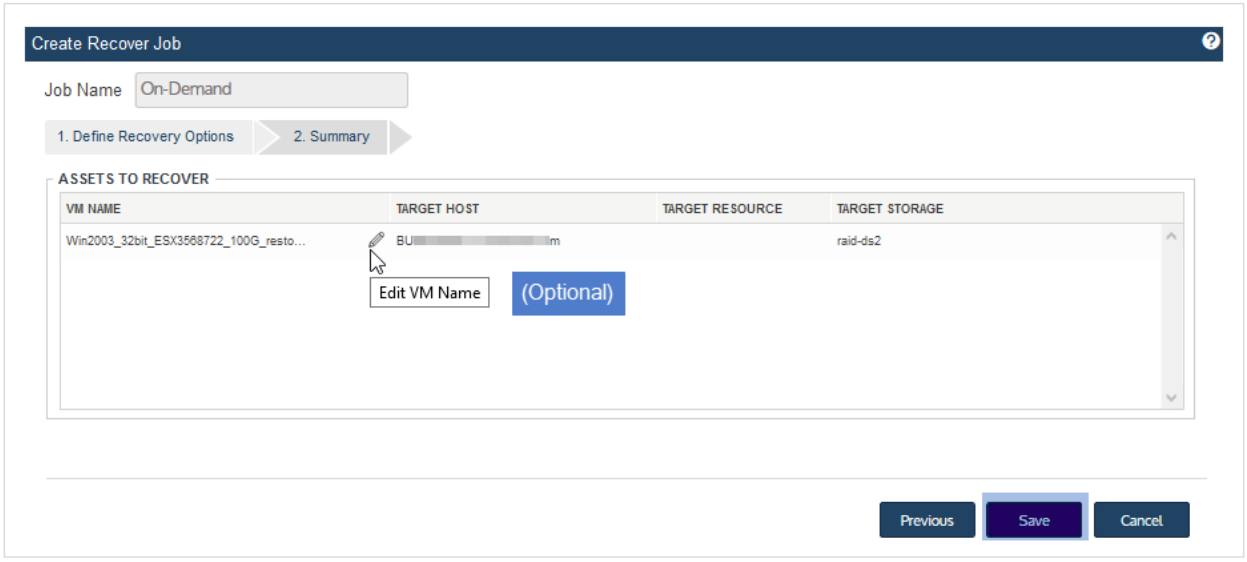

10 Click **OK** to close the Information message.

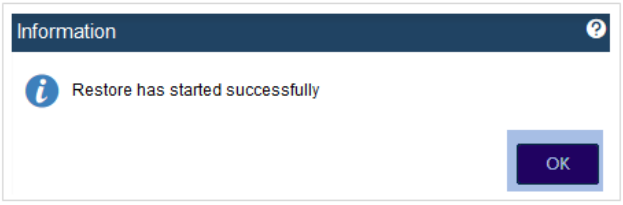

11 To view job progress, select **Jobs > Active Jobs**. Recovery is complete when the job status changes to Success.

Notes:

- For VMware, Hyper-V, and XenServer, recovery consists of these tasks: first the VM's configuration files (metadata) are recovered; next the VM's data is recovered. In our example, the VM has multiple disks. Data for each disk is recovered as a separate task.
- For AHV, recovery is done in one task.
- The recovered VM is created in a powered off state.

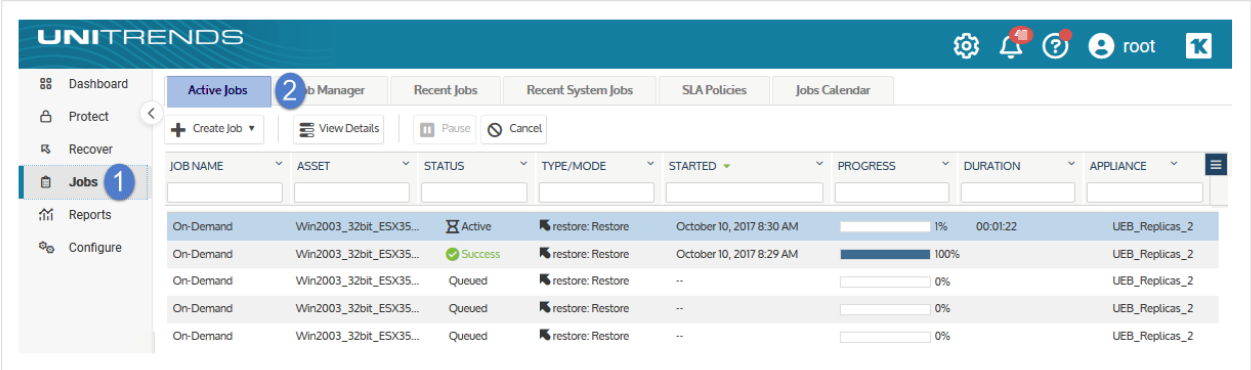

12 After the recovery job completes, go to the hypervisor and power on the recovered virtual machine.

13 Modify VM settings and backup schedules as needed.

- A recovered VM may not have the same network settings as the original. Check network settings and modify if needed.
- The recovered VM has the same username/password credentials as the original VM. Access the VM and verify that it is functioning as expected in production.
- Create or edit backup schedules to begin protecting the recovered VM. The next backup of the recovered VM is promoted to a full.

Notes:

- Windows server VMs - In rare instances, after a restore is performed for a Windows server VM, a disk may be inaccessible because it has been placed in an offline state. To bring disks into an online state, login to the VM, go to Disk Management, right-click on the offline disk, and select **Online** from the drop-down menu.
- Debian VMs - In some instances, Gnome might not start after a Debian VM is recovered. You can resolve this issue by rebooting the VM or restarting Gnome from the console. To access the console, enter *Ctl+Alt+F1* and log in as root. Then run *startx*.

To recover a VM by using the Backup Browser

Use this procedure to recover an entire virtual machine by using the Backup Browser.

Note: Backup copies and imported backups cannot be viewed in the Backup Browser. To recover from a backup copy or imported backup, use the Backup Catalog procedure above.

1 Log in to the backup appliance .

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

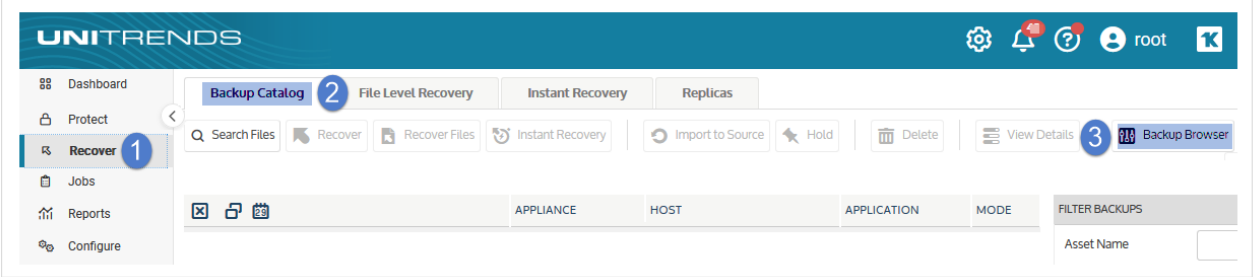

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display:

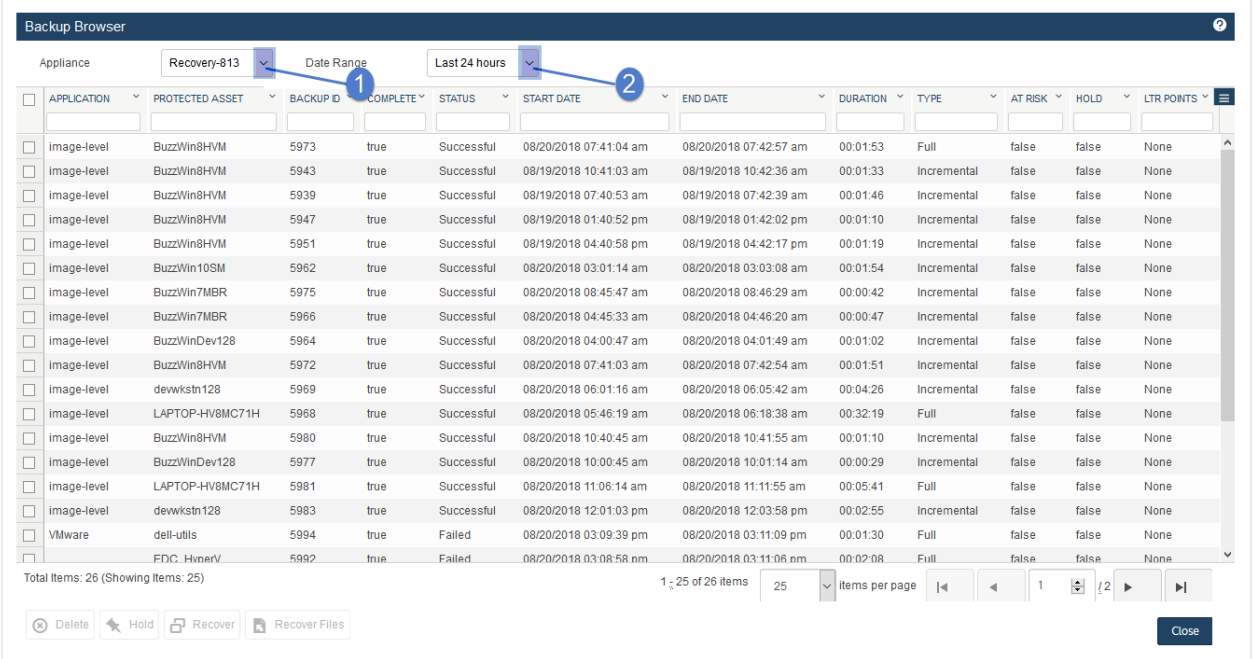

APPLICATION	PROTECTED ASSET	BACKUP ID	COMPLETE	STATUS	START DATE	END DATE	DURATION	TYPE	AT RISK	HOLD	LTR POINTS
image-level	BuzzWin8HVM	5973	true	Successful	08/20/2018 07:41:04 am	08/20/2018 07:42:57 am	00:01:53	Full	false	false	None
image-level	BuzzWin8HVM	5943	true	Successful	08/19/2018 10:41:03 am	08/19/2018 10:42:36 am	00:01:33	Incremental	false	false	None
image-level	BuzzWin8HVM	5939	true	Successful	08/19/2018 07:40:53 am	08/19/2018 07:42:39 am	00:01:46	Incremental	false	false	None
image-level	BuzzWin8HVM	5947	true	Successful	08/19/2018 01:40:52 pm	08/19/2018 01:42:02 pm	00:01:10	Incremental	false	false	None
image-level	BuzzWin8HVM	5951	true	Successful	08/19/2018 04:40:58 pm	08/19/2018 04:42:17 pm	00:01:19	Incremental	false	false	None
image-level	BuzzWin10SM	5962	true	Successful	08/20/2018 03:01:14 am	08/20/2018 03:03:08 am	00:01:54	Incremental	false	false	None
image-level	BuzzWin7MBR	5975	true	Successful	08/20/2018 08:45:47 am	08/20/2018 08:46:29 am	00:00:42	Incremental	false	false	None
image-level	BuzzWin7MBR	5966	true	Successful	08/20/2018 04:45:33 am	08/20/2018 04:46:20 am	00:00:47	Incremental	false	false	None
image-level	BuzzWinDev128	5964	true	Successful	08/20/2018 04:00:47 am	08/20/2018 04:01:49 am	00:01:02	Incremental	false	false	None
image-level	BuzzWin8HVM	5972	true	Successful	08/20/2018 07:41:03 am	08/20/2018 07:42:54 am	00:01:51	Incremental	false	false	None
image-level	dewkstn128	5969	true	Successful	08/20/2018 06:01:16 am	08/20/2018 06:05:42 am	00:04:26	Incremental	false	false	None
image-level	LAPTOP-HV8MC71H	5968	true	Successful	08/20/2018 05:46:19 am	08/20/2018 06:18:38 am	00:32:19	Full	false	false	None
image-level	BuzzWin8HVM	5980	true	Successful	08/20/2018 10:40:45 am	08/20/2018 10:41:55 am	00:01:10	Incremental	false	false	None
image-level	BuzzWinDev128	5977	true	Successful	08/20/2018 10:00:45 am	08/20/2018 10:01:14 am	00:00:29	Incremental	false	false	None
image-level	LAPTOP-HV8MC71H	5981	true	Successful	08/20/2018 11:06:14 am	08/20/2018 11:11:55 am	00:05:41	Full	false	false	None
image-level	dewkstn128	5983	true	Successful	08/20/2018 12:01:03 pm	08/20/2018 12:03:58 pm	00:02:55	Incremental	false	false	None
VMware	dell-utils	5994	true	Failed	08/20/2018 03:09:39 pm	08/20/2018 03:11:09 pm	00:01:30	Full	false	false	None
	EDC_HyperV	5992	true	Failed	08/20/2018 03:08:58 pm	08/20/2018 03:11:06 pm	00:02:08	Full	false	false	None

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Select the host-level backup.

6 Click **Recover**.

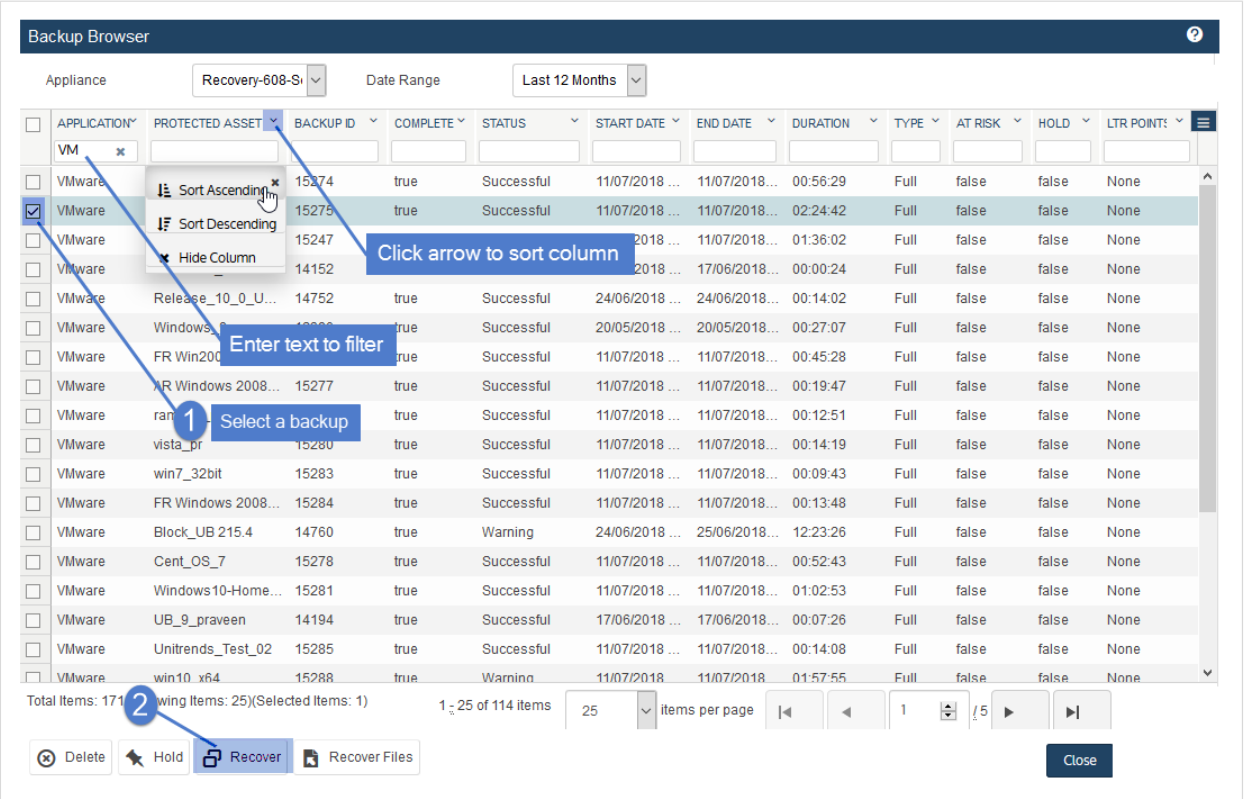

7 Select these Recovery Options:

Recovery Options	Description
Target Location	Select the host where the VM will be recovered. The list contains hosts that have been added to the appliance and are compatible with the VM being recovered. Incompatible hosts do not display in the list. To add a host, see "Adding a virtual host" on page 308. For Hyper-V clusters – To create a clustered VM, you must select a cluster as the target Hyper-V host.
Target Resource	(Optional) Select a resource pool. This field displays only if the Target Location is an ESXi host that has resource pools.
Target Storage	Select a datastore (ESXi host), a volume (Hyper-V host), a storage container (AHV host), or a Storage Repository (XenServer host).

Recovery Options	Description
Recover excluded disk metadata	(Optional) Check this box to recover the metadata for disks that were excluded from the backup.

8 Click **Next**.

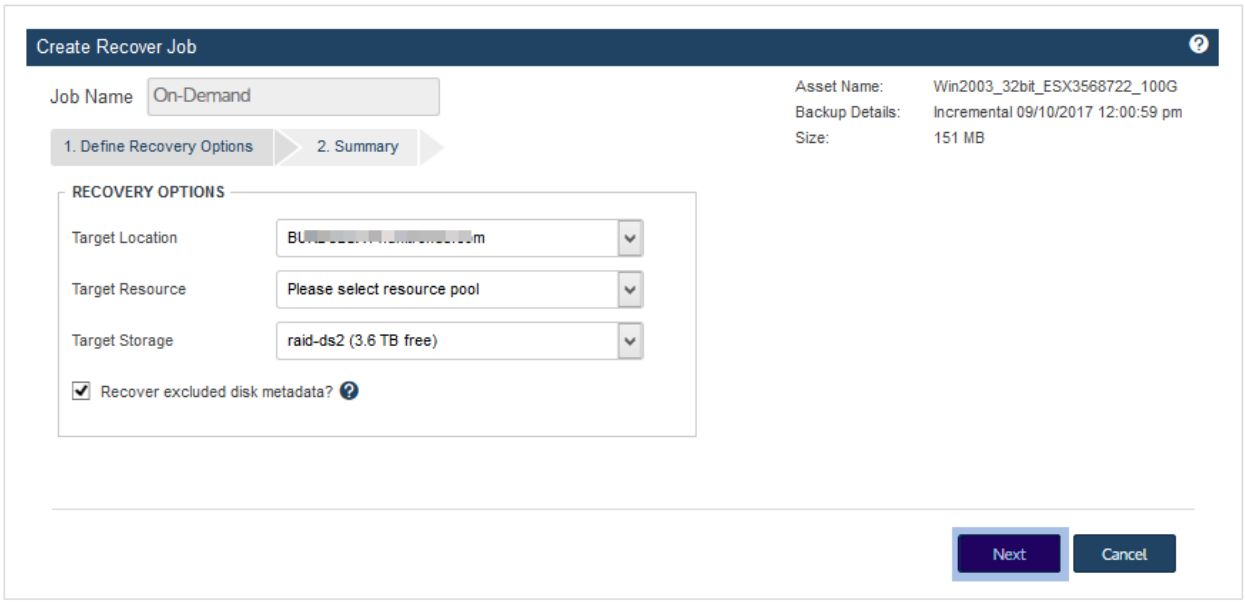

A summary of the selected recovery options display.

9 (Optional) Modify the VM Name by clicking it in the Assets to Recover list and entering a new name. (Supported for VMware, AHV, and XenServer only.)

10 Click **Save**. The job is queued immediately.

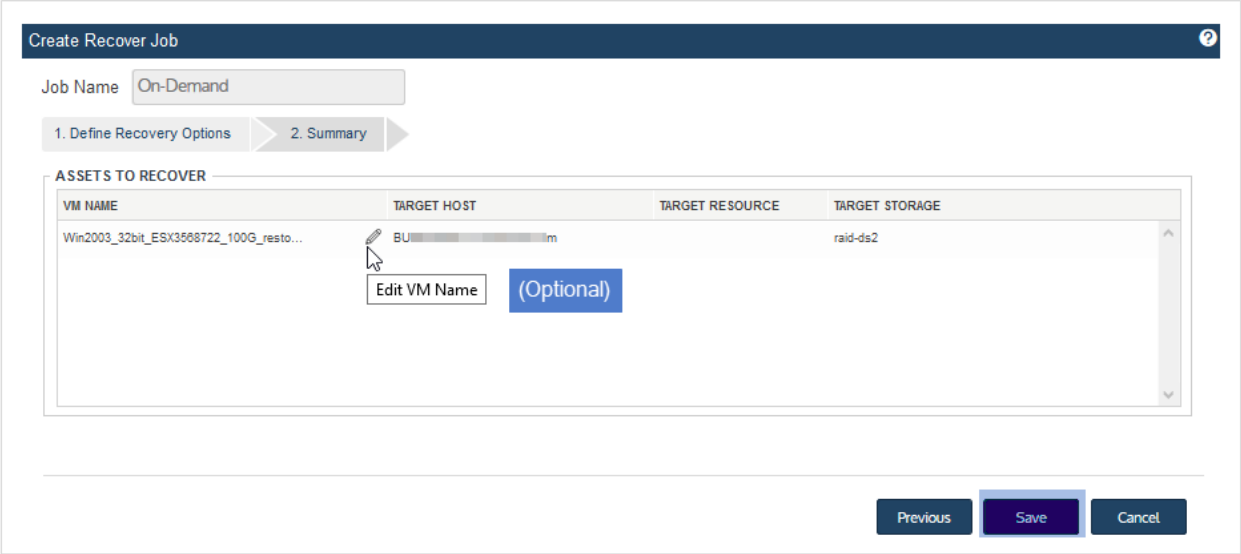

11 Click **OK** to close the Information message.

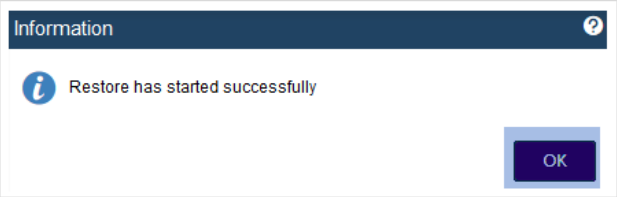

12 To view job progress, select **Jobs > Active Jobs**. Recovery is complete when the job status changes to Success.

Notes:

- For VMware, Hyper-V, and XenServer, recovery consists of these tasks: first the VM's configuration files (metadata) are recovered; next the VM's data is recovered. In our example, the VM has multiple disks. Data for each disk is recovered as a separate task.
- For AHV, recovery is done in one task.
- The recovered VM is created in a powered off state.

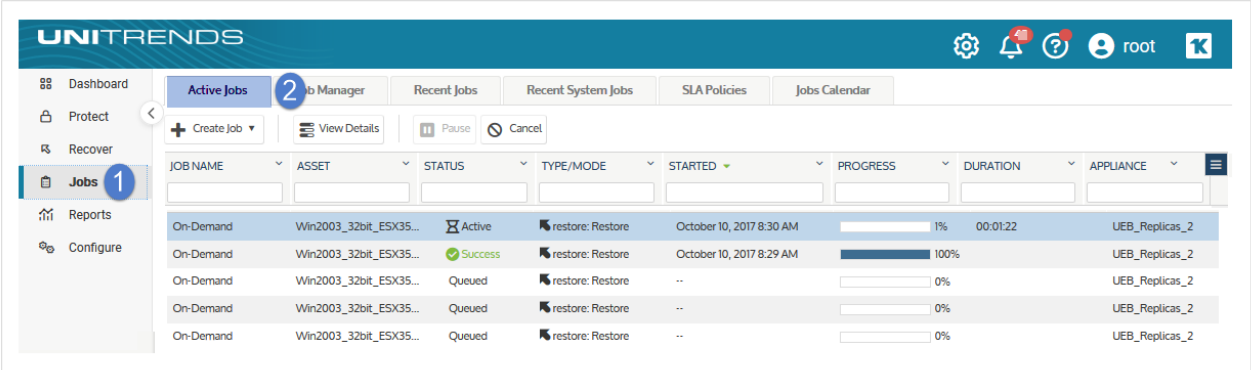

13 After the recovery job completes, go to the hypervisor and power on the recovered virtual machine.

14 Modify VM settings and backup schedules as needed.

- A recovered VM may not have the same network settings as the original. Check network settings and modify if needed.
- The recovered VM has the same username/password credentials as the original VM. Access the VM and verify that it is functioning as expected in production.
- Create or edit backup schedules to begin protecting the recovered VM. The next backup of the recovered VM is promoted to a full.

Notes:

- Windows server VMs - In rare instances, after a restore is performed for a Windows server VM, a disk may be inaccessible because it has been placed in an offline state. To bring disks into an online state, login to the VM, go to Disk Management, right-click on the offline disk, and select **Online** from the drop-down menu.
- Debian VMs - In some instances, Gnome might not start after a Debian VM is recovered. You can resolve this issue by rebooting the VM or restarting Gnome from the console. To access the console, enter *Ctl+Alt+F1* and log in as root. Then run *startx*.

Recovering files from virtual machine backups

You can recover files from host-level backups and host-level backup copies of Windows and Linux VMs. File recovery is supported for VMs that reside on VMware, Hyper-V, AHV, and XenServer hosts. A single file-level recovery can be performed on multiple VMs simultaneously.

Recovery procedures can be run from the backup appliance or from the backup copy target appliance:

- Any host-level backup of a Windows or Linux VM – Recovering from a local backup or imported backup copy involves creating a recovery object on the backup appliance that contains files from the backup. You recover files by downloading directly from this object or by mounting the object to a recovery target machine.
- Any host-level backup of a Windows or Linux VM – Recovering directly from a backup copy that resides on a target appliance or resides in the Unitrends Cloud involves creating a recovery object on the target appliance or in the

Unitrends Cloud. You recover files by downloading directly from this object or by mounting the object to a recovery target machine.

- Indexed VMware host-level backup of a Windows VM – Run the recovery on the local backup appliance to search a VMware virtual machine's indexed backups for files that meet specified criteria and recover selected items to an agent-based Windows asset.

The method you use to recover files is determined by the following:

- Where you run the procedure. Procedures run on a backup appliance differ from those run on a backup copy target appliance.
- Whether you are recovering from a backup, imported backup copy, or hot backup copy.
- The operating system (OS) and configuration of the protected VM.
- Whether you are running indexed host-level backups of a VMware Windows virtual machine.

If file recovery is not supported for your Windows or Linux VM, recover the virtual machine instead as described in "Recovering a virtual machine" on page 796.

Recovery procedures overview

Use the "Recovering from an indexed VMware backup of a Windows VM by using Search Files" on page 810, "Windows file-level recovery" on page 817, or "Linux file-level recovery" on page 834 procedures to recover files. The recovery procedure you use is determined by the following: whether the backup is an indexed VMware backup, where you are running the procedure and whether you are recovering from a local backup, from an imported backup copy, or directly from a hot backup copy. See the following table for a description of the options you can use in each case:

Backup type and location	Run from	Use this procedure to create the recovery object
Host-level backup on the local appliance	Backup appliance	Run the "Windows file-level recovery" on page 817 or "Linux file-level recovery" on page 834 procedure from the backup appliance to recover files from a local backup or from an imported backup copy. • For Windows - In "Step 2: Create the recovery object", use this option: "To create the recovery object and recover from a backup or imported backup copy". • For Linux - In "Step 2: Create the recovery object", use this option: "To create the recovery object and recover from a backup or imported backup copy". To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.
Indexed backup of a VMware Windows VM on the local appliance	Backup appliance	Run the "Recovering from an indexed VMware backup of a Windows VM by using Search Files" on page 810 procedure from the backup appliance to search indexed backups.

Backup type and location	Run from	Use this procedure to create the recovery object
		Note: This recovery procedure searches indexed local backups only. Search Files cannot be used to search imported backup copies.
Host-level backup copy in the Unitrends Cloud or on a backup copy target appliance (release 9.1 or later only)	Source backup appliance	Run the "Windows file-level recovery" on page 817 or "Linux file-level recovery" on page 834 procedure from the source backup appliance to recover files directly from a backup copy that resides on a target appliance or resides in the Unitrends Cloud. Both the source appliance and the target appliance must be running release 9.1 or later. • For Windows - In "Step 2: Create the recovery object", use this option: "To create the recovery object and recover from a hot backup copy by using the source backup appliance". • For Linux - In "Step 2: Create the recovery object", use this option: "To create the recovery object and recover from a hot backup copy by using the source backup appliance".
Host-level backup copy on a backup copy target appliance	Backup copy target appliance	Run the "Windows file-level recovery" on page 817 or "Linux file-level recovery" on page 834 procedure from the backup copy target appliance to recover files from a hot backup copy that resides on that target appliance. • For Windows - In "Step 2: Create the recovery object", use this option: "To create the recovery object and recover from a hot backup copy by using the target appliance". • For Linux - In "Step 2: Create the recovery object", use this option: "To create the recovery object and recover from a hot backup copy by using the target appliance".

Recovering from an indexed VMware backup of a Windows VM by using Search Files

Use this procedure to search a VMware Windows VM's indexed backups for files/folders that meet specified criteria and recover selected items from the search results.

Notes:

- The backup appliance must be running release 10.5 or higher to run this procedure.

- This procedure can only be used for local backups that were run with the Edit Asset > Index Backups option. For details on configuring this option, see "To edit a virtual machine asset" on page 317.
- During the procedure, you will select an agent-based Windows asset where files will be recovered. Choose from the Windows assets that have been added to the appliance. To recover files to the original location or to another VM, install the Unitrends agent on the Windows VM (see "Installing the Windows agent" on page 362), then add the VM to the backup appliance as an agent-based asset (see "To add an agent-based asset" on page 289).
- File search of imported backup copies is not supported. Recover by browsing the imported backup instead (see "Windows file-level recovery" on page 817 or "Linux file-level recovery" on page 834).
- File search is not supported for recovery of ReFS filesystems. Recover by browsing the backup instead (see "Windows file-level recovery" on page 817).

1 Log in to the backup appliance.

2 Click **Recover > Backup Catalog > Search Files**.

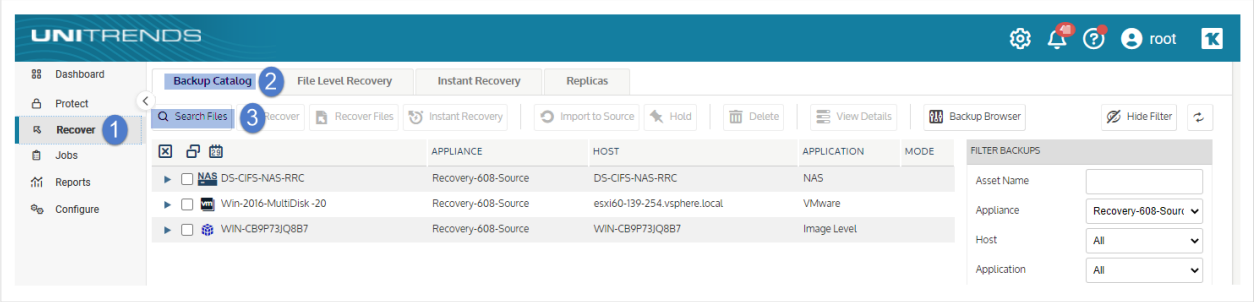

3 In the **Type** list, select **Backup**.

4 Select the VMware **Asset** whose backups will be searched. VMware assets are listed as *VM_name* (*host_name*).

5 Enter one or more search options:

Search Options	Description
String	Enter text to search. The wildcard character * is supported.
Match Case	Select to match the letter case of the entered string.
From/To	Use to search for files that were last modified within the specified date range. Results do not include files modified on the From or To date.
Size	Use to search for files that meet this size criteria.
Advanced	Click to search using a regular expression.

6 Click **Search**.

All indexed backups of this VM stored on the appliance are searched for matching files.

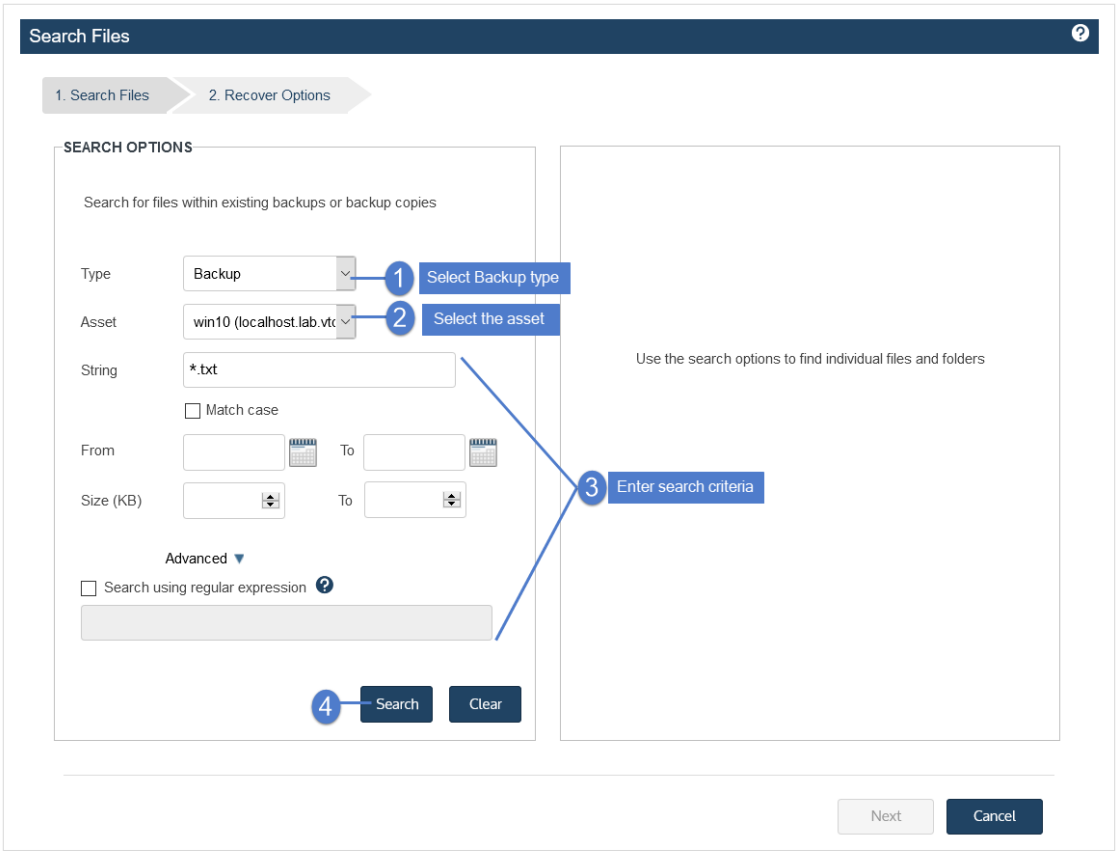

7 In the results list, click to select files/folders to recover.

Notes:

- All items you select must be from a single backup. Check the Backup ID to determine an item's backup. If you select items from multiple backups, the Next button becomes disabled.
- Softlinks cannot be downloaded and are not included in the search results.

8 Click **Next**.

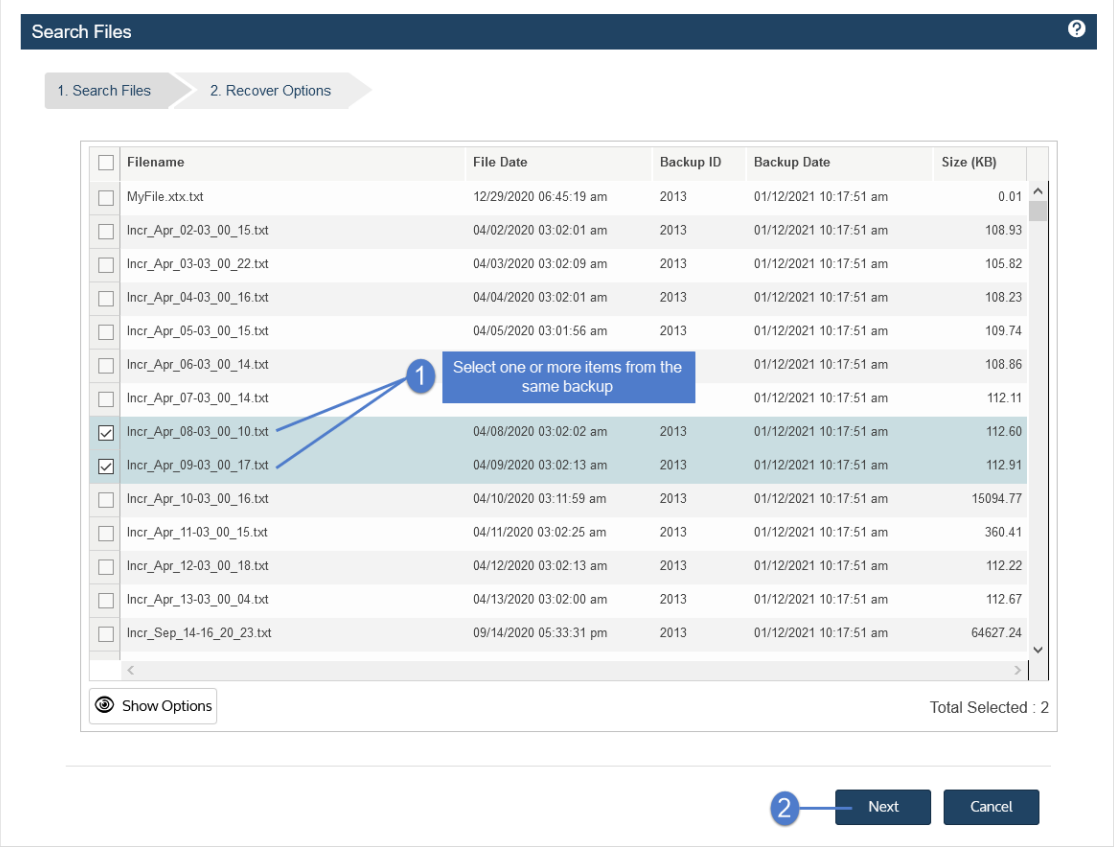

9 Select the Windows **Asset** where the files will be recovered.

The list contains agent-based assets that have been added to the appliance. To add a target asset, install the Unitrends agent on the asset (see "Installing the Windows agent" on page 362), then add the asset (see "To add an agent-based asset" on page 289).

10 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

11 (Optional) Specify Exclusions.

12 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.

Option	Description
Commands to run post-restore	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Preserve directory structure	Check this box to preserve the existing file structure within the target directory. **Note:** To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

Overwrite existing files and *Restore newer files only* options

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	**Non-Windows backup**
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
		• If the file to recover is older than the one in the Target Directory, does not recover the file.
Overwrite existing files = Yes Restore newer files only = No	Recovers the file and overwrites the existing file.	Recovers the file and overwrites the existing file.
Overwrite existing files = No Restore newer files only = Yes	Does not recover the file.	Does not recover the file.

13 Click **Save**.

14 Click **OK** to close the Notice message.

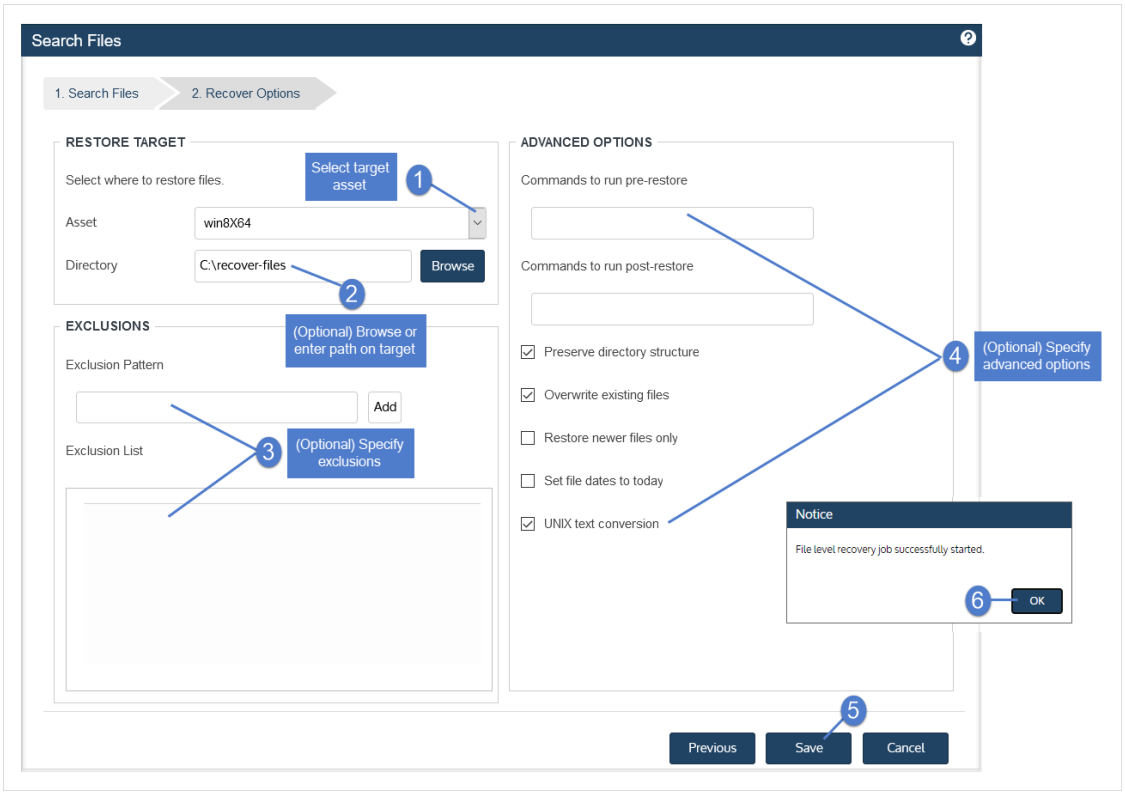

15 To monitor the recovery job:

- Select **Jobs > Active** Jobs.
- Select the job in the list and click **View Details**.

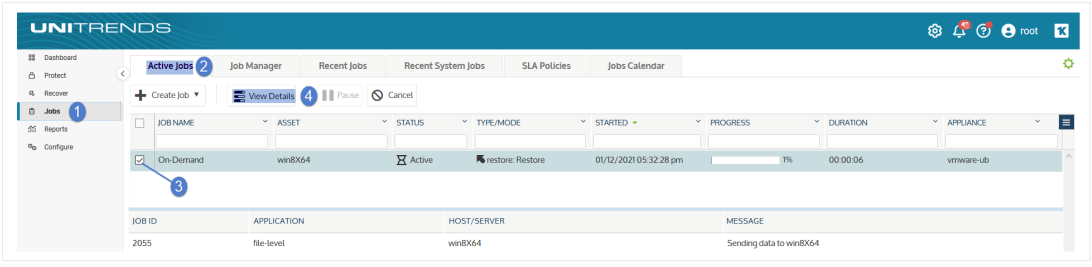

- The recovery is complete when the job's status changes to Success.

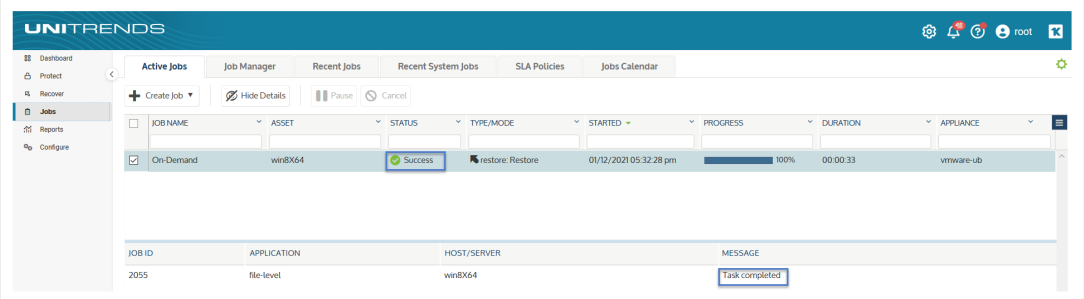

16 Access the recovered files on the recovery target.

Note: Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

Windows file-level recovery

Use the following procedures to recover files from a backup, imported backup copy, or hot backup copy of a Windows VM.

Notes:

- Hyper-V – Recovery of selected files or pathnames that contain non-UTF-8 compatible characters is not supported.
- VMware – Do not use these procedures to recover files from an indexed VMware backup. Instead, see "Recovering from an indexed VMware backup of a Windows VM by using Search Files" on page 810.

Step 1: Ensure prerequisites have been met

The following requirements and considerations apply to recovering files from a host-level backup or host-level backup copy of a Windows VM by creating a recovery object:

Prerequisite or consideration	Description
Supported recovery methods	To recover files from a host-level backup or copy, the appliance creates a recovery object that contains the backup's files. For some Windows VMs, this object is also

Prerequisite or consideration	Description
	exposed as a CIFS (Samba) share and/or an iSCSI LUN on the backup appliance. After you create the recovery object, you will view it on the File Level Recovery tab to see whether the CIFS and iSCSI options are available. You can recover files from this object in several ways. Options include: • Browse the recovery object and download selected files to a *.zip* file. This is the simplest method. • Mount the CIFS share on a recovery target machine. From the target machine, select files to recover. • Mount the iSCSI LUN on a recovery target machine. From the target machine, select files to recover. (You must use an iSCSI LUN in some cases. For details, see "When to use an iSCSI LUN" on page 818.)
Recovery target requirements	The target can be configured with basic, GUID Partition Table (GPT), or dynamic disks. All configured disks must have unique names. To use a CIFS share for the recovery, the target Windows asset must be able to access the appliance's Samba share: • SMB 2.0 – The SMB 2.0 security option is enabled by default on Unitrends appliances that were originally imaged or deployed with version 10.4.8 or higher. SMB 2.0 must be enabled on the target Windows asset. • SMB 1.0 – The SMB 1.0 security option is enabled by default on appliances that were originally imaged or deployed with a pre-10.4.8 version. SMB 1.0 must be enabled on the target Windows asset. **Note:** Upgrading from a pre-10.4.8 version does not change the SMB 1.0 setting. (To configure your appliance to use SMB 2.0, see How Unitrends supports SMBv2.)
When to use an iSCSI LUN	You must recover by mounting the iSCSI LUN to perform the following tasks: • Recover access control information on files and folders. • Recover New Technology File System (NTFS) encrypted files. • Recover Resilient File System (ReFS) files.

Prerequisite or consideration	Description
	Note: ReFS limitation - ReFS file system versions are not compatible with all Windows operating system versions. To avoid compatibility issues, recover ReFS files by mounting the iSCSI LUN on a machine whose operating system version is the same or later than that of the machine where the backup was taken. • Recover files on dynamic disks. If the dynamic volumes are still in use on the original VM, you must mount the recovery object on a different machine. **Note:** For the recovery, iSCSI disks are writable and a 1 GB write limit is enforced. Errors display on the recovery target machine if more than 1 GB is required. In this case, you cannot recover by using iSCSI. Recover files by downloading to a *.zip* file or by mounting the CIFS share, or perform a VM recovery.

Step 2: Create the recovery object

Use one of these procedures to create the recovery object:

Note: If a previously-created recovery object is still mounted for the VM, you must remove it before creating a new one.

- "To create the recovery object and recover from a backup or imported backup copy" on page 819 – Run on the backup appliance to recover from a backup or imported backup copy.
- "To create the recovery object and recover from a hot backup copy by using the source backup appliance" on page 821 – Run on the backup appliance to recover from a backup copy that resides on a target appliance or in the Unitrends Cloud .
- "To create the recovery object and recover from a hot backup copy by using the target appliance" on page 823 – Run on the target appliance to recover from a backup copy that resides on that target appliance.

To create the recovery object and recover from a backup or imported backup copy

Run this procedure on the backup appliance to recover from a backup or imported backup copy.

1 Log in to the backup appliance.

2 Select **Recover** and click the **Backup Catalog** tab.

(Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

3 Expand the asset and select the backup or imported backup copy from which you want to recover files.

(To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)

4 Click **Recover Files**.

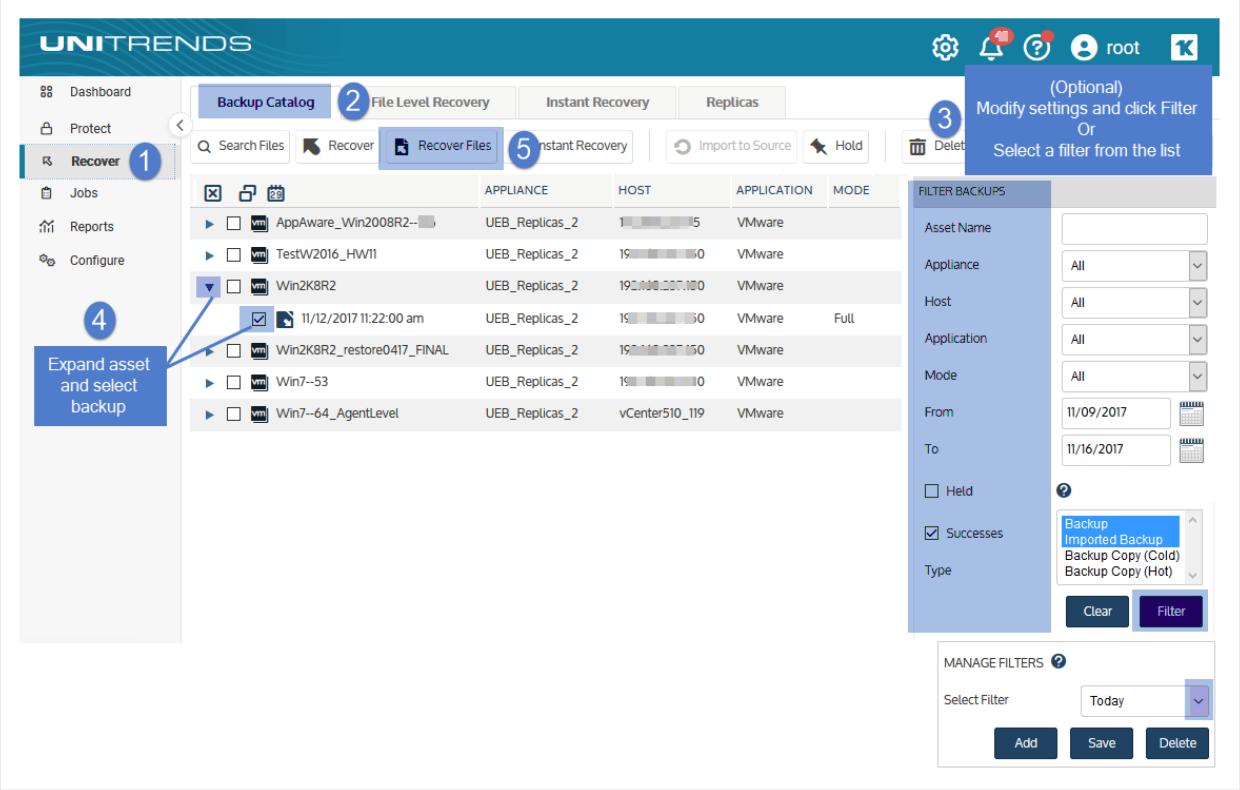

5 Click **Confirm** to continue. The appliance creates the recovery object.

Note: If you receive an error on a Unitrends Backup appliance while creating the recovery object, increase the memory allocation for the Unitrends Backup VM by using the host that manages it.

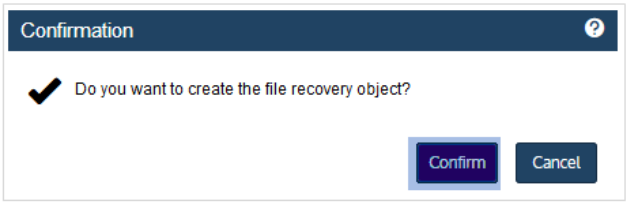

6 Click **View FLR**.

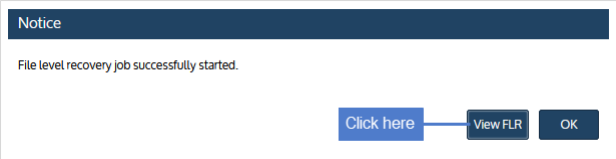

Proceed to "Recover files" on page 825.

To create the recovery object and recover from a hot backup copy by using the source backup appliance

Run this procedure on the backup appliance to recover from a backup copy that resides on a target appliance or in the Unitrends Cloud.

1. Log in to the source backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.
3. Use Filter Backups to the right to display hot backup copies:
   - In the Type box, select **Backup Copy (Hot)**.
   - (Optional) Select other filter options. For details, see "Working with custom filters" on page 67.
   - Click **Filter**.
4. Expand the asset and select the hot backup copy from which you want to recover files.
5. Click **Recover Files**.

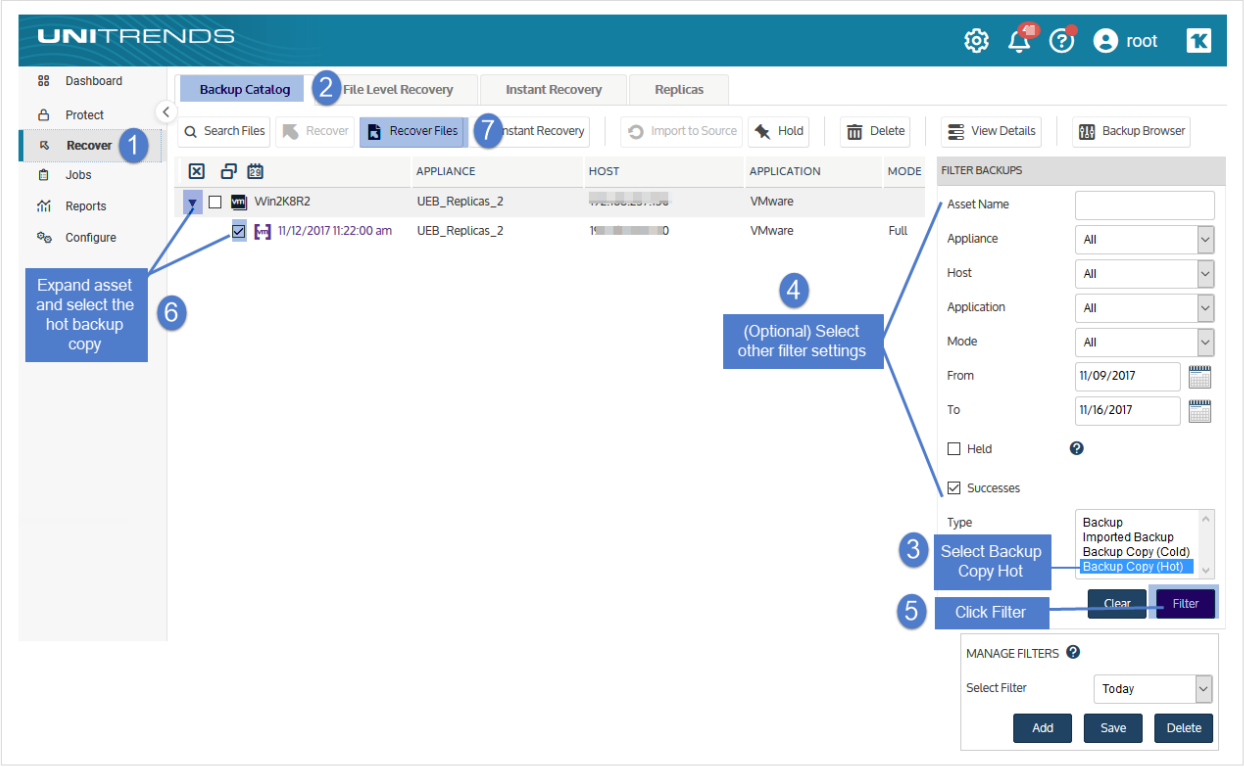

6 Click **Confirm** to continue. The appliance creates the recovery object in the Cloud or on the target appliance.

Notes:

- If your appliance is a Unitrends Backup virtual appliance and you receive an error while creating the recovery object, increase the memory allocation for the Unitrends Backup VM using the host that manages it.
- Recovery objects created in the Unitrends Cloud are automatically removed after 96 hours.

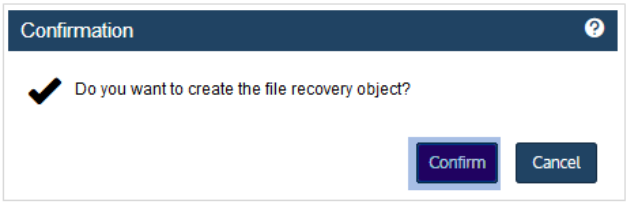

7 Click **View FLR** to view the recovery object on the File Level Recovery tab. The recovery object Name is *AssetName (on the target)*.

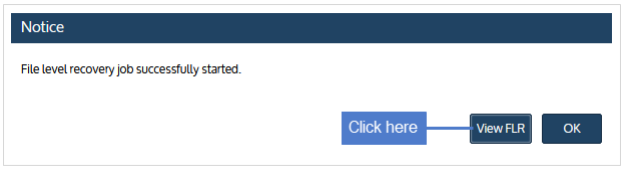

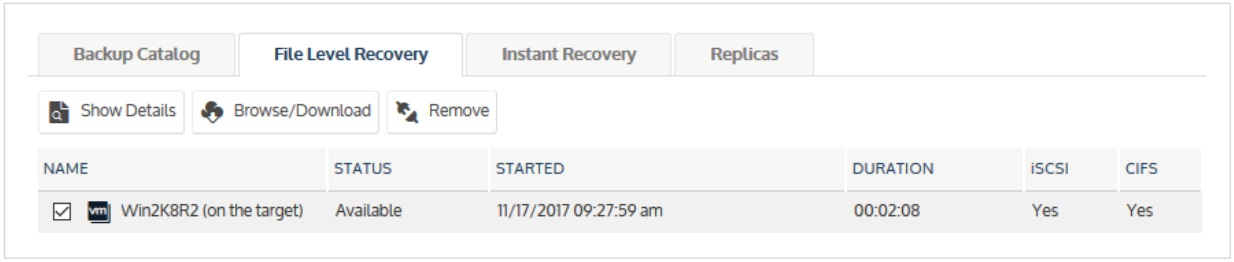

Proceed to "Step 3: Recover files" on page 825.

To create the recovery object and recover from a hot backup copy by using the target appliance

Run this procedure on the target appliance to recover from a backup copy that resides on that target appliance.

1. Log in to the backup copy target appliance.
2. Select **Recover** and click the **Backup Catalog** tab.
3. Use Filter Backups to the right to display hot backup copies:
   - In the Type box, select **Backup Copy (Hot)**.
   - (Optional) Select other filter options. For details, see "Working with custom filters" on page 67.
   - Click **Filter**.
4. Expand the asset and select the hot backup copy from which you want to recover files.
5. Click **Recover Files**.

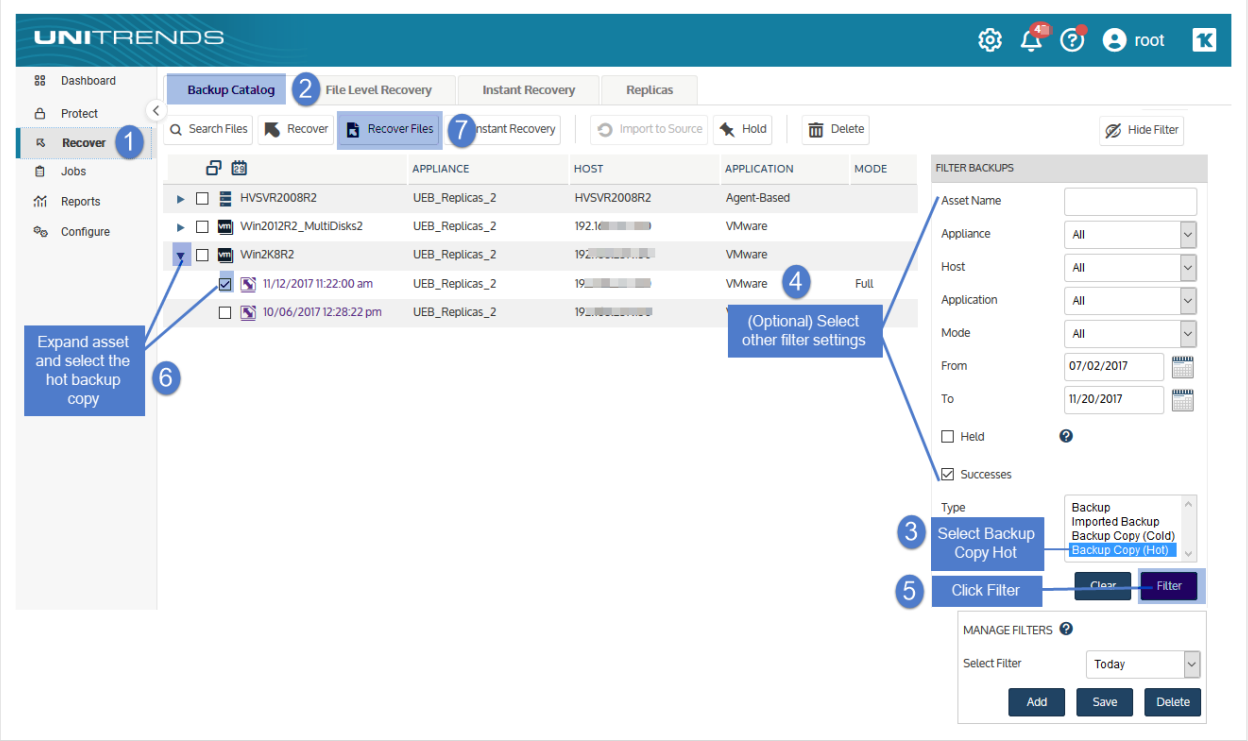

6 Click **Confirm** to continue. The appliance creates the recovery object on the target appliance.

Note: If your appliance is a Unitrends Backup virtual appliance and you receive an error while creating the recovery object, increase the memory allocation for the Unitrends Backup VM using the host that manages it.

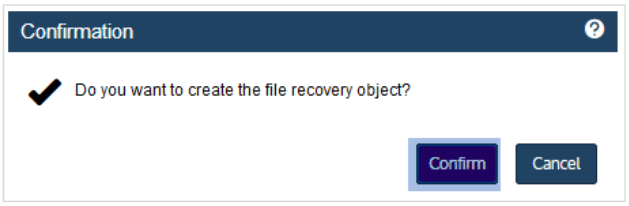

7 Click **View FLR** to view the recovery object on the File Level Recovery tab.

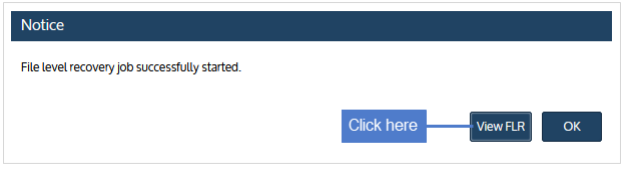

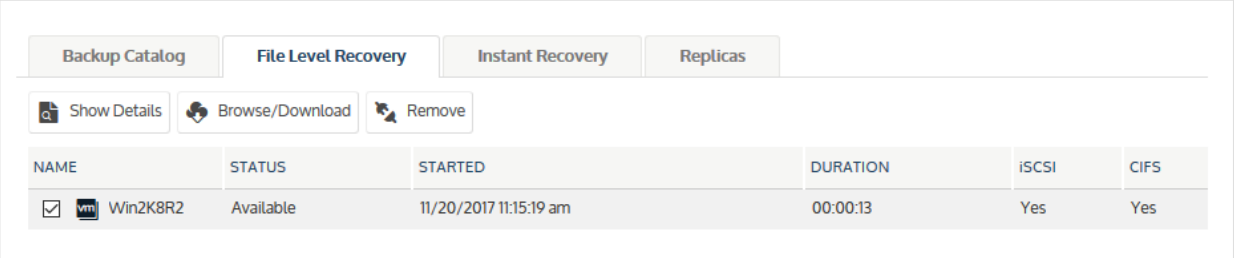

8 Proceed to "Step 3: Recover files".

Step 3: Recover files

View the recovery object on the File Level Recovery tab to see which recovery options are supported for the VM you selected. Use one of the following procedures to recover files. For a description of each method, see "Recovery procedures overview" on page 809.

- "To recover files by browsing and downloading to a .zip file" on page 825
- "To recover files by mounting the CIFS share" on page 827
- "To recover files by mounting the iSCSI LUN" on page 829

To recover files by browsing and downloading to a *.zip* file

1 On the **File Level Recovery** tab, locate the recovery object.

Recovery objects display on the tab with the following details: the name of the VM asset for which the object was created, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

2 Select the recovery object and click **Browse/Download**.

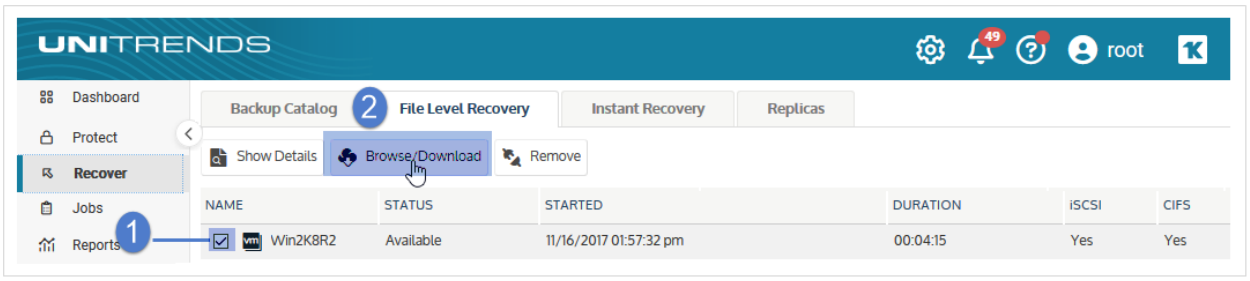

3 In the File Browser, select or drag files and/or directories to recover.

Note: Softlinks (also called *symbolic links*) are excluded from download. If you select a directory that contains files and softlinks, only the files are downloaded.

4 Click **Download**.

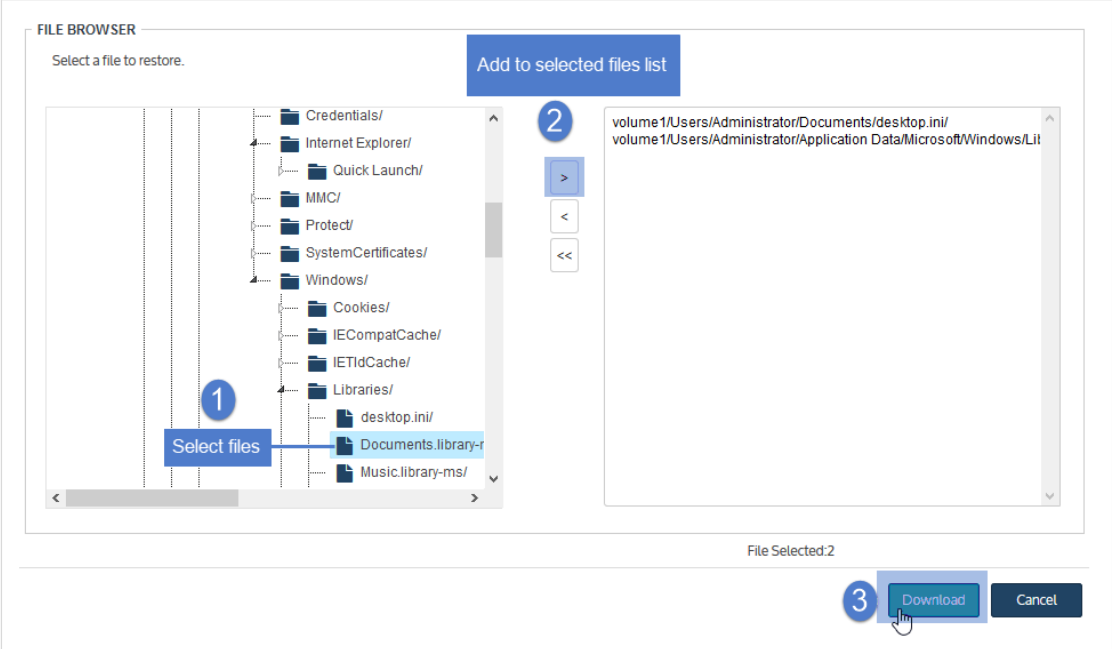

5 Click **Confirm** to download the selected files to a *.zip* file. The *.zip* file is downloaded to your browser's default location.

Notes:

- Volumes are assigned numbers during recovery that do not necessarily match the numbers from the original disk.
- The duration of the download is impacted by various factors, such as the size of the files, bandwidth, and download speed.
- Persistent browser and UI sessions are required to create the *.zip* file in the browser's default download location. If you close the browser or UI session during the recovery, you must run a new job.

6 When the download completes, the *Unitrends-Restore.zip* file displays in the browser. Select whether to open or save the file.

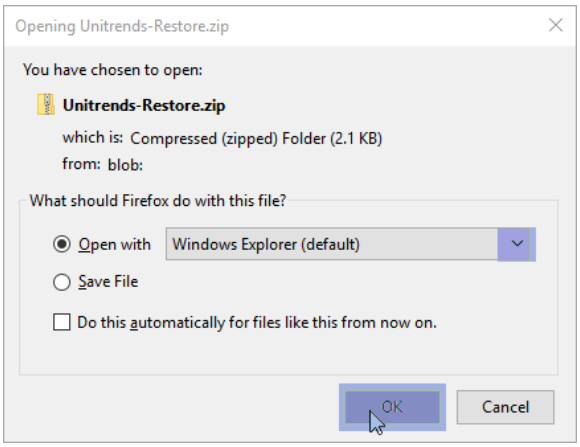

7 Access the recovered files in the download location and move them to another location on the local machine.

> **Note:** The Windows file explorer contains a setting to hide protected/system files. Turn off this setting to access all files.

Proceed to "Step 4: Remove the recovery object from the appliance" on page 833.

To recover files by mounting the CIFS share

1 Select **Recover** and click the **File Level Recovery** tab.

Recovery objects display with the following details: the name of the VM asset for which the object was created, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

2 Select the recovery object and click **Show Details**.

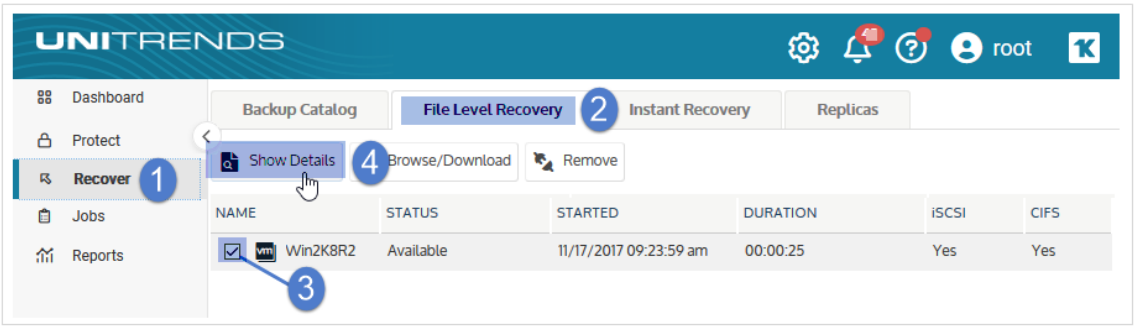

3 Note the CIFS path that displays in the File Level Recovery Details window. You will need this path to mount the CIFS share on the target machine.

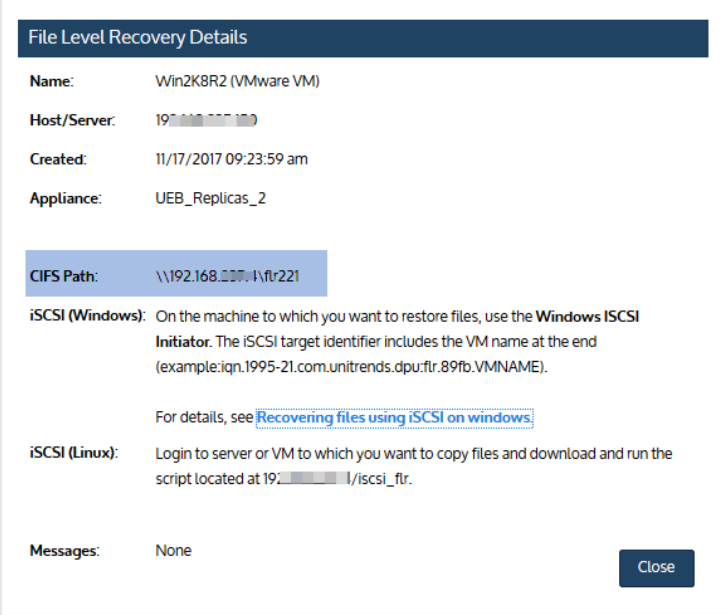

4 Log in to the recovery target workstation.

5 Enter the CIFS path into a file browser on the recovery target.

6 Browse the share to locate the files you want to recover.

Notes:

- Volumes are assigned numbers during recovery that do not necessarily match the letters from the original disks.
- The Windows file explorer contains a setting to hide protected/system files. Turn off this setting to access all files.

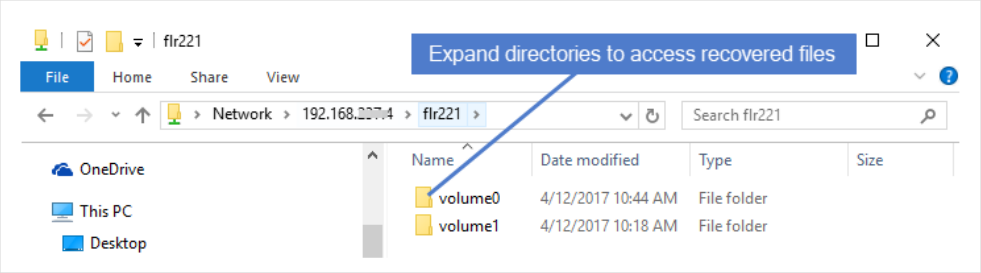

7 Move selected files to another location on the local machine.

8 Disconnect the network share by right-clicking the share and selecting **Disconnect**.

9 Proceed to "Step 4: Remove the recovery object from the appliance" on page 833.

To recover files by mounting the iSCSI LUN

1 Log in to the recovery target workstation.

2 Launch the iSCSI Initiator from **Administrative Tools** in the **Control Panel**.

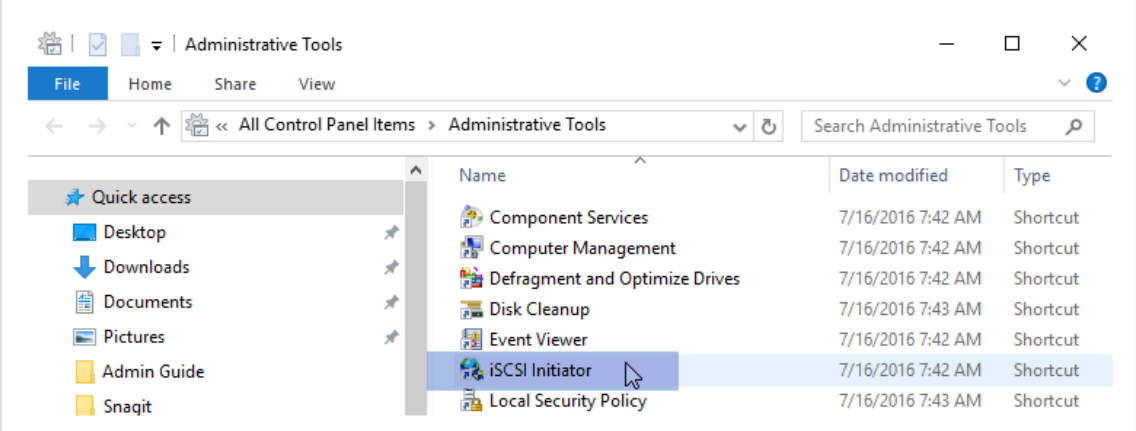

3 In the **Target** field, enter the appliance IP address and click **Quick Connect...**.

The **Discovered targets** field populates with a list of iSCSI LUN targets.

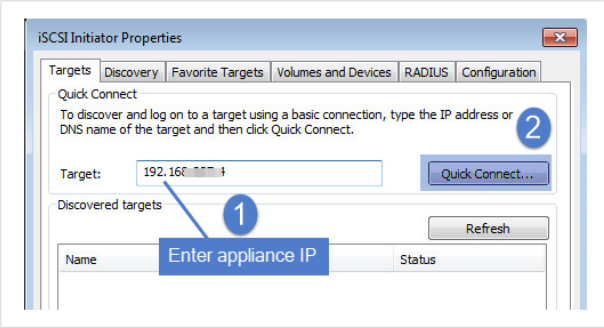

4 Select the iSCSI target from the list.

5 The iSCSI target is discovered and connected to the local machine. Click **Done**.

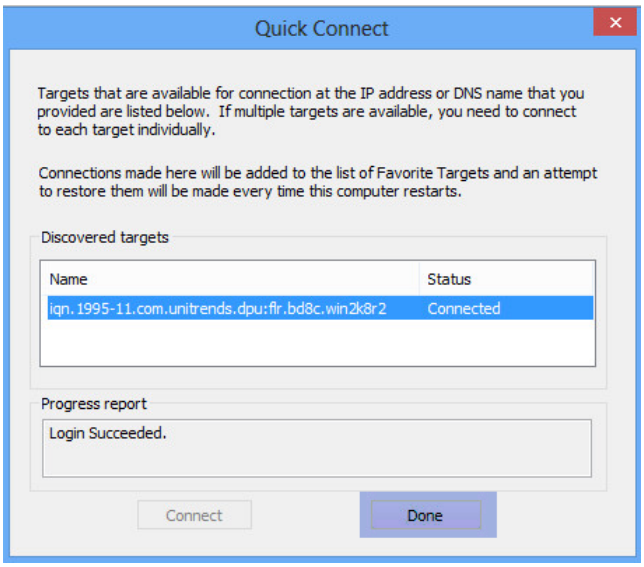

6 Use Disk Manager or diskpart to verify that the mounted iSCSI disk is online. If not, bring the drive online.

7 Return to the iSCSI Initiator. On the Volumes and Devices tab, click **Auto Configure** to map drives from the iSCSI target to the local machine (or map them manually if you prefer).

> **Note:** Volumes are assigned letters during recovery that do not necessarily match the letters from the original disks.

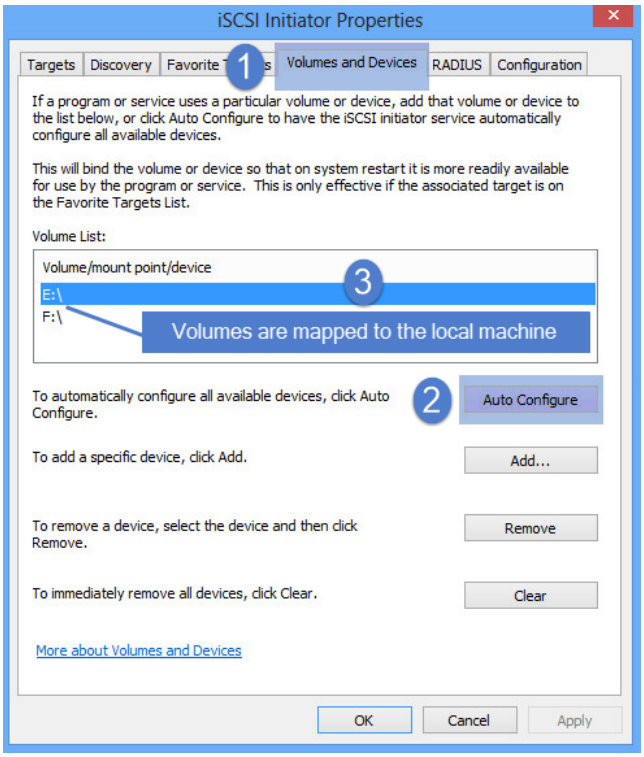

8 Access the files under the mapped drives and move them to another location on the local machine.

Note: The Windows file explorer contains a setting to hide protected/system files. Turn off this setting to access all files.

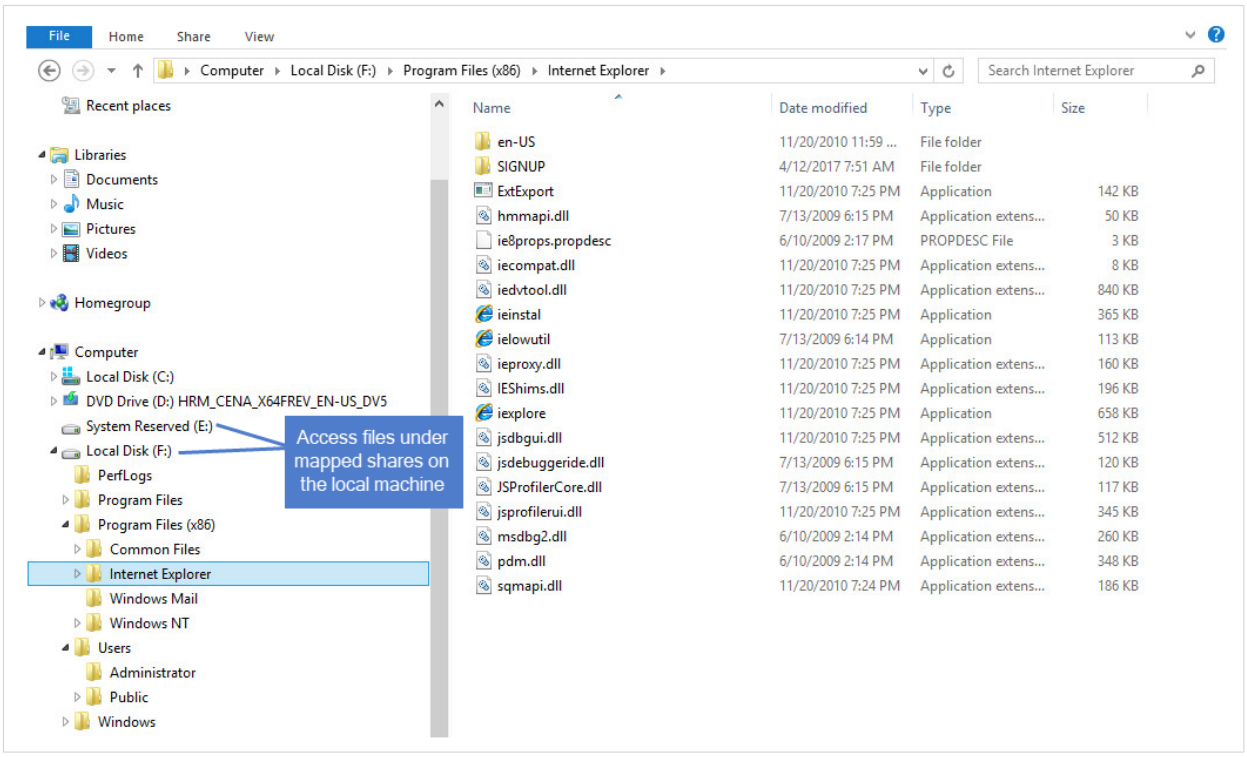

9 Use the iSCSI Initiator to disconnect from the LUN.

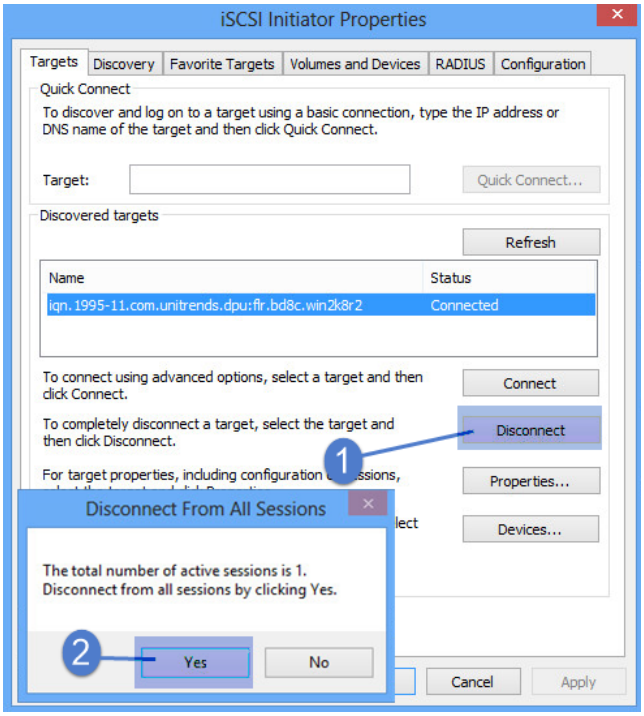

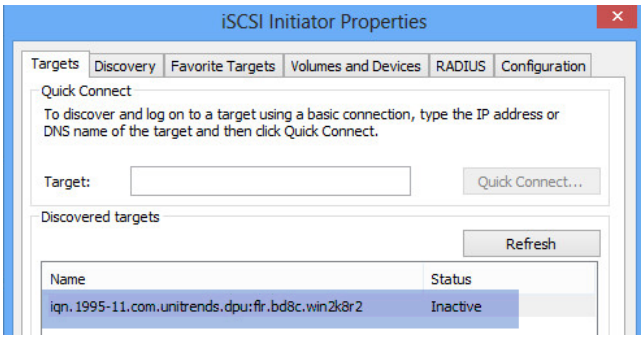

10 Proceed to "Step 4: Remove the recovery object from the appliance".

Step 4: Remove the recovery object from the appliance

To ensure optimal performance, remove the recovery object from the appliance.

WARNING! If you mounted the CIFS share or iSCSI LUN, be sure to unmount it from the target before you remove the recovery object. Removing the recovery object while the target is still connected causes undesired results and errors on the target machine.

To remove a file-level recovery object

1 Select **Recover** and click the **File Level Recovery** tab.

2 Select the recovery object.

3 Click **Remove**.

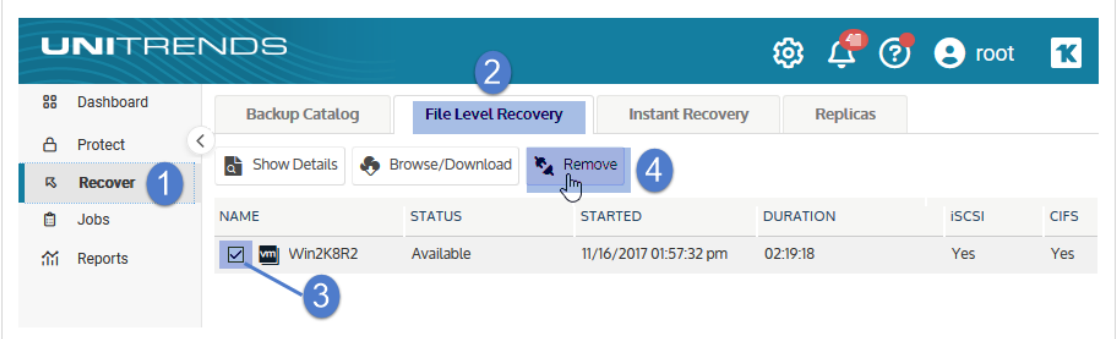

4 Click **Confirm** to continue. The object is removed and no longer displays on the File Level Recovery tab.

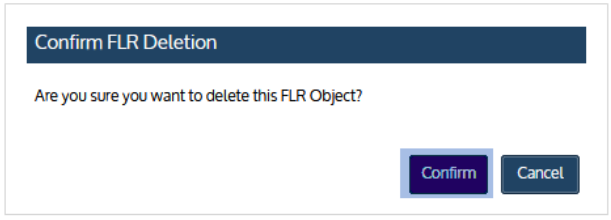

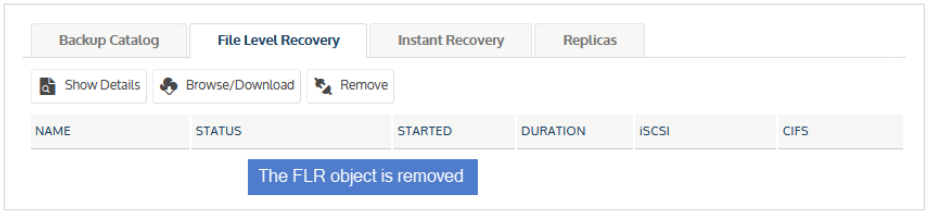

Linux file-level recovery

Use the following procedures to recover files from a backup, imported backup copy, or hot backup copy of a Linux VM.

- "Step 1: Ensure prerequisites have been met"
- "Step 2: Create the recovery object"
- "Step 3: Recover files"
- "Step 4: Remove the recovery object from the appliance"

Step 1: Ensure prerequisites have been met

The following requirements and considerations apply to recovering files from a host-level backup or host-level backup copy of a Linux VM by creating a recovery object:

Prerequisite or consideration	Description
Supported recovery methods	To recover files from a host-level backup or copy, the appliance creates a recovery object that contains the backup's files. For some Linux VMs, this object is also exposed as a CIFS (Samba) share and/or an iSCSI LUN on the backup appliance. After you create the recovery object, you will view it on the File Level Recovery tab to see whether the CIFS and iSCSI options are available. You can recover files from this object in several ways. Options include: • Browse the recovery object and download selected files to a *.zip* file. This is the simplest method. • Mount the CIFS share on a recovery target machine. From the target machine,

Prerequisite or consideration	Description
	select files to recover. • Mount the iSCSI LUN on a recovery target machine. From the target machine, select files to recover.
Linux kernel version	The Linux VM cannot be running a higher Linux kernel version than the Unitrends appliance. If your Linux VM is running a higher Linux kernel version, recover the virtual machine instead as described in "Recovering a virtual machine" on page 796.
Configuration of the protected Linux VM	These requirements apply to the original Linux VM whose backup or backup copy will be used for the recovery: • Software RAID (mdraid) configurations are not supported. If the VM is configured with software raid, you cannot recover files. Recover the entire VM instead, as described in "Recovering a virtual machine" on page 796. • For NTFS, FAT32, ext2, ext3, ext4, or xfs Linux file systems, you can recover by downloading to a *.zip* file or by mounting the CIFS share. • For other file systems, including Linux mounted volumes, you must mount the iSCSI LUN to access and recover files. For iSCSI requirements, see "Requirements for recovery by mounting the iSCSI LUN".
Requirements for recovery by mounting the iSCSI LUN	To recover by mounting the iSCSI LUN, the following prerequisites and considerations apply: • The iscsi-initiator-utils package must be installed on the recovery target. • For the recovery, iSCSI disks are writable and a 1 GB write limit is enforced. Errors display on the recovery target machine if more than 1 GB is required. In this case, you cannot recover by using iSCSI. Recover files by downloading to a *.zip* file or by mounting the CIFS share, or perform a VM recovery.

Step 2: Create the recovery object

Use one of these procedures to create the recovery object:

Note: If a previously-created recovery object is still mounted for the VM, you must remove it before creating a new one.

- "To create the recovery object and recover from a backup or imported backup copy" on page 836 – Run on the backup appliance to recover from a backup or imported backup copy.

- "To create the recovery object and recover from a hot backup copy by using the source backup appliance" on page 838 – Run on the backup appliance to recover from a backup copy that resides in the Unitrends Cloud or resides on the target appliance.
- "To create the recovery object and recover from a hot backup copy by using the target appliance" on page 839 – Run on the target appliance to recover from a backup copy that resides on that target appliance.

To create the recovery object and recover from a backup or imported backup copy

Run this procedure on the backup appliance to recover from a backup or imported backup copy.

1. Log in to the backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the asset and select the backup or imported backup copy from which you want to recover files.

   (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
4. Click **Recover Files**.

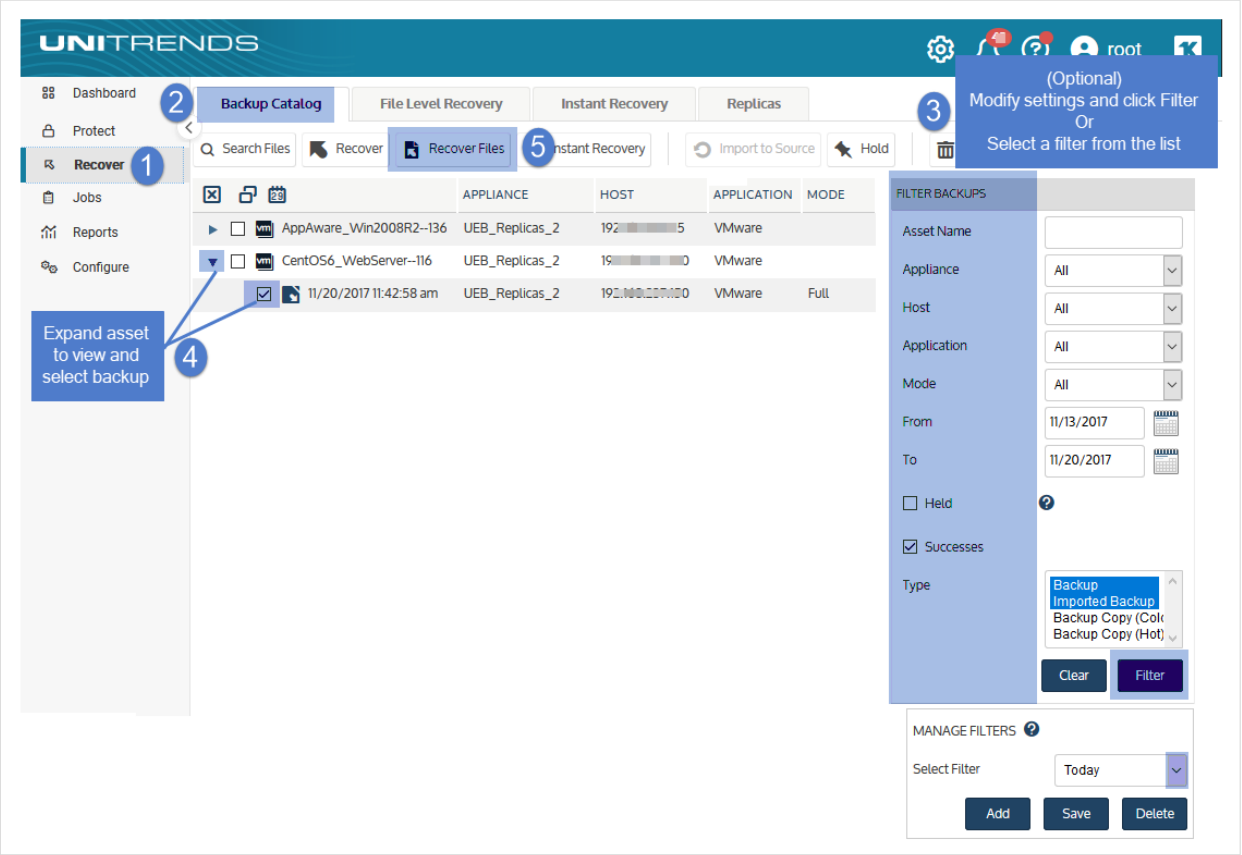

5 Click **Confirm** to continue. The appliance creates the recovery object.

Note: If you receive an error on a Unitrends Backup appliance while creating the recovery object, increase the memory allocation for the Unitrends Backup VM by using the host that manages it.

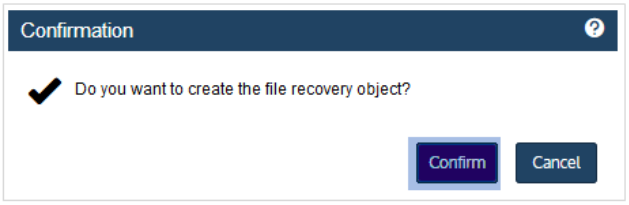

6 Click **View FLR**.

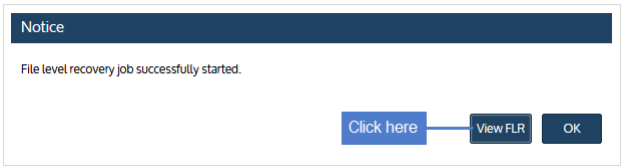

Proceed to "Step 3: Recover files" on page 841.

To create the recovery object and recover from a hot backup copy by using the source backup appliance

Run this procedure on the backup appliance to recover from a backup copy that resides in the Unitrends Cloud or resides on the target appliance.

1 Log in to the source backup appliance.

2 Select **Recover** and click the **Backup Catalog** tab.

3 Use Filter Backups to the right to display hot backup copies:

- In the Type box, select **Backup Copy (Hot)**.
- (Optional) Select other filter options. For details, see "Working with custom filters" on page 67.
- Click **Filter**.

4 Expand the asset and select the hot backup copy from which you want to recover files.

5 Click **Recover Files**.

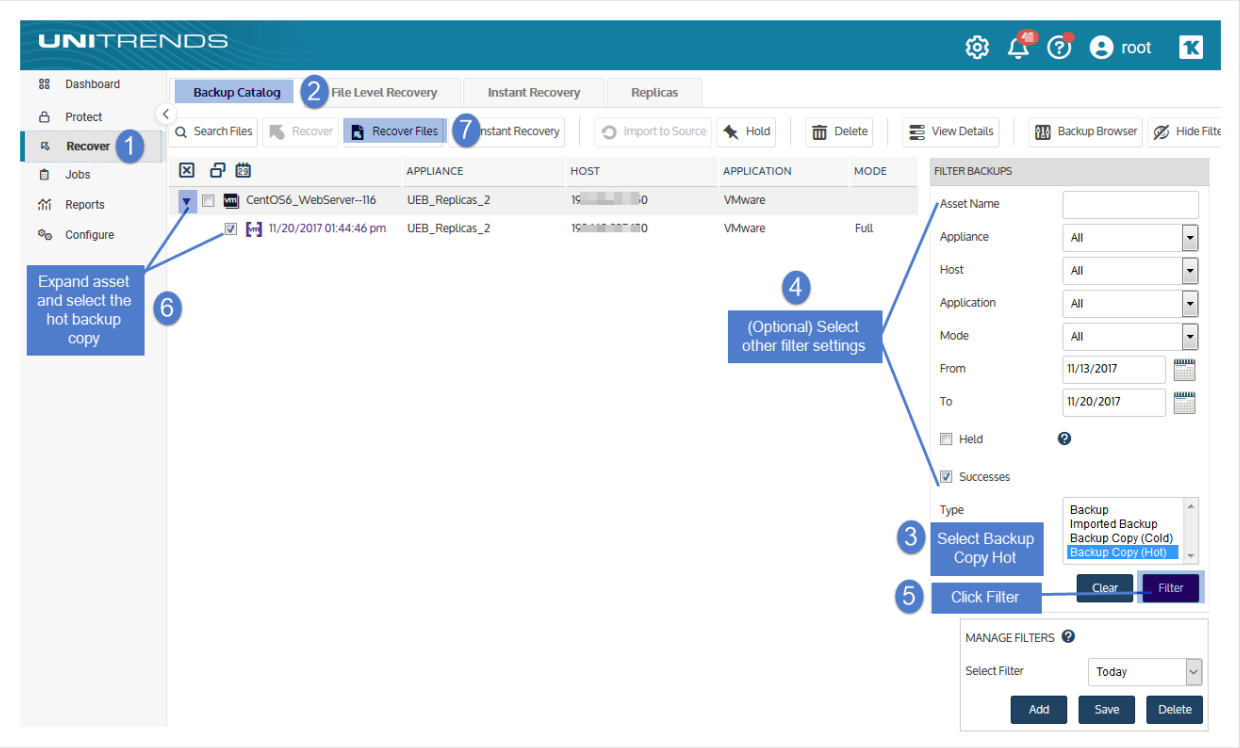

6 Click **Confirm** to continue. The appliance creates the recovery object in the Cloud or on the target appliance.

Notes:

- If your appliance is a Unitrends Backup virtual appliance and you receive an error while creating the recovery object, increase the memory allocation for the Unitrends Backup VM using the host that manages

it.

- Recovery objects created in the Unitrends Cloud are automatically removed after 96 hours.

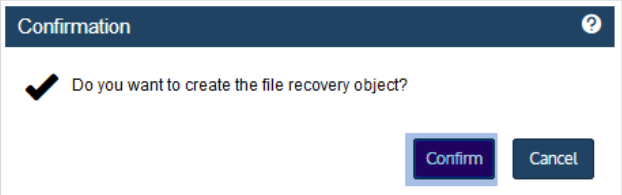

7 Click **View FLR** to view the recovery object on the File Level Recovery tab. The recovery object Name is *AssetName (on the target)*.

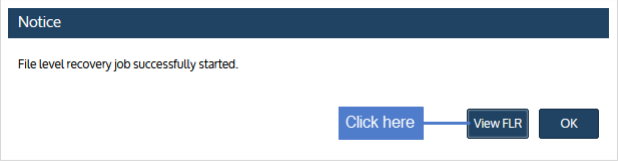

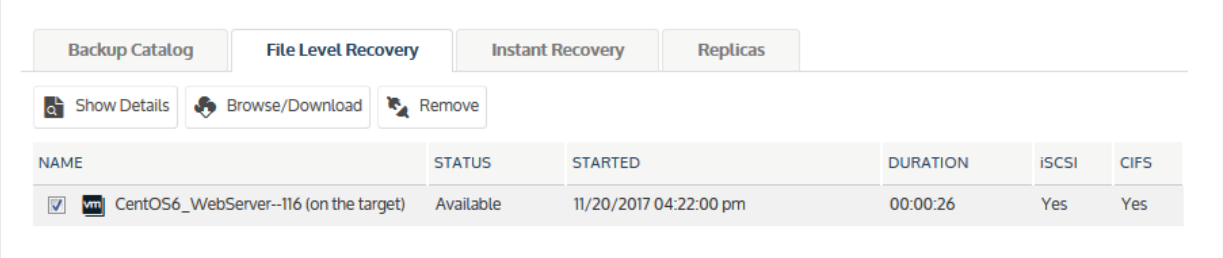

8 Proceed to "Step 3: Recover files" on page 841.

To create the recovery object and recover from a hot backup copy by using the target appliance

Run this procedure on the target appliance to recover from a backup copy that resides on that target appliance.

1 Log in to the backup copy target appliance.

2 Select **Recover** and click the **Backup Catalog** tab.

3 Use Filter Backups to the right to display hot backup copies:

- In the Type box, select **Backup Copy (Hot)**.
- (Optional) Select other filter options. For details, see "Working with custom filters" on page 67.
- Click **Filter**.

4 Expand the asset and select the hot backup copy from which you want to recover files.

5 Click **Recover Files**.

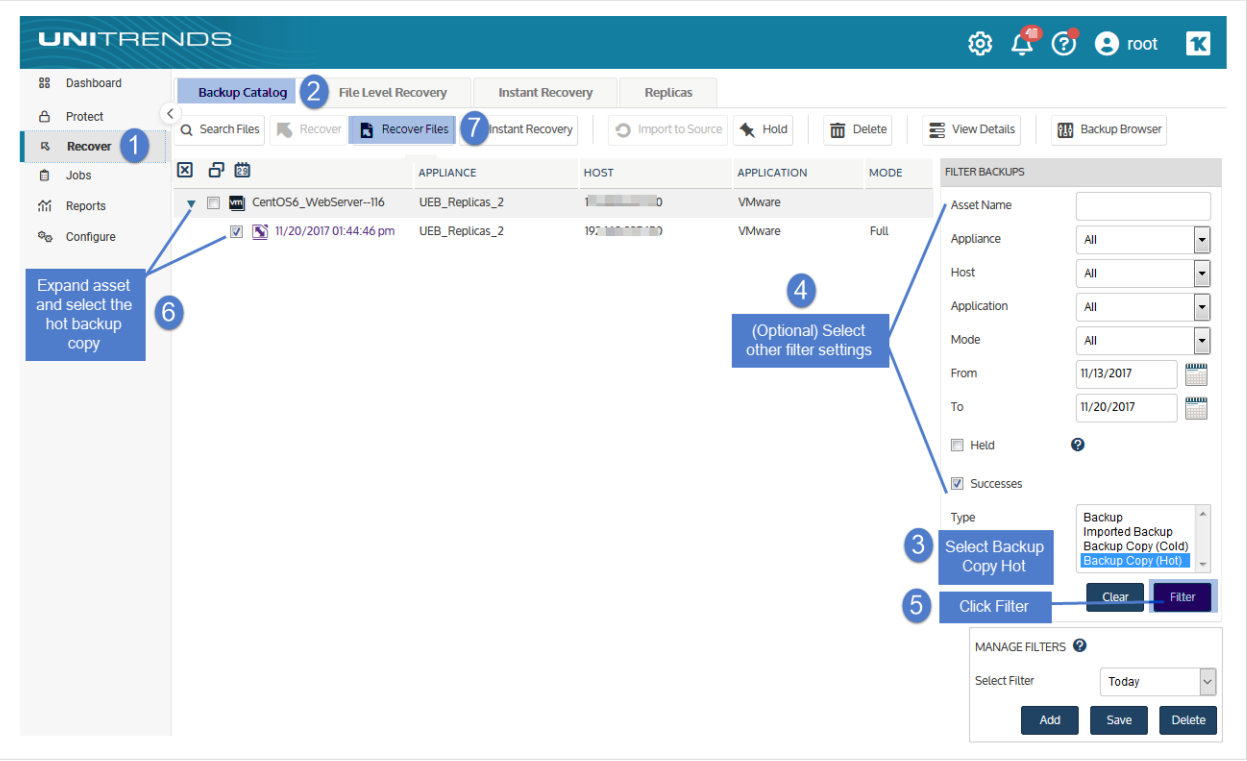

6 Click **Confirm** to continue. The appliance creates the recovery object on the target appliance.

Note: If your appliance is a Unitrends Backup virtual appliance and you receive an error while creating the recovery object, increase the memory allocation for the Unitrends Backup VM using the host that manages it.

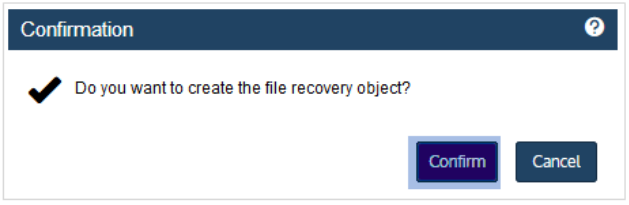

7 Click **View FLR** to view the recovery object on the File Level Recovery tab.

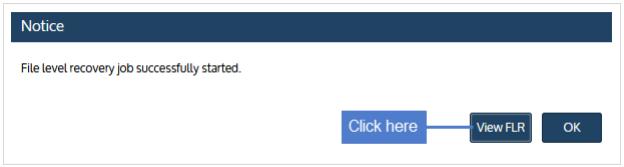

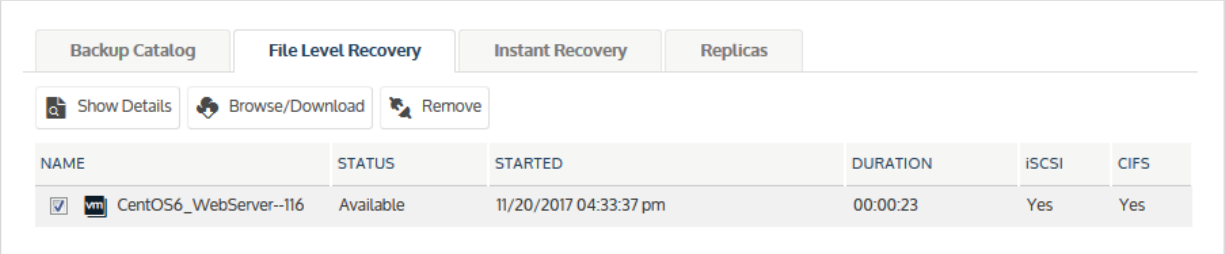

Proceed to "Step 3: Recover files".

Step 3: Recover files

Use one of the following procedures to recover files. For a description of each method, see "Recovery procedures overview" on page 809.

- "To recover files by browsing and downloading to a .zip file" on page 841
- "To recover files by mounting the CIFS share" on page 843
- "To recover files to a Linux machine by mounting the iSCSI LUN" on page 845

To recover files by browsing and downloading to a *.zip* file

1 On the **File Level Recovery** tab, locate the recovery object.

Recovery objects display on the tab with the following details: the name of the VM asset for which the object was created, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

2 Select the recovery object and click **Browse/Download**.

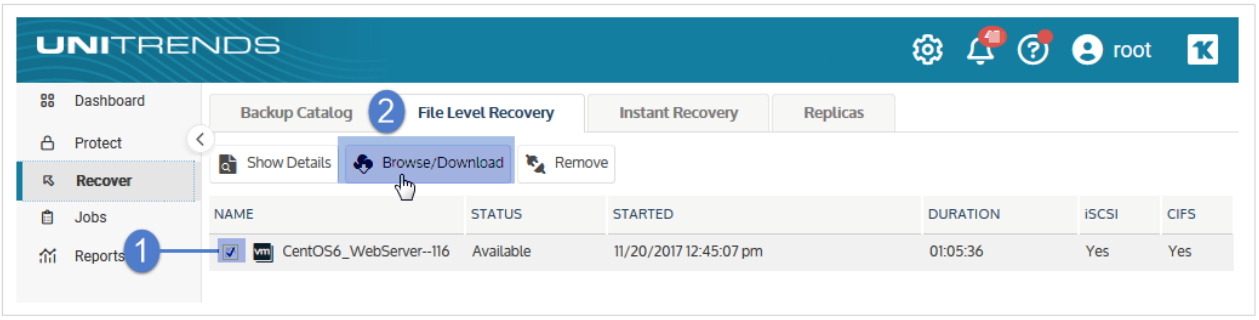

3 In the File Browser, select or drag files and/or directories to recover.

> **Note:** Softlinks (also called *symbolic links*) are excluded from download. If you select a directory that contains files and softlinks, only the files are downloaded.

4 Click **Download**.

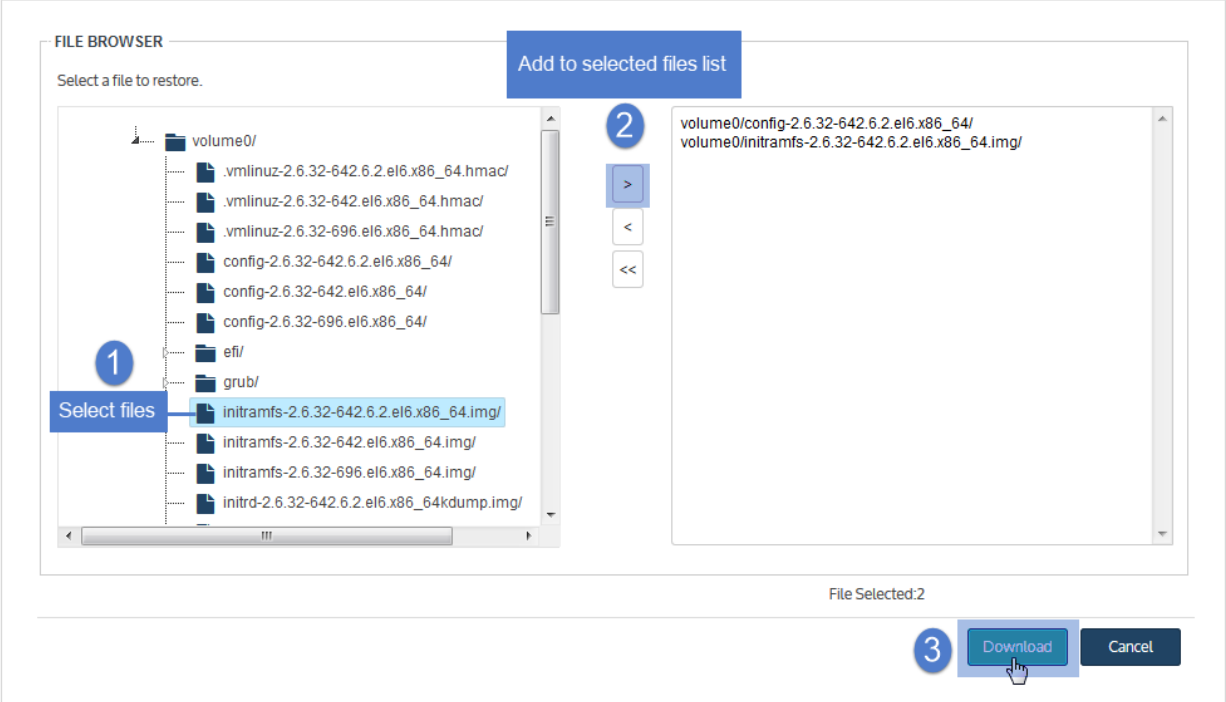

5 Click **Confirm** to download the selected files to a *.zip* file. The *.zip* file is downloaded to your browser's default location.

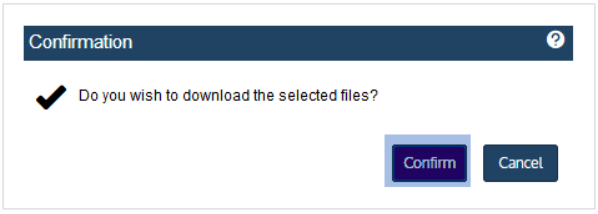

Notes:

- Volumes are assigned numbers during recovery that do not necessarily match the numbers from the original disk.
- The duration of the download is impacted by various factors, such as the size of the files, bandwidth, and download speed.
- Persistent browser and UI sessions are required to create the *.zip* file in the browser's default download location. If you close the browser or UI session during the recovery, you must run a new job.

6 When the download completes, the *Unitrends-Restore.zip* file displays in the browser. Select whether to open or save the file.

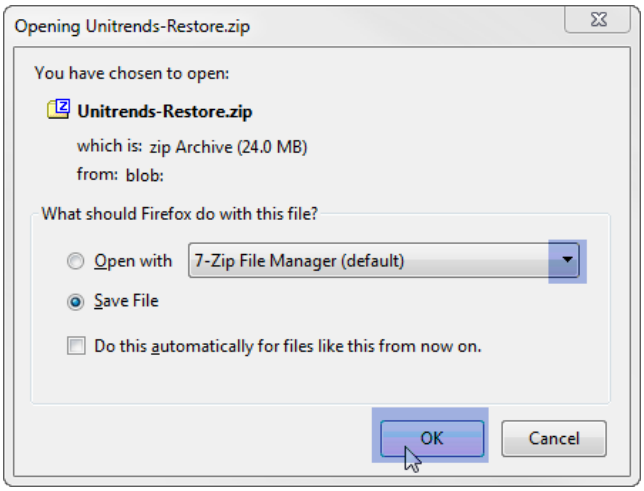

7 Access the recovered files in the download location and move them to another location on the local machine.

> **Note:** The Windows file explorer contains a setting to hide protected/system files. Turn off this setting to access all files.

8 Proceed to "Step 4: Remove the recovery object from the appliance" on page 847.

To recover files by mounting the CIFS share

1 Select **Recover** and click the **File Level Recovery** tab.

Recovery objects display with the following details: the name of the VM asset for which the object was created, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

2 Select the recovery object and click **Show Details**.

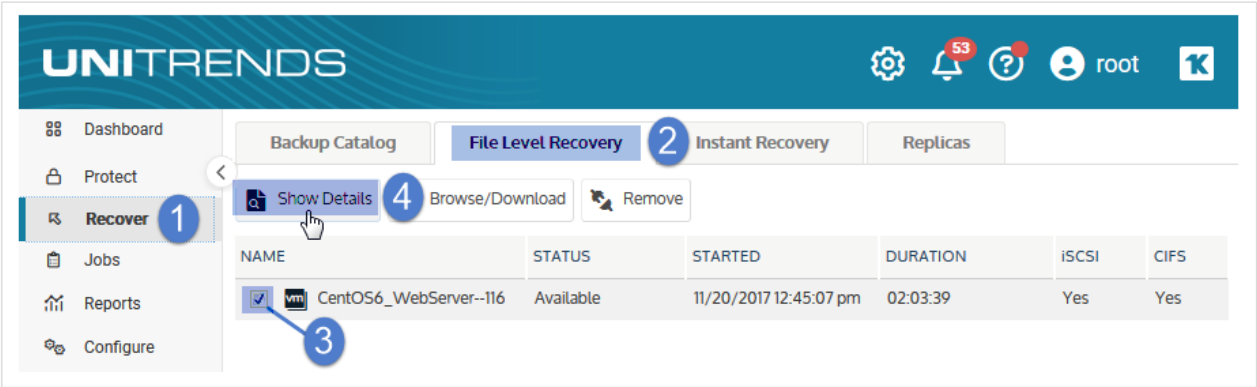

3 Note the CIFS path that displays in the File Level Recovery Details window. You will need this path to mount the CIFS share on the target machine.

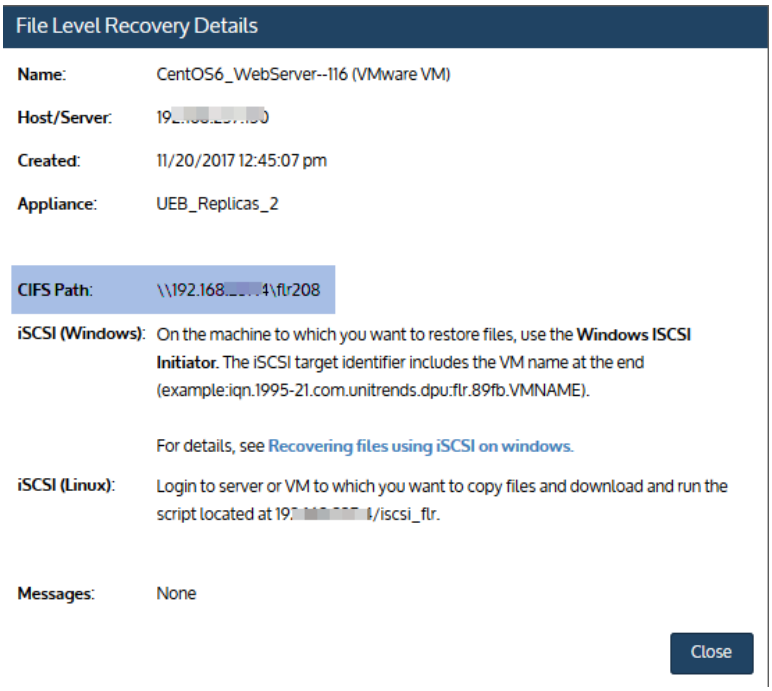

4 Log in to the recovery target workstation.

5 Enter the CIFS path into a file browser on the recovery target.

6 Browse the share to locate the files you want to recover.

Note: Volumes are assigned numbers during recovery that do not necessarily match the numbers from the original disks.

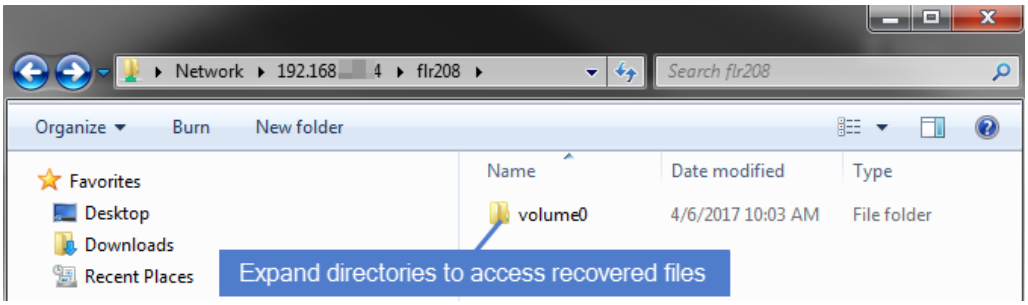

7 Move selected files to another location on the local machine.

8 Disconnect the network share by right-clicking the share and selecting **Disconnect**.

9 Proceed to "Step 4: Remove the recovery object from the appliance" on page 833.

To recover files to a Linux machine by mounting the iSCSI LUN

Use these steps to mount the iSCSI LUN to the target machine and copy the files.

1 In the appliance UI, select **Recover** and click the **File Level Recovery** tab.

Recovery objects display with the following details: the name of the VM asset for which the object was created, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

2 Select the recovery object and click **Show Details**.

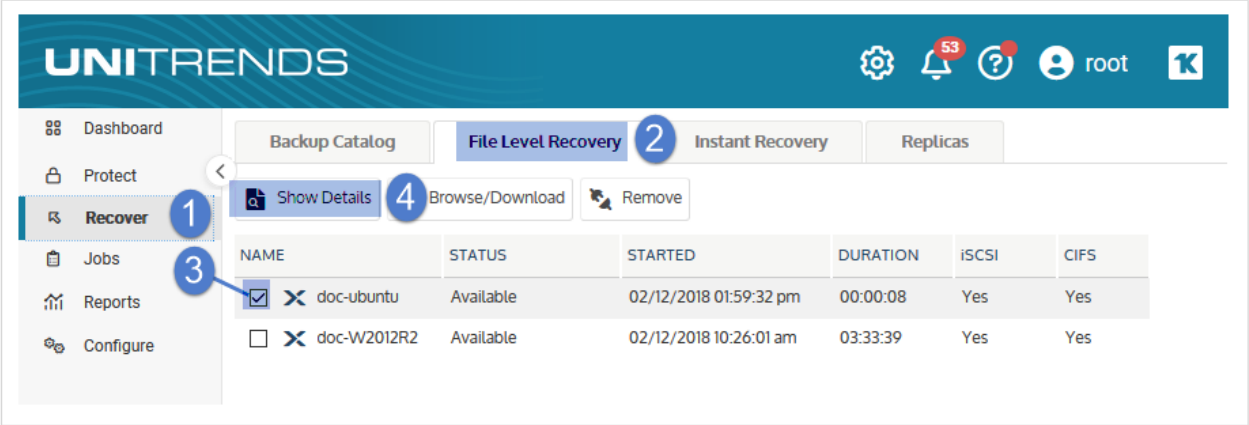

3 Note the full path of the iSCSI mount point directory that displays in the File Level Recovery Details window. You will need this path to mount the iSCSI object on the target machine. The mount point is normally: /iscsi_flr.

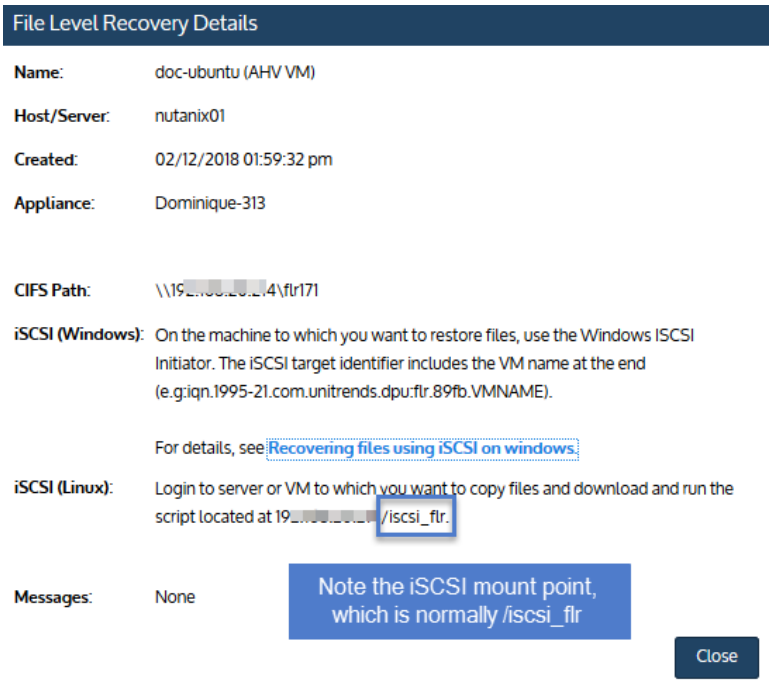

4 Log in to the recovery target.

5 Enter the following command to change to the */tmp* directory:

```
# cd /tmp
```

6 Run the following command to copy the *iscsi_flr* script from the backup appliance:

```
# wget http://<appliance IP>/iscsi_flr
```

7 After the script downloads, add the execute permission:

```
# chmod +x iscsi_flr
```

8 Run the following command to mount the recovery object:

```
# ./iscsi_flr mount
```

9 Enter the appliance IP address:

```
# Enter address of the Unitrends backup system: <appliance IP>
```

10 Enter the full path of the mount point directory. The full path is likely: /iscsi_flr. This procedure uses /iscsi_flr as an example. Be sure to enter the actual mount point that was displayed in the appliance UI.

```
# Enter mount point directory (full path): /iscsi_flr
```

11 Discovered iSCSI targets display. Choose the target that contains the appliance IP by entering its number. In this example, session 1 is the appliance target:

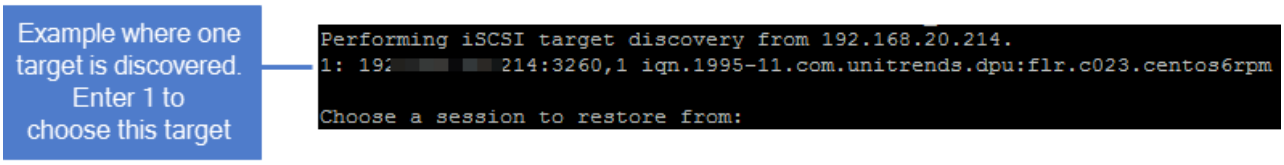


```
# Choose a session to restore from: <sessionNumber>
```

12 Verify that the mounted iSCSI disk is online. If not, bring the drive online.

13 Change to the mount point directory to access the files. For example:

```
#  cd /iscsi_flr
```

14 Move selected files to another location on the local machine.

15 Run the following command from the */tmp* directory to disconnect from the LUN:

```
# ./iscsi_flr unmount
```

16 Proceed to "Step 4: Remove the recovery object from the appliance".

Step 4: Remove the recovery object from the appliance

To ensure optimal performance, remove the recovery object from the appliance.

WARNING! If you recovered by mounting a LUN, be sure to unmount the LUN from the target before you remove the recovery object. Removing the recovery object while the target is still connected causes undesired results and errors on the target machine.

To remove a file-level recovery object

1 Select **Recover** and click the **File Level Recovery** tab.

2 Select the object to remove from the appliance.

3 Click **Remove**.

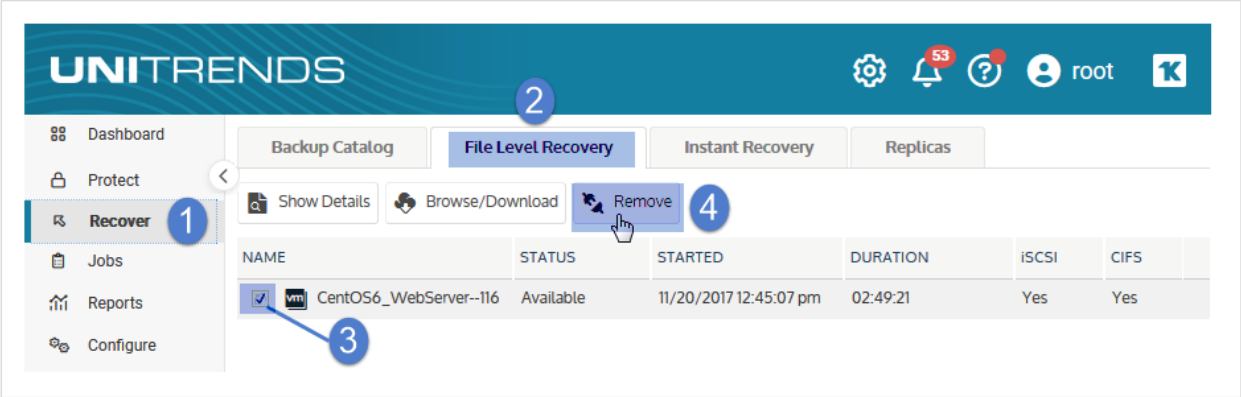

4 Click **Confirm** to continue. The object is removed and no longer displays on the File Level Recovery tab.

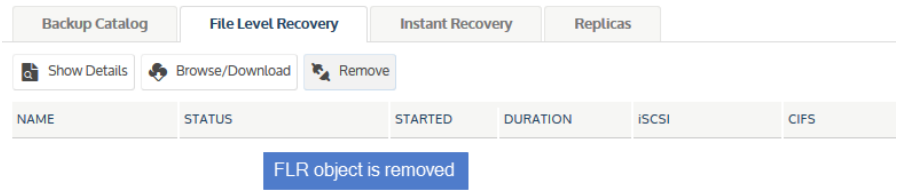

Viewing a file recovery object

After creating a recovery object, you can view it on the Recover page to see whether it is available, details for accessing it, and other information about the object.

To view a file recovery object

1 Select **Recover** and click the **File Level Recovery** tab.

Recovery objects display with the following details: the name of the asset for which the object was created, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

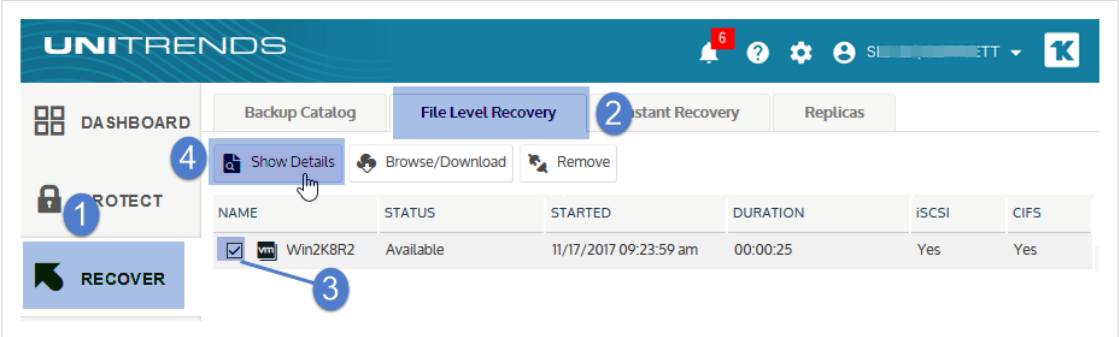

2 To view more information, check the recovery object and click **Show Details**.

The following additional details display, if applicable:

- Path to the CIFs share
- Messages
- Name of the appliance on which the object resides

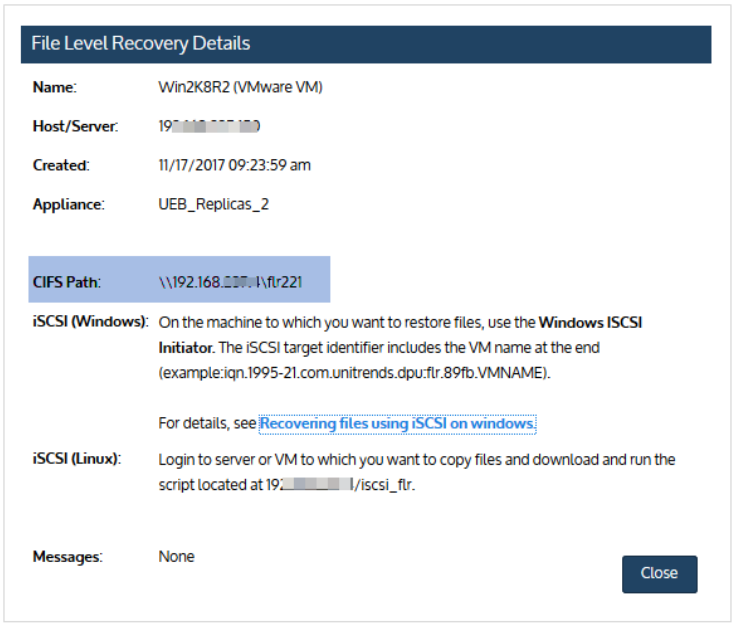

Recovering SQL, Exchange, or SharePoint items from virtual machine backups with Ontrack® PowerControls™

To recover items from a host-level backup or copy, you will use the backup appliance to create a recovery object containing the backup's files, then use Ontrack PowerControls to view and recover application items.

See these procedures for details:

- "To recover SQL items from a backup or imported backup copy"
- "To recover Exchange items from a backup or imported backup copy"
- "To recover SharePoint items from a backup or imported backup copy"

To recover SQL items from a backup or imported backup copy

1. Log in to the backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the asset and select the backup or imported backup copy from which you want to recover items.

   (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
4. Click **Recover Files**.

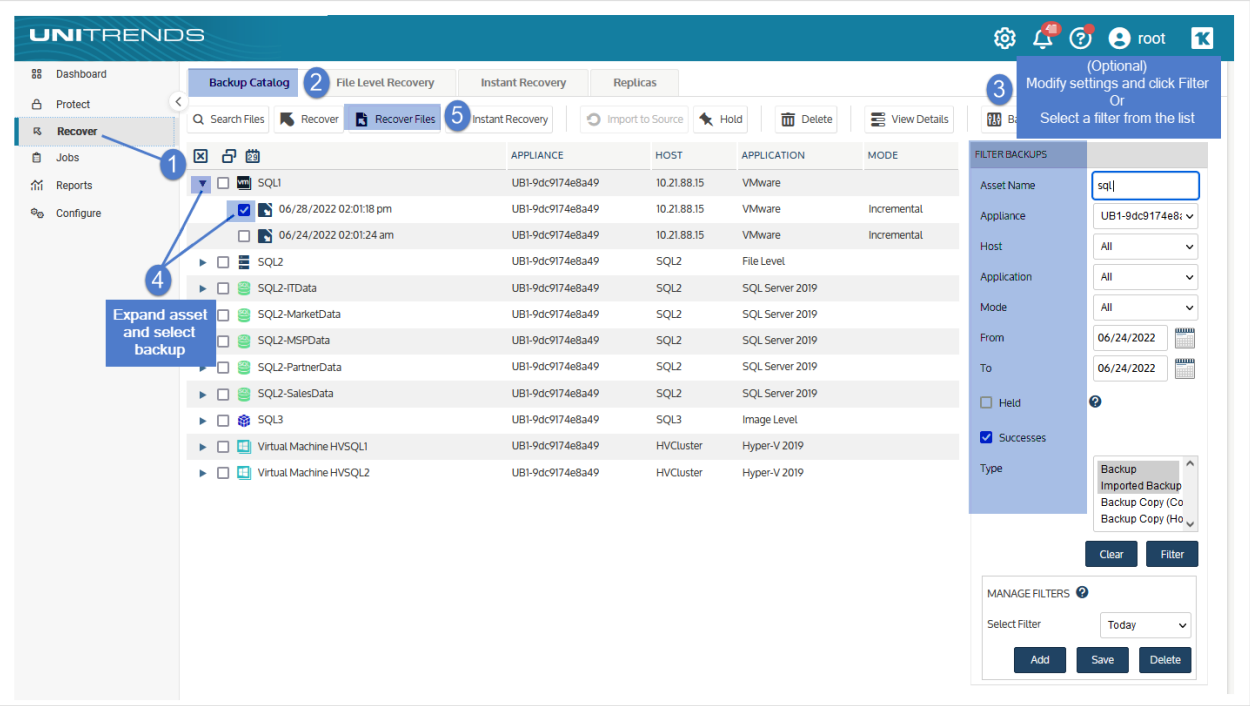

5 Click **Confirm** to continue. The appliance creates the recovery object.

Note: If you receive an error on a Unitrends Backup appliance while creating the recovery object, increase the memory allocation for the Unitrends Backup VM by using the host that manages it.

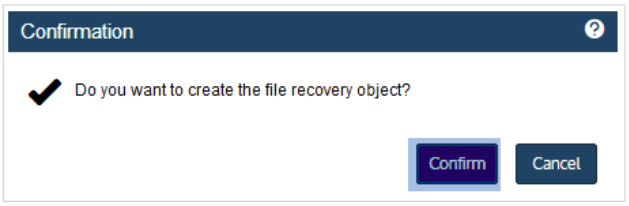

6 Click **View FLR**.

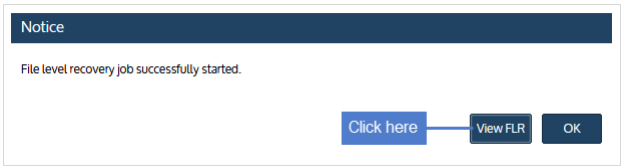

7 On the **File Level Recovery** tab, locate the recovery object.

Recovery objects display on the tab with the following details: the name of the VM asset for which the object was created, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

8 Select the recovery object and click **Browse/Download**.

Note: In most cases you can browse the recovery object directly in the appliance UI by clicking **Browse/Download**. If this method is not supported in your environment, the Browse/Download button is disabled and you will need to mount a CIFS share or iSCSI LUN on a target machine instead, as described in "To recover files by mounting the CIFS share" or "To recover files by mounting the iSCSI LUN" on page 829.

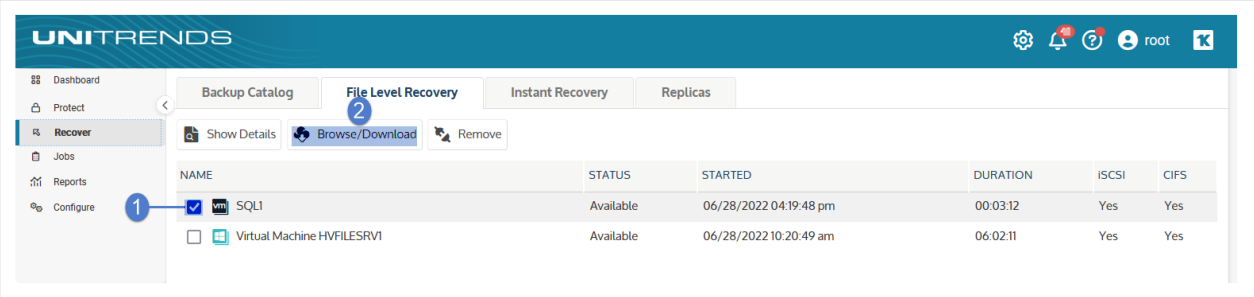

9 In the File Browser, browse to the SQL instance whose items you will recover. Note the filepath as you will browse to it in Ontrack PowerControls® for SQL™.

In our example, the path is *volume1\ProgramFiles\Microsoft SQL Server\MSSQL 15.MSSQLSERVER\MSSQL\DATA*.

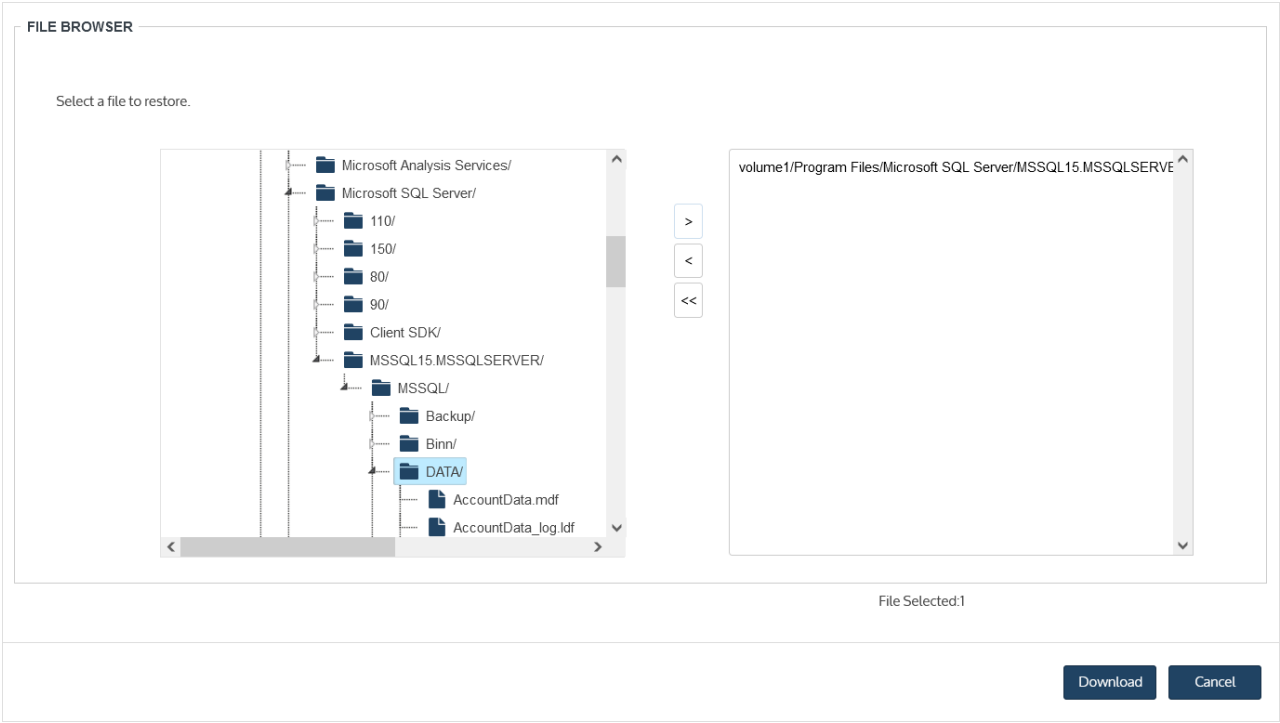

10 Launch Ontrack PowerControls® for SQL™.

11 Select **File > Open Source**. Click **Add**.

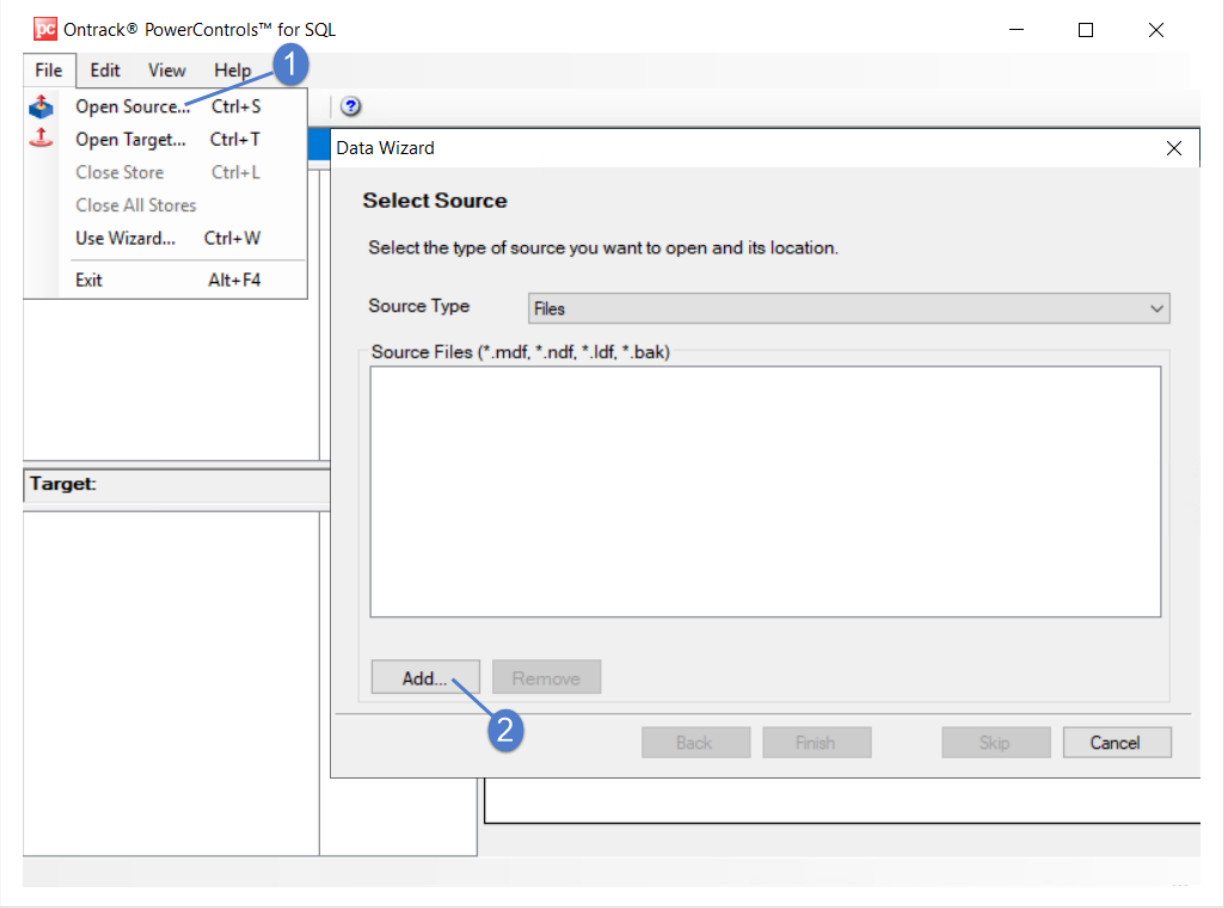

12 Enter \\ followed by the appliance IP address to create a network mapping to your Unitrends appliance. For example, *192.168.225.24*.

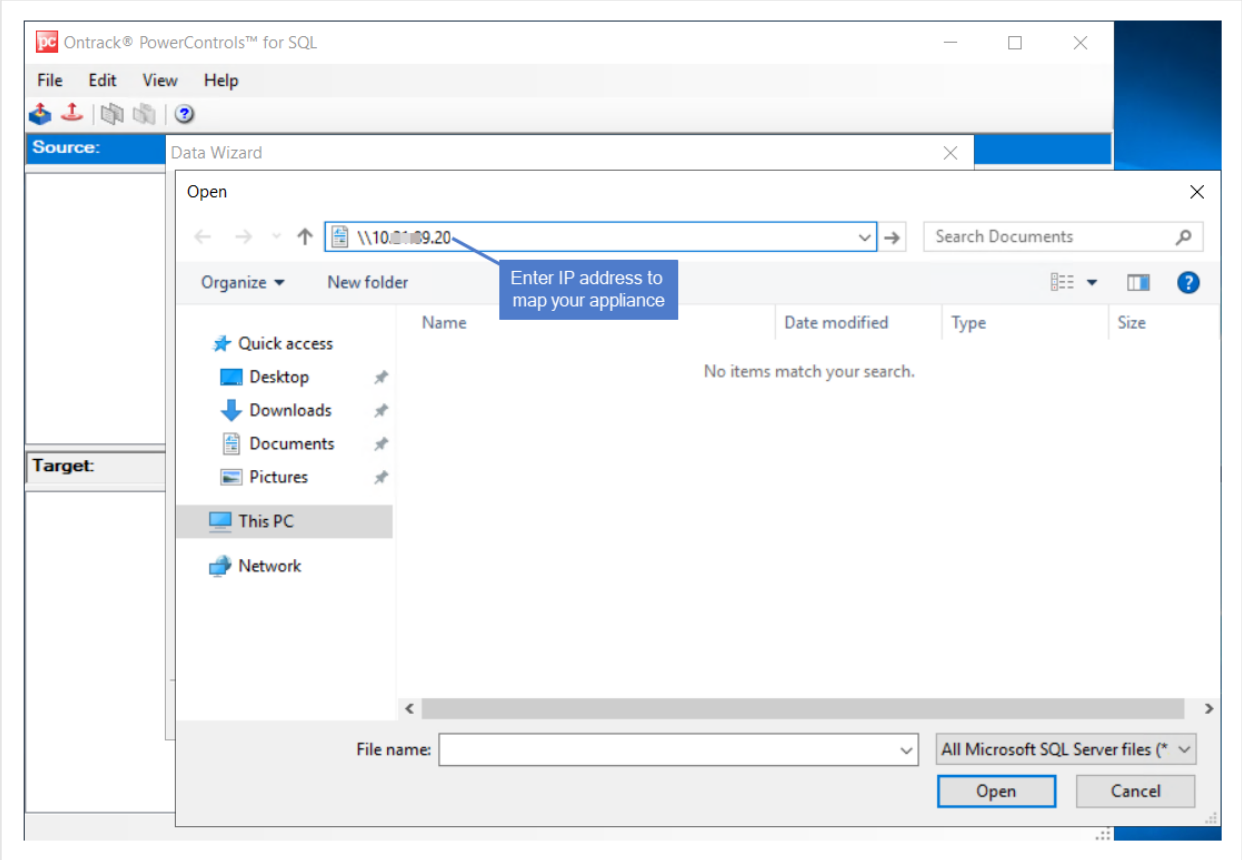

13 Browse under the *flr* folder to the SQL data path you noted above in step 9 .

Note: An *flr* folder exists for each recovery object that displays on the File Level Recovery tab in the appliance UI. In our example, there are two recovery objects: *flr15* and *flr226*.

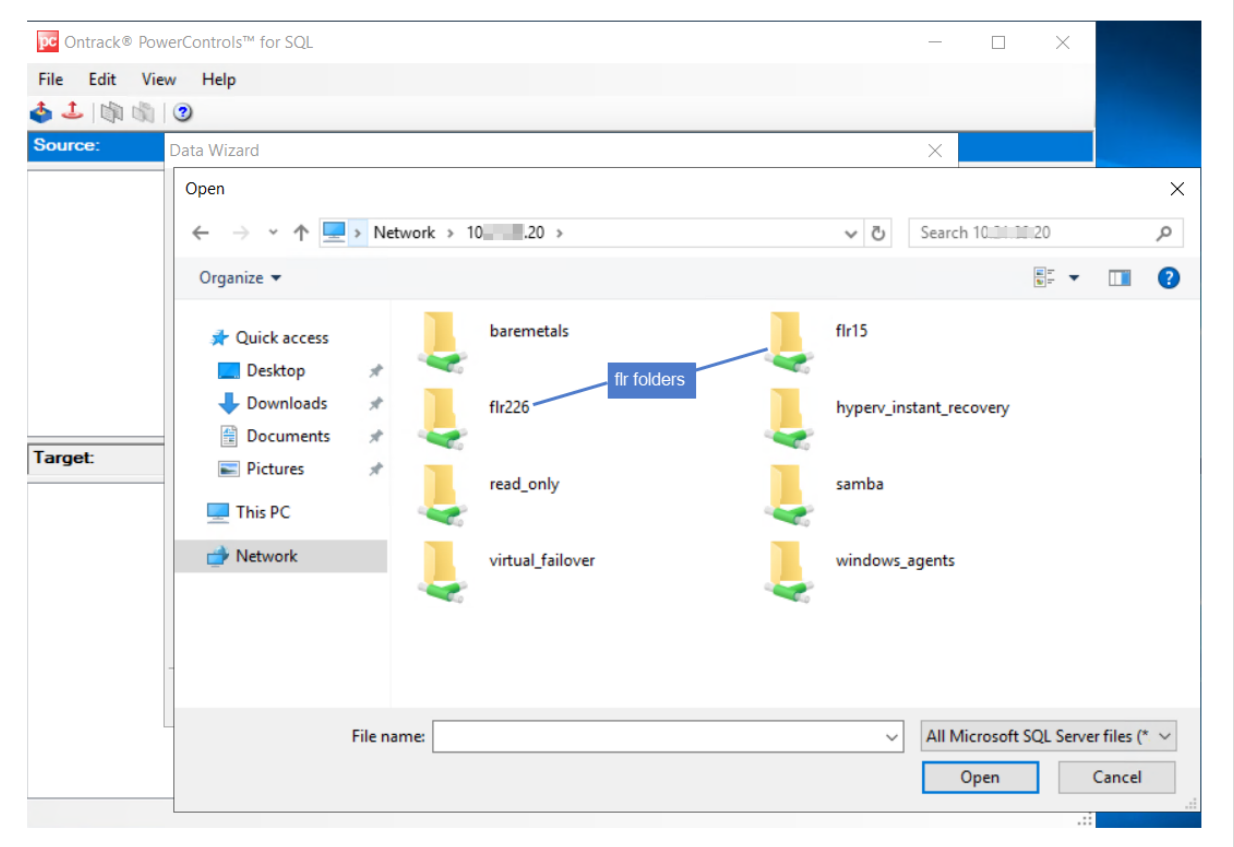

14 Select databases and logs. Click **Open**.

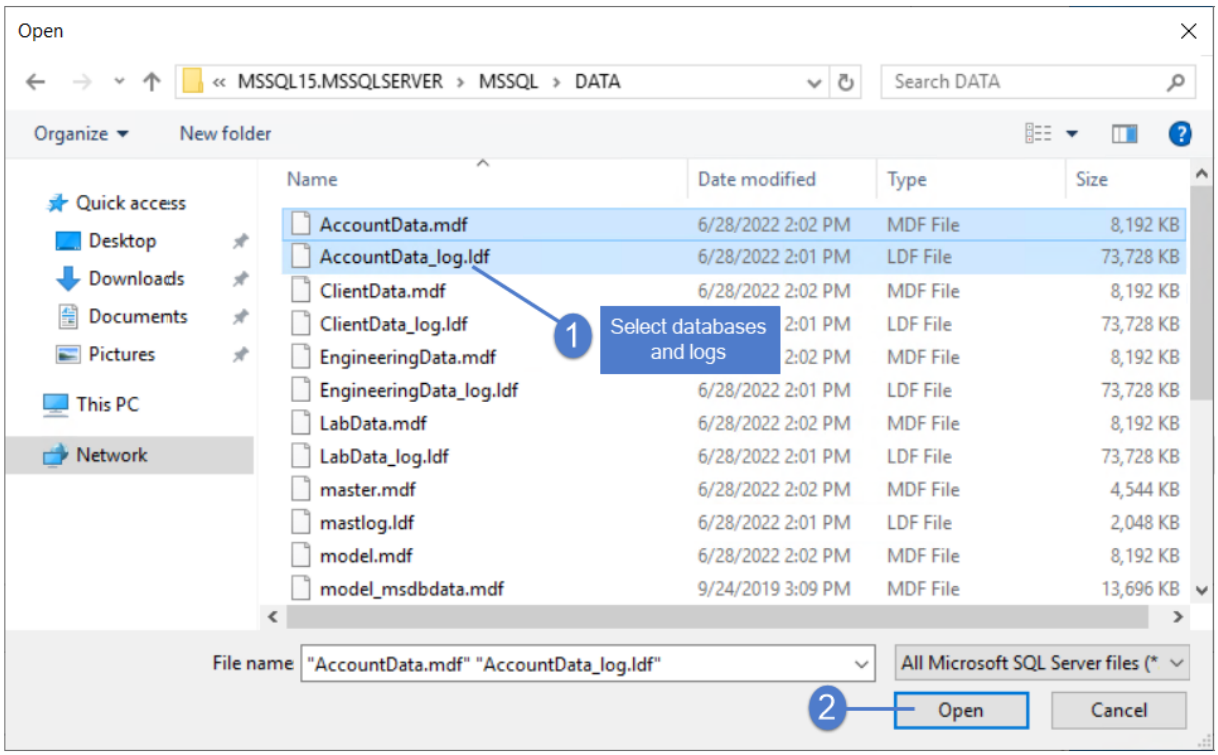

15 Click **Next**.

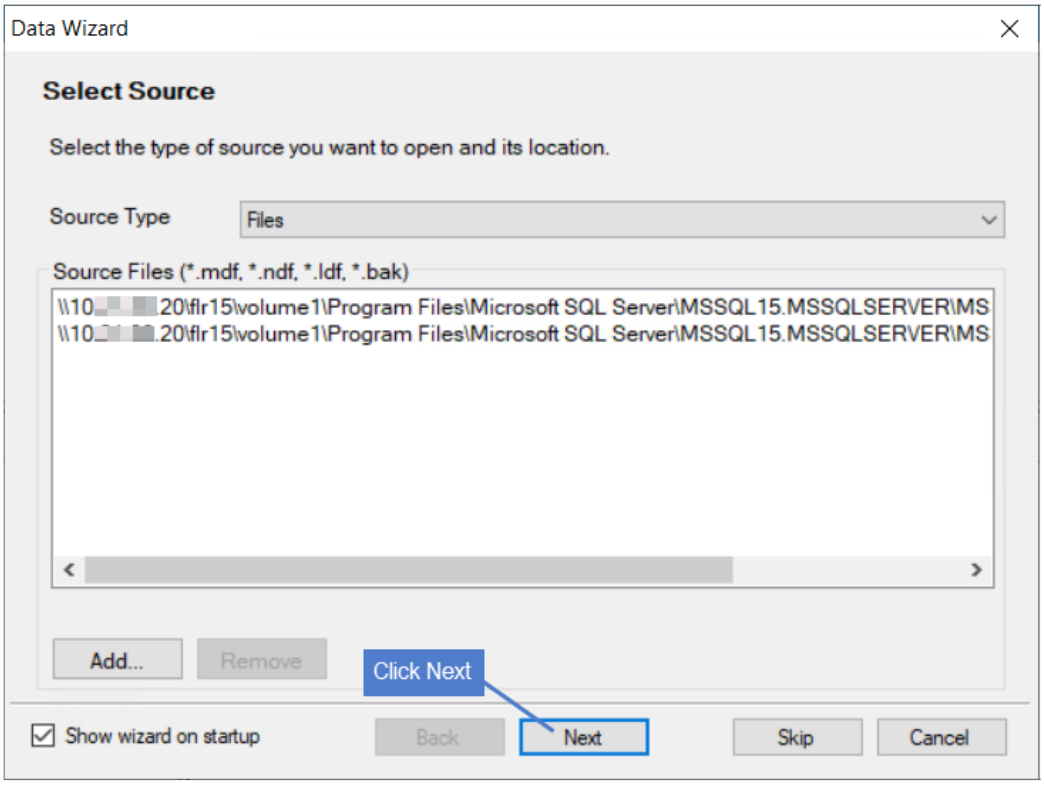

16 Select the target server, authentication method/credentials, and database. Click **Finish**.

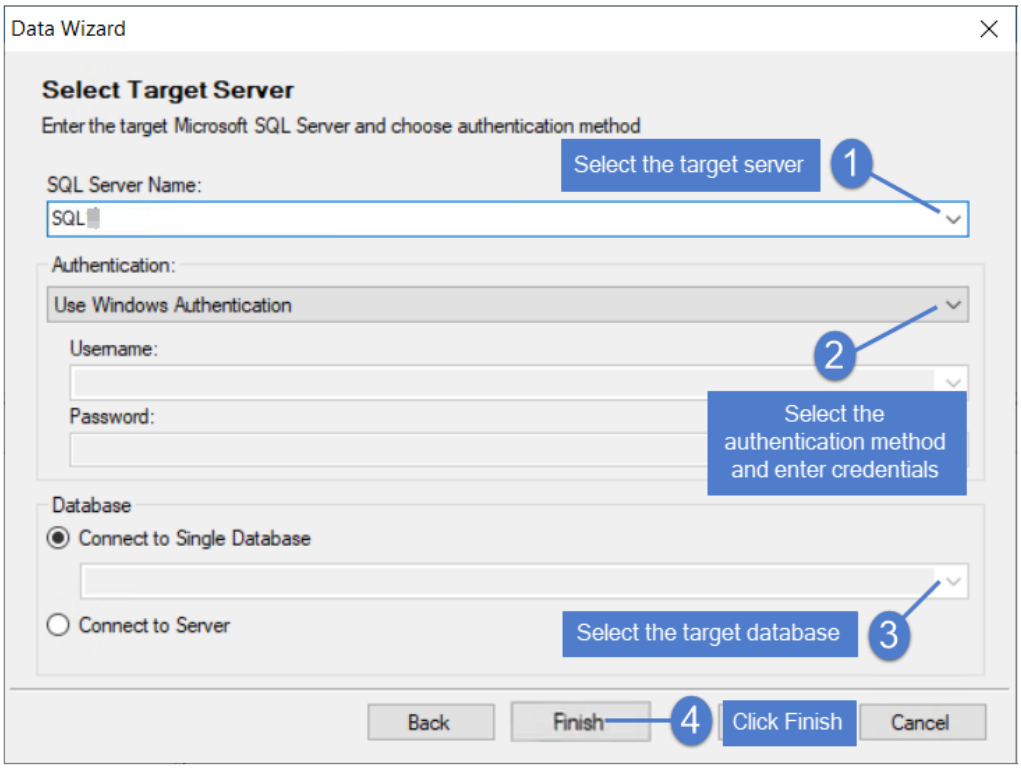

17 Selected databases/logs display in the Source area. If needed, select **File > Open Source > Add** to add more databases/logs.

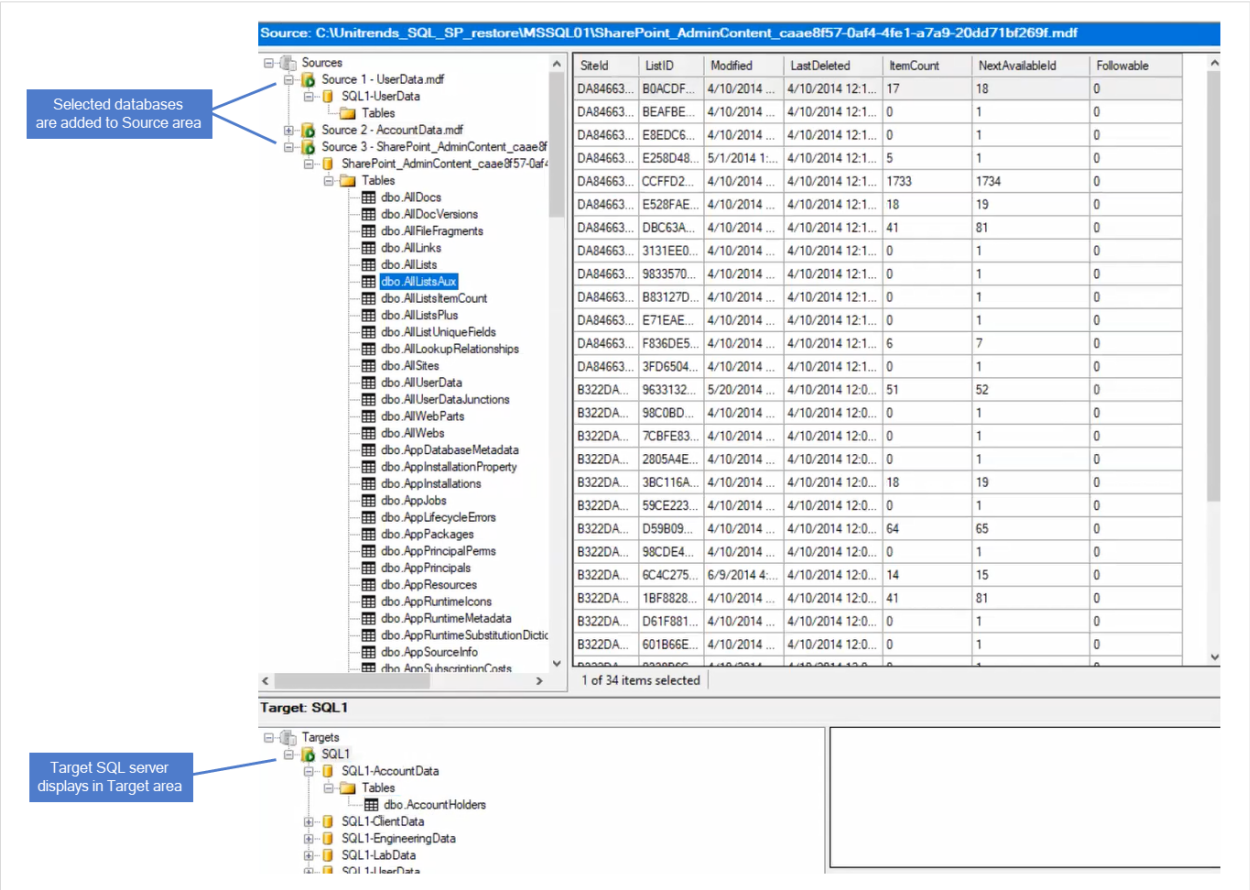

18 Locate tables to restore by browsing databases in the Source area. Copy and paste (or drag) a table from a SQL database backup (source) into a live SQL database (target). For additional detail, see *Restoring Microsoft SQL Server Tables* in the Ontrack® PowerControls™ User Guide for SQL.

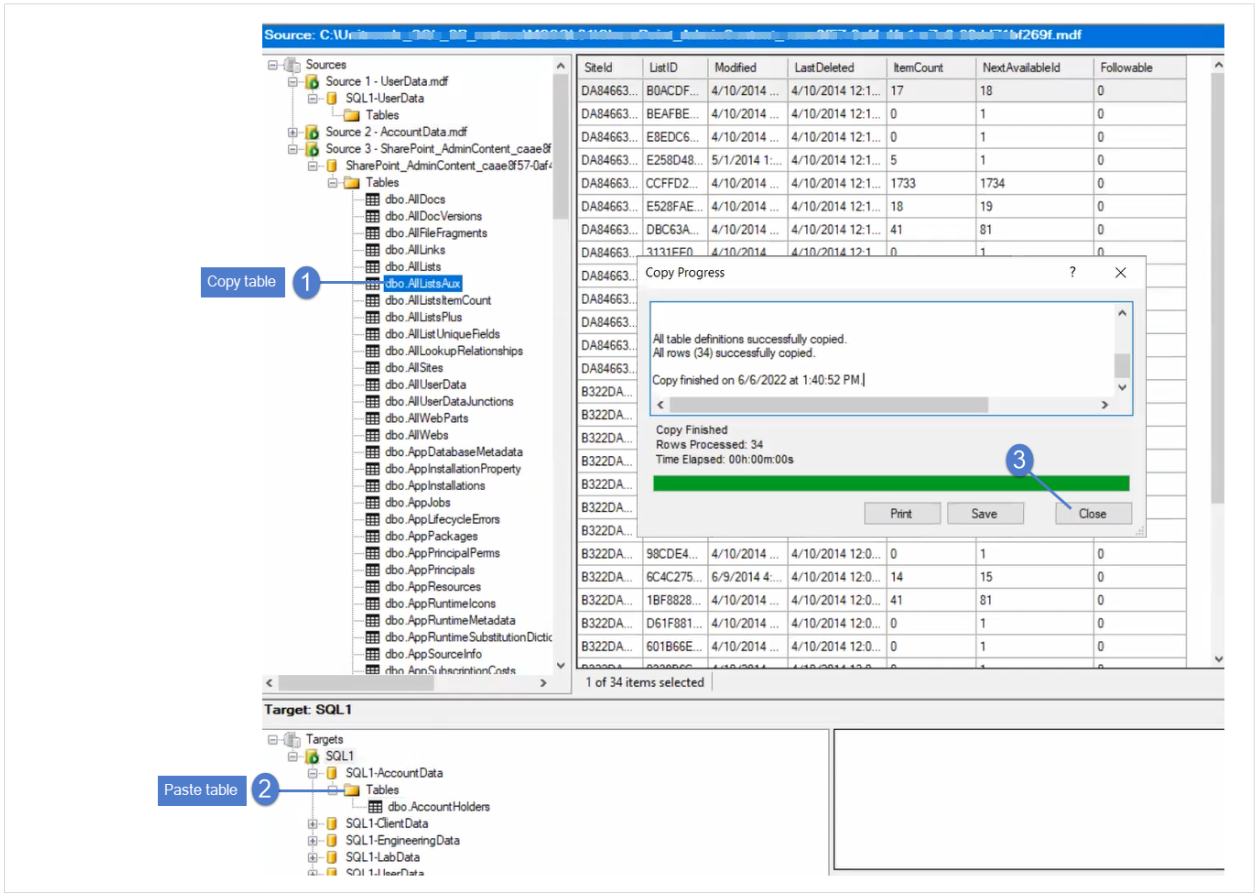

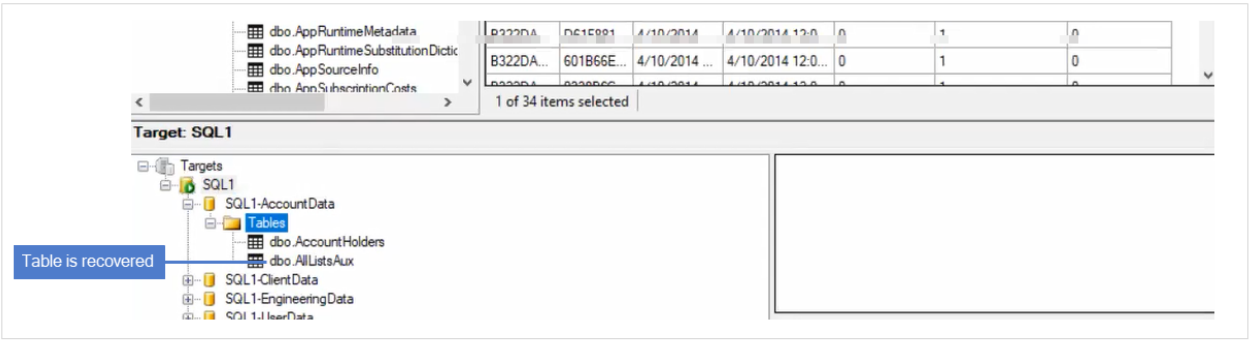

19 When you have finished recovering items, remove the recovery object:

- Return to the appliance UI. On the **File Level Recovery** tab, select the recovery object and click **Remove**. Click **Confirm**.

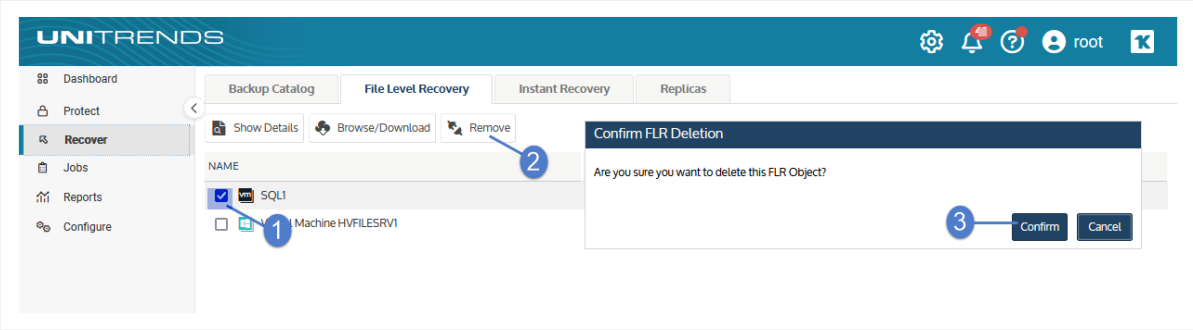

- The recovery object is removed:

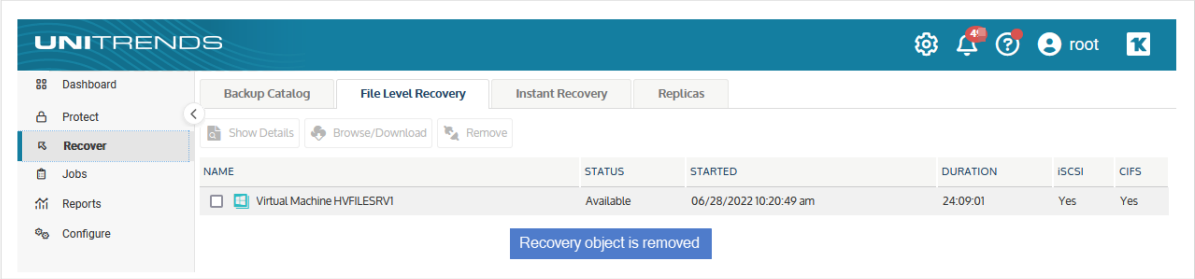

To recover Exchange items from a backup or imported backup copy

1. Log in to the backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the asset and select the backup or imported backup copy from which you want to recover items.

   (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
4. Click **Recover Files**.

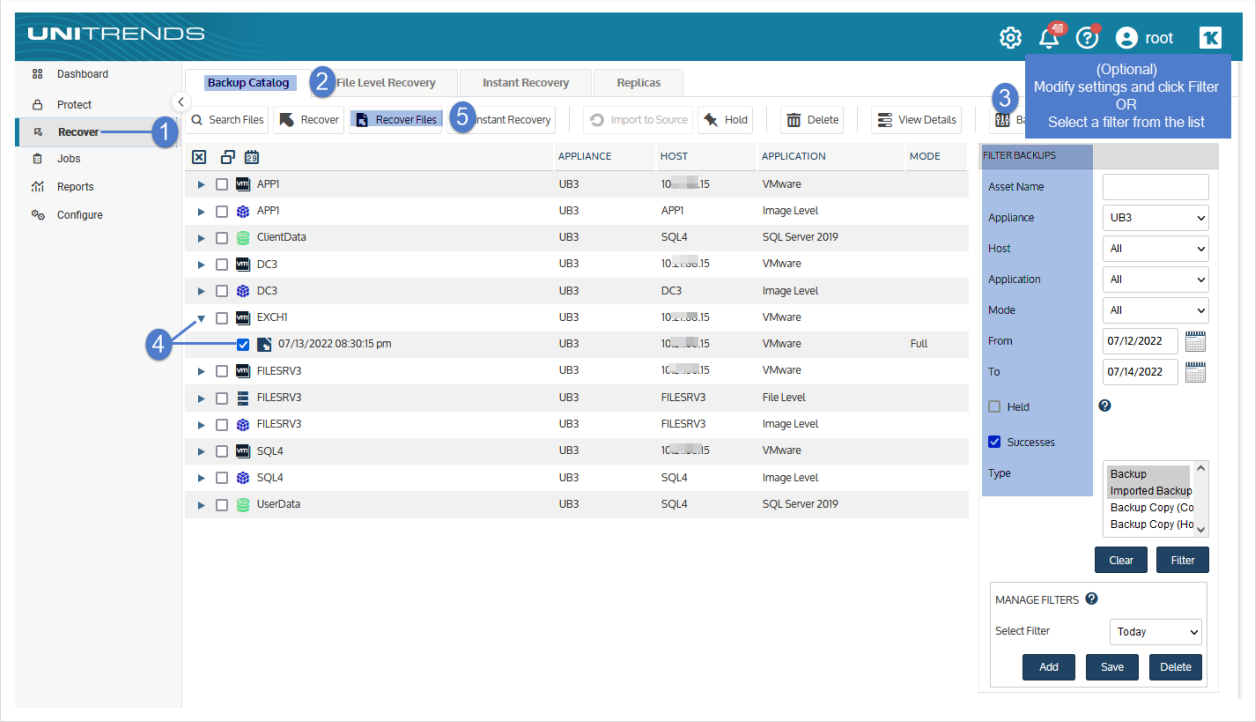

5 Click **Confirm** to continue. The appliance creates the recovery object.

Note: If you receive an error on a Unitrends Backup appliance while creating the recovery object, increase the memory allocation for the Unitrends Backup VM by using the host that manages it.

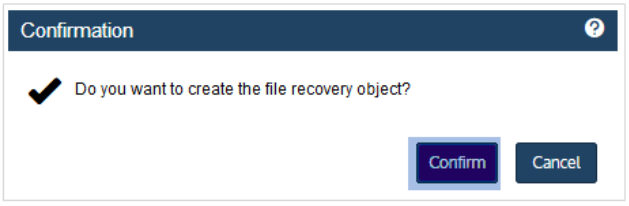

6 Click **View FLR**.

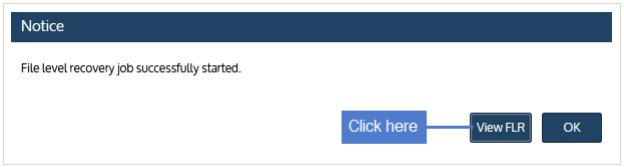

7 On the **File Level Recovery** tab, locate the recovery object.

Recovery objects display on the tab with the following details: the name of the VM asset for which the object was created, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

8 Select the recovery object and click **Show Details**. The CIFS path shows the appliance IP and flr folder (*flr9* in our example). Note the name of the *flr* folder. You will need it to access the EDB file in Ontrack PowerControls® for Exchange™.

Note: An *flr* folder exists for each recovery object that displays on the File Level Recovery tab in the appliance UI.

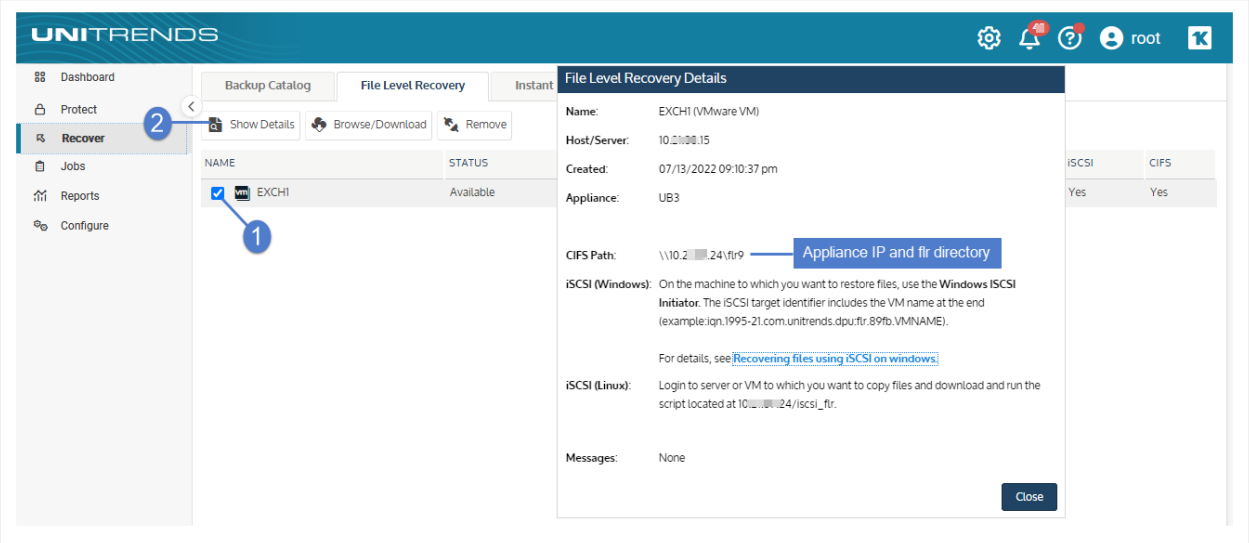

9 Select the recovery object and click **Browse/Download**.

Note: In most cases you can browse the recovery object directly in the appliance UI by clicking **Browse/Download**. If this method is not supported in your environment, the Browse/Download button is disabled and you will need to mount a CIFS share or iSCSI LUN on a target machine instead, as described in "To recover files by mounting the CIFS share" or "To recover files by mounting the iSCSI LUN" on page 829.

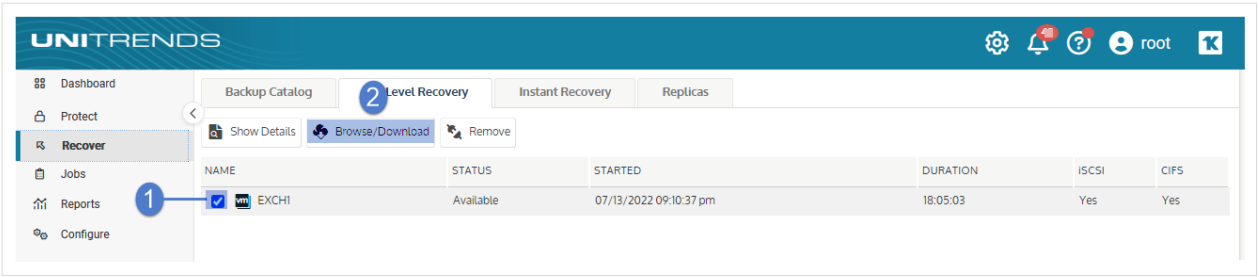

10 In the File Browser, browse to the mailbox database whose items you will recover. Note the filepath as you will browse to it in Ontrack PowerControls® for Exchange™.

In our example, the path is *volume1\Program Files\Microsoft\Exchange Server\V15\Mailbox\Group2 Database*.

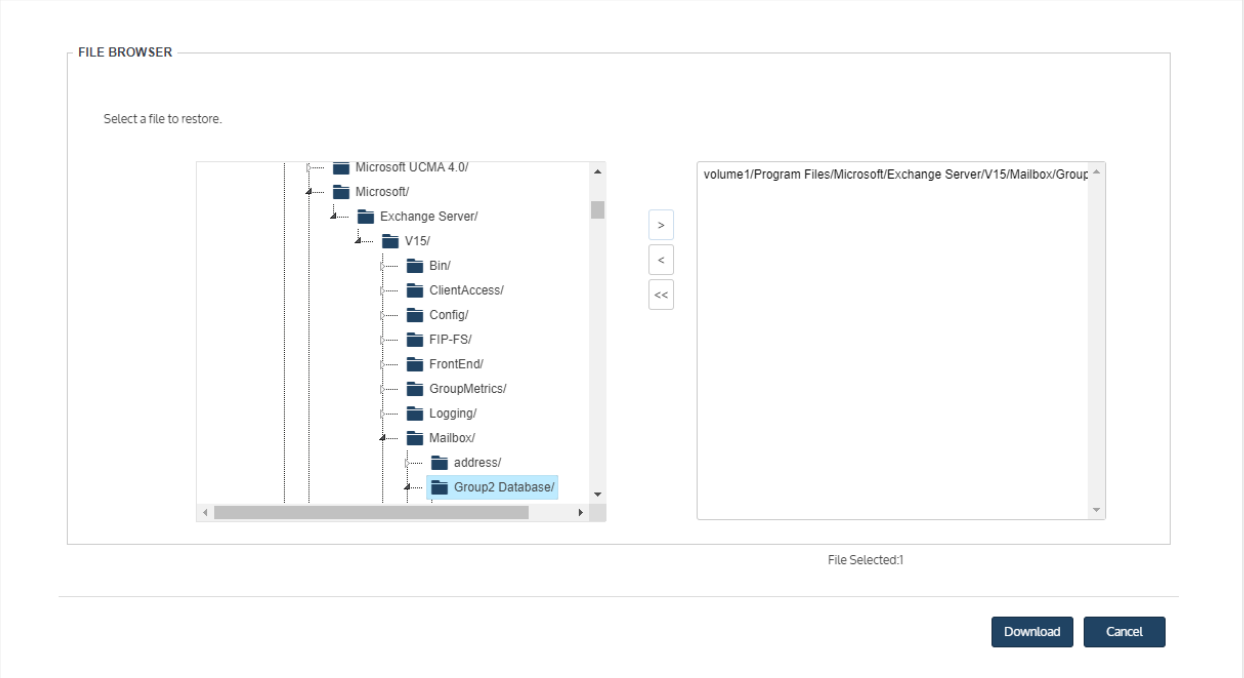

11 Launch Ontrack PowerControls® for Exchange™.

12 Enter the following in the Data Wizard Source Selection dialog:

> **Note:** If the Data Wizard does not display, select **File > Open Source**.

- Select Source – Select **EDB File** in the Select Source area.
- EDB Path – Browse to the EDB file or enter the full UNC path to the EDB file. In our example, the full UNC path is *\\[ApplianceIP]\flr9\volume1\Program Files\Microsoft\Exchange Server\V15\Mailbox\Group2 Database\Group2 Database.edb*.
- Log File Path – Browse to or enter the full UNC path to the log file directory. In our example, the full UNC path is *\\[ApplianceIP]\flr9\volume1\Program Files\Microsoft\Exchange Server\V15\Mailbox\Group2 Database*.

13 Click **Next**.

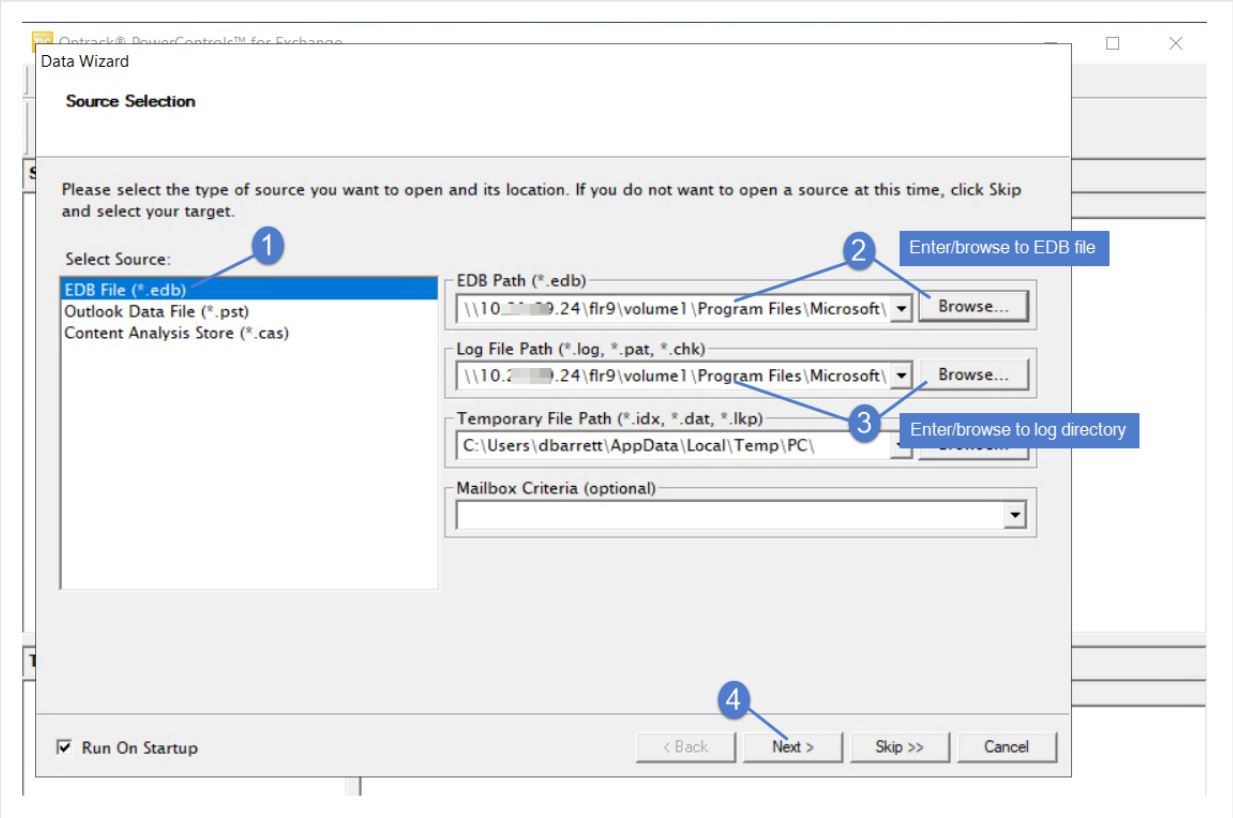

14 In the Data Wizard Target Selection dialog, specify the location where items will be recovered. In our example we are recovering to an Outlook Data File.

15 Click **Next**.

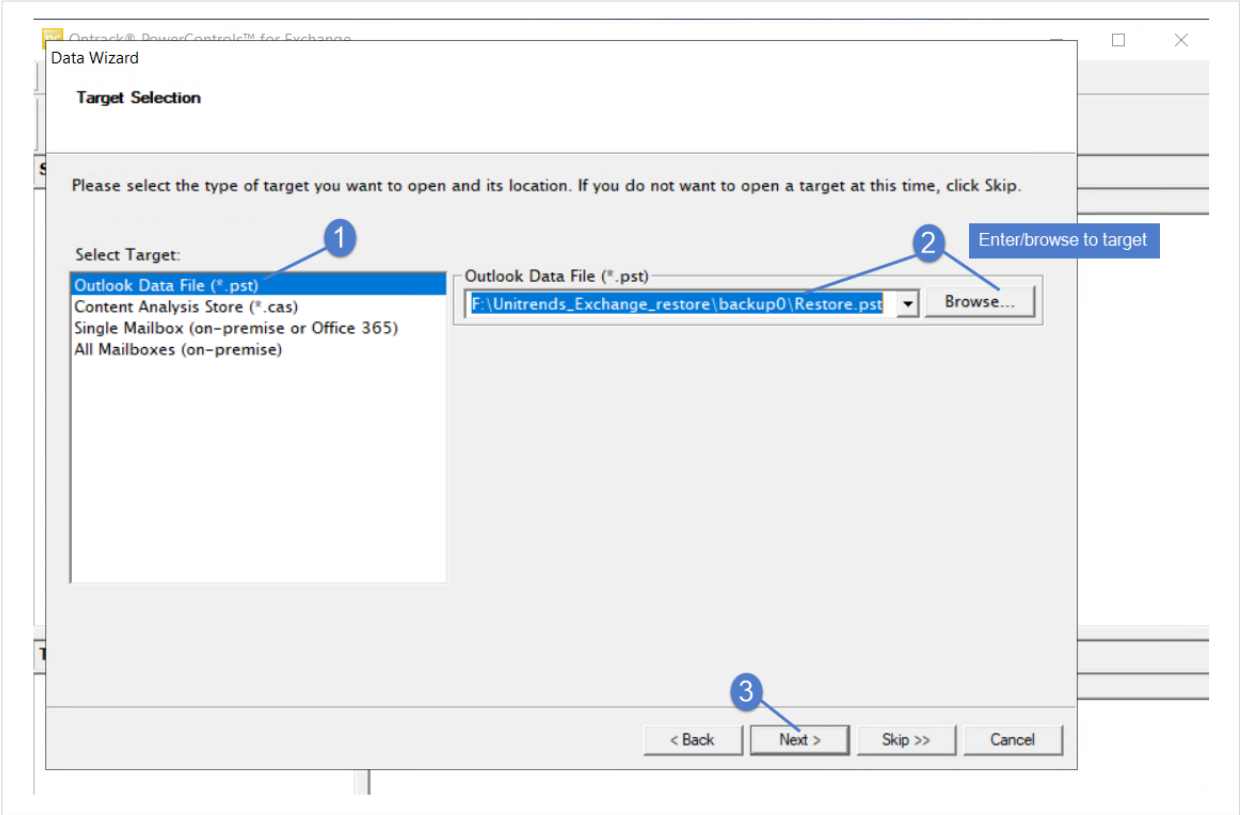

16 To recover items, do one of the following:

- To recover items to a PST file or live Exchange environment, navigate to the items you want to recover in the Source pane on top. Copy and paste (or drag) them to the Target pane.
- To recover items to a network location, navigate to the items you want to recover in the Source pane on top, select them, right click and select **Export**, select a message format and recovery location, and click **Export**.

For additional detail, see the Ontrack® PowerControls™ User Guide for Exchange.

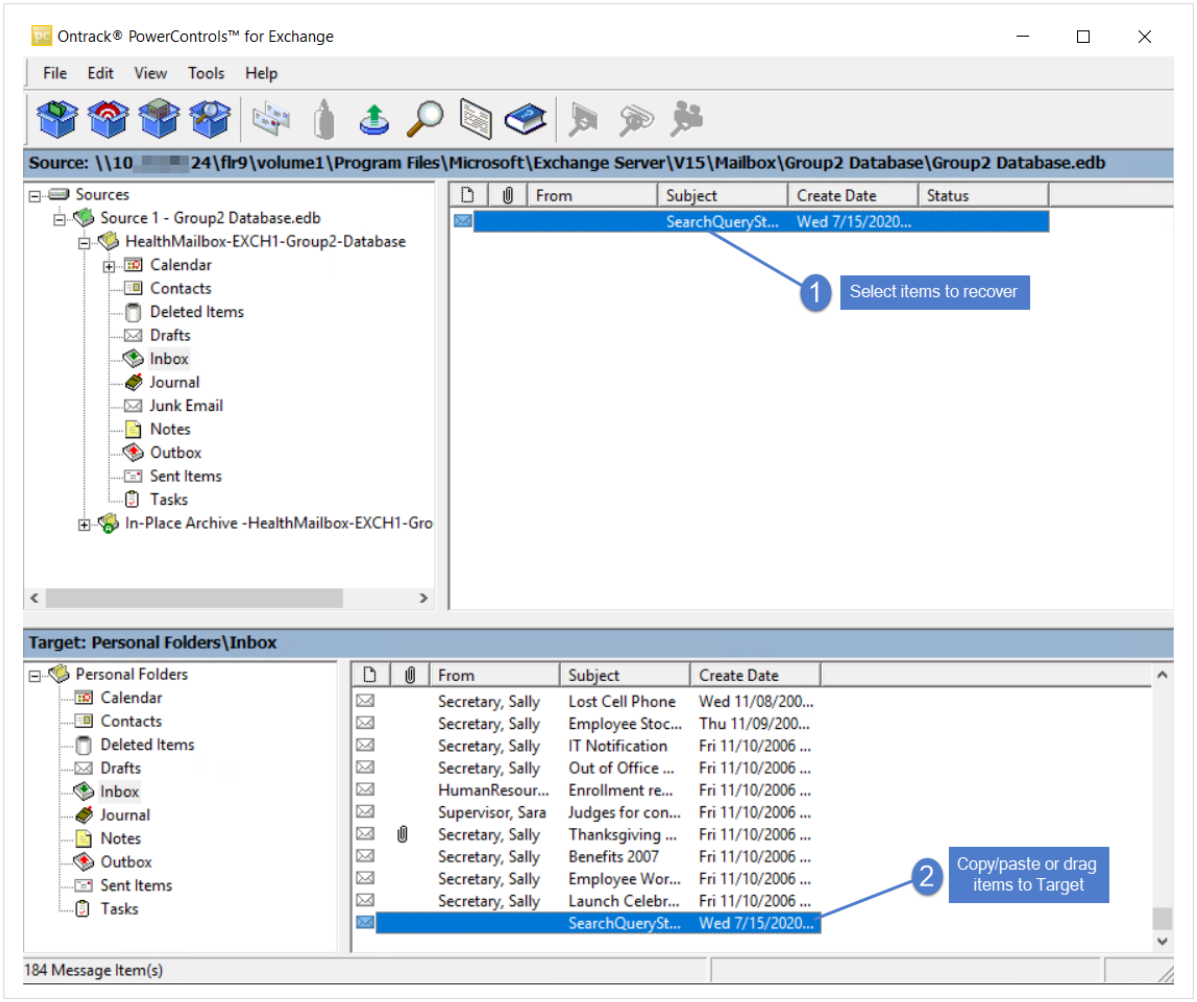

17 When you have finished recovering items, remove the recovery object:

- Return to the appliance UI. On the **File Level Recovery** tab, select the recovery object and click **Remove**. Click **Confirm**.

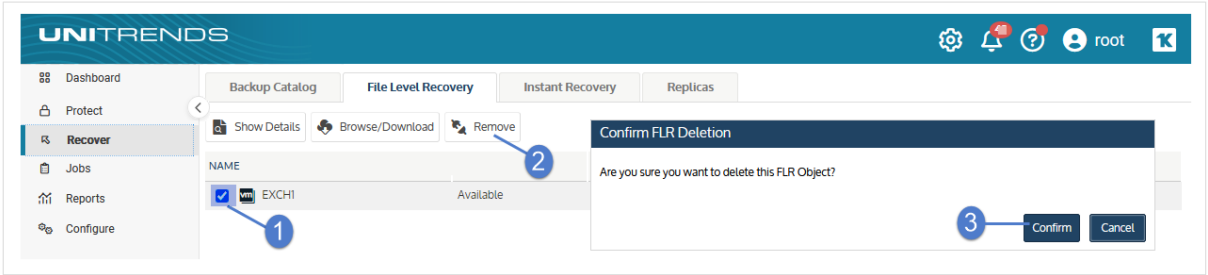

- The recovery object is removed:

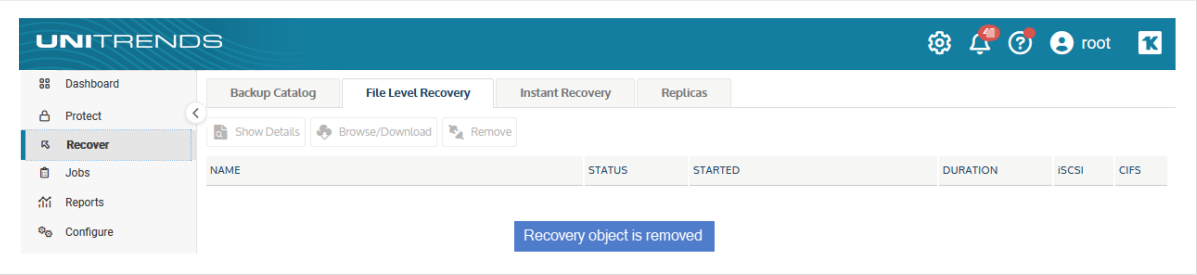

To recover SharePoint items from a backup or imported backup copy

1. Log in to the backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the asset and select the backup or imported backup copy from which you want to recover items.

   (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
4. Click **Recover Files**.

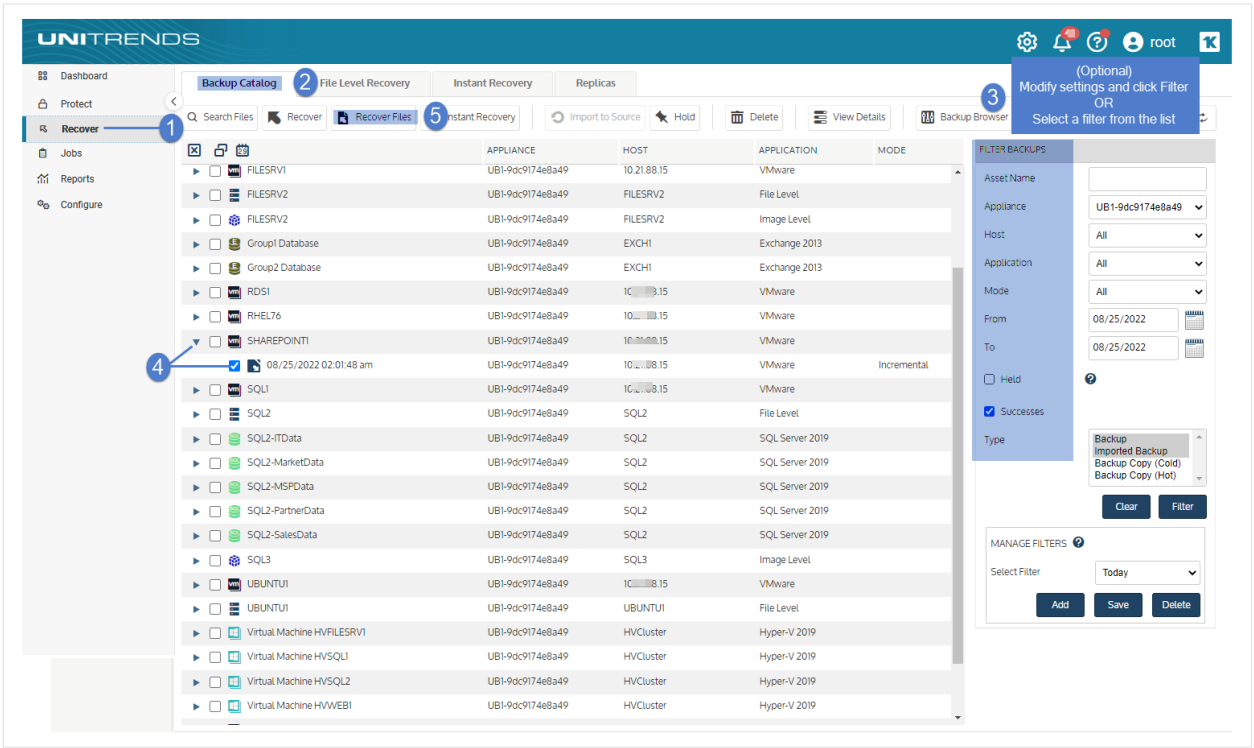

5 Click **Confirm** to continue. The appliance creates the recovery object.

Note: If you receive an error on a Unitrends Backup appliance while creating the recovery object, increase the memory allocation for the Unitrends Backup VM by using the host that manages it.

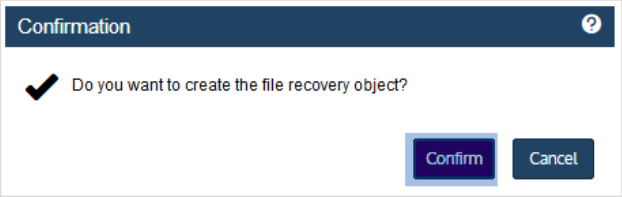

6 Click **View FLR**.

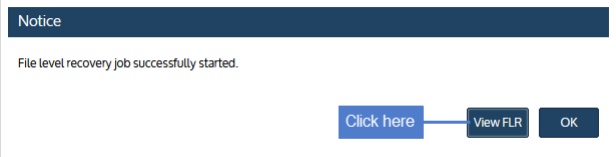

7 On the **File Level Recovery** tab, locate the recovery object.

Recovery objects display on the tab with the following details: the name of the VM asset for which the object was created, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

8 Select the recovery object and click **Browse/Download**.

Note: In most cases you can browse the recovery object directly in the appliance UI by clicking **Browse/Download**. If this method is not supported in your environment, the Browse/Download button is disabled and you will need to mount a CIFS share or iSCSI LUN on a target machine instead, as described in "To recover files by mounting the CIFS share" or "To recover files by mounting the iSCSI LUN" on page 829.

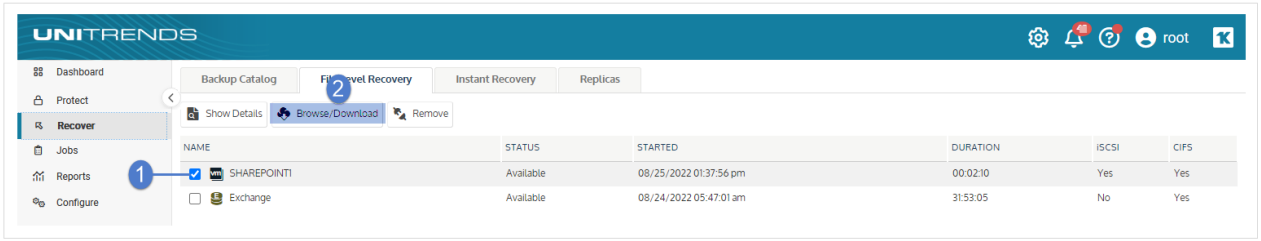

9 In the File Browser, browse to the SharePoint instance whose items you will recover. Note the filepaths to the *MDF* and *LDF* files, as you will browse to them in Ontrack PowerControls® for SharePoint™.

In our example, the paths are *volume1/Program Files/Microsoft SQL Server/MSSQL13.MSSQLSERVER/MSSQL/DATA/SharePoint_AdminContent_bd2a51ae-feac-4a66-aa42-967b825c50a3.mdf* and *volume1/Program Files/Microsoft SQL Server/MSSQL13.MSSQLSERVER/MSSQL/DATA/SharePoint_AdminContent_bd2a51ae-feac-4a66-aa42-*

967b825c50a3_log.ldf.

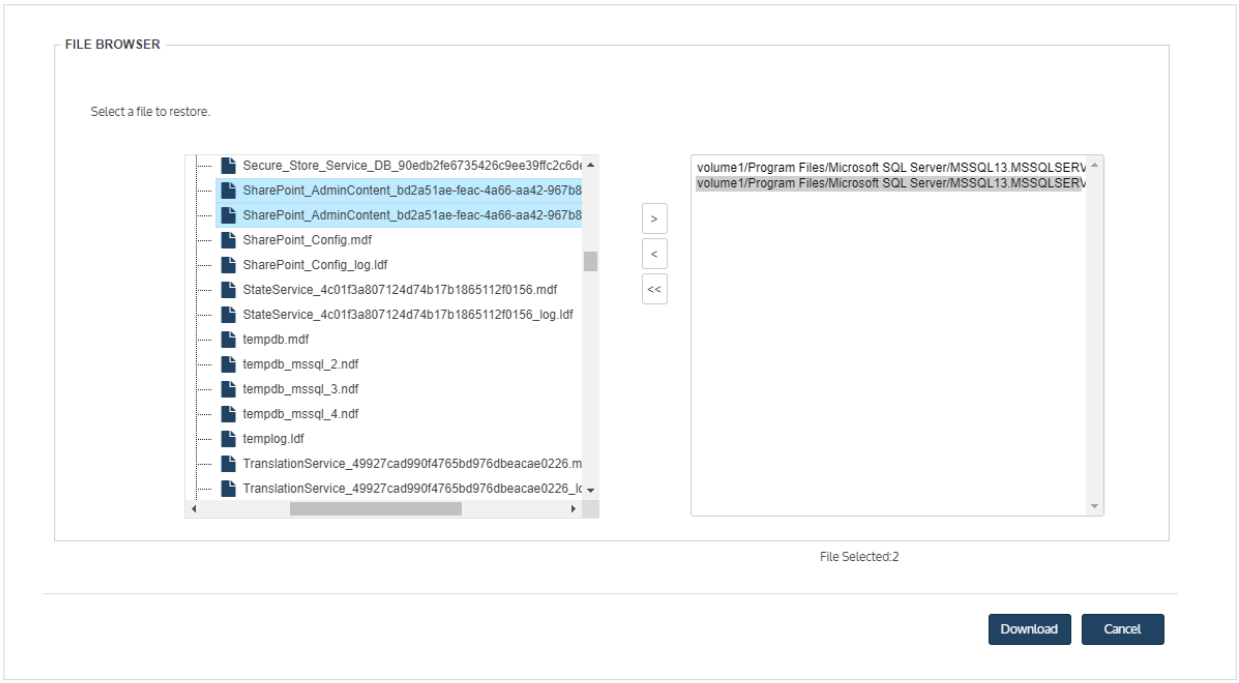

10 Launch Ontrack PowerControls® for SharePoint™.

11 Select **File > Open Source**. Click **Add**.

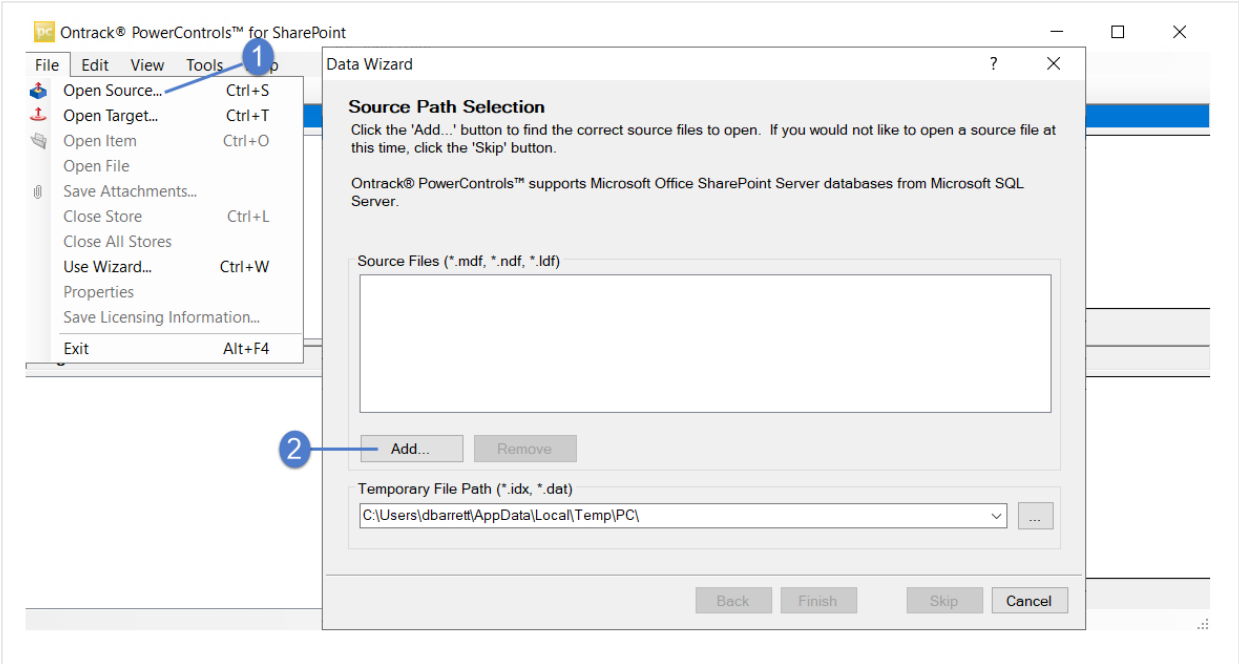

12 Enter \\ followed by the appliance IP address to create a network mapping to your Unitrends appliance. For example, *192.168.225.24*.

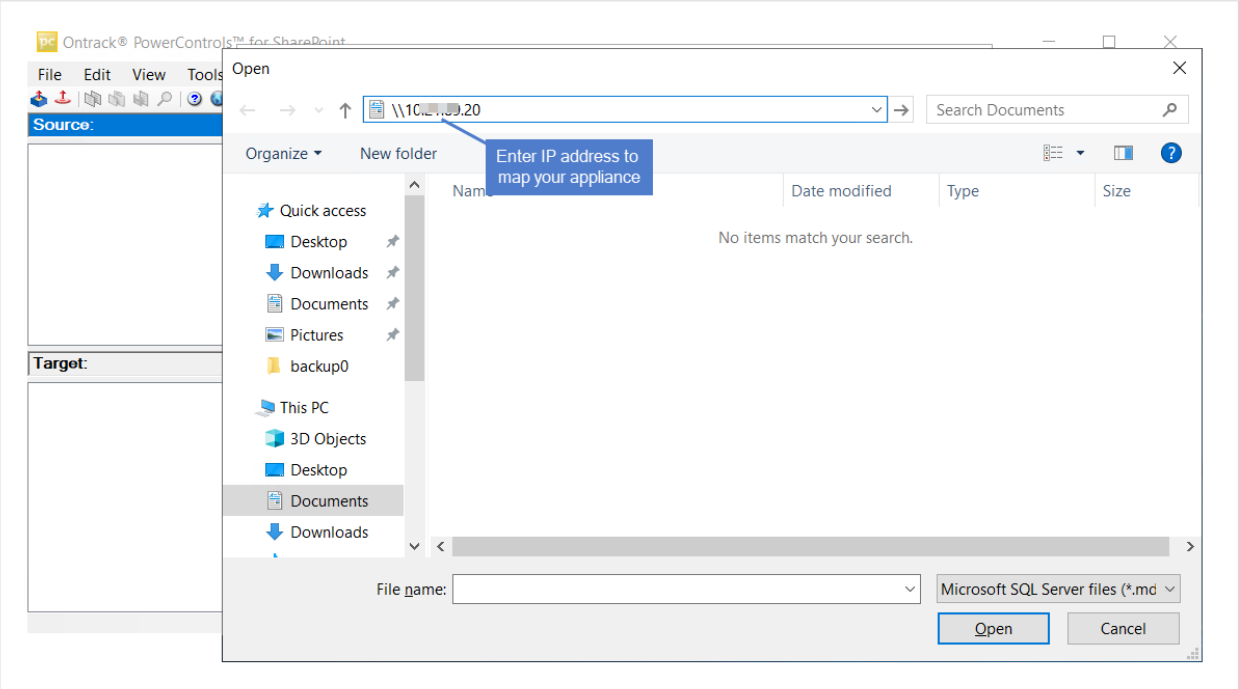

13 Browse under the *flr* folder to the SharePoint data folder you noted above in step 9 .

Note: An *flr* folder exists for each recovery object that displays on the File Level Recovery tab in the appliance UI. To see the name of the SharePoint *flr* folder, select the SharePoint recovery object on the File Level Recovery tab and click **Show Details**.

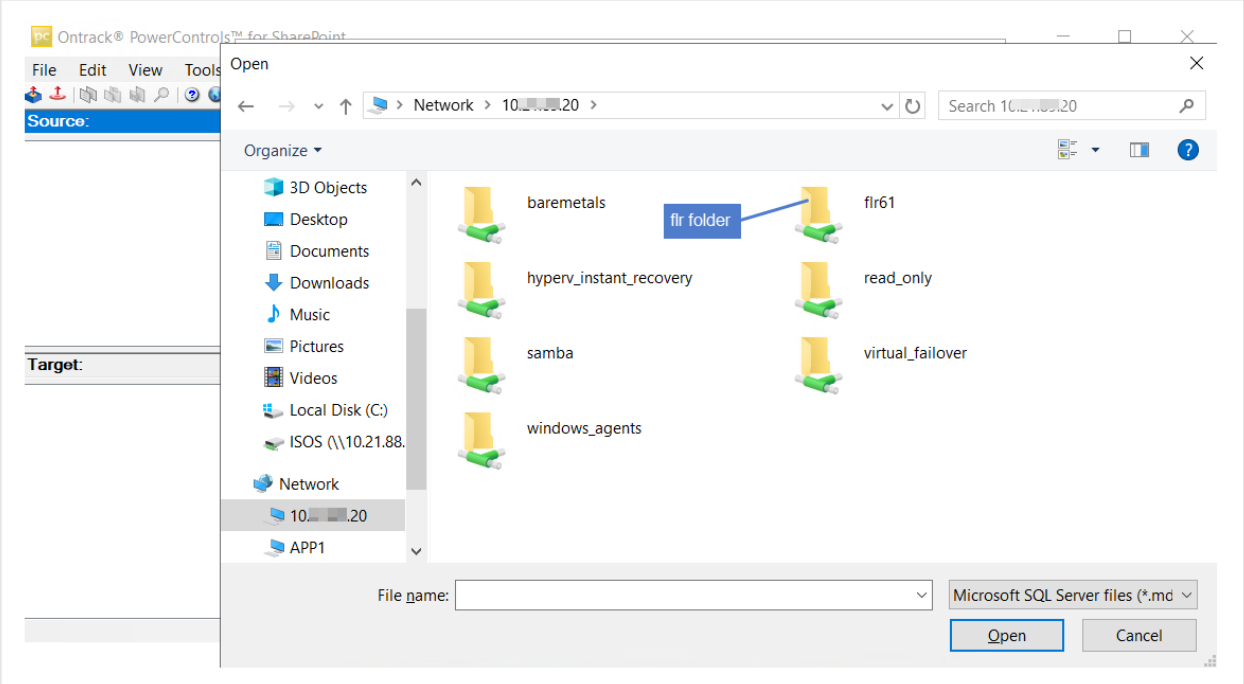

14 Select databases and logs. Click **Open**.

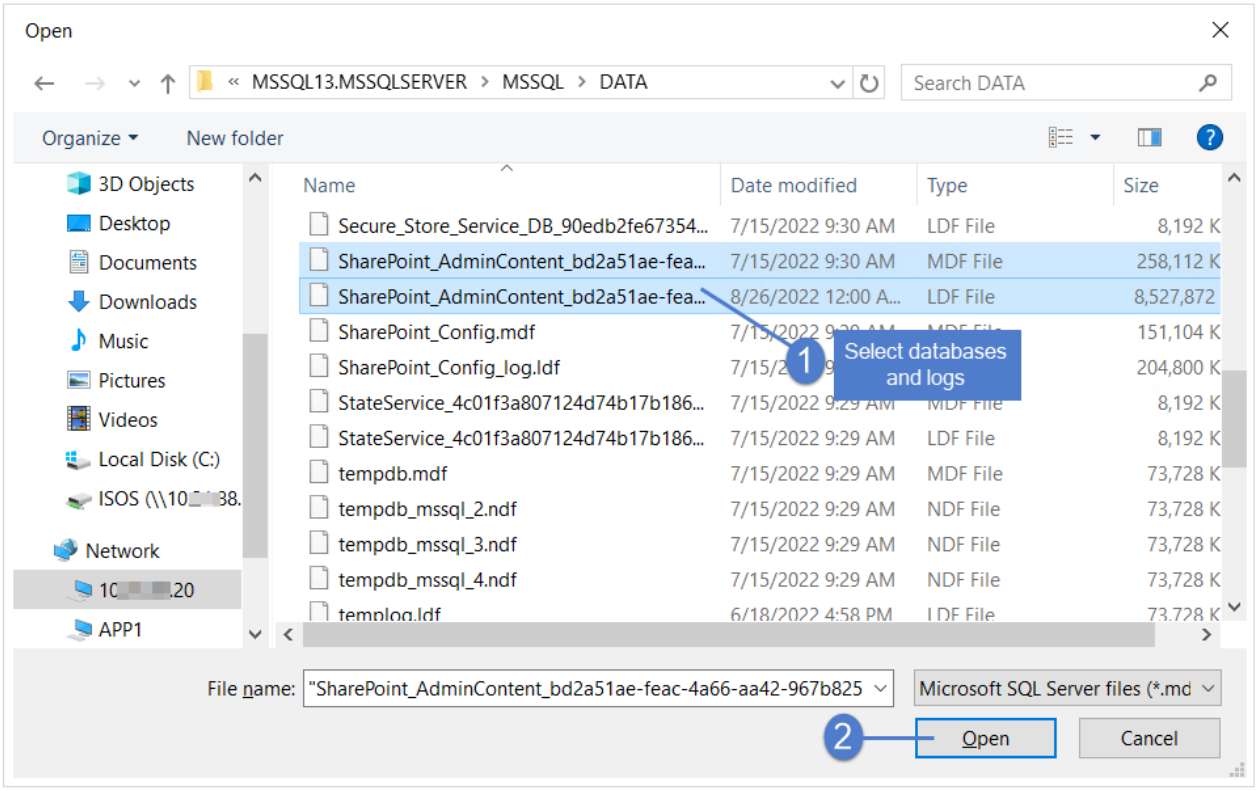

15 Click **Finish**.

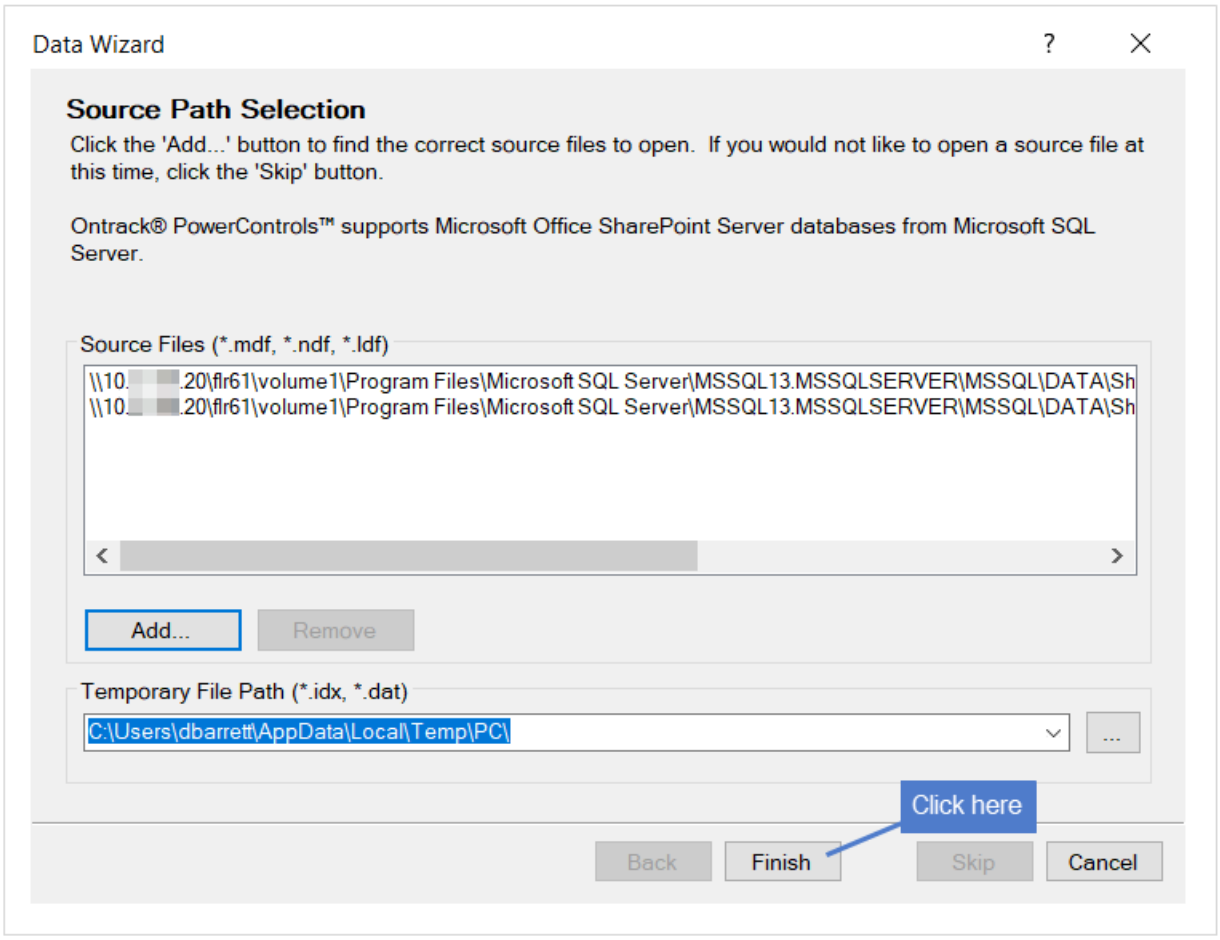

16 Selected databases/logs display in the Source area. If needed, select **File > Open Source > Add** to add more databases/logs.

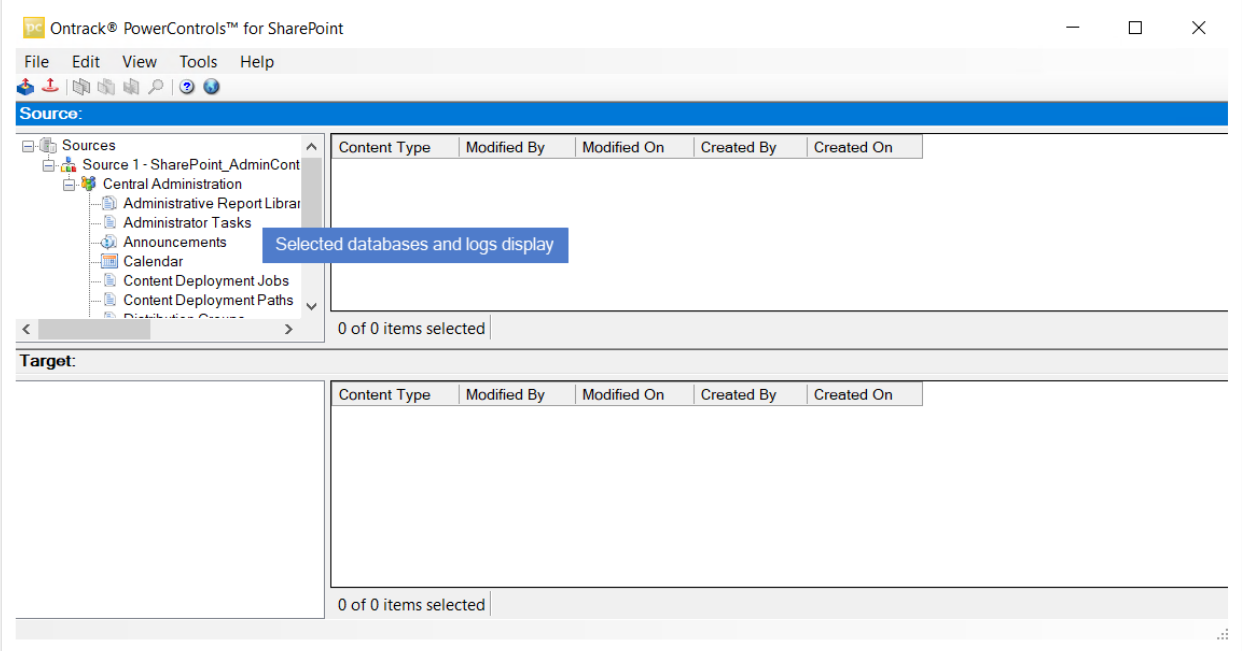

17 Select **File > Open Target.**

18 Select the target server and enter authentication information. Click **Next.**

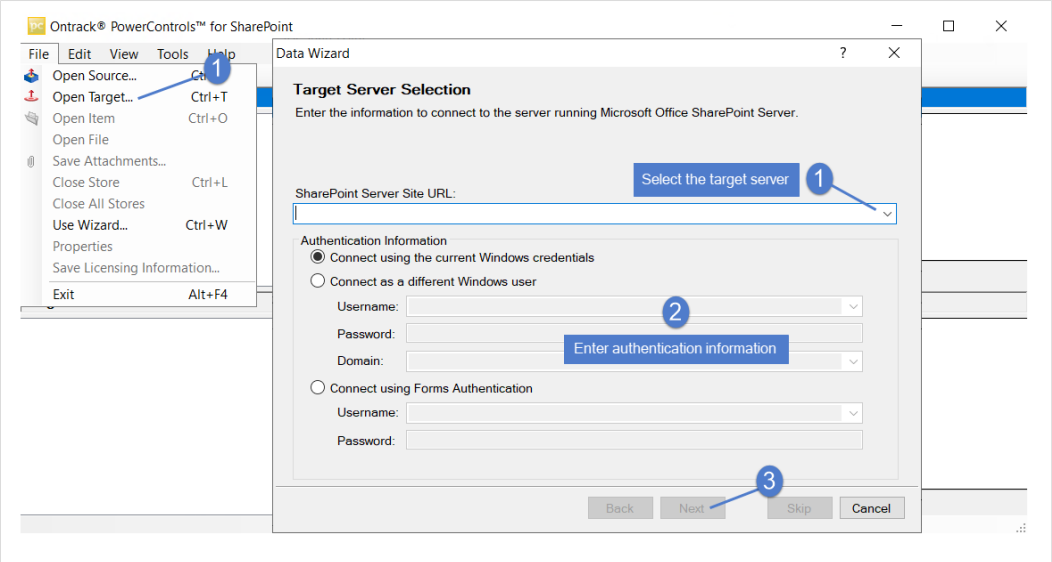

19 The target SharePoint server displays in the Target area.

20 To recover items, select files in the Source area and copy/paste (or drag) them to a library on the target. For additional detail, see the Ontrack® PowerControls™ User Guide for SharePoint.

21 When you have finished recovering items, remove the recovery object:

- Return to the appliance UI. On the **File Level Recovery** tab, select the recovery object and click **Remove**. Click **Confirm**.

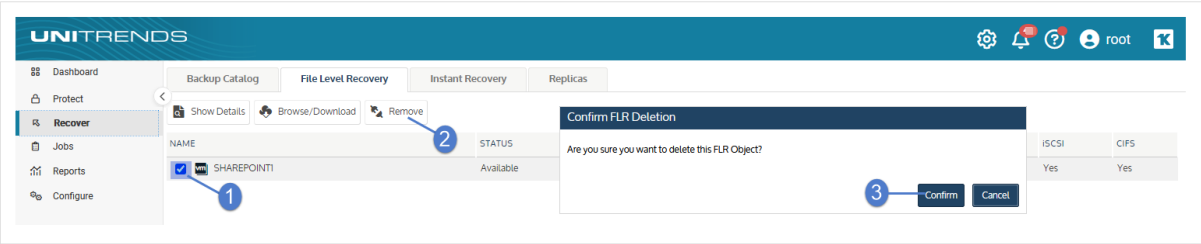

- The recovery object is removed:

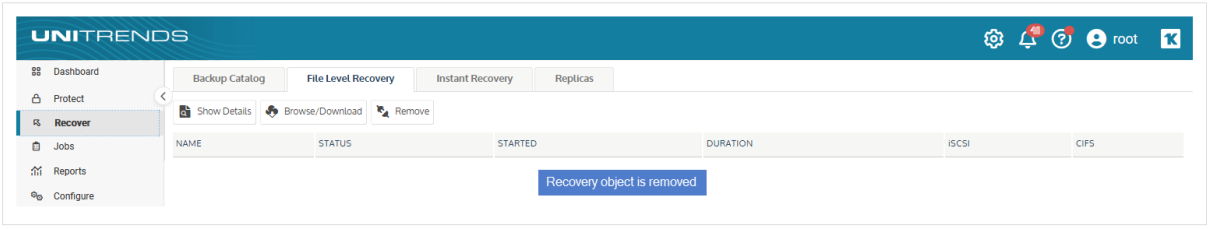

VM replicas

The VM replica feature provides a quick way to recover a failed VMware VM. It creates a virtual machine replica of the original VM, then keeps this replica up-to-date by applying backups of the original VM as they run. In the event of a disaster, you can bring this replica online to immediately assume the role of the failed VM. You can use the replica to replace the original VM temporarily or use it as a permanent replacement.

To use the feature, simply set up the replica by using the Create Replica VMs dialog. The appliance then creates the replica from the most recent backup of the original VM, and automatically applies all subsequent backups. Because the replica is continually updated, it is ready for production use at any time.

While creating the replica, you specify the ESXi host location where the replica VM will reside. The replica VM is created as a cold stand-by in the specified location. The replica is powered off and has no network connectivity. Because the replica remains powered off even as backups are applied, it consumes no compute resources.

After the first backup has been applied, replica creation is complete. You can then do the following as needed:

- Audit the replica to verify the integrity of the machine and its data and applications. In audit mode, the replica runs on a private network (inaccessible from the production network). This enables you to check the replica VM while the original VM is still operating in production. It is recommended that you periodically audit the replica to ensure it functions as expected.
- Bring the replica into live mode to assume the role of the original VM. Once the replica has booted into live mode, simply configure network settings to bring the VM online in your production environment. Live mode is intended as a temporary solution. You should exit live mode (by tearing down the replica) as soon as you have verified that the live replica is functioning as expected in production. (You can keep using the replica VM after it exits live

mode by selecting the *Delete the VM replica from the appliance only* option in the Tear Down VM replica dialog. For considerations, see "Live mode recommendations" on page 895.)

See the following topics for details on using the VM replicas feature:

- "Replica restore jobs"
- "VM replica requirements" on page 882
- "Creating VM replicas" on page 885
- "Working with VM replicas" on page 891

Replica restore jobs

To keep a replica up-to-date, the appliance automatically queues a replica restore job each time a backup successfully completes. The restore job runs as soon as possible to apply the backup to the replica. (For example, the job cannot run if the replica is in audit mode or if the job is queued behind other active jobs.)

The restore job applies the backup to the replica VM and then consolidates VM snapshots. During the restore, the replica is in restore mode and the replica restore job displays on the Active Jobs tab:

Note: Do not cancel a replica restore job by clicking Cancel on the Active Jobs tab. Instead, bring the replica into audit mode to temporarily stop applying backups. (For details, see "Do not cancel an active replica restore job" on page 881.)

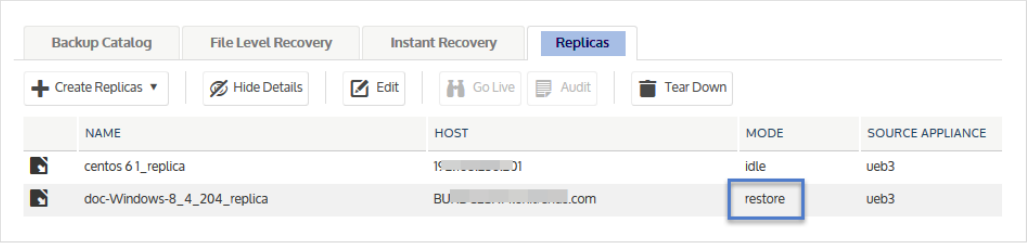

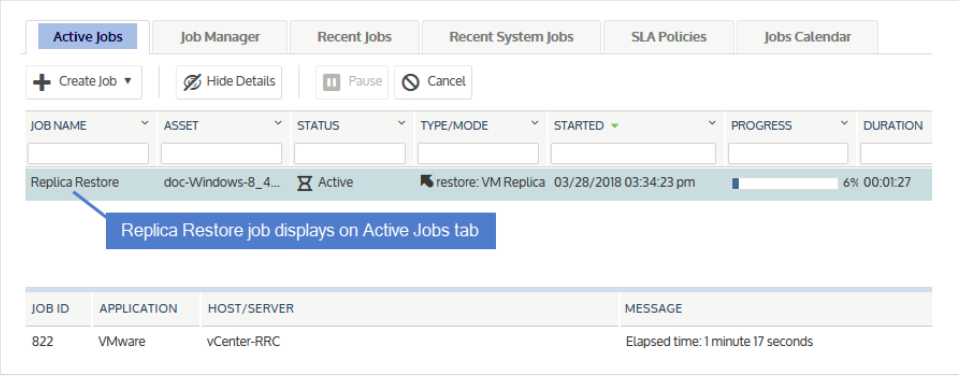

Entering live mode while a restore is in progress

Beginning in release 10.1.2, the replica can quickly enter live mode while a restore is in progress. The appliance boots the VM into live mode immediately and cancels the running replica restore job. Upon entering live mode, the replica is marked *invalid* because the replica role no longer applies. After configuring network settings for the live replica VM and verifying that it is functioning as expected in production, you tear down the replica on the Unitrends appliance. You can then opt to create a new replica for the live VM.

Entering audit mode while a restore is in progress

Beginning in release 10.1.2, the replica can quickly enter audit mode while a restore is in progress. The appliance boots the VM into audit mode immediately and cancels the running replica restore job. See these topics for details:

- "If the backup had not yet been applied to the replica"
- "If the backup has been applied to the replica" on page 880

If the backup had not yet been applied to the replica

If the restore job did not finish applying the backup to the replica:

- The canceled replica restore job moves to Error status on the Active Jobs tab:

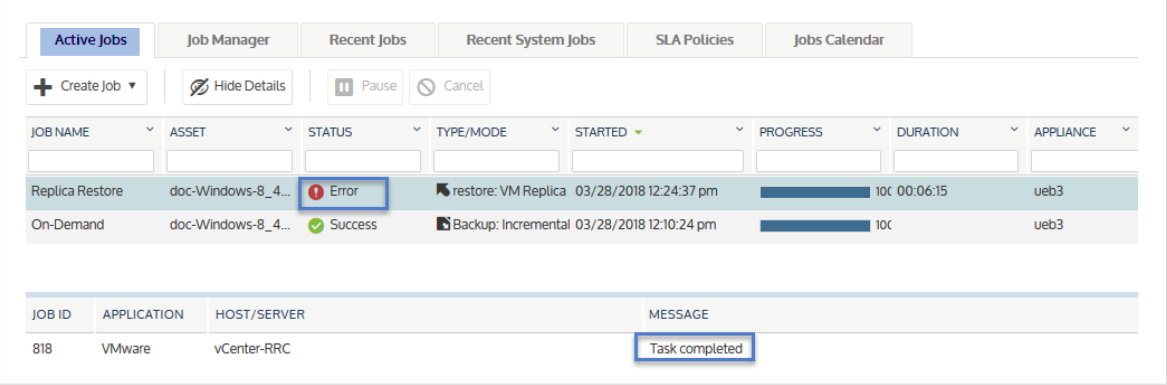

- The canceled replica restore job displays in Failure status on the Replicas Report. Click the job to view details in the Replica History Status dialog. Details contain the message *Canceled and failed. Next replica restore will retry*:

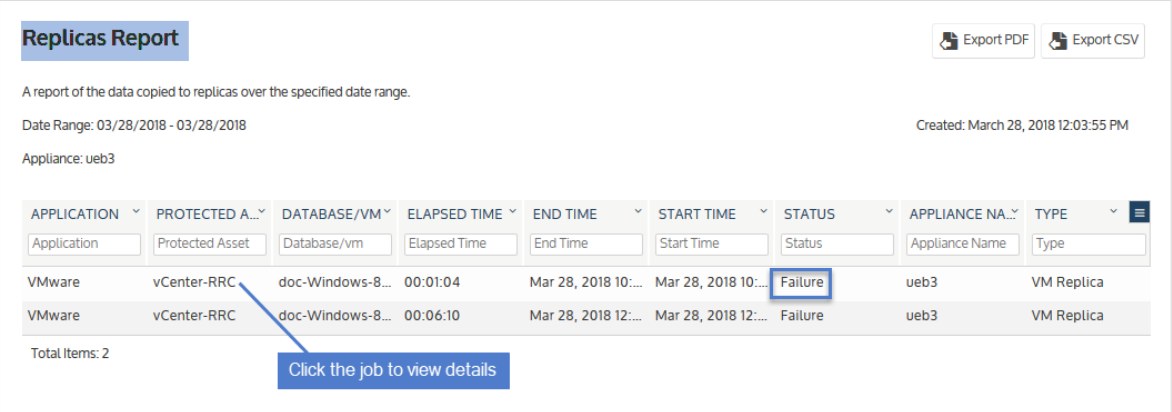

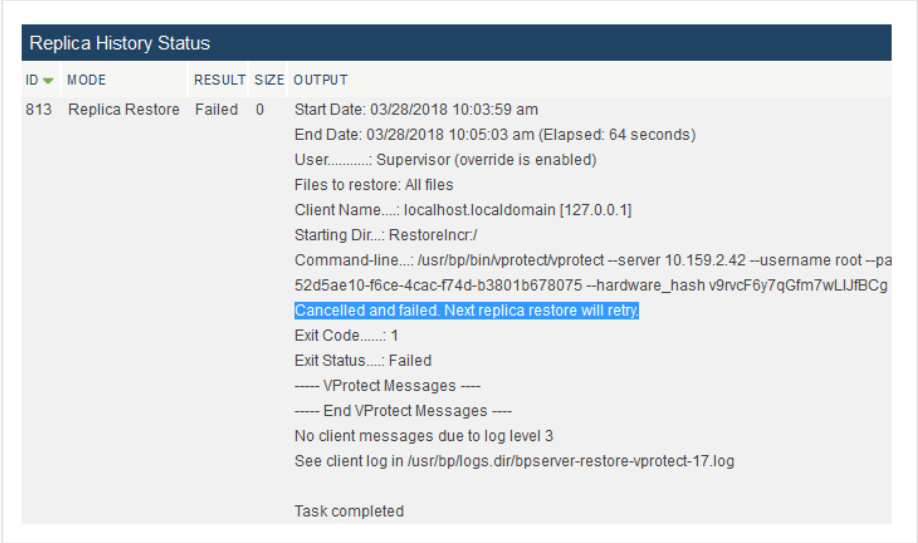

- The canceled job remains in the replica backlog. When the replica exits audit mode, the appliance queues a new restore job to replace the one that was canceled.

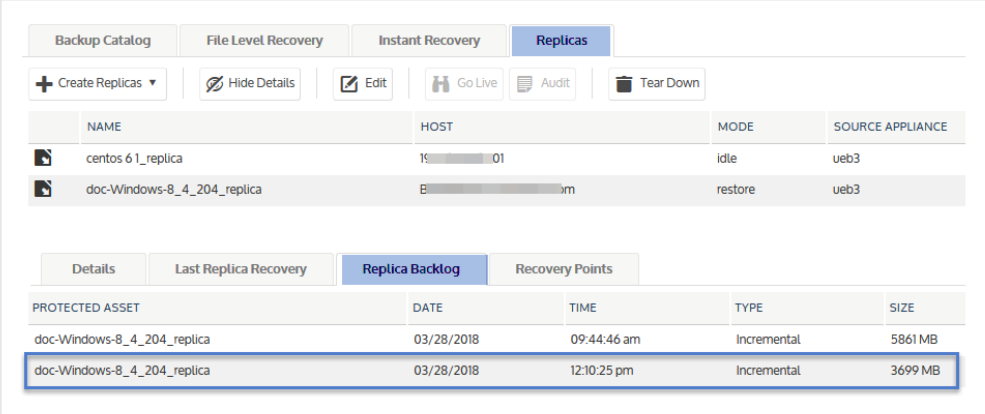

If the backup has been applied to the replica

If the restore job finished applying the backup to the replica:

- The canceled replica restore job moves to Warning status on the Active Jobs tab:

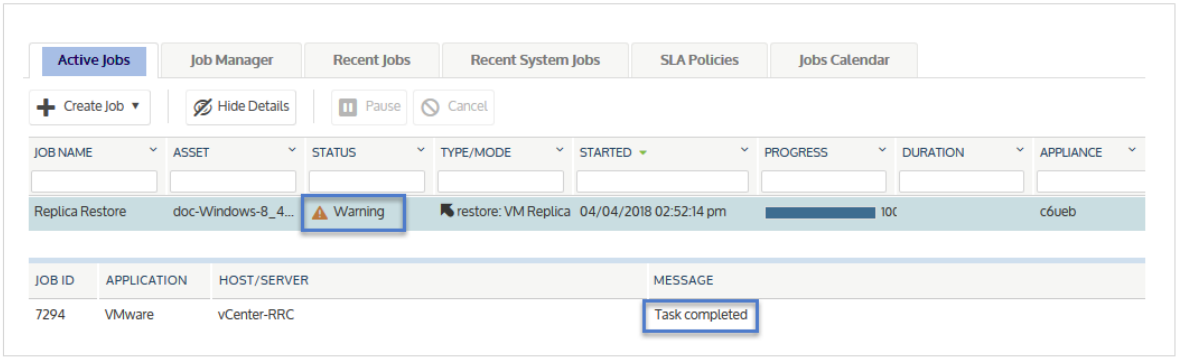

- The canceled replica restore job displays in Warning status on the Replicas Report. Click the job to view details in the Replica History Status dialog. Details contain the message *Canceled but successful. Some cleanup may occur on the next replica restore*:

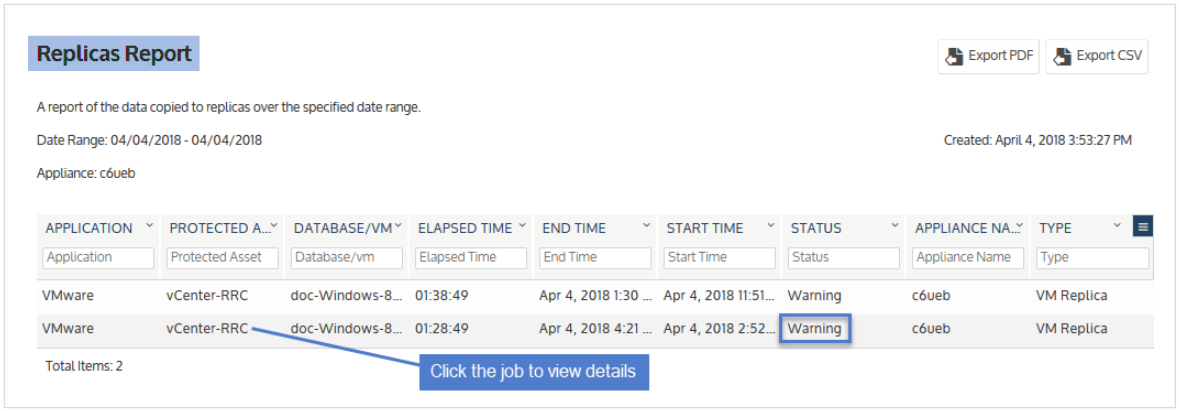

- If needed, snapshot consolidation for the canceled job is performed during the next replica restore. Note that the next restore may take extra time. Do not interrupt the next restore. Repeatedly interrupting restore jobs significantly degrades performance.

Do not cancel an active replica restore job

Do not cancel an active replica restore job. Instead, bring the replica into audit mode.

Each successful backup of the original VM is applied to the replica as soon as the backup completes. The appliance applies the backup by running a replica restore job, which displays on the Active Jobs tab as shown here:

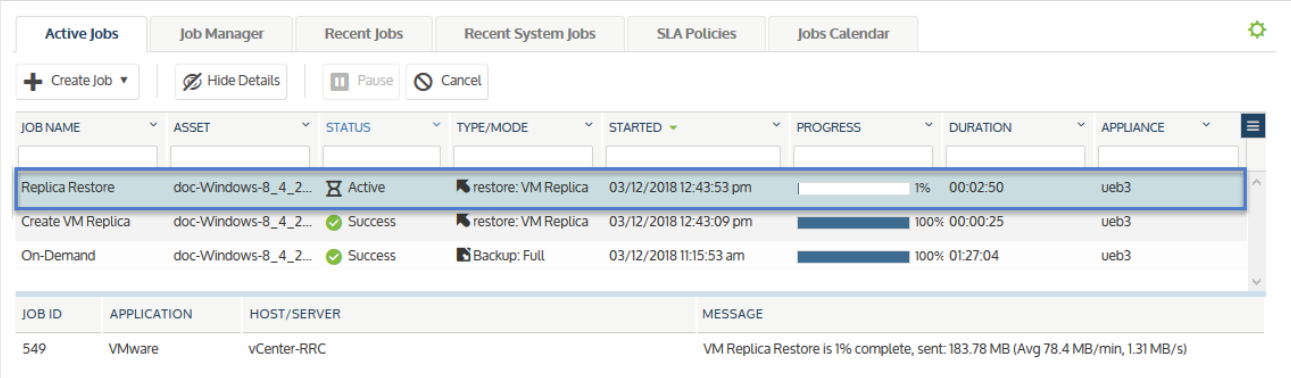

If you cancel a replica restore job by using the Cancel button on the Active Jobs page, the replica may become invalid and need to be recreated (for details, see *Halted* mode in "VM replica modes" on page 902). To temporarily stop applying backups, bring the replica into audit mode instead (as described in "Auditing a VM replica" on page 893). Use the procedure "To exit audit mode" on page 895 to start applying backups again. Note that all backups that ran while the replica was in audit mode will be applied to the replica upon exiting audit mode. You cannot skip applying a specific backup to a replica.

VM replica requirements

The following requirements must be met to use the VM replica feature:

Requirement	Description
Unitrends appliance	The Unitrends appliance used to create the VM replica must meet these requirements: • Must be running version 10.1 or higher. • Must be the backup appliance that is protecting the original VM or the target appliance that is receiving copies of the original VM's backups. • License Feature String must contain *MKT=2* or higher. To add this feature, contact your Sales Representative. 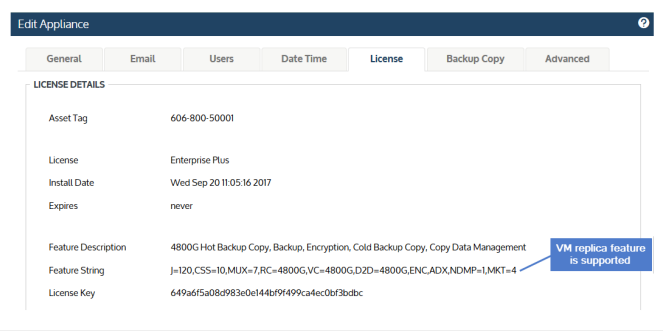

Requirement	Description
Host-level backup or hot backup copy	A successful host-level backup or host-level hot backup copy of a VMware virtual machine is required to create the replica VM. (Hyper-V, AHV, and XenServer VMs are not yet supported). The replica VM is kept up to date by applying subsequent backups as they run. Be sure to capture changes on the original VM by running backups at regular intervals. **Notes:** • The replica VM is based on the latest backup or hot copy of the virtual machine. The replica VM does not include any VM disks that were excluded from the backup. • If no backup or hot copy exists, you can still set up the replica by running the create replica procedure. In this case, the appliance creates a "shell" replica in *uninitialized(pending: create)* mode. The replica remains in this mode until a successful VM backup or hot backup copy is created. (An uninitialized replica cannot assume the role of a failed VM. At least one backup must be applied before the replica can replace the original VM.)
Hypervisor version	The ESXi server that hosts the replica VM must meet these requirements: • Must be running a paid ESXi 5.0 or a higher version listed in the Compatibility and Interoperability Matrix. • Essentials and Essentials Plus editions are supported for ESXi servers that are managed by a vCenter only. (Essentials and Essentials Plus editions running on stand-alone ESXi servers are not supported.) • Must support the operating system (OS) of the VM. (See the VMware documentation for details.) For example, a replica of a Windows 2016 VM cannot reside on an ESXi 5.1 host. • Can be the host where the original VM resides or an alternate host running the same software version as the original or a later version.
Virtual host asset	The ESXi server must be added to the appliance as an asset. See "Adding a virtual host" on page 308 for details.
One replica per VM	An appliance can create only one replica for each virtual machine instance. To create another replica for a given VM, you must first tear down any existing replica. If you retain the replica VM on the hypervisor during teardown, you can create another as long as you specify a different replica name.
Compute and storage	The ESXi server that hosts the replica must have adequate compute and storage resources:

Requirement	Description
	• The replica's compute resources (processors and memory) match those of the original VM. • The replica's virtual disks are based on the original VM, but are always thin provisioned. Any virtual-mode raw device mapping (RDM) disks recover as standard virtual disks. • The maximum disk size is capped by what the hypervisor supports. For VMs with disks larger than 2 TB, the ESXi server must be running ESXi 5.5 or a higher version listed in the Compatibility and Interoperability Matrix. If the virtual host does not have or support the specified resources, the replica cannot be created and an error displays in the Create Replica VMs dialog. **Notes:** • The replica VM is created as a cold standby. No compute resources are used until you boot the replica in live or audit mode. • Once you create the replica, you can edit its compute resource settings as needed. For details, see "Editing a VM replica" on page 892.
Replica VM hardware version	The replica VM is configured with the highest hardware version that the hypervisor supports.
Replica VM changeability	Once you have created the replica, do not make any changes to the replica VM outside of the Edit VM Replica Details dialog. Any alteration to the replica VM may invalidate the replica. **Notes:** • If the replica becomes invalid due to VM configuration changes made outside of the Edit VM Replica Details dialog, it enters *halted* mode and can no longer be used. You must delete the halted replica and create a new one. • When a replica enters *live* mode, it becomes invalid because the replica role is no longer applicable. You can make changes to a live replica VM by using the ESXi hypervisor.
Appliance DR limitation	Appliance disaster recovery does not recover VM replicas. You must set up new replicas after you recover the appliance.

Creating VM replicas

The appliance creates a replica by using the most recent backup of the original VM. The replica's compute resources (processors and memory) match those of the original VM. The replica's virtual disks are also based on those of the most recent backup, but are always thin provisioned. During setup, you select the virtual host and storage destination where the replica will be created. (You can choose any eligible virtual host that has been added to the Unitrends appliance as an asset. To add a host, see "Adding a virtual host" on page 308.)

The virtual host must have adequate compute and storage resources to create the replica. The replica VM is created as a cold standby, so no compute resources are used until you boot the replica (to run an audit or to bring it live in production). But if the virtual host does not have or support the specified resources, the replica is not created and an error displays. Ensure adequate resources are available before running the create replica procedure.

During replica setup, you select a virtual host server, a datastore, and the number of recovery points to retain. You can create replicas for a group of VMs or for a single VM by using these procedures:

- "To create one VM replica" – Use this option to create a replica for a single VM. With this procedure, you can opt to edit the default replica name (<VMname>_replica).
- "To create multiple VM replicas" – Use this option to quickly create replicas for a group of VMs. All VMs are created on the same virtual host by using the same datastore. The number of recovery points applies to all VMs in the group. Each replica is named <VMname>_replica.

Note: You do not need to back up a VM before doing the create VM replica procedure. If no backup exists, the appliance creates a "shell" replica. When a backup completes, the appliance automatically applies it to the shell to create the full VM replica. This way you can easily set up replicas for groups of VMs without checking for backups or waiting for backups to run.

To create one VM replica

1. Log in to the backup appliance.
2. Select **Recover**, then click the **Replicas** tab.
3. Click **Create Replicas** and select **VMware**.

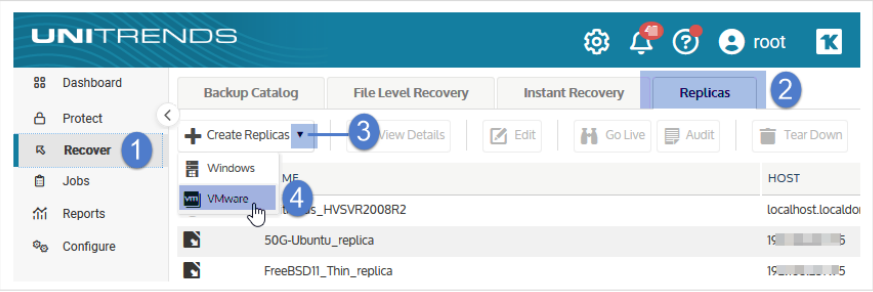

4. In the Inventory tree, expand the virtual host and select the VM.
   - To locate a VM by name, use the **Search** field below.
   - To view individual VMs, expand the virtual host and any vApps and resource pools.

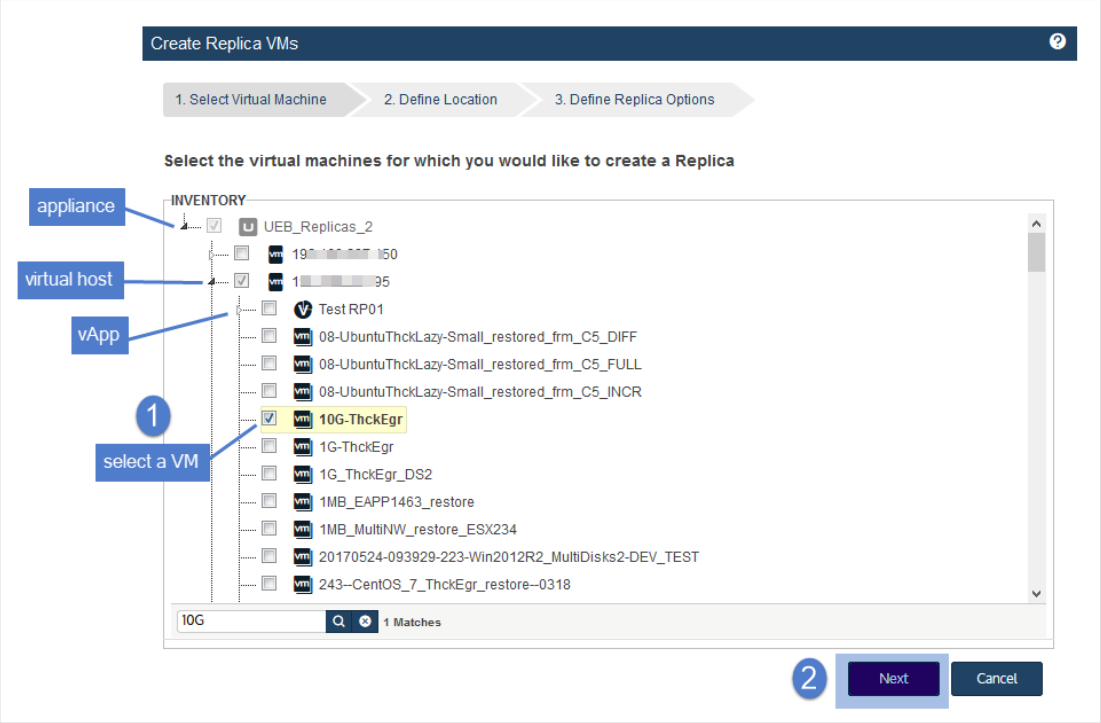

5 Click **Next**.

6 Enter Location Details to specify the location where the replica will reside.

- Location – Select a VMware host from the list. The list contains all eligible virtual host assets that have been added to the appliance. (For details on adding a virtual host, see "To add a virtual host asset" on page 311.)
- Resource Pool – (Optional) If your VMware environment has resource pools or vApps, you can opt to select one in the list.
- Storage – Select the datastore that will be used to create the replica VM disks.

7 Click **Next**.

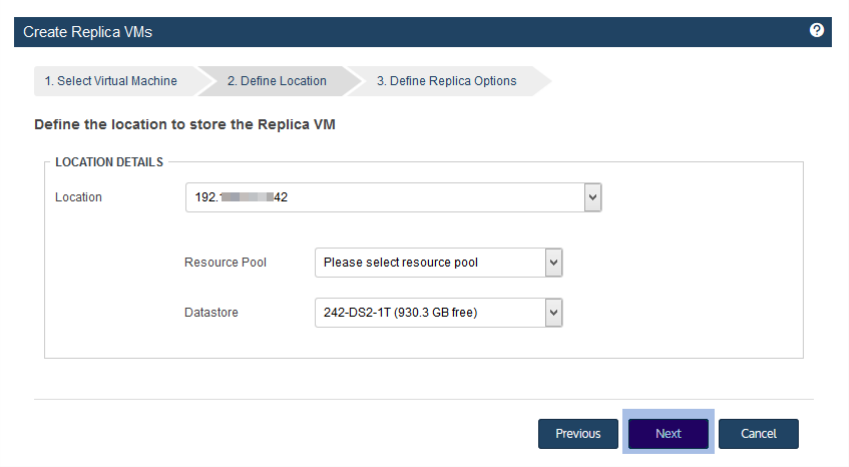

8 Enter Replica Details:

- Name – Replica name. By default, the replica is named <VMname>_replica. You can opt to edit this name.
- Number of recovery points – Maximum number of recovery points stored with the replica. By default, only the most recent recovery point is saved. You can opt to increase this value, up to the maximum number configured for the appliance. (To modify the maximum allowed on the appliance, see "VM replica configuration" on page 177.)

9 Click **Save**.

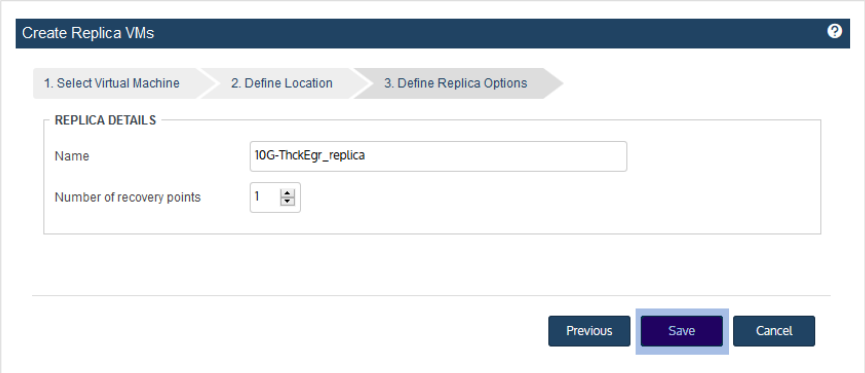

10 The appliance does one of the following:

- If a backup exists, creates a replica for the selected VM, then applies the latest backup. The replica's mode changes to *restore* while the backup is being applied, then to *idle*. After the replica enters idle mode, you can audit the replica or bring it 'live' as needed. We recommend that you audit the replica soon after it enters idle mode, to verify its integrity. See "Auditing a VM replica" on page 893 for details.
- If no backup is found, creates a "shell" replica in *uninitialized(pending: create)* mode. The replica remains in this mode until a successful VM backup is created. Run a backup of the original VM as soon as possible. (An

uninitialized replica cannot assume the role of a failed VM. At least one backup must be applied before the replica can replace the original VM.)

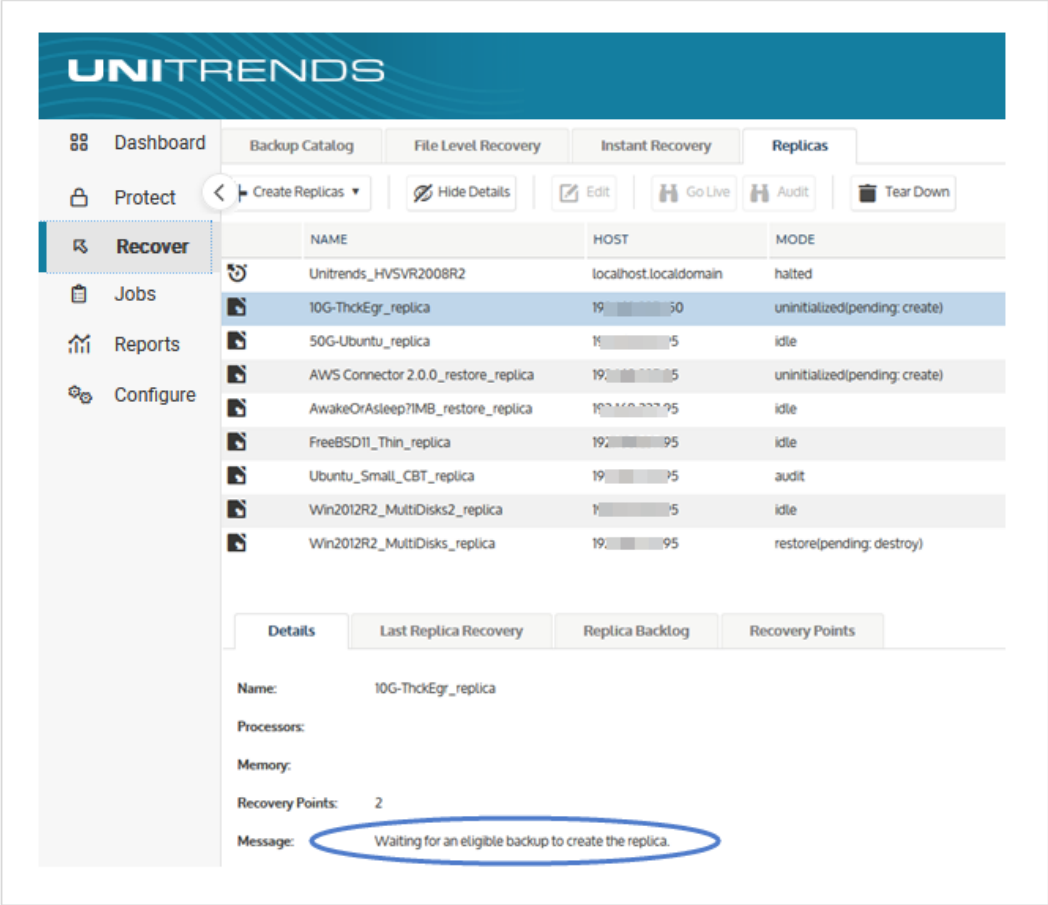

To create multiple VM replicas

1. Log in to the backup appliance.
2. Select **Recover**, then click the **Replicas** tab.
3. Click **Create Replicas** and select **VMware**.

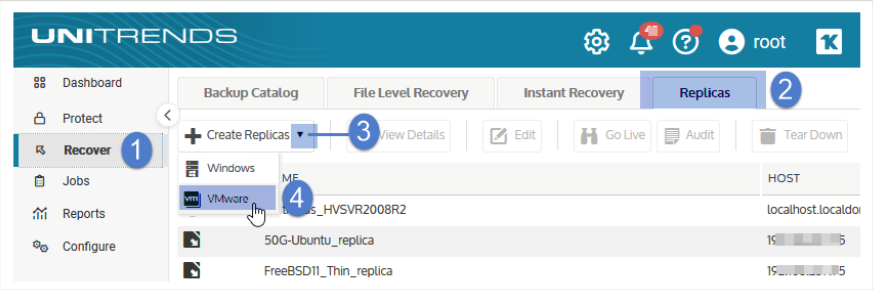

4 In the Inventory tree, expand the virtual host and select the VMs to protect.

- To locate a VM by name, use the **Search** field below.
- To view individual VMs, expand the virtual host and any vApps and resource pools.
- To quickly select multiple VMs, click a virtual host, vApp, or resource pool checkbox.
- To select one VM, click its checkbox.

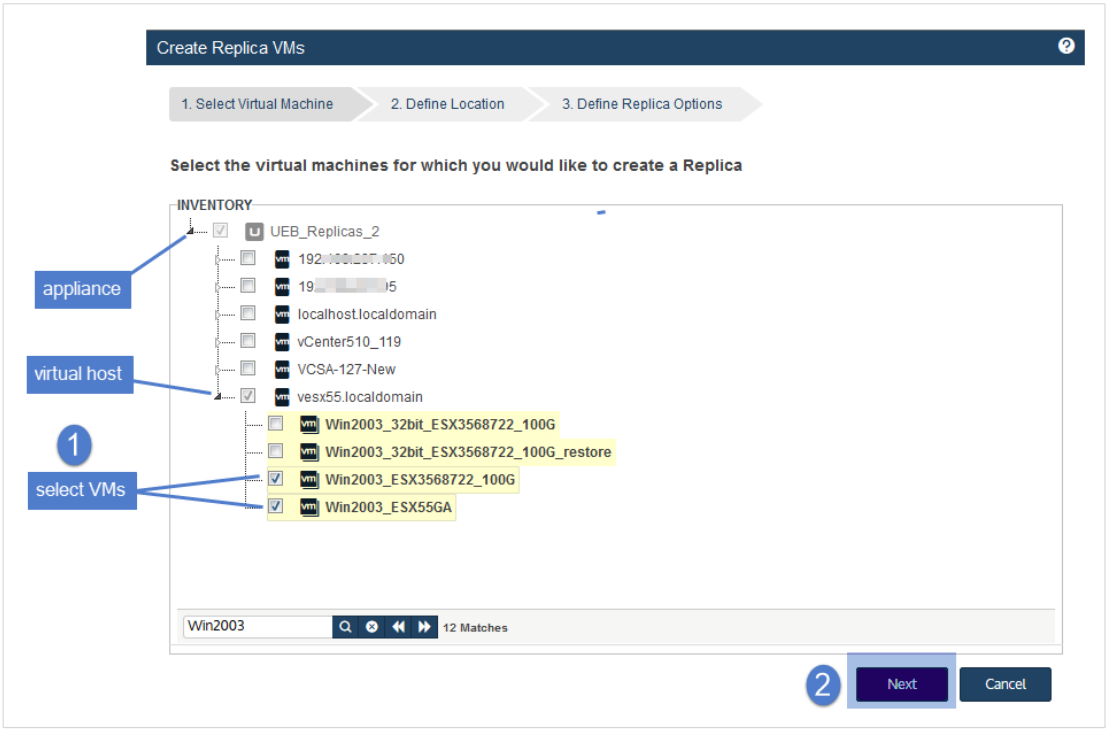

5 Click **Next**.

6 Enter Location Details to specify the location where the replicas will reside:

- Location – Select a VMware host from the list. The list contains all eligible virtual host assets that have been added to the appliance. (For details on adding a virtual host, see "To add a virtual host asset" on page 311.)

- Resource Pool – (Optional) If your VMware environment has resource pools or vApps, you can opt to select one in the list.
- Storage – Select the datastore that will be used to create the virtual disks for all of the replica VMs.

7 Click **Next**.

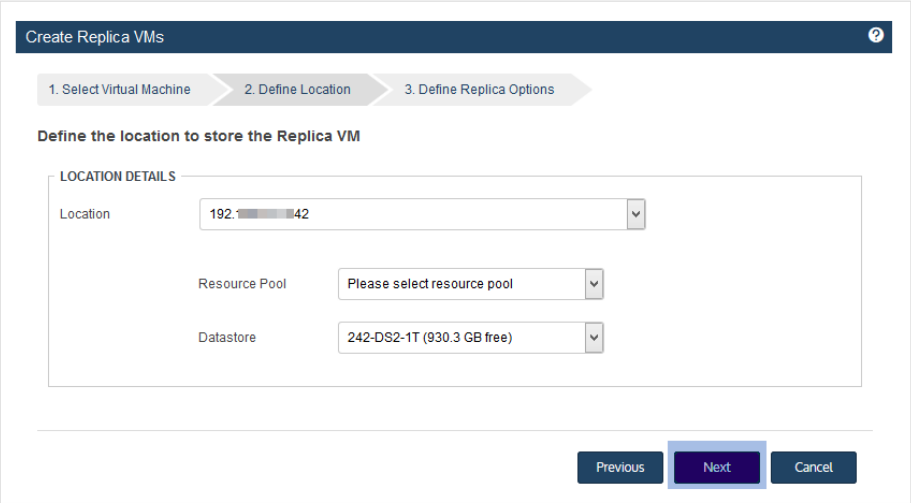

8 Review Replica Details:

- Replica Name – Replicas are named <VMname>_replica.
- Number of recovery points – Maximum number of recovery points stored with each replica. By default, only the most recent recovery point is saved. You can opt to increase this value, up to the maximum number configured for the appliance. (To modify the maximum allowed on the appliance, see "VM replica configuration" on page 177.).

9 Click **Save**.

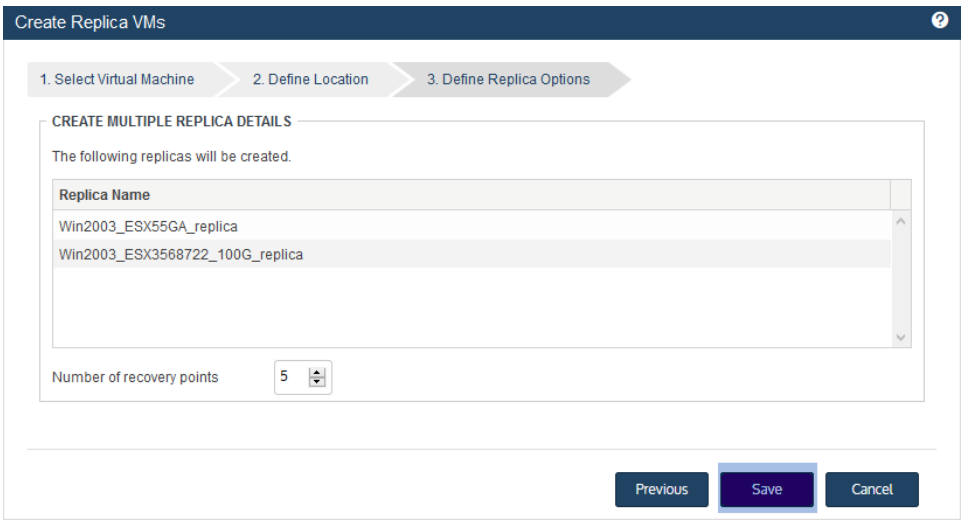

10 The appliance does one of the following for each VM:

- If a backup exists, creates a replica, then applies the latest backup. The replica's mode changes to *restore* while the backup is being applied, then to *idle*. After the replica enters idle mode, you can audit the replica or bring it 'live' as needed. We recommend that you audit the replica soon after it enters idle mode, to verify its integrity. See "Auditing a VM replica" on page 893 for details.
- If no backup is found, creates a "shell" replica in *uninitialized(pending: create)* mode. The replica remains in this mode until a successful VM backup is created. Run a backup of the original VM as soon as possible. (An uninitialized replica cannot assume the role of a failed VM. At least one backup must be applied before the replica can replace the original VM.)

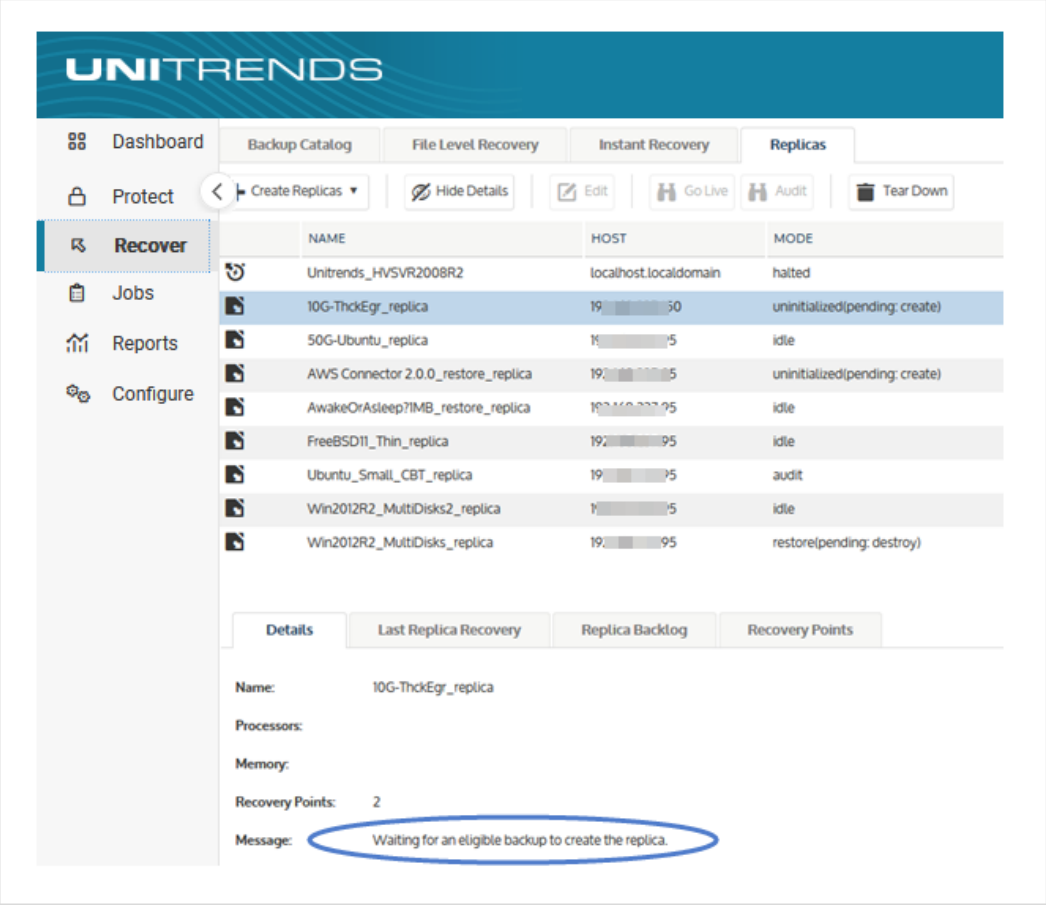

Working with VM replicas

After creating VM replicas, use the following procedures as needed:

- "Editing a VM replica"
- "Auditing a VM replica"

- "Bringing the VM replica live in production" on page 895
- "Tearing down a VM replica" on page 898
- "Monitoring VM replicas" on page 899

Editing a VM replica

After creating a replica, you can modify the following settings at any time:

- The number of processors connected to the replica.
- The amount of memory attached to the replica.
- The number of recovery points to store with the replica.
- Whether to apply new backups to the replica (*Enable restores* checkbox).

To edit a VM replica

1. Log in to the backup appliance.
2. Select **Recover**, then click the **Replicas** tab.
3. Select the replica, then click **Edit**.

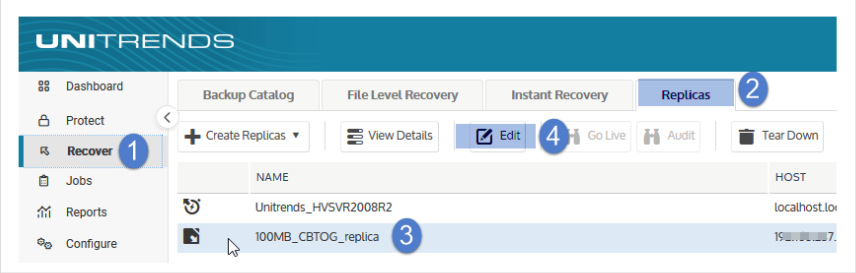

4. Modify settings as desired.
5. Click **Save**.

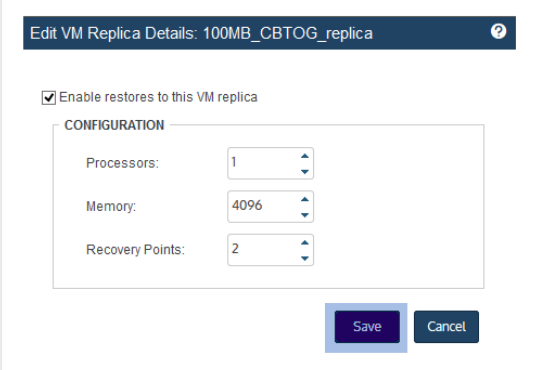

Auditing a VM replica

Audit mode enables you to run the replica while the original VM is still operating in production. A replica running in audit mode boots with no network interface. Auditing the replica with the original VM still online does not result in network conflicts or impact the original VM in any way. However, applications on the replica that require network access do not function fully in audit mode.

It is recommended that you audit each newly created replica to ensure it functions as expected, and that you perform additional audits at regular intervals to check subsequent recovery points.

A newly created replica cannot be audited until at least one backup has been applied. During the audit, no subsequent backups are applied. Upon exiting audit mode, the appliance applies any backups that completed during the audit to bring the replica up to date.

Auditing the replica is a two-part process where you bring the replica into audit mode and then access the replica to verify that it is functioning as expected. During the audit, you should verify the following:

- The replica boots successfully and is operational.
- The replica contains the expected data and applications.

After you have finished auditing the replica, you must take it out of audit mode so the appliance can resume applying backups. (Note that any changes made to the replica VM during the audit are lost upon exiting audit mode.)

Audit mode procedures

Use these procedures to audit the VM replica:

> **Note:** In audit mode, the replica uses hypervisor resources. By default, the replica's compute resources (processors and memory) match those of the original VM. If needed, you can view and modify these settings before entering audit mode by using the "Editing a VM replica" on page 892 procedure.

- "To bring the VM replica into audit mode"
- "To access a VM replica in audit mode" on page 895
- "To exit audit mode" on page 895

To bring the VM replica into audit mode

1. Log in to the backup appliance.
2. Select **Recover**, then click the **Replicas** tab.
3. Select the replica, then click **View Details**.

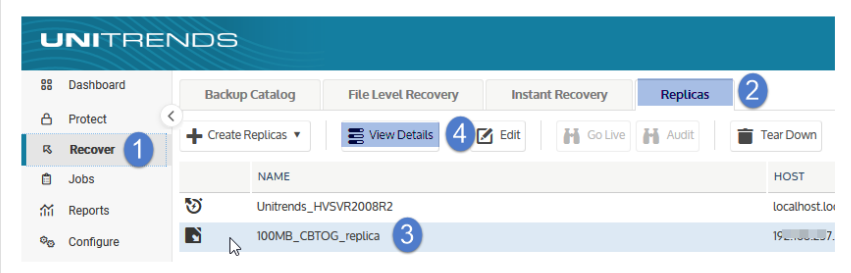

4 Select a recovery point, then click **Audit**.

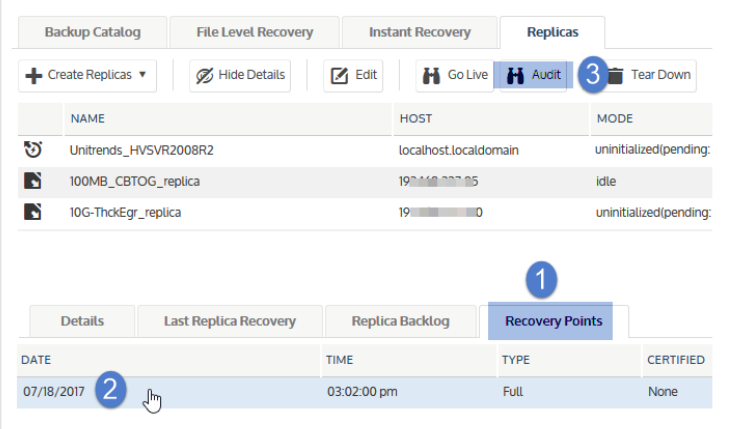

5 Click **Confirm**.

The replica's mode changes to *idle (pending audit)*, then to *audit*.

Note: If a restore is in progress - Beginning in release 10.1.2, the VM boots into audit mode immediately and the running replica restore job is canceled. Note that the next restore may take extra time. If you have performed an audit while a restore was running, do not interrupt the next restore. Repeatedly interrupting restore jobs significantly degrades performance. (For details, see "Entering audit mode while a restore is in progress" on page 878.)

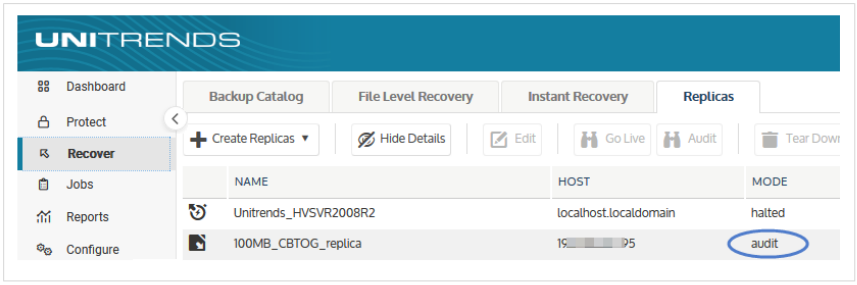

6 After the replica is in audit mode, you can connect to the replica to verify that it is functioning as expected. See "To access a VM replica in audit mode" for details.

To access a VM replica in audit mode

After a VM replica has entered audit mode, use this procedure to access the replica:

1 Connect to your hypervisor manager.

2 Locate the replica in the list of virtual machines, and access it the same way you access all VMs on the hypervisor.

3 After verifying that the replica is running with its recovered data, proceed to "To exit audit mode".

To exit audit mode

1 Log in to the backup appliance.

2 Select **Recover**, then click the **Replicas** tab.

3 Select the replica, then click **End Audit**.

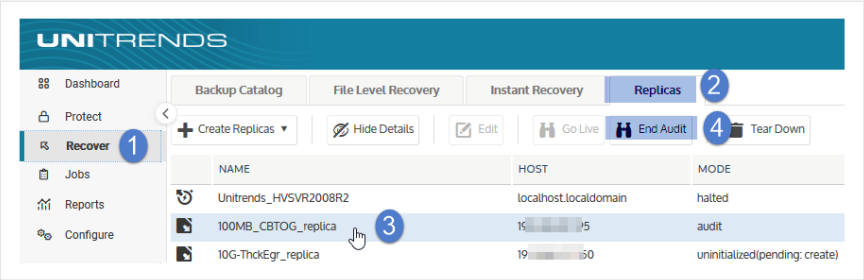

4 The replica exits *audit* mode. Its mode changes to *audit (pending: off)*, then to one of the following:

- *Restore* – One or more backups successfully completed during the audit. The appliance is currently applying a backup.
- *Idle* - The replica is idle (there are no backups to apply or a replica restore job is not yet running).

Bringing the VM replica live in production

If disaster strikes and the original VM fails, you can replace it with the replica by booting into *live* mode. Because the replica is continually updated with the original VM's data, it can immediately assume the role of the original VM.

See these topics for details:

- "Live mode recommendations"
- "To bring the VM replica into live mode" on page 896

Live mode recommendations

Review these recommendations before going into live mode:

- The replica can replace the original VM temporarily or be used as a permanent replacement.

- Beginning in release 10.1.2, restore enhancements enable you to quickly enter live mode while a restore is in progress. If a restore is in progress, the VM boots into live mode and cancels the running replica restore job.
- Upon entering live mode, the replica is marked *invalid* because the replica role no longer applies. After configuring network settings for the live replica VM and verifying that it is functioning as expected in production, you tear down the replica on the Unitrends appliance. You can then opt to create a new replica for the live VM.
- In live mode, the replica uses hypervisor resources. By default, the replica's compute resources (processors and memory) match those of the original VM. If needed, you can view and modify these settings before entering live mode by using the "Editing a VM replica" on page 892 procedure.
- Entering live mode enables the replica VM to be discovered by the appliance. Once in live mode, you can protect the replica VM with Unitrends backups.
- Live mode should be used for a short time only. You should exit live mode (as described in "Tearing down a VM replica" on page 898) as soon as you have verified that the replica is functioning as expected in production. Select the *Delete the VM replica from the appliance only* option so you can retain the replica VM. This way, you can keep using the replica for as long as needed and capture any changes by running backups of the replica VM.

To bring the VM replica into live mode

This procedure provides instructions for booting a replica in live mode. Be sure to shut down the original VM before running this procedure.

1. Log in to the backup appliance.
2. Select **Recover**, then click the **Replicas** tab.
3. Select the replica, then click **View Details**.

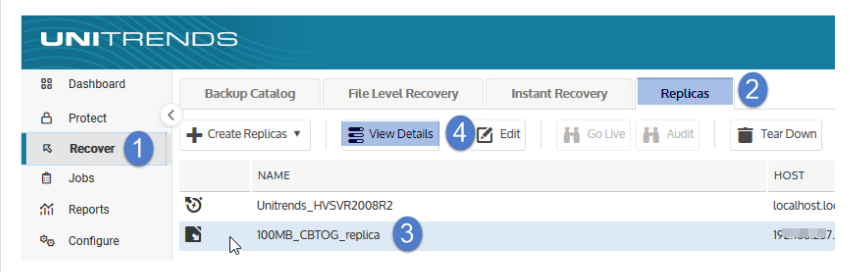

4. On the Recovery Points tab below, select a recovery point. Then click **Go Live**.

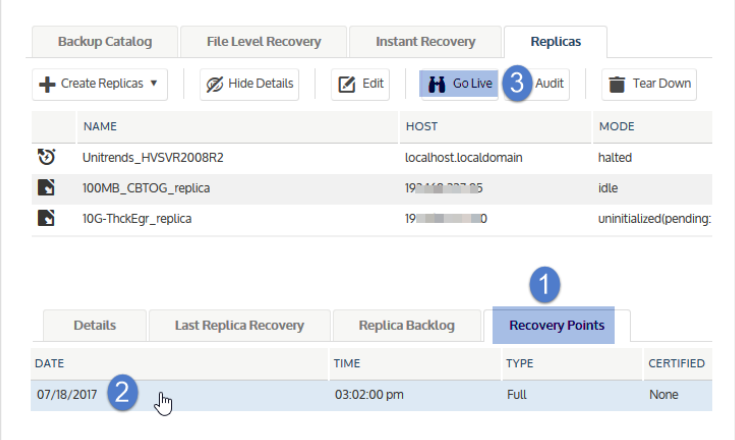

5 Click **Confirm**. The replica's mode changes to *idle(pending: live)* and then to *live*.

Note: Upon entering live mode, the replica is marked *invalid* because the replica role no longer applies.

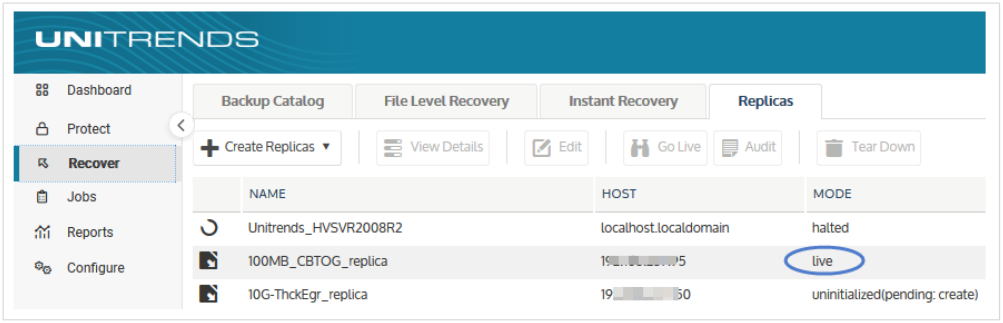

6 Access the replica VM by using the hypervisor manager and verify that the replica is functioning as expected.

7 Configure the replica VM's network settings to bring it onto the production network.

8 Verify that the VM is functioning as expected in production.

At this point, the replica has assumed the identity of the original VM.

9 Delete the replica *from the appliance only*, by using the "Tearing down a VM replica" on page 898 procedure. Be sure to select the *Delete the VM replica from the appliance only* option to retain the replica VM on the ESXi host.

10 Create a backup schedule to protect the replica VM. For details, see "Backup Administration and Procedures" on page 425.

Notes:

It can take several minutes for the replica VM to show up in the list of VMs to protect with VMware backups. To refresh the list of discovered VMs, click the Gear icon in the upper-right of the UI and select **Inventory Sync**.

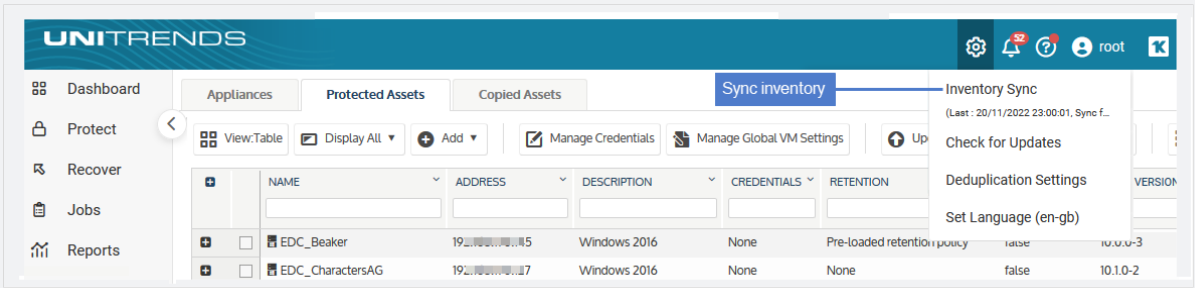

Tearing down a VM replica

This section provides instructions for deleting a VM replica. You have these options:

Tear down option	Description and considerations
Delete the VM replica from the appliance only	Select this option to retain the replica VM on the hypervisor. **Note:** To optimize backups, you must bring the replica into live mode before you do the tear down procedure. (For details, see "Bringing the VM replica live in production" on page 895.) Upon entering live mode, the appliance merges existing snapshots to enable optimized (sparse) backups. Be sure to enter live mode before doing the tear down if you will be retaining the replica VM on the hypervisor.
Delete the VM replica from the hypervisor and appliance	Select this option if you no longer need to use the replica VM. (For example, you have recovered the original VM to a different virtual host and will not be using the replica VM as a permanent replacement.)

To tear down a VM replica

IMPORTANT! If you are running a live replica as a replacement for a VM that has failed, do not tear down the replica until you have verified that the replica VM is functioning as expected in production.

1. Log in to the backup appliance.
2. Select **Recover**, then click the **Replicas** tab.
3. Select the replica, then click **Tear Down**.

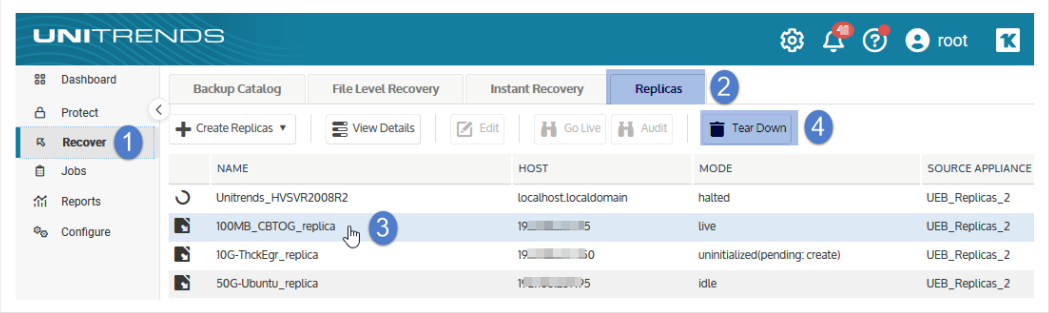

4 Select the desired option and click **Tear Down**.

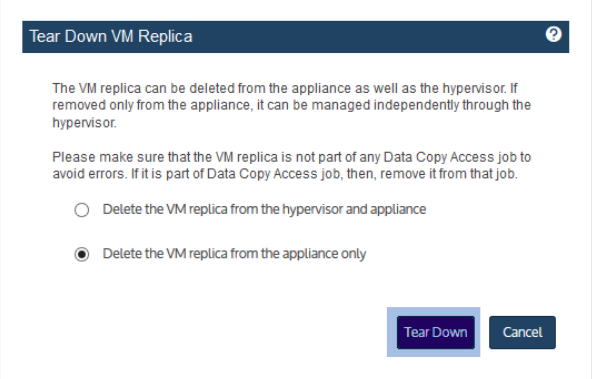

It can take several minutes for the appliance to purge all information about the replica. If you need to create a new replica for the original VM, you must wait for this information to purge. If it has not yet purged, the original VM does not display in the list of discovered VMs in the Create Replica VMs dialog.

Monitoring VM replicas

Use these procedures to check the status and details of existing VM replicas:

- "To view all VM replicas"
- "To view VM replica details" on page 900
- "VM replica modes" on page 902

To generate a report of backups that have been applied to replicas, see "Replicas History report" on page 1360.

To view all VM replicas

1 Log in to the backup appliance.

2 Select **Recover**, then click the **Replicas** tab.

3 All VM replicas and Windows replicas display in a list on the Replicas tab.

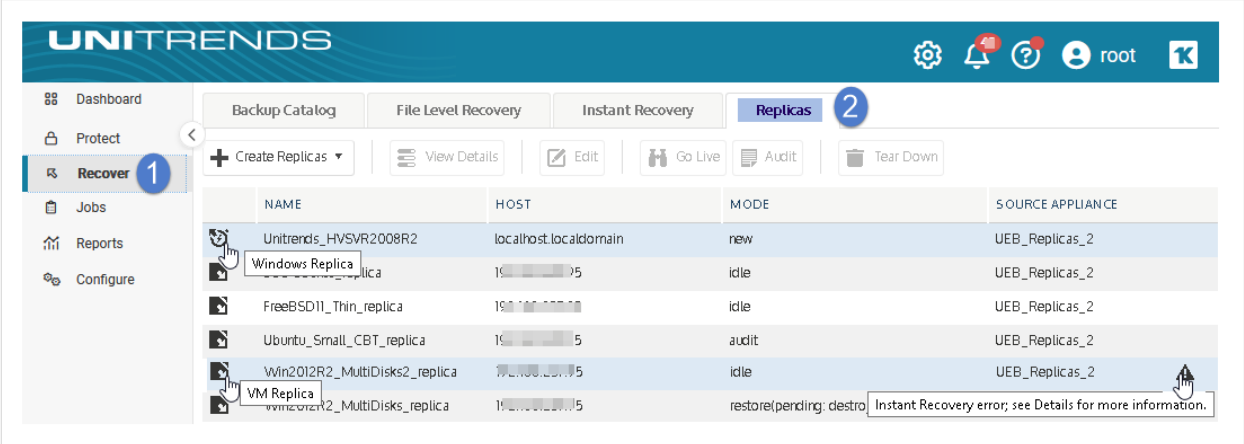

The following information is given for each VM replica:

Column	Description
Replica type icon	Indicates the replica type: VM (or Windows).
Name	Replica name. By default, the replica is named <VMName>_replica.
Host	Virtual host where the VM replica resides.
Mode	Replica mode. Examples: new, audit, restore, or idle. See "VM replica modes" on page 902 for additional details.
Source Appliance	Appliance where the replica was created. **Note:** If the replica was created on a backup copy target by using a hot backup copy, this field displays the source appliance where the original backup was taken and the target appliance where the hot backup copy resides, in this format: *BackupAppliance --> BackupCopyTargetAppliance*.
Alert icon	Indicates that an alert has been generated for the replica. Hover over the icon for details.

To view VM replica details

1 Log in to the backup appliance.

2 Select **Recover**, then click the **Replicas** tab.

3 Select the replica, then click **View Details**.

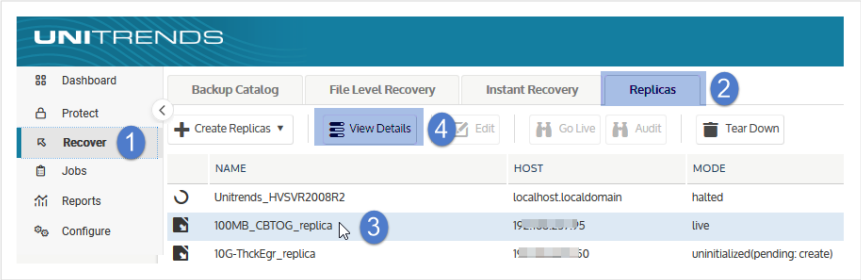

4 Tabs containing replica details display below:

Details tab	Description
Details	Shows the replica name, number of processors, amount of memory, number of recovery points, and status message. In this example, the replica entered live mode on September 22 by using a recovery point snapshot from July 18: 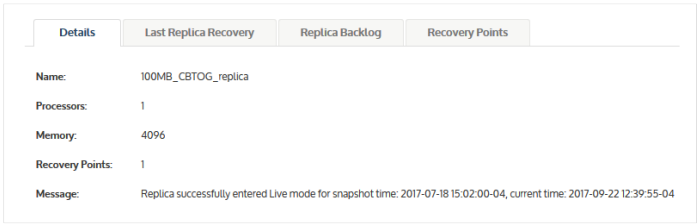
Last Replica Recovery	Each time a backup is taken of the original VM, the appliance applies it to the replica (by running a recovery operation). This tab shows the date and time of the backup that was applied to the replica in the most recent recovery operation. In this example, a backup taken on July 18 is the most recent recovery point that has been applied to the replica: 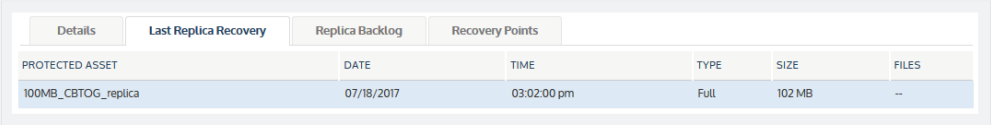
Replica Backlog	List of backups that have not yet been applied to the replica. No backups are listed if the replica is up to date. Example: 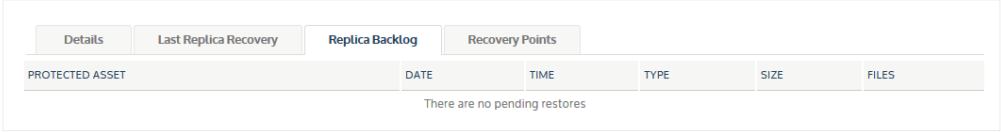

Details tab	Description
Recovery Points	Shows the recovery point snapshots that can be used to boot the replica in audit or live mode. (To modify the maximum number of recovery points retained for the replica, see "Editing a VM replica" on page 892.) In this example, the five most recent recovery points are available: **Note:** *Certified* indicates whether the recovery point has been certified by Copy Data Management. For details, see "Recovery Assurance" on page 1263.

Details | Last Replica Recovery | Replica Backlog | **Recovery Points**

DATE	TIME	TYPE	CERTIFIED
09/09/2017	12:00:33 am	Incremental	None
09/09/2017	12:00:37 pm	Incremental	None
09/10/2017	12:00:40 am	Incremental	None
09/10/2017	12:00:33 pm	Incremental	None
09/11/2017	12:00:36 am	Incremental	None

5 (Optional) Click **Hide Details** to stop displaying details for the selected replica.

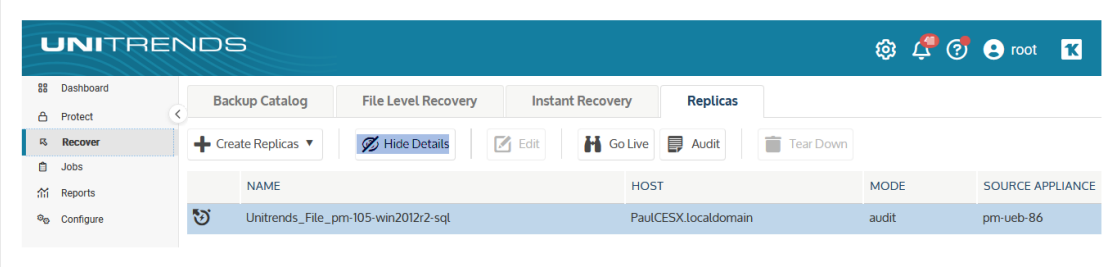

VM replica modes

You can monitor a replica by checking its mode on the Replicas tab. The mode indicates what is currently happening with the replica (for example, whether it is newly created, whether a backup is being applied, or whether it is in audit mode.)

VM replica modes are described in the following table:

Mode	Description
New	The replica is new and no backup has been applied. A replica in *new* mode cannot be audited or booted into live mode.
Restore	A backup has completed, and the appliance is applying it to the replica. The replica remains

Mode	Description
	in *restore* mode until the restore completes, or until you enter audit or live mode. **Note:** Beginning in release 10.1.2, restore enhancements enable you to quickly enter audit or live mode while a restore is in progress. If a restore is in progress, the VM boots into audit or live mode immediately and the running replica restore job is canceled. For details, see "Replica restore jobs" on page 877.
Idle	At least one backup has been applied to the replica, but currently no action is occurring.
Halted	The appliance has attempted to apply a backup to the replica but the restore could not be performed. The failed restore job has caused the replica to go into *halted* mode. The following can occur when a replica is in this mode: • If the restore could not be performed because the appliance could not reach the replica, it tries again after several minutes, and the mode changes from *halted* to *idle*. After three failed attempts, the replica becomes invalid, and it remains in *halted* mode until a user deletes it. • If the restore could not be performed because a configuration change was made to the original VM, the replica becomes invalid, and it remains in *halted* mode until a user deletes it. • If the restore could not be performed because the recovery job was canceled by a user (by clicking **Cancel** on the Active Jobs tab) and the Unitrends appliance is running a pre-10.1.2 release, the replica becomes invalid, and it remains in *halted* mode until a user deletes it. **Note:** For VM replica restores where at least one backup has been successfully applied to the replica and the appliance is running release 10.1.2 or later, the replica is not invalidated and does not enter halted mode. For details, see "Replica restore jobs" on page 877. Once a replica is invalidated and enters *halted* mode, you must tear down the replica and create a new one. (For details, see "Tearing down a VM replica" on page 898 and "Creating VM replicas" on page 885.)

Mode	Description
	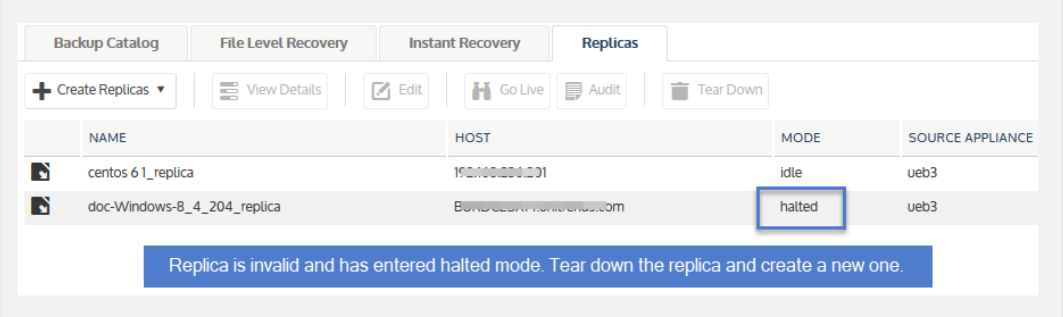
Audit	A user is performing a manual audit and the replica has booted in *audit* mode. For details, see "Auditing a VM replica" on page 893.
Live	A user has booted the replica in *live* mode to replace the original VM. For details, see "Bringing the VM replica live in production" on page 895.

Virtual machine instant recovery

Instant recovery enables you to recover a failed or corrupted VMware or Hyper-V VM and access it in minutes.

To perform instant recovery, you specify a recovery point (by selecting a backup or backup copy) and a target location where the recovered VM will reside (by selecting a virtual host). Instant recovery then creates a disk image *recovery object* on the Unitrends appliance and a new VM on the target host. This takes just a few minutes. The new VM is created with the same network settings as the original and can immediately assume production operations for the failed VM.

Upon creating the recovered VM, instant recovery migrates data from the on-appliance recovery object to the new VM. The recovered VM remains fully operational during the migration.

After the data has been migrated, instant recovery is complete and the recovery object is no longer needed. You then tear down the instant recovery session. This deletes the recovery object, freeing the appliance resources and reserved space.

Instant recovery modes

You can choose to perform the recovery in *audit* mode or *instant recovery* mode. Use audit mode to verify recovery points for VMs that are still running in production. Use instant recovery mode to replace a failed or corrupted VM. Descriptions of each mode are given in the following table:

Mode	Description
Audit	Enables you to verify that a VM can be created from a backup or backup copy. The appliance uses data from the selected backup or backup copy to create a disk image on the appliance and a new VM on the target host. Although the VM resides on the host, it runs from the disk image on the appliance. All other resources, such as the processors and memory, reside on the host. A VM in audit mode is not intended for production use. It does not have network connectivity, and changes made to the VM in audit mode are not backed up on the Unitrends appliance. Recovering a VM in audit mode has no impact on the original VM. It is not necessary to shut down the original VM during the audit, even if you use the original host as the recovery target. After verifying that the VM has booted and its data is accessible, you tear down the recovery session. This deletes the VM from the target host and deletes the on-appliance recovery object, freeing the appliance resources.
Instant recovery	Enables you to replace a failed or corrupted VM. When you select a backup or backup copy, the appliance uses data from this backup to create a disk image recovery object on the appliance and a new VM on the target host. The new VM is available for use immediately. The Unitrends appliance uses Storage vMotion (VMware) or Storage Live Migration (Hyper-V) to copy the data from the disk image to the target host. Once the data migration is complete, you tear down the recovery session to free appliance resources.

Preparing for instant recovery

Unitrends recommends planning for instant recovery before a VM fails. Following is a summary of the steps needed to set up instant recovery for your virtual machines. Steps include links to detailed instructions for each procedure.

Step 1: Ensure that all requirements have been met. For details, see "Prerequisites for VMware instant recovery" on page 906 or "Prerequisites for Hyper-V instant recovery" on page 907.

Step 2: Run host-level backups of the VMs:

- To create a job manually, see "To create a VMware backup job" on page 455 or "To create a Hyper-V backup job" on page 462.
- To create a job by using an SLA policy, see "To create an SLA policy for VMware assets" on page 551 or "To create an SLA policy for Hyper-V assets" on page 556.
- For a comparison of the manual and SLA policy job creation methods, see "About creating backup and backup copy jobs" on page 426.

Step 3: Reserve space on the appliance for instant recovery. For details, see "Allocate storage for instant recovery" on page 911.

Step 4: Add target virtual hosts to the Unitrends appliance as needed.

While running the IR procedure, you select a virtual host where the recovered VM will be created. You can recover to the host where the original VM resides or to a different location. If needed, add additional virtual

hosts to the appliance to make them available for IR. For details, see "Adding a virtual host" on page 308.

Note: For VMware, a vCenter is required. Add the ESXi hosts and the vCenter managing the hosts.

Step 5: Run IR in audit mode to check the recovered VMs. For details, see "Perform instant recovery in audit mode" on page 913.

Prerequisites for VMware instant recovery

The following table describes prerequisites and considerations for VMware instant recovery.

Prerequisite or consideration	Description
vCenter version and license	To perform instant recovery, the ESXi server used as the instant recovery target must be managed by a vCenter that meets the following requirements: • Must be running one of the following: vCSA 5, vCenter version 5, or a higher vCenter version listed in the Compatibility and Interoperability Matrix. • Must have a license that supports Storage vMotion. • Must be a added to the Unitrends appliance from which you are performing the instant recovery. (Both the vCenter and the ESXi server must be added to the backup appliance as a virtual host asset. For details, see "Adding a virtual host" on page 308.) **Note:** You can perform the audit process using a stand-alone ESXi server, but instant recovery mode is not supported.
ESXi server	The ESXi server used as the instant recovery target must meet the following requirements: • Must be managed by a vCenter that meets the version and license criteria above. • Must be running ESXi version 5 or a higher version listed in the Compatibility and Interoperability Matrix. • Must be added to the Unitrends appliance from which you are performing the instant recovery. (See "Adding a virtual host" on page 308.) • Must be running the same ESXi version as the original server hosting the virtual machine, or a higher version listed in the Compatibility and Interoperability Matrix. It is highly recommended that you recover to an ESXi version that matches the original. • Must support the operating system (OS) of the VM you are recovering. (See the VMware documentation for details.) For example, you cannot recover a Windows

Prerequisite or consideration	Description
	2016 VM to ESXi 5.1. • Must have sufficient space and compute resources for the new VM.
Backup	You must have a backup to recover a virtual machine. The backup used for the instant recovery must be: • A successful host-level backup or backup copy of the virtual machine. This can be a full, incremental, or differential backup. (To run a backup, see "To create a VMware backup job" on page 455.) • A local backup, an imported backup copy, or a hot backup copy (supported if performing IR on the target appliance where the hot copy resides). The backup used for instant recovery cannot contain a 4096 sector disk. 4096 sector disks are not supported for instant recovery.
VM configuration	The following configuration settings apply to the recovered VM: • The Mac address and network settings of the recovered VM match those of the source VMware backup used for the recovery. • The recovered VM is configured with the latest hardware version that is supported by the target hypervisor. For example, if a hardware version 8 VM is recovered to an ESXi 5.5 server, the recovered VM is hardware version 10. • The recovered VM's disks are provisioned as Thick Eager Zeroed.

Prerequisites for Hyper-V instant recovery

The following table describes prerequisites and considerations for Hyper-V instant recovery.

Prerequisite or consideration	Description
Hyper-V server	• The Hyper-V server used as the instant recovery target must be one of the following: – A Windows Server with the Hyper-V role enabled, running Windows 2012 or a higher version listed in the Compatibility and Interoperability Matrix. – A Hyper-V Server running 2012 or a higher version listed in the Compatibility and Interoperability Matrix.

Prerequisite or consideration	Description
	• The Unitrends Windows agent must be installed on the Hyper-V server hosting the protected VM and on the Hyper-V server used as the recovery target. Unitrends recommends reloading the list of VMs on the appliance after installing the agent on the host Hyper-V servers. To reload the list of VMs, select **Options > Inventory Sync**. • The restore target can be the original Hyper-V server or an alternate Hyper-V server. • The restore target must be running the same version as the Hyper-V server hosting the original VM, or a higher version listed in the Compatibility and Interoperability Matrix. It is recommended that you restore to a Hyper-V server whose version matches that of the original, where possible. **Note:** Backups from older versions of Hyper-V can be used for Hyper-V instant recovery as long as the target server is running 2012 or a higher version listed in the Compatibility and Interoperability Matrix. • The target server must be added to the Unitrends appliance from which you are performing the instant recovery. (See "Adding a virtual host" on page 308.) • The target server must support the operating system (OS) of the VM you are recovering. (See the Microsoft documentation for details. For Windows, see this Microsoft article: Should I create a generation 1 or 2 virtual machine in Hyper-V?) For example, you cannot recover a Windows 2016 VM to Hyper-V 2008 R2. • The target server must have sufficient space and compute resources for the new VM. • The Hyper-V host must be configured for adequate simultaneous storage migrations. The host must be able to run simultaneous storage migrations for all sessions that are currently running in instant recovery mode. (Storage migration is not needed to run IR in audit mode). Inadequate simultaneous storage migrations leads to undesirable results. To modify the number of simultaneous storage migrations, open Hyper-V manager, right-click the host server and select **Properties**, select **Storage Migrations** and modify the **Simultaneous storage migrations** setting:

Prerequisite or consideration	Description
	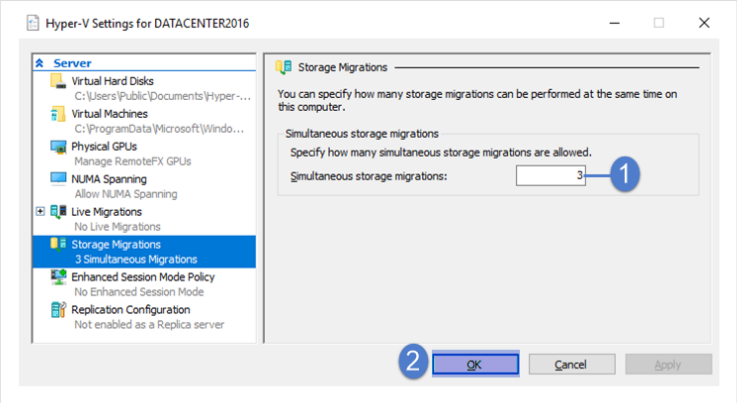
Backup	You must have a backup to recover a virtual machine. The backup used for the instant recovery must be: • A successful host-level backup or backup copy of the virtual machine. This can be a full or incremental backup. (To run a backup, see "To create a Hyper-V backup job" on page 462.) • A local backup, an imported backup copy, or a hot backup copy (supported if performing IR on the target appliance where the hot copy resides). • Hyper-V Instant Recovery is not supported for backups with files or pathnames that contain non-UTF8 characters. • Hyper-V Instant Recovery is not supported for Host component (AzMan Security Database for Hyper-V) backups of the original VM.
VM configuration	The following VM configuration settings and requirements apply to the recovered VM: • A recovered VM's disk names match those of the source Hyper-V backup used for the recovery. • All disks created on the Hyper-V VM must have unique names. • The network settings of the recovered VM match those of the source Hyper-V backup used for the recovery. • The recovered VM is configured with the latest hardware generation version that is supported by the target Hyper-V server. • A recovered VM's configuration version matches that of the source Hyper-V

Prerequisite or consideration	Description
	backup used for the recovery. • Disks excluded during backup, including independent and physical RDM disks, are not restored with instant recovery for Hyper-V. **Note:** Not all VM settings are preserved during instant recovery. Review the list of preserved and non-preserved settings to determine if instant recovery is the best approach for your VM. **VM settings preserved during instant recovery:** • Firmware type (BIOS or UEFI) • Processor count • Memory – Startup memory, minimum memory, maximum memory, and dynamic memory • Number of network adapters **VM settings NOT preserved during instant recovery:** • BIOS startup order (the position of the boot disk is preserved) • UEFI – – Boot order (the position of the boot disk is preserved) – Secure boot if disabled – Encryption support enable trusted platform module – Security policy shielding • Processor resource control • Network Adapter VLAN ID and Bandwidth Management • COM ports • Integration Services Offerred • Checkpoint File Location • Smart Paging File Location • Automatic Start Action

Prerequisite or consideration	Description
	• Automatic Stop Action
Cluster requirements	Hyper-V instant recovery supports clustered and non-clustered virtual machines. The following requirements must be met to recover a clustered VM: • The target Hyper-V server determines the cluster status of the new VM. To create a clustered VM with instant recovery, you must select a cluster as the target Hyper-V server. If you select an individual member node, the resulting VM is not clustered. • When creating a clustered VM with instant recovery, you must select the network switch common to all nodes in the cluster. If you do not select this switch, the new VM will lose network connectivity if it fails over to another cluster node.

Allocate storage for instant recovery

Instant recovery can be performed at any time, as long as there are backups or backup copies for your VMs and sufficient backup storage is allocated as instant recovery space. Unitrends strongly recommends reserving space for instant recovery soon after initial deployment and before a VM fails. If necessary, you can allocate instant recovery space later, but doing so may require the appliance to purge local backups to make room for the newly allocated instant recovery space.

Because the disks for a recovered VM reside on the appliance until storage migration completes, you must allocate a portion of your backup storage for instant recovery. You must allocate at least 20 percent of the space used on the VM's original disks. Once disks have been migrated to the recovered VM, appliance instant recovery storage is no longer needed.

Notes:

- Appliance storage that is allocated to instant recovery can also be used for the Windows replica feature.
- Storage allocated for instant recovery cannot be used for backups.
- Recovery Series and Recovery MAX physical appliances come with a set amount of backup storage. Backup storage allocation can be modified to increase the amount used for instant recovery. Backup storage cannot be added to the appliance.
- Unitrends Backup virtual appliances are deployed as virtual machines. After initial deployment, you can add more backup storage as desired. For details, see "About adding backup storage to a Unitrends Backup appliance" on page 199.

To allocate storage for instant recovery

1. On the **Configure > Appliances** page, select the appliance.
2. On the **Storage** tab, select the **Internal** storage and click **Edit**.

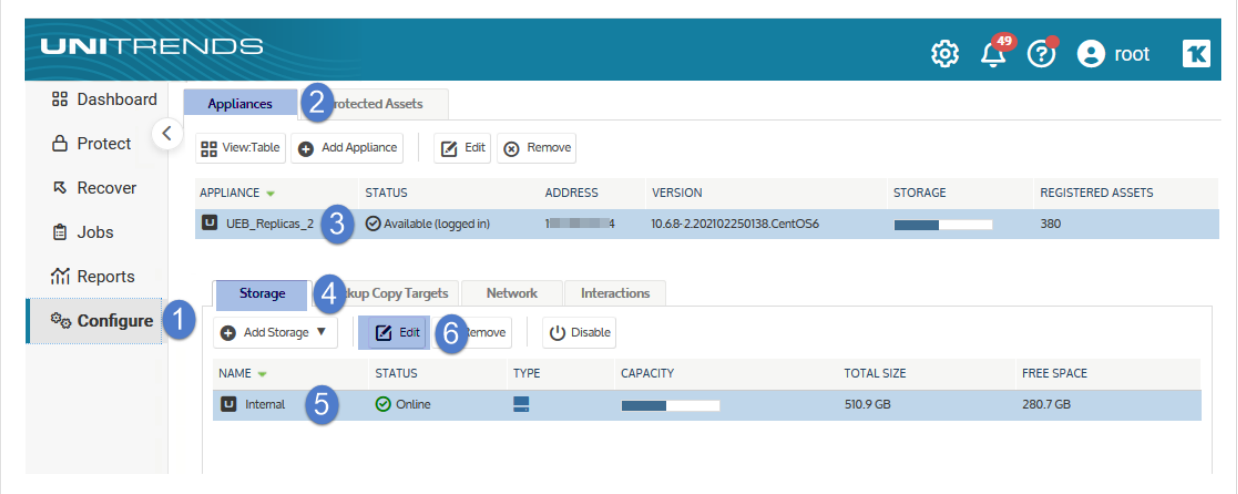

3 Modify the percentages used for backups versus instant recovery, and click **Save**.

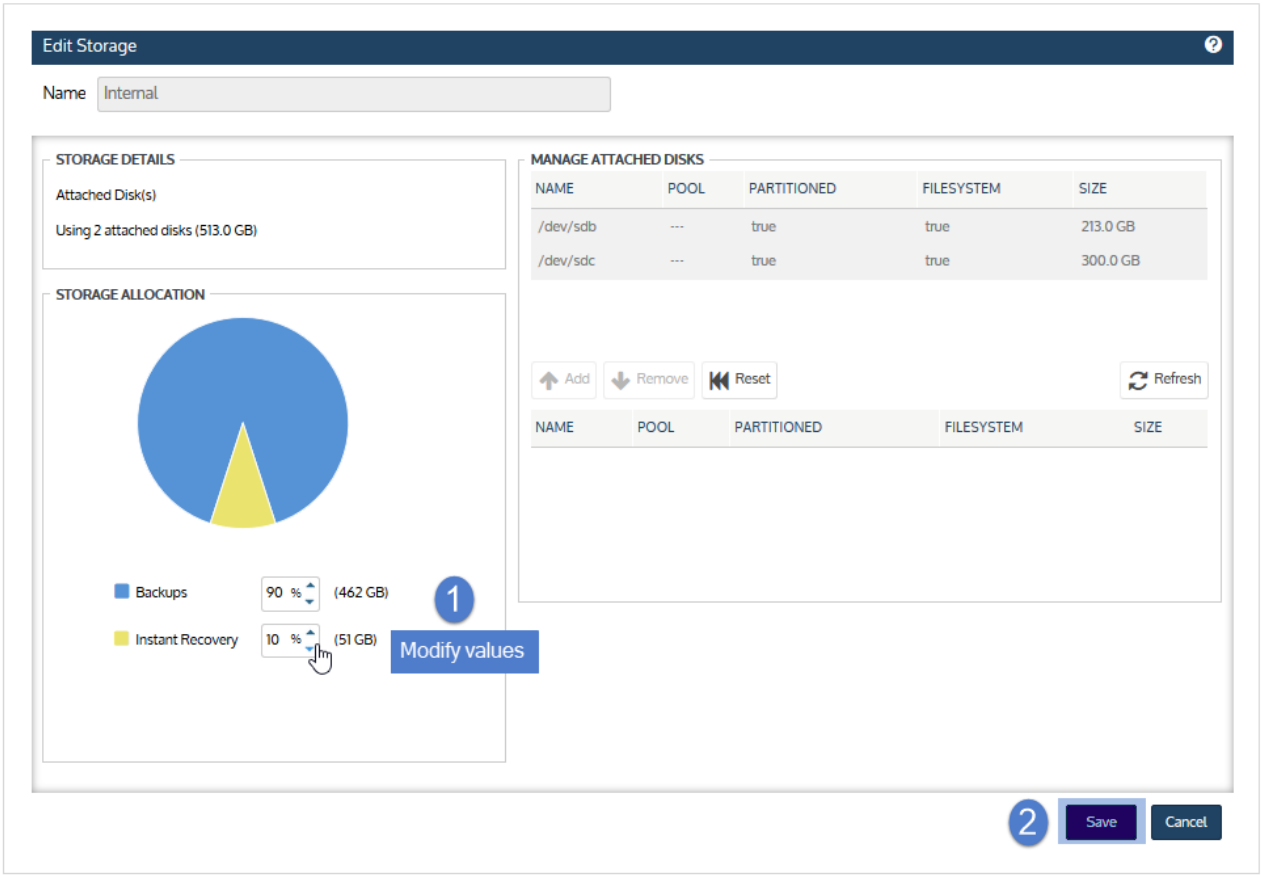

Perform instant recovery in audit mode

Perform instant recovery in audit mode to verify that backups and backup copies can be used to recover the VM in the event of a disaster. Repeat this procedure as needed to test new backups.

To perform instant recovery in audit mode

Run this procedure from the backup appliance to recover from a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover from a hot backup copy.

1 Log in to the backup appliance or target appliance (if recovering from a hot backup copy).

2 Select **Recover** and click the **Backup Catalog** tab.

(Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

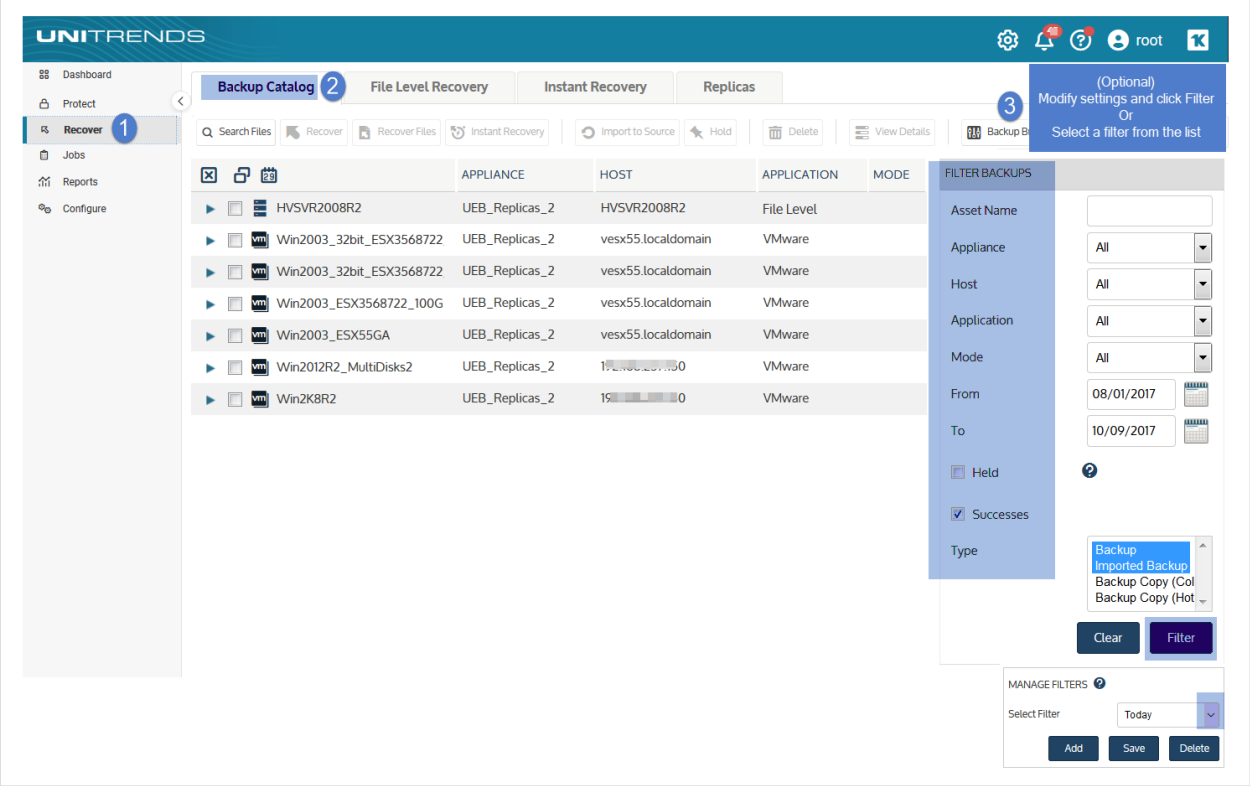

3 Expand the VM asset and select one of the following to use for the recovery:

- A host-level backup.
- An imported host-level backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
- A hot backup copy (supported if performing IR on the target appliance where the hot copy resides).

4 Click the **Instant Recovery** button.

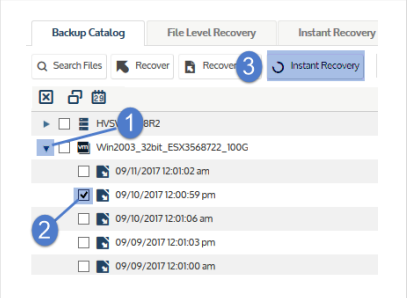

5 Check the **Recover this VM in Audit Mode** box.

6 Select the following Recovery Options:

Recovery Options	Description
Target Location	Select the host where the VM will be recovered. The list contains hosts that have been added to the appliance and are compatible with the VM being recovered. Incompatible hosts do not display in the list. • To add a host, see "Adding a virtual host" on page 308. • For Hyper-V clusters – To create a clustered VM, you must select a cluster as the target Hyper-V host.
Target Resource	(Optional) Select a resource pool. This field displays only if the Target Location is an ESXi host that has resource pools.
Target Storage	Select a datastore (ESXi host) or volume (Hyper-V host).
Target Appliance Network	Select a network adapter on the appliance to use for the recovery. *eth0* is selected by default. If your appliance is configured with multiple adapters, you can opt to select a different adapter from the list. The appliance uses the selected adapter for communication with the hypervisor during the storage migration.

7 Click **Next**.

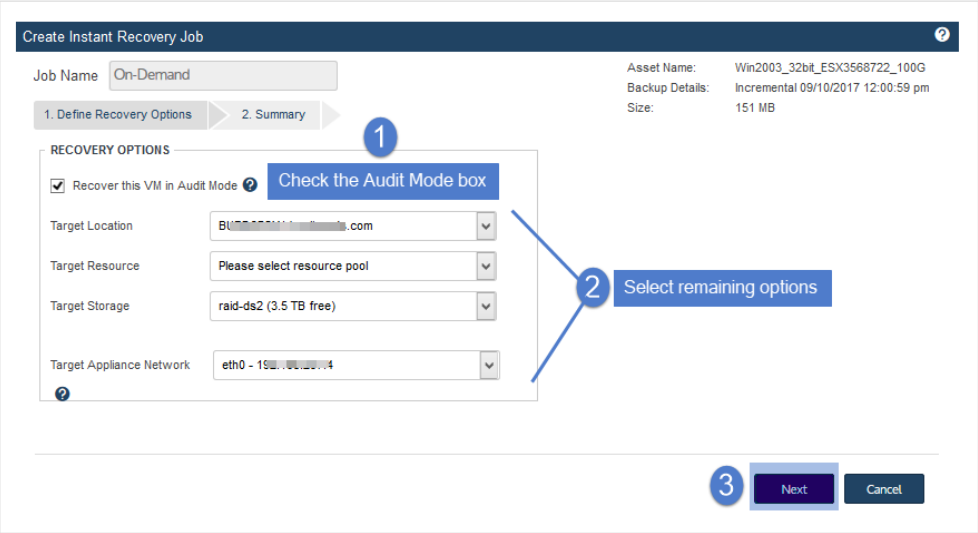

8 Review the recovery settings.

- A recovered VMware VM is created with the following default name: *<original_VM_name>_restore*. To change this name, click the pencil icon.
- A recovered Hyper-V VM is created with the same name as the original VM and no suffix. Due to Hyper-V limitations, it is not possible to rename the VM during the recovery, and the original VM is overwritten by the recovery operation if it resides on the recovery target.

9 Click **Save** to start the recovery.

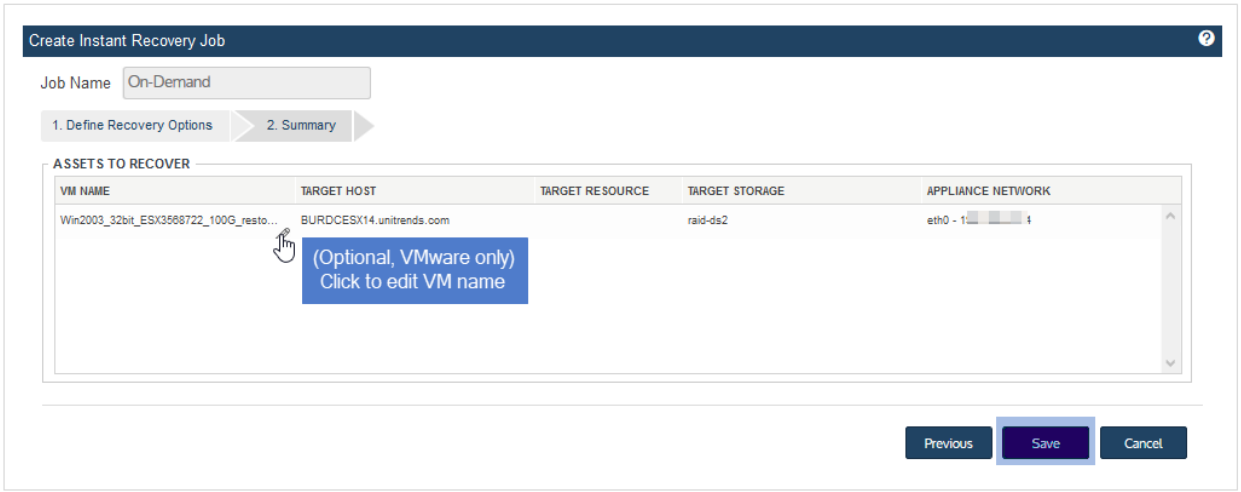

10 Click **OK** to close the Information message.

The recovered VM is created on the target host. The VM has no network settings.

11 Check IR progress by viewing details on the Instant Recovery tab.

Note: If you are recovering an incremental that has a long chain of dependent incrementals, it can take extra time to create the recovered VM.

- Click the **Instant Recovery** tab. Active IR sessions display by VM name.
- Click to select the IR session you created.
- Click **View Details** above.

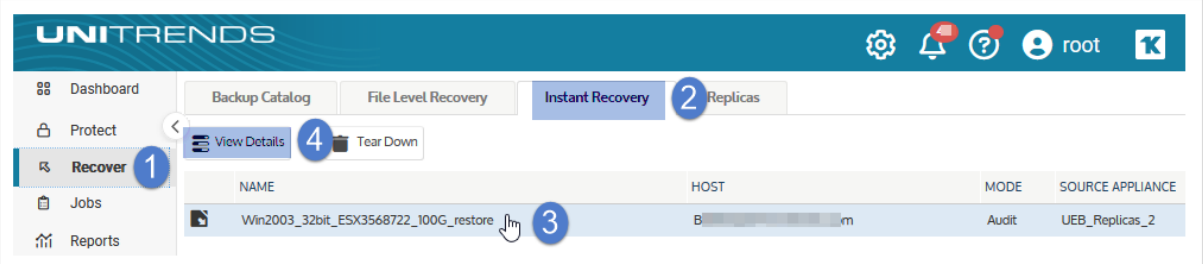

- Status messages display in the Details tab below. IR moves through several phases. The VM can be accessed when you see this message: *VM is available for use*.

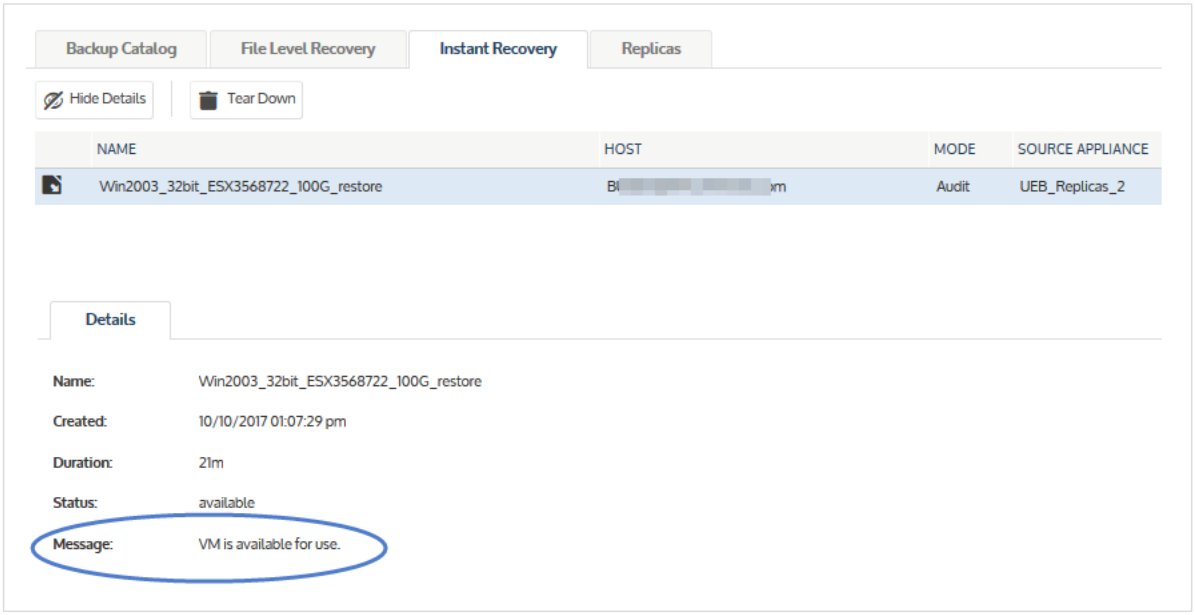

12 Access the VM directly from the target VMware or Hyper-V host. Log in to the VM console and verify that the recovered VM is functioning as expected.

Notes:

- Any features or applications that require network access do not have full functionality in audit mode.

- Any changes made to the recovered VM in audit mode are lost when you tear down the IR session. These changes are not applied to the production VM that you are auditing.
- Windows server VMs - In rare instances, after a restore is performed for a Windows server VM, a disk may be inaccessible because it has been placed in an offline state. To bring disks into an online state, login to the VM, go to Disk Management, right-click on the offline disk, and select **Online** from the drop-down menu.
- Debian VMs - In some instances, Gnome might not start after a Debian VM is recovered. You can resolve this issue by rebooting the VM or restarting Gnome from the console. To access the console, enter *Ctl+Alt+F1* and log in as root. Then run *startx*.

13 Tear down the IR session:

- On the Instant Recovery tab, click to select the IR session.
- Click **Tear Down**.

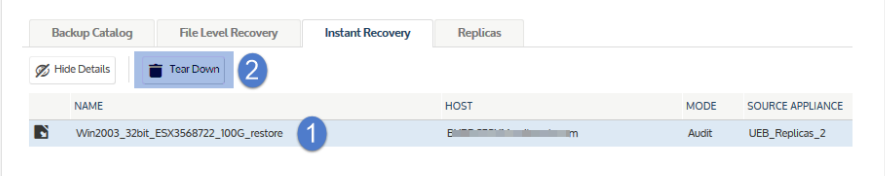

- Click **Confirm**.

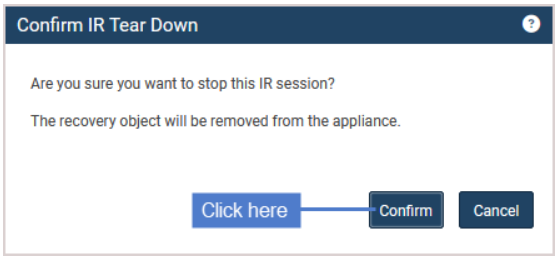

The recovery object is removed from the appliance and the recovered VM is removed from the target host.

Performing instant recovery

Perform instant recovery after a VM fails. If the VM has not failed, you can use audit mode to verify that the VM can be created from a backup. (See "Perform instant recovery in audit mode" on page 913 for details.)

To perform instant recovery

Be sure to shut down the original VM before running this procedure.

Run this procedure from the backup appliance to recover from a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover from a hot backup copy.

1 Log in to the backup appliance or target appliance (if recovering from a hot backup copy).

2 Select **Recover** and click the **Backup Catalog** tab.

(Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

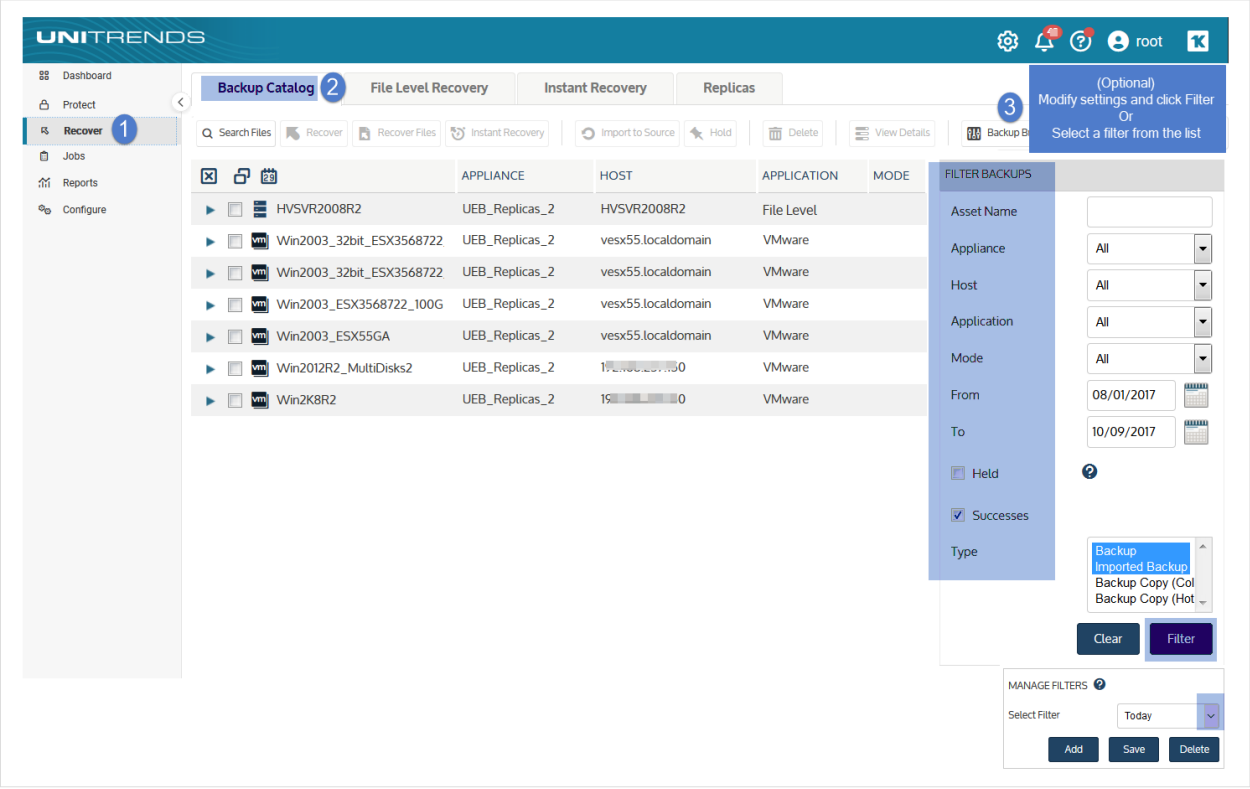

3 Expand the VM asset and select one of the following to use for the recovery:

- A host-level backup.
- An imported host-level backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
- A hot backup copy (supported if performing IR on the target appliance where the hot copy resides).

4 Click the **Instant Recovery** button.

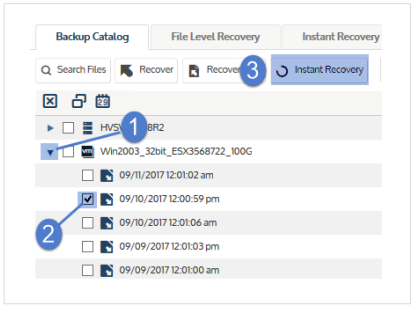

5 Do *not* check the **Recover this VM in Audit Mode** box.

This box must be unchecked to recover the failed VM with its original network settings.

Note: If the VM has not failed, you can do the recovery in audit mode by checking this box. See "Perform instant recovery in audit mode" on page 913 for details.

6 Select these recovery options:

Recovery Options	Description
Target Location	Select the host where the VM will be recovered. The list contains hosts that have been added to the appliance and are compatible with the VM being recovered. Incompatible hosts do not display in the list. • To add a host, see "Adding a virtual host" on page 308. • For Hyper-V clusters – To create a clustered VM, you must select a cluster as the target Hyper-V host.
Target Resource	(Optional) Select a resource pool. This field displays only if the Target Location is an ESXi host that has resource pools.
Target Storage	Select a datastore (ESXi host) or volume (Hyper-V host).
Target Appliance Network	Select a network adapter on the appliance to use for the recovery. *eth0* is selected by default. If your appliance is configured with multiple adapters, you can opt to select a different adapter from the list. The appliance uses the selected adapter for communication with the hypervisor during the storage migration.
Target Network Switch	For Hyper-V hosts only, select the target network switch. • This field displays only when recovering to a Hyper-V host. • For Hyper-V clusters – If recovering a clustered VM, be sure to select a switch that is common to all nodes in the cluster.

7 Click **Next**.

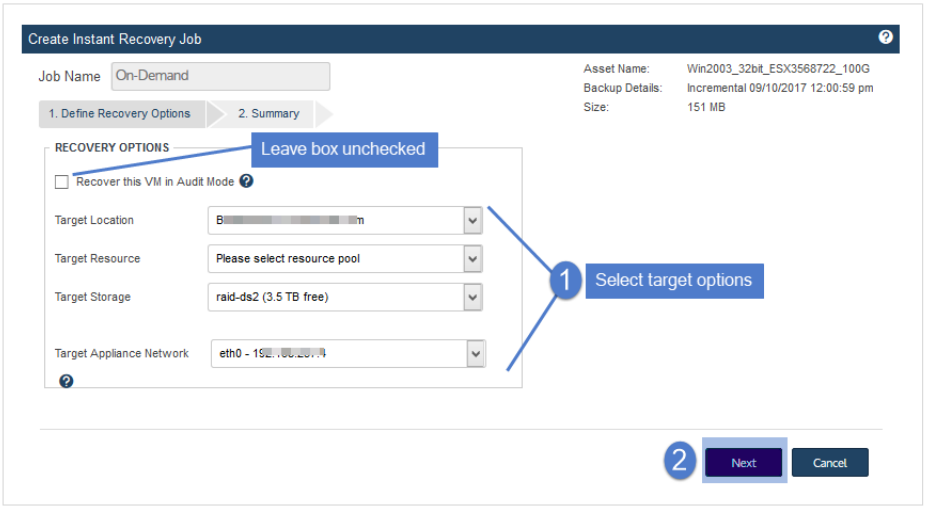

8 Review the recovery settings.

- A recovered VMware VM is created with the following default name: *<original_VM_name>_restore*. To change this name, click the pencil icon.
- A recovered Hyper-V VM is created with the same name as the original VM and no suffix. Due to Hyper-V limitations, it is not possible to rename the VM during the recovery, and the original VM is overwritten by the recovery operation if it resides on the recovery target.

9 Click **Save** to start the recovery.

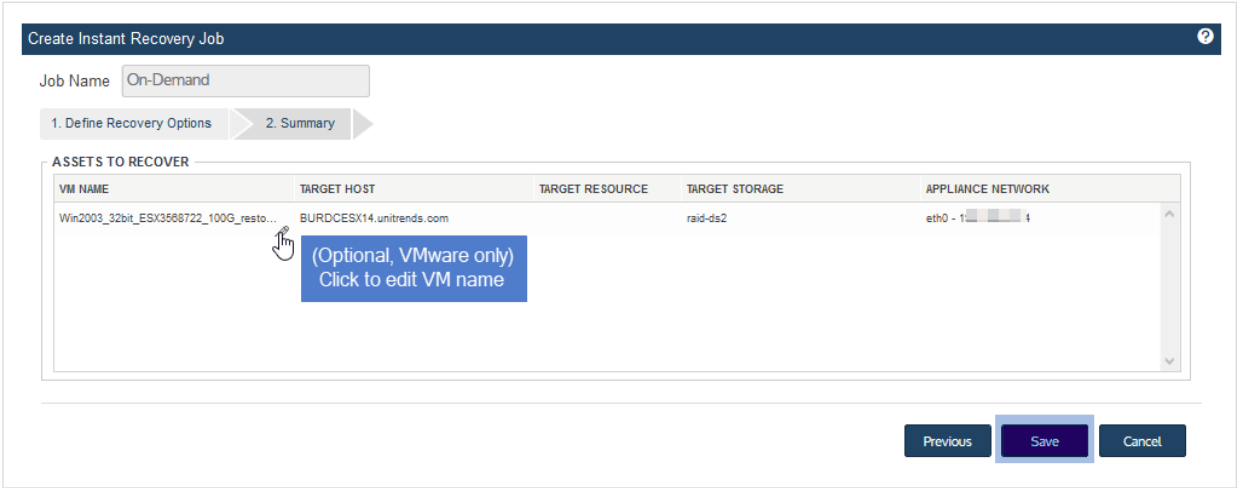

10 Click **OK** to close the Information message.

The recovery object and new VM are created, and data migration begins. During this time, the VM is fully operational.

IMPORTANT! Do not tear down the recovery session or modify the VM in any way until the data migration completes.

11 Check IR progress on the Instant Recovery tab:

- Click the **Instant Recovery** tab. Active IR sessions display by VM name.
- Click to select the IR session you created.
- Click **View Details** above.

- Status messages display in the Details tab below. IR moves through several phases. Data migration is complete when you see this message: *VM is available for use*.

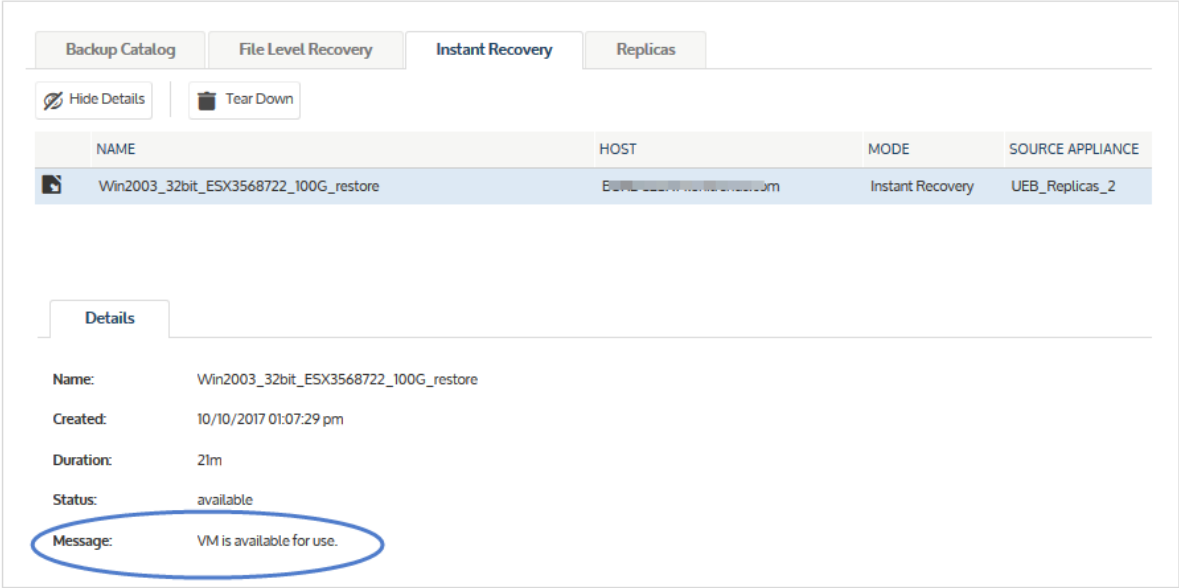

12 When you see the *VM is available for use* message, all data has been migrated and IR is complete.

- The recovered VM has the same network settings and username/password credentials as the original VM. For details on other settings, see "VM configuration" on page 907 (VMware) or "VM configuration" on page 909 (Hyper-V).
- Access the VM and verify that it is functioning as expected in production.

Notes:

- Windows server VMs - In rare instances, after a restore is performed for a Windows server VM, a disk may be inaccessible because it has been placed in an offline state. To bring disks into an online state, login to the VM, go to Disk Management, right-click on the offline disk, and select **Online** from the drop-down menu.
- Debian VMs - In some instances, Gnome might not start after a Debian VM is recovered. You can resolve this issue by rebooting the VM or restarting Gnome from the console. To access the console, enter *Ctl+Alt+F1* and log in as root. Then run *startx*.

13 Tear down the IR session:

- On the Instant Recovery tab, click to select the IR session.
- Click **Tear Down**.

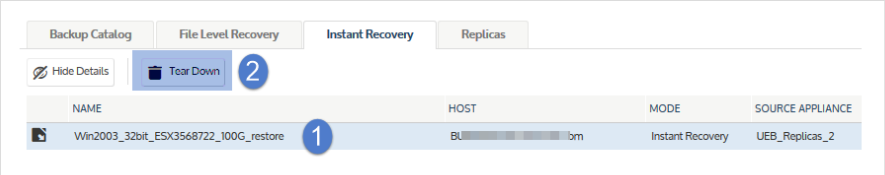

- Click **Confirm**.

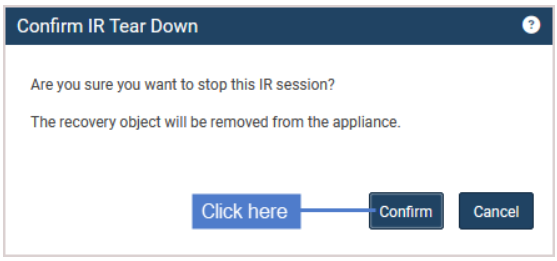

The session is removed from the appliance and no longer displays on the Instant Recovery tab. The IR VM is retained on the hypervisor.

14 Modify VM settings and backup schedules as needed to begin protecting the recovered VM. The next backup of the recovered VM is promoted to a full.

Tearing down the instant recovery session

After performing an instant recovery, tear down the IR session. This removes the recovery object from the appliance, freeing resources that were being used for the recovery. If the recovery was done in audit mode, this also deletes the recovered VM from the virtual host.

To tear down the instant recovery session

Note: If the recovery was done in instant recovery mode, do not perform this procedure until all data has been migrated and the VM is ready for use. Tearing down too early invalidates the recovered VM (and you must perform IR again to create a new one).

1 On the appliance used for the recovery, select **Recover** and click the **Instant Recovery** tab.

Active IR sessions display by VM name.

2 Click to select the applicable instant recovery session.

3 Click **Tear Down**.

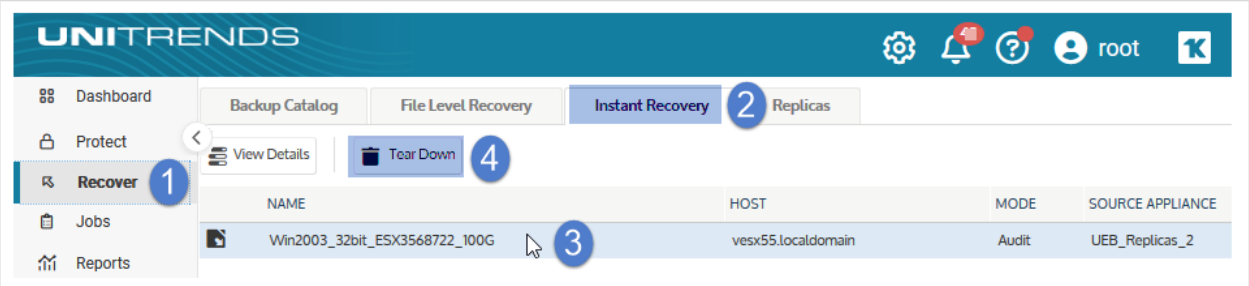

4 Click **Confirm**.

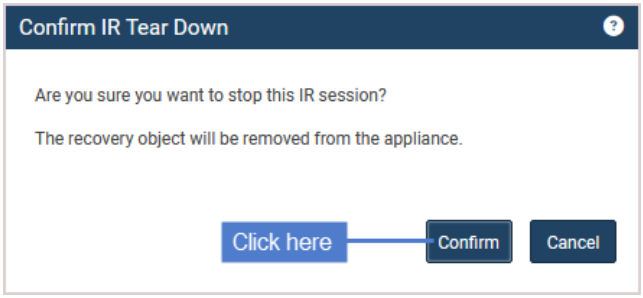

The session is removed from the appliance and no longer displays on the Instant Recovery tab. If the IR VM was running in live mode, it is retained on the hypervisor. If the IR VM was running in audit mode, it is removed from the hypervisor.

This page is intentionally left blank.

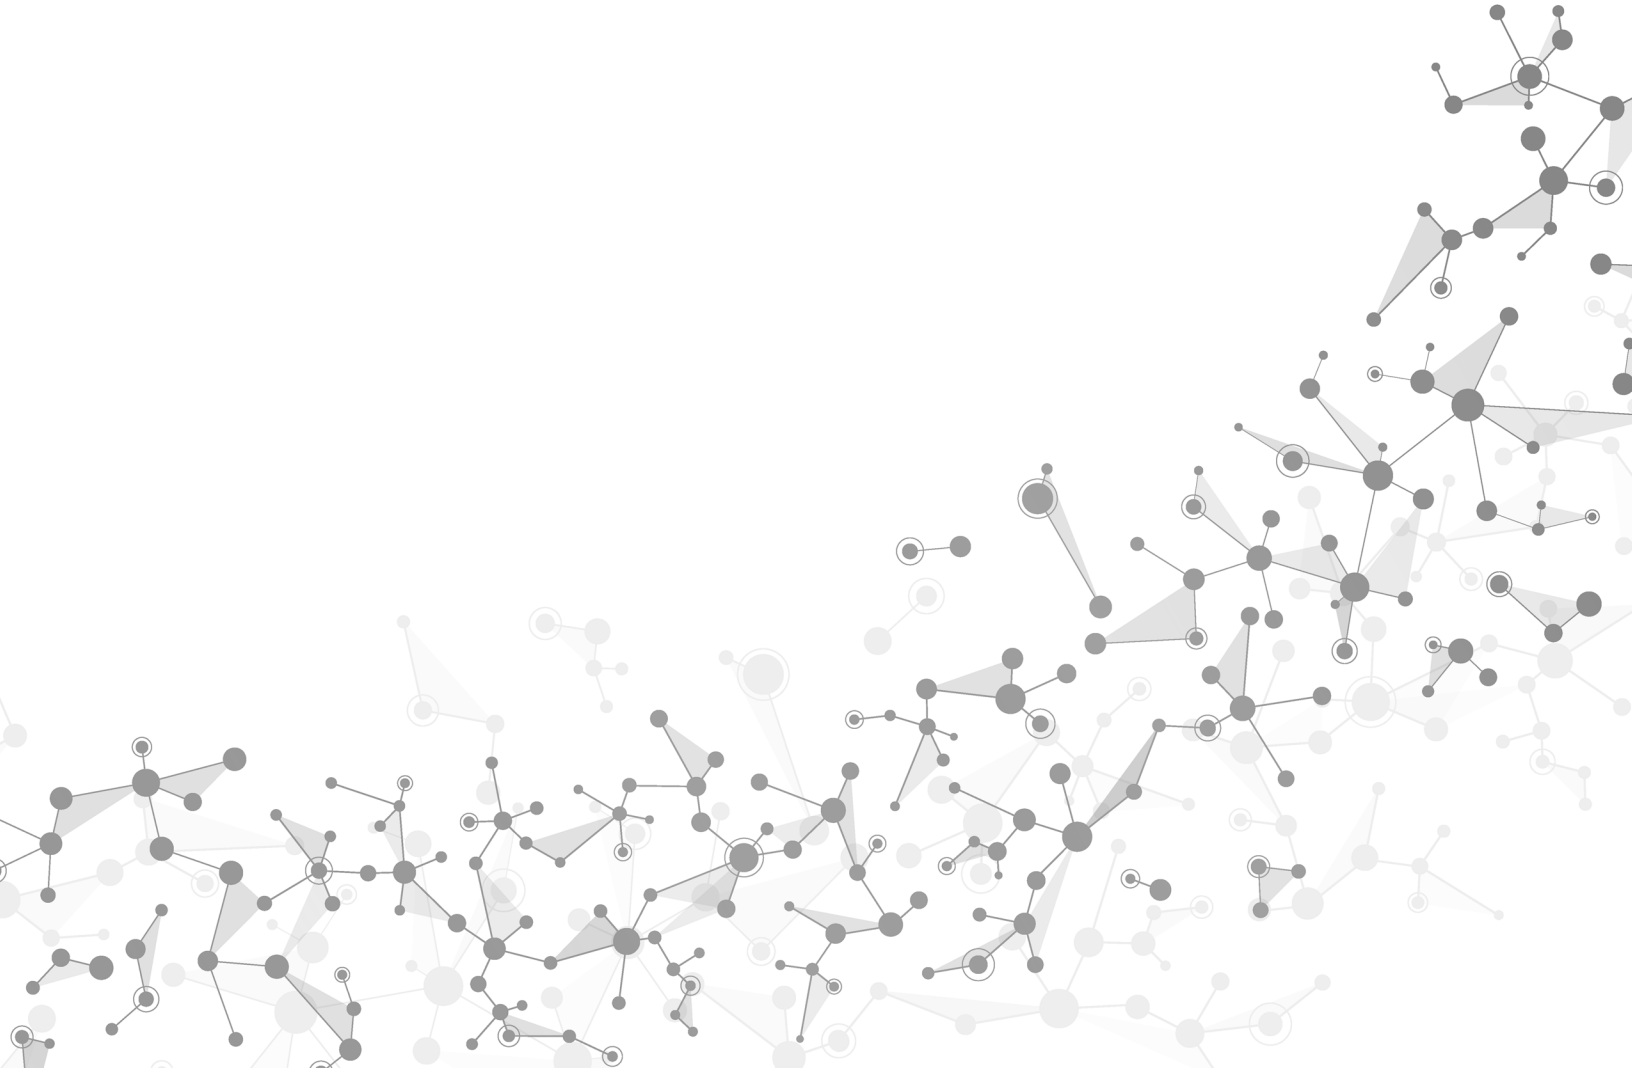

Chapter 15: Recovering File-level Backups

This chapter describes recovery procedures for file-level backups (backups that were run using a Unitrends agent). Assets must have an eligible backup or backup copy before running these procedures.

Notes:

- The procedures in this chapter are used to recover file-level backups for all operating systems. In some cases, additional limitations and steps apply. For example, recovering Windows Active Directory requires additional steps. For additional considerations, see "File-level Backups Overview" on page 703.
- The Windows agent is used for file-level and image-level protection. If you are using image-level protection, see the recovery procedures in "Recovering Windows Image-level Backups" on page 1031.

Each backup creates a recovery point of the entire asset. It contains the data captured during the backup job, plus all dependent data in the backup group. You can recover entire backups or pick a subset of files. Recovering an entire backup recovers all dependent data in the group. For example, recovering an incremental also recovers its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

You can recover from backups or imported backup copies that reside on a Unitrends backup appliance, from hot backup copies that reside on a Unitrends backup copy target appliance, or from hot backup copies that reside in the Unitrends Cloud.

To perform the recovery, you will start by selecting a backup or backup copy. For backups, you can do this in the Backup Catalog or in the Backup Browser. For backup copies and imported backups, you must use the Backup Catalog.

In the Backup Catalog, backups and backup copies display under the protected asset. You can modify the display by entering filter criteria. Expand an asset to view its backups and backup copies:

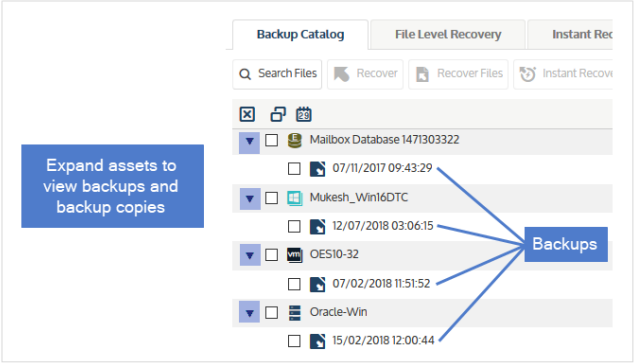

The Backup Browser provides advanced search and filter options. Backups are not grouped under the protected asset. Search for backups by selecting an appliance and date range. Filter the display by entering text in the column fields. Click an arrow to sort by column:

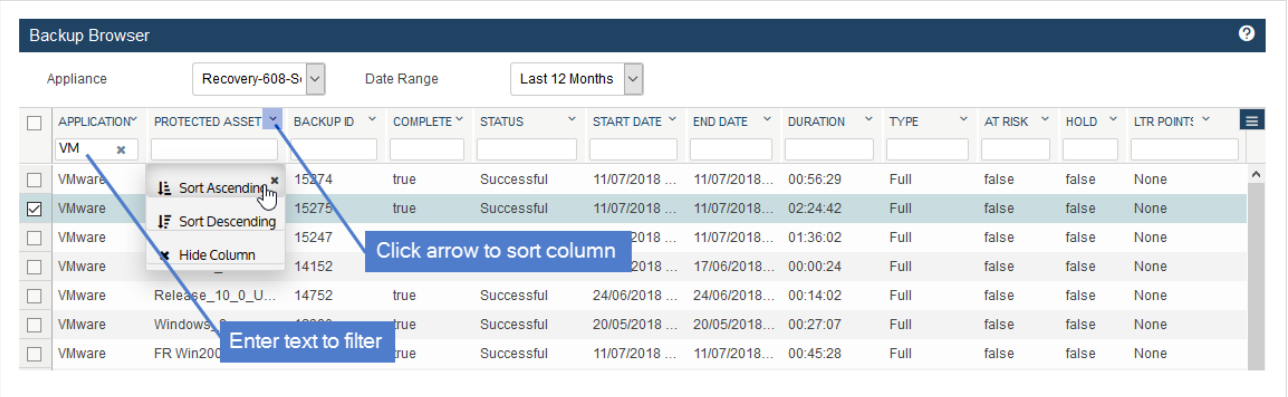

For more on working with these features, see "Backup Catalog tab" on page 60.

Supported procedures vary by whether you are running them from the source or target appliance. See the following topics for details:

- "Recover from backups or imported backup copies" on page 926
- "Recover files from cold backup copies" on page 957
- "Recover from hot backup copies by running procedures on the target appliance" on page 974
- "Recover from hot backup copies by running procedures on the source appliance" on page 985
- "Recover Windows Active Directory information" on page 992
- "Recover a Windows cluster database" on page 993

Recover from backups or imported backup copies

Before you start, be sure the following prerequisites have been met:

- Recovery target is available – You must recover to an agent-based asset that has been added to the backup appliance. (The asset must have a Unitrends agent installed and must display on the appliance's Protected Assets tab.) If necessary, add the asset as described in "To add an agent-based asset" on page 289.
- Backup copy has been imported – To recover from a backup copy, you must first import the backup copy before running these recovery procedures. For details, see "To import a hot backup copy" on page 780 or "To import a cold backup copy" on page 786.

  Note: For hot backup copies (copies that reside in the Unitrends Cloud or on another Unitrends appliance), you can recover files directly (without first importing the backup copy). For details, see "Recover from hot backup copies by running procedures on the source appliance" on page 985.

Run these procedures from the backup appliance to recover from backups or imported backup copies:

- "To recover from a file-level backup by using Search Files" on page 927

- "To browse one file-level backup and recover files by using the Backup Catalog" on page 933
- "To browse one file-level backup and recover files by using the Backup Browser" on page 939
- "To recover an entire file-level backup by using the Backup Catalog" on page 945
- "To recover an entire file-level backup by using the Backup Browser" on page 951

To recover from a file-level backup by using Search Files

Use this procedure to search an asset's backups and imported backup copies for files that meet specified criteria and recover selected files from the search results.

1 Log in to the backup appliance.

2 Click **Recover > Backup Catalog > Search Files**.

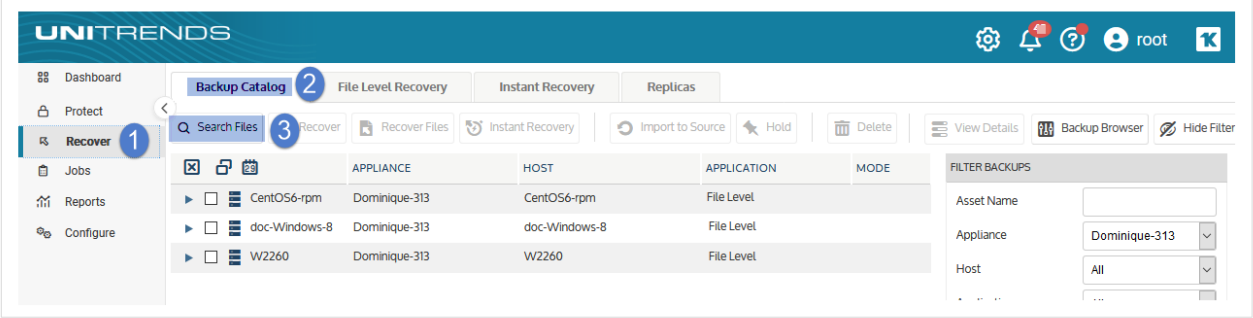

3 In the **Type** list, select **Backup**.

4 Select the **Asset** whose backups and imported backup copies will be searched.

5 Enter one or more search options:

Search Options	Description
String	Enter text to search. The wildcard character * is supported.
Match Case	Select to match the letter case of the entered string.
From/To	Use to search for files that were last modified within the specified date range. Results do not include files modified on the From or To date.
Size	Use to search for files that meet this size criteria.
Advanced	Click to search using a regular expression.

6 Click **Search**.

All backups of this asset stored on the appliance are searched for matching files.

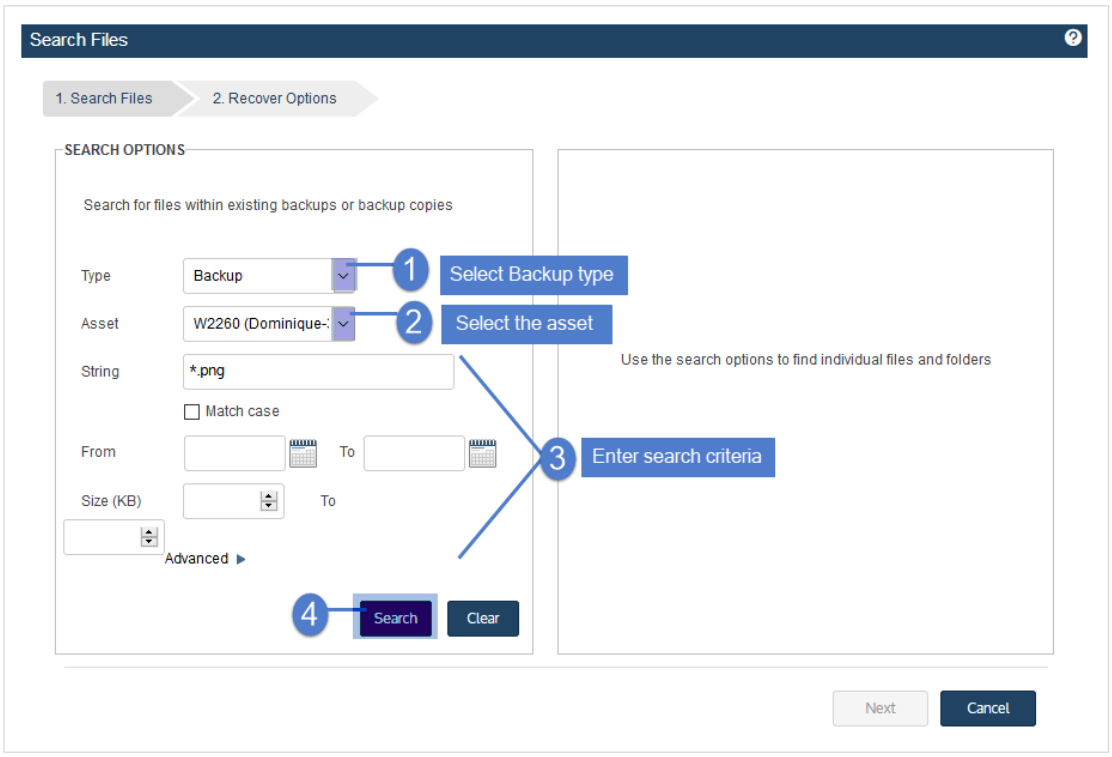

7 In the results list, click to select files to recover.

Notes:

- All files you select must be from a single backup. Check the Backup ID to determine a file's backup. If you select files from multiple backups, the Save button becomes disabled.
- Softlinks cannot be downloaded and are not included in the search results.

8 Click **Next**.

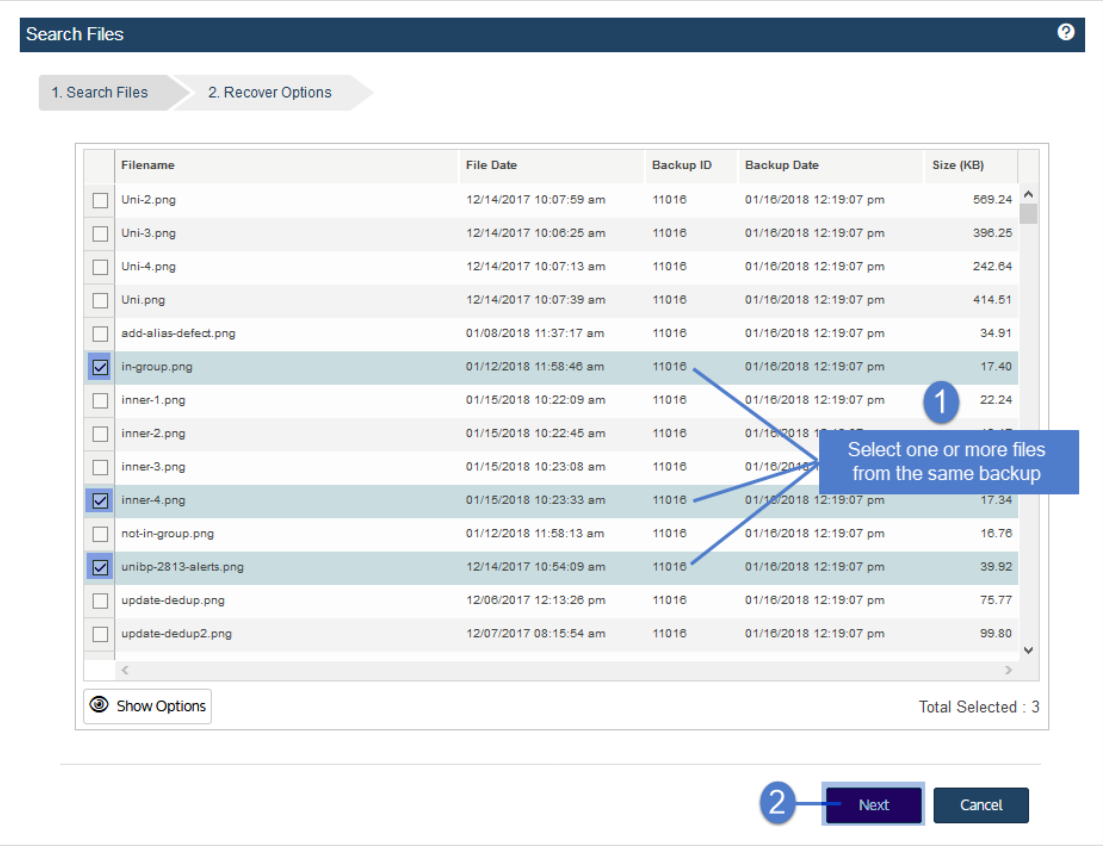

9 Select the **Asset** where the files will be recovered.

10 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

11 (Optional) Specify Exclusions.

12 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Commands to run post-restore	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.

Option	Description
Preserve directory structure	Check this box to preserve the existing file structure within the target directory. **Note:** To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

Overwrite existing files and *Restore newer files only* options

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file. • If the file to recover is older than the one in the Target Directory, does not recover the file.
Overwrite	Recovers the file and	Recovers the file and overwrites the existing file.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
existing files = Yes Restore newer files only = No	overwrites the existing file.	
Overwrite existing files = No Restore newer files only = Yes	Does not recover the file.	Does not recover the file.

13 Click **Save**.

14 Click **OK** to close the Notice message.

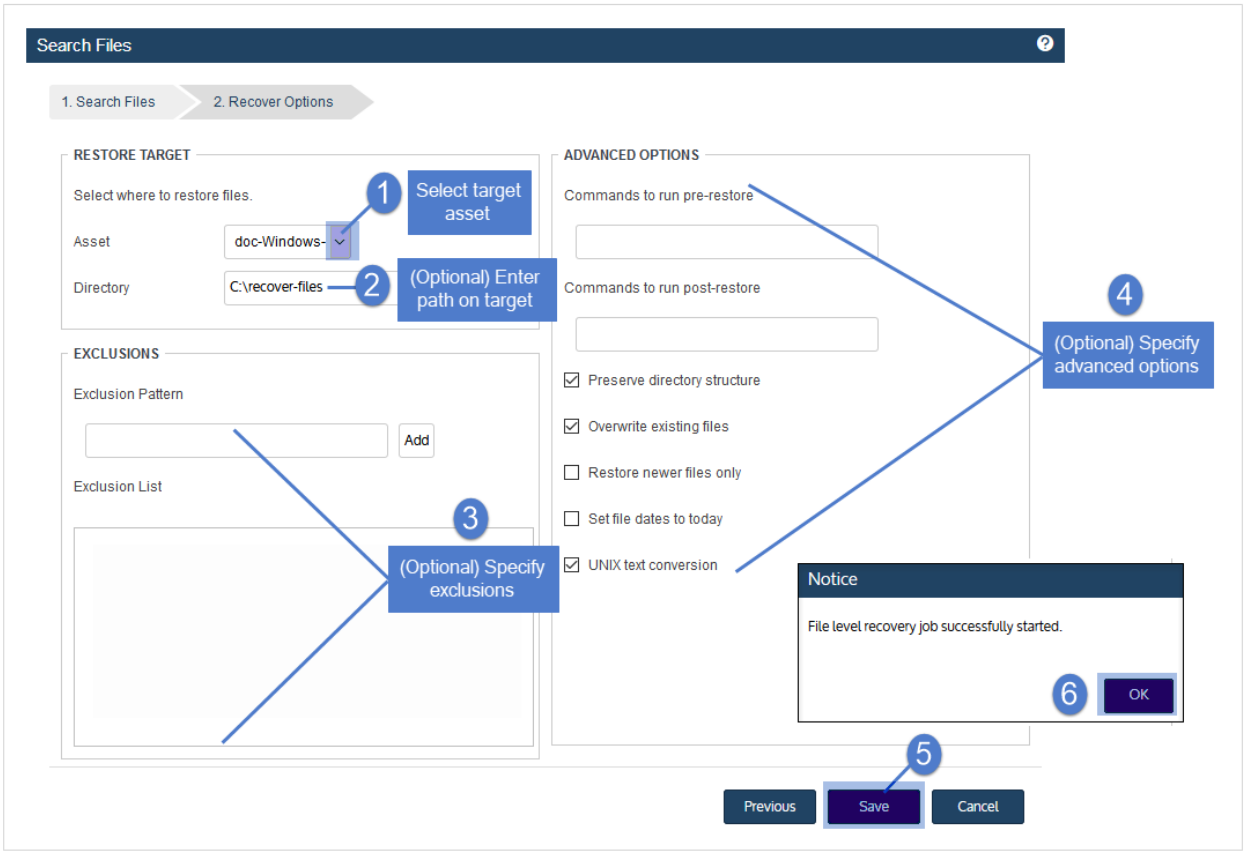

15 To monitor the recovery job:

- Select **Jobs > Active** Jobs.
- Select the job in the list and click **View Details**.

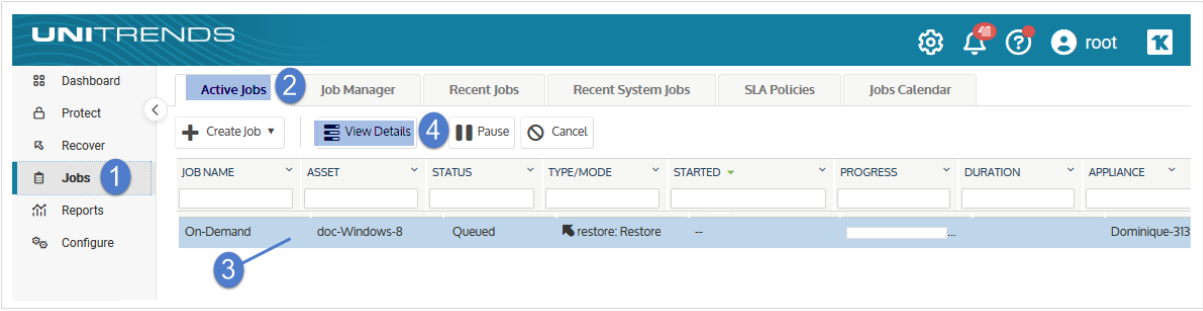

- The recovery is complete when the job's status changes to Success.

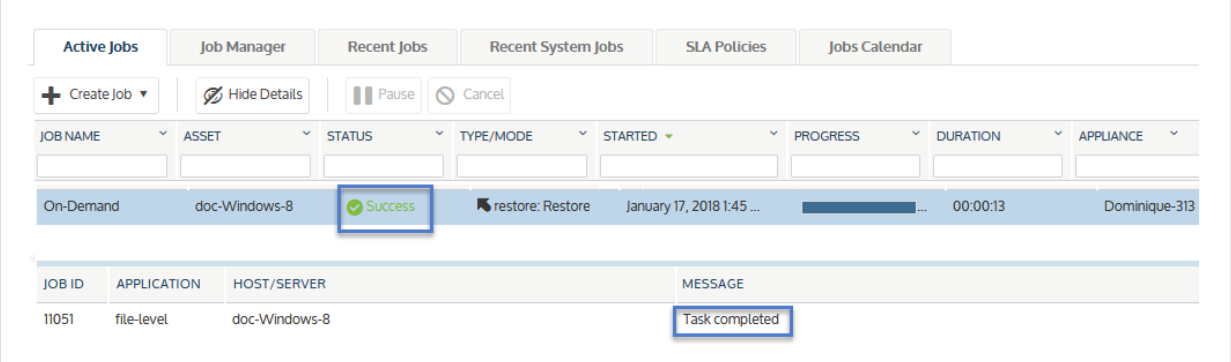

16 Access the recovered files on the recovery target.

Note: Windows assets – Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

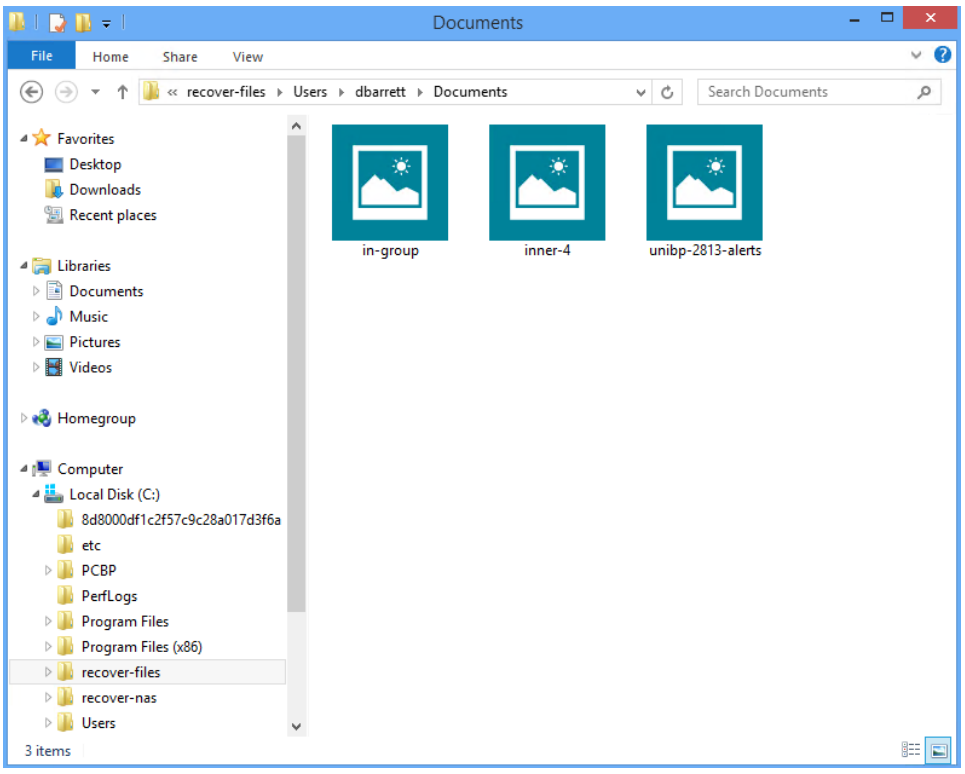

To browse one file-level backup and recover files by using the Backup Catalog

Use this procedure to browse the contents of one backup or imported backup copy and recover selected files.

Note: The file browser contains the backup you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog**.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the asset and select a backup or imported backup copy.

   (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
4. Click **Recover Files**.

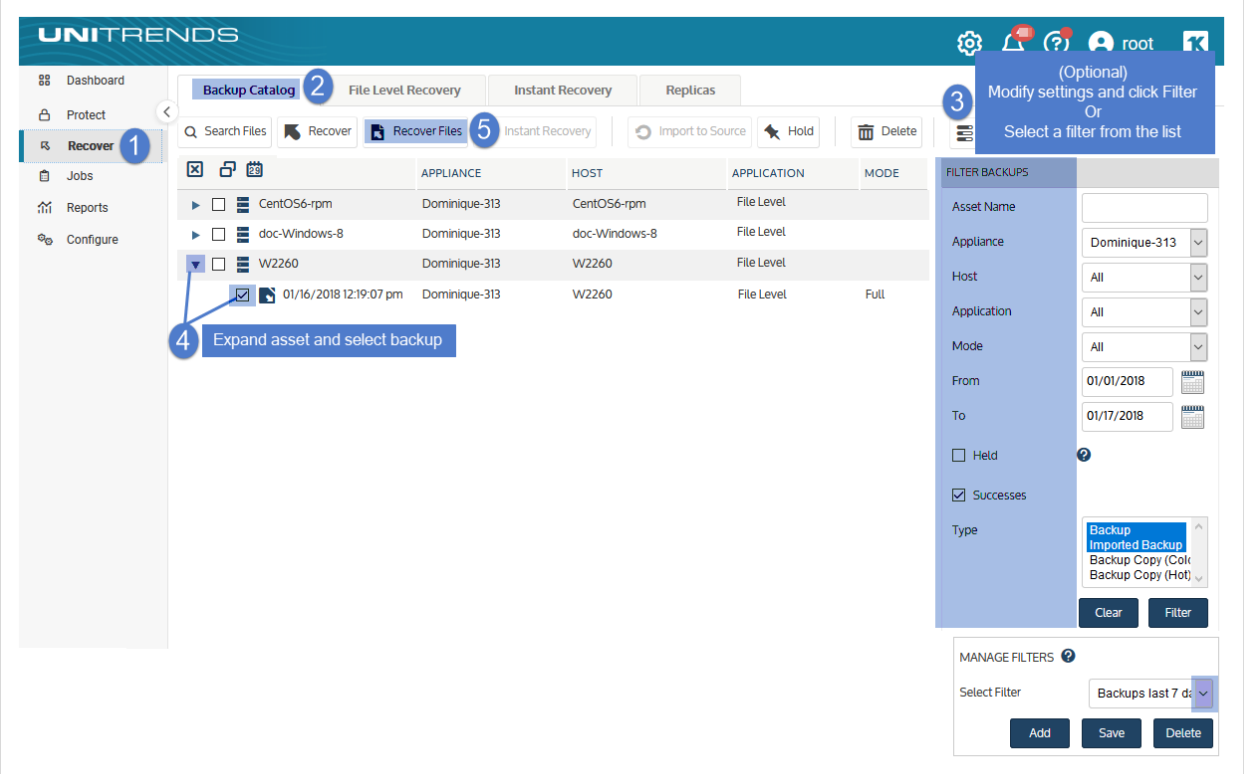

5. In the File Browser, expand folders to view items in the backup.
6. Select or drag files and/or folders to recover.
7. Click **Next**.

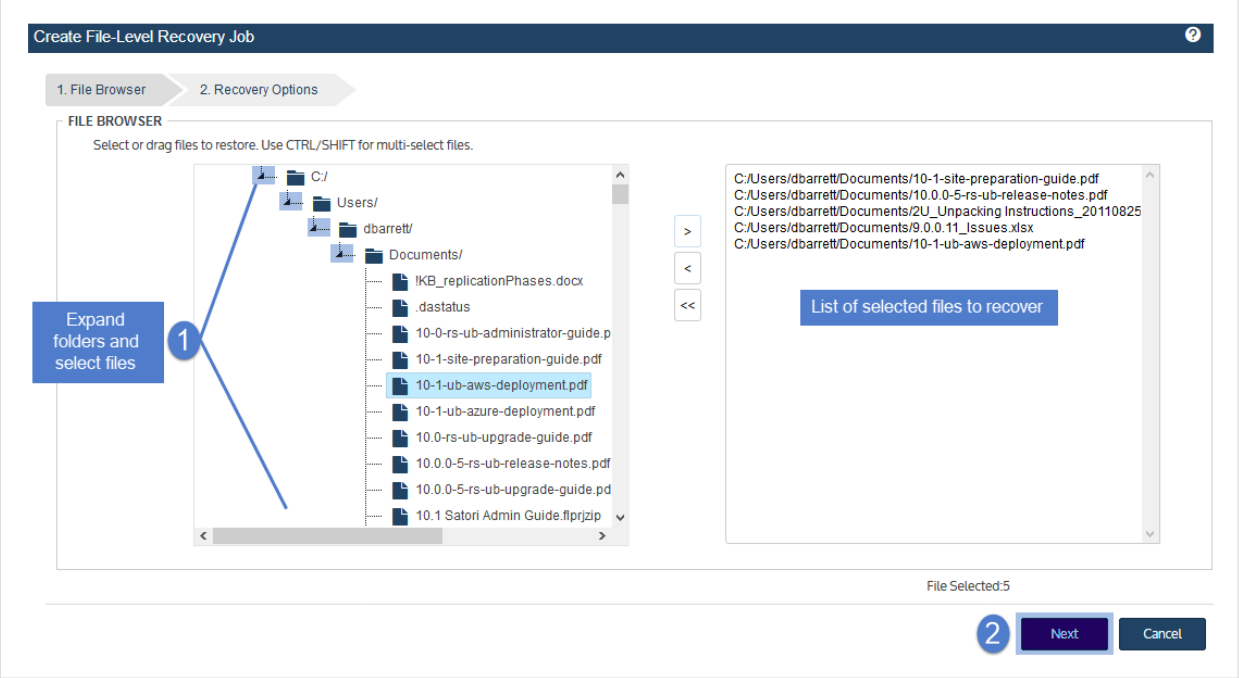

8 Select an **Asset**. Choose from the agent-based assets that have been added to this appliance.

9 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

10 (Optional) Specify Exclusions.

11 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Commands to run post-restore	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Preserve directory structure	Check this box to preserve the existing file structure within the target directory.

Option	Description
	Note: To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

Overwrite existing files and *Restore newer files only* options

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	**Non-Windows backup**
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file. • If the file to recover is older than the one in the Target Directory, does not recover the file.
Overwrite existing files = Yes	Recovers the file and overwrites the existing file.	Recovers the file and overwrites the existing file.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
Restore newer files only = No		
Overwrite existing files = No Restore newer files only = Yes	Does not recover the file.	Does not recover the file.

12 Click **Save**.

13 Click **OK** to close the Notice message.

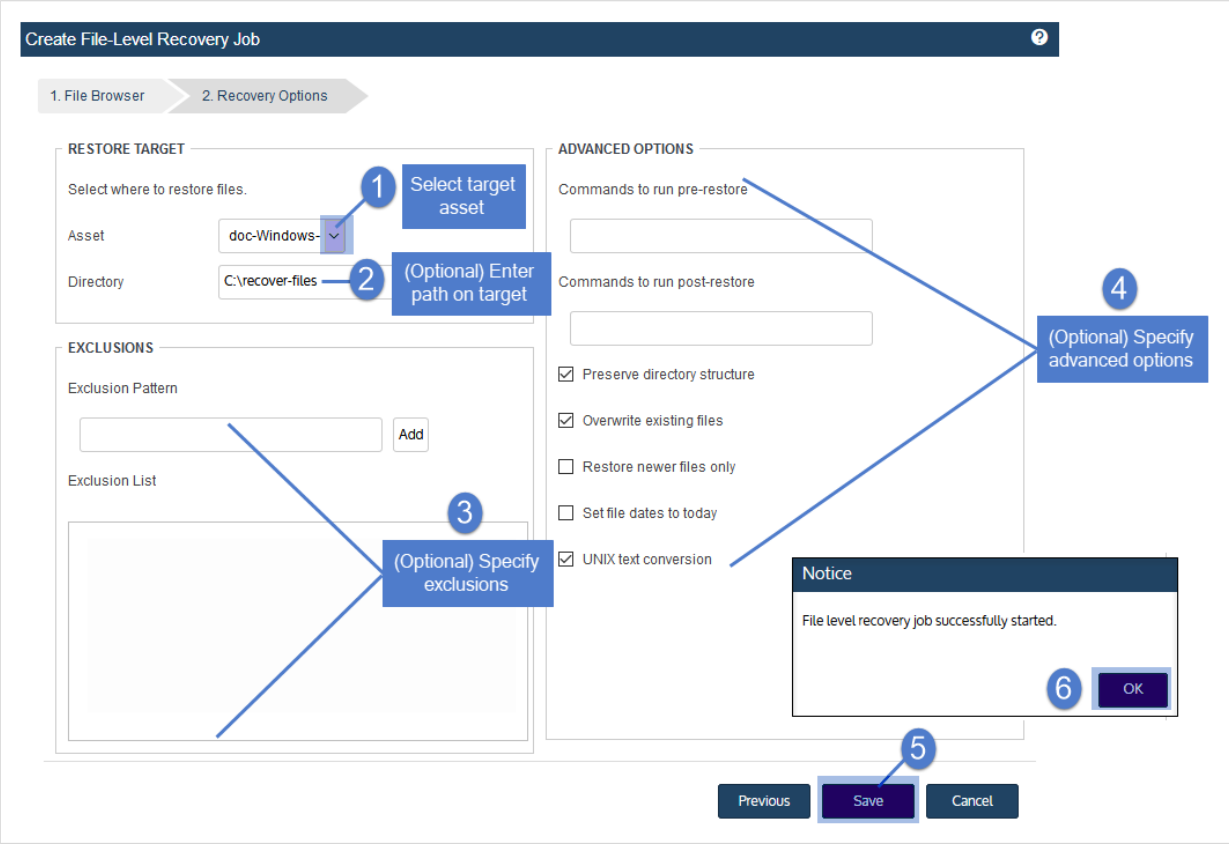

14 To monitor the recovery job:

- Select **Jobs** > **Active** Jobs.
- Select the job in the list and click **View Details**.

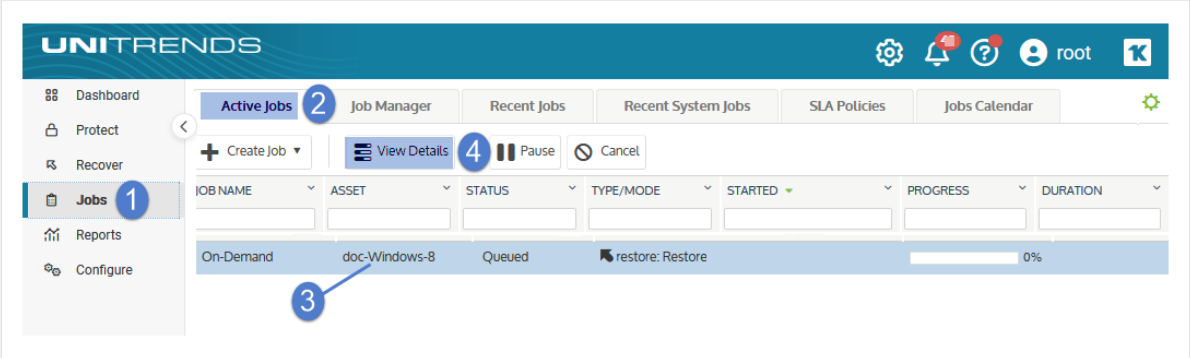

- The recovery is complete when the job's status changes to Success.

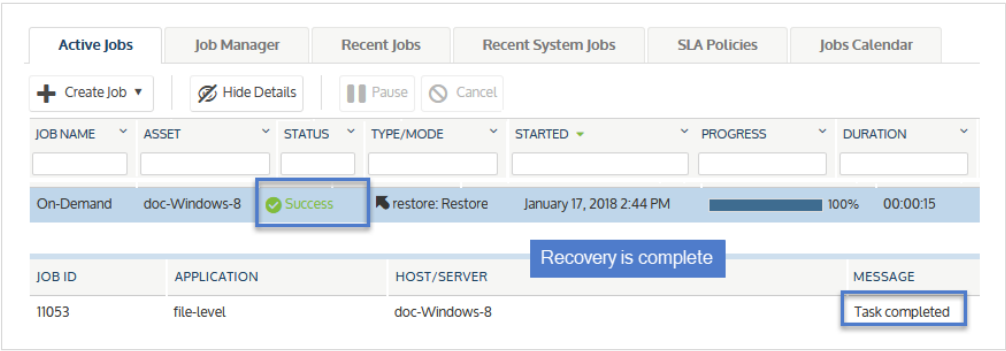

15 Access the recovered files on the recovery target.

Note: Windows assets – Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

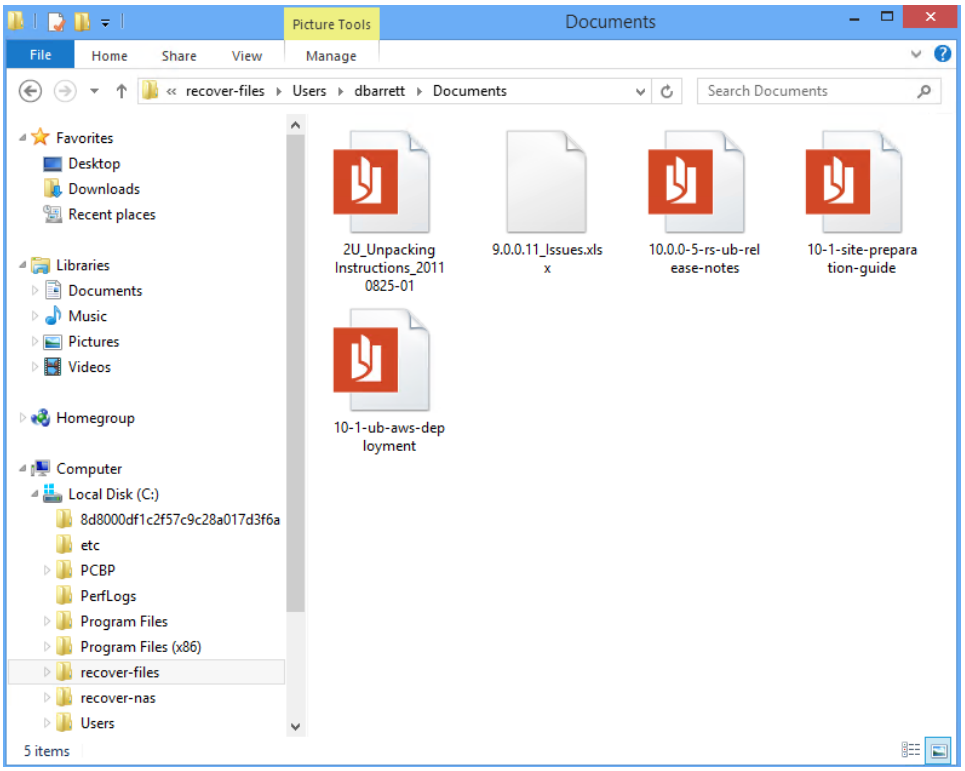

To browse one file-level backup and recover files by using the Backup Browser

Use this procedure to browse the contents of one backup and recover selected files.

Notes:

- You must use the Backup Catalog to recover from an imported backup. (See the procedure above.)
- The file browser contains the backup you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

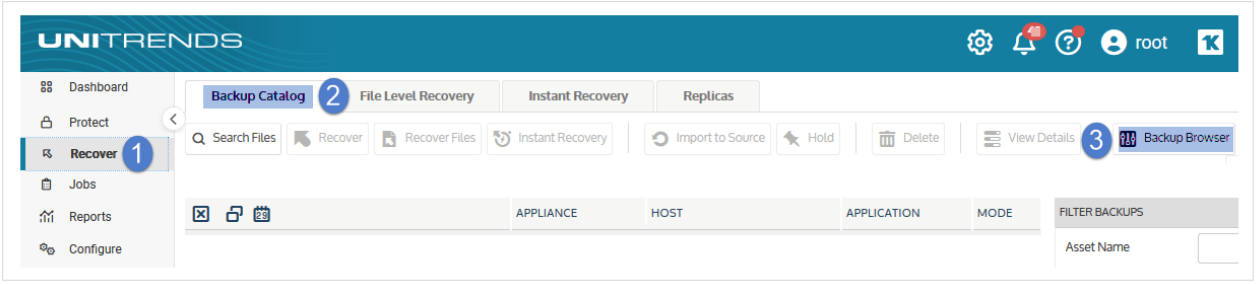

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display:

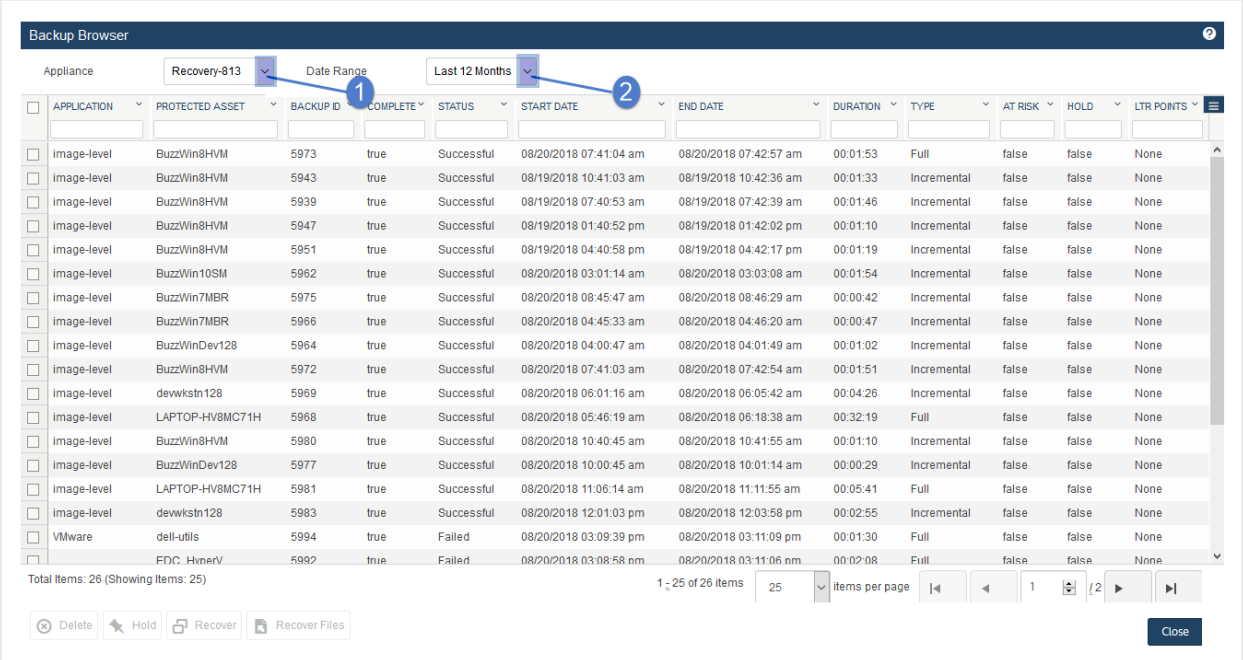

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Select the file-level backup and click **Recover Files**.

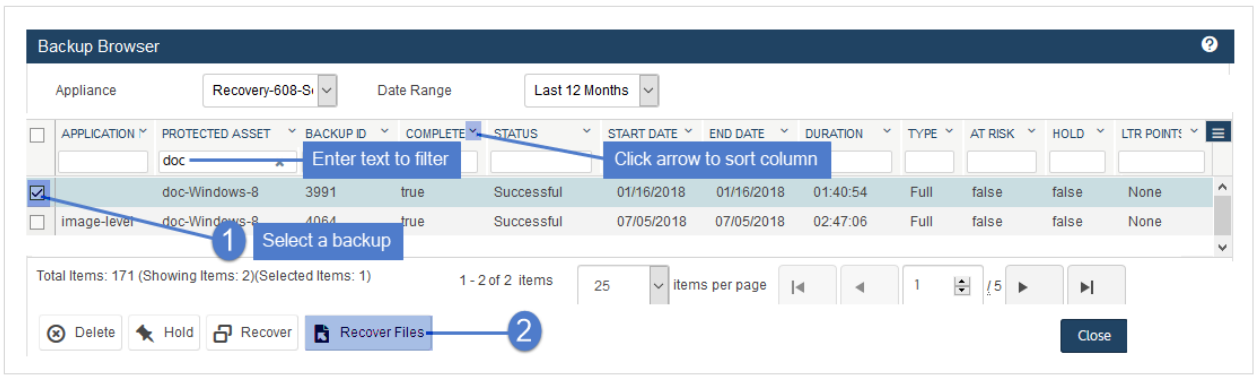

6 In the File Browser, expand folders to view items in the backup.

7 Select or drag files and/or folders to recover.

8 Click **Next**.

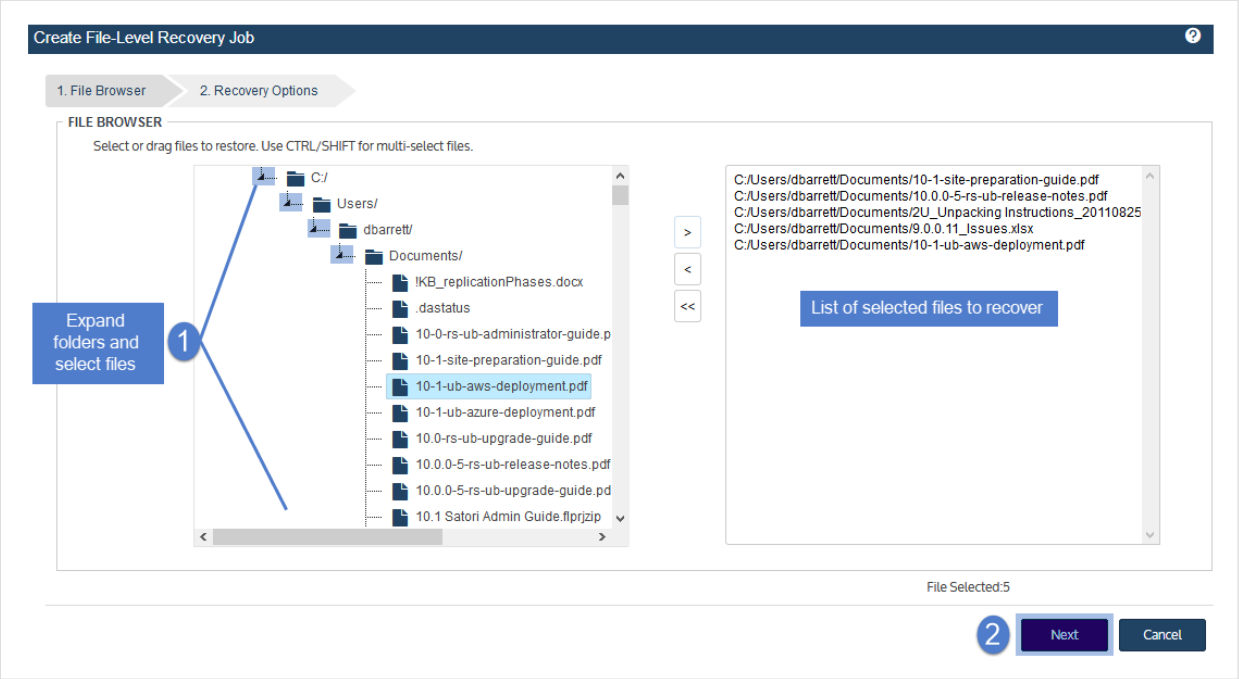

9 Select an **Asset**. Choose from the agent-based assets that have been added to this appliance.

10 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

11 (Optional) Specify Exclusions.

12 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Commands to run post-restore	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.

Option	Description
Preserve directory structure	Check this box to preserve the existing file structure within the target directory. **Note:** To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

Overwrite existing files and *Restore newer files only* options

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	**Non-Windows backup**
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file. • If the file to recover is older than the one in the Target Directory, does not recover the file.
Overwrite	Recovers the file and	Recovers the file and overwrites the existing file.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
existing files = Yes Restore newer files only = No	overwrites the existing file.	
Overwrite existing files = No Restore newer files only = Yes	Does not recover the file.	Does not recover the file.

13 Click **Save**.

14 Click **OK** to close the Notice message.

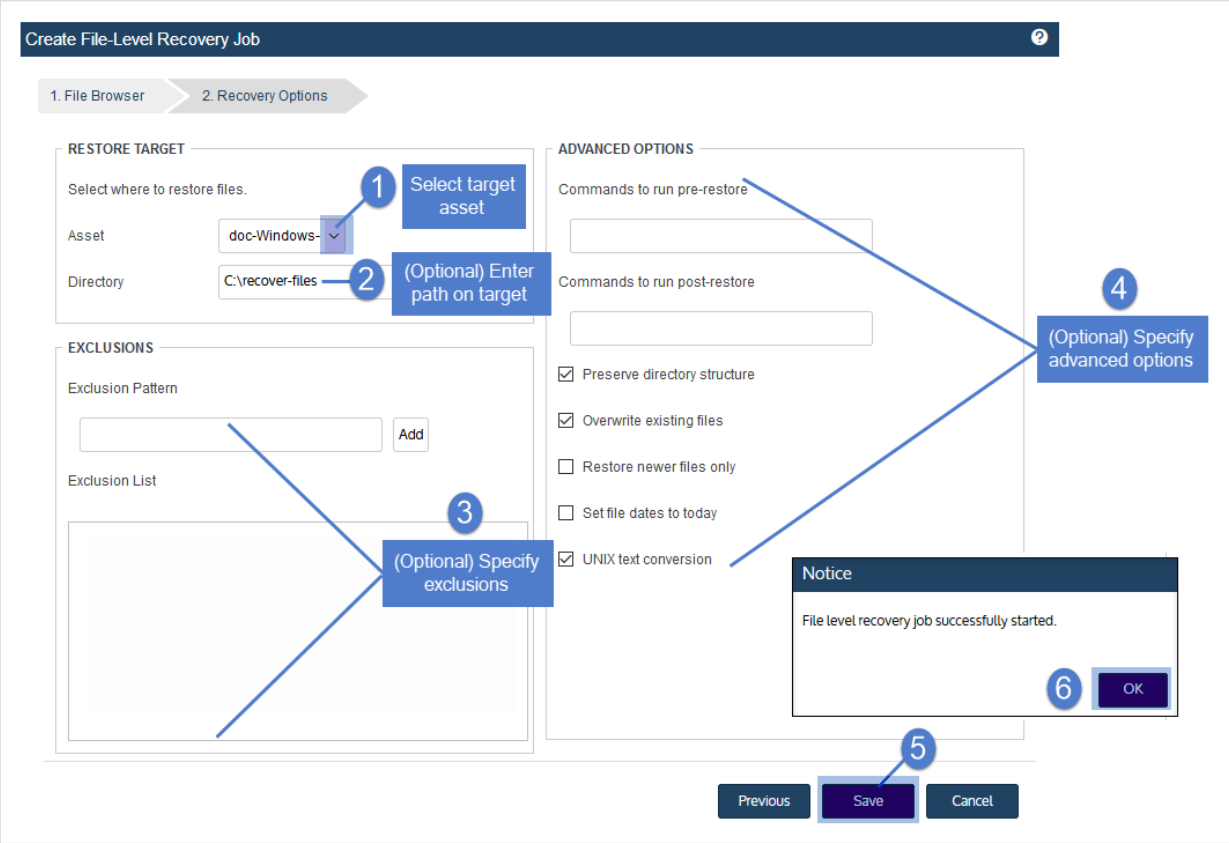

15 To monitor the recovery job:

- Select **Jobs** > **Active** Jobs.
- Select the job in the list and click **View Details**.

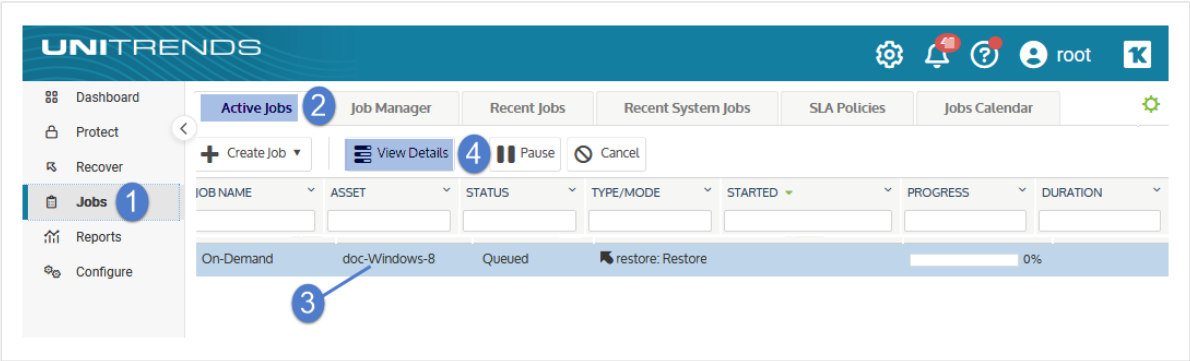

- The recovery is complete when the job's status changes to Success.

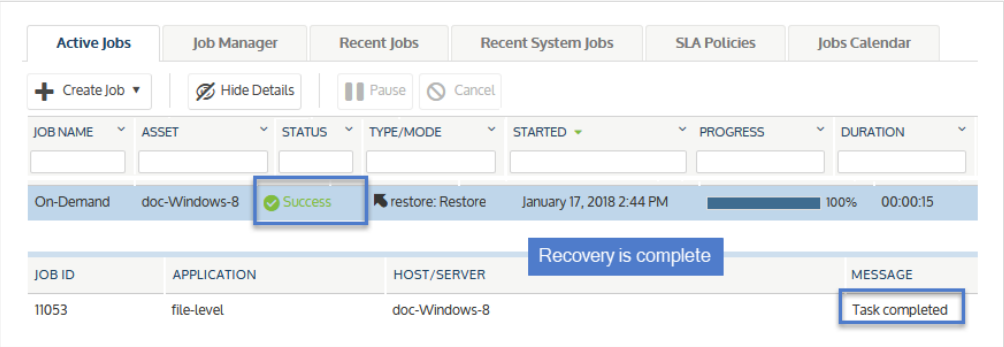

16 Access the recovered files on the recovery target.

Note: Windows assets – Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

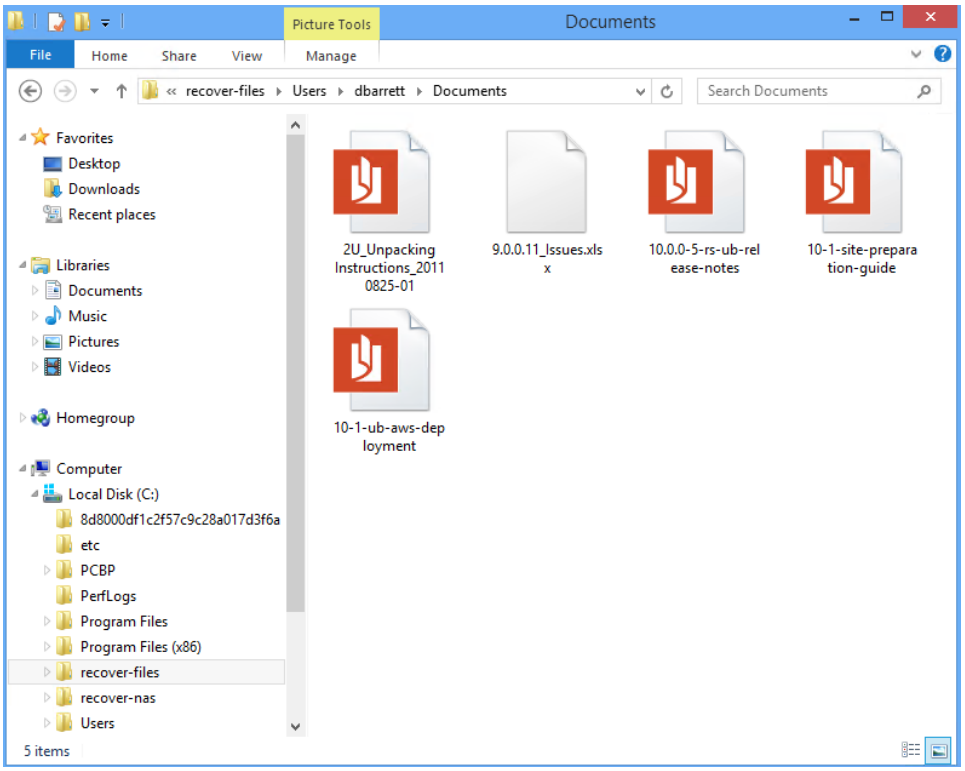

To recover an entire file-level backup by using the Backup Catalog

Use this procedure to recover a backup or imported backup copy.

Note: This procedure recovers the backup you select, plus all dependent data in the backup group. For example, recovering an incremental also recovers its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog**.

(Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

3 Expand the asset and select a backup or imported backup copy.

(To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)

4 Click **Recover**.

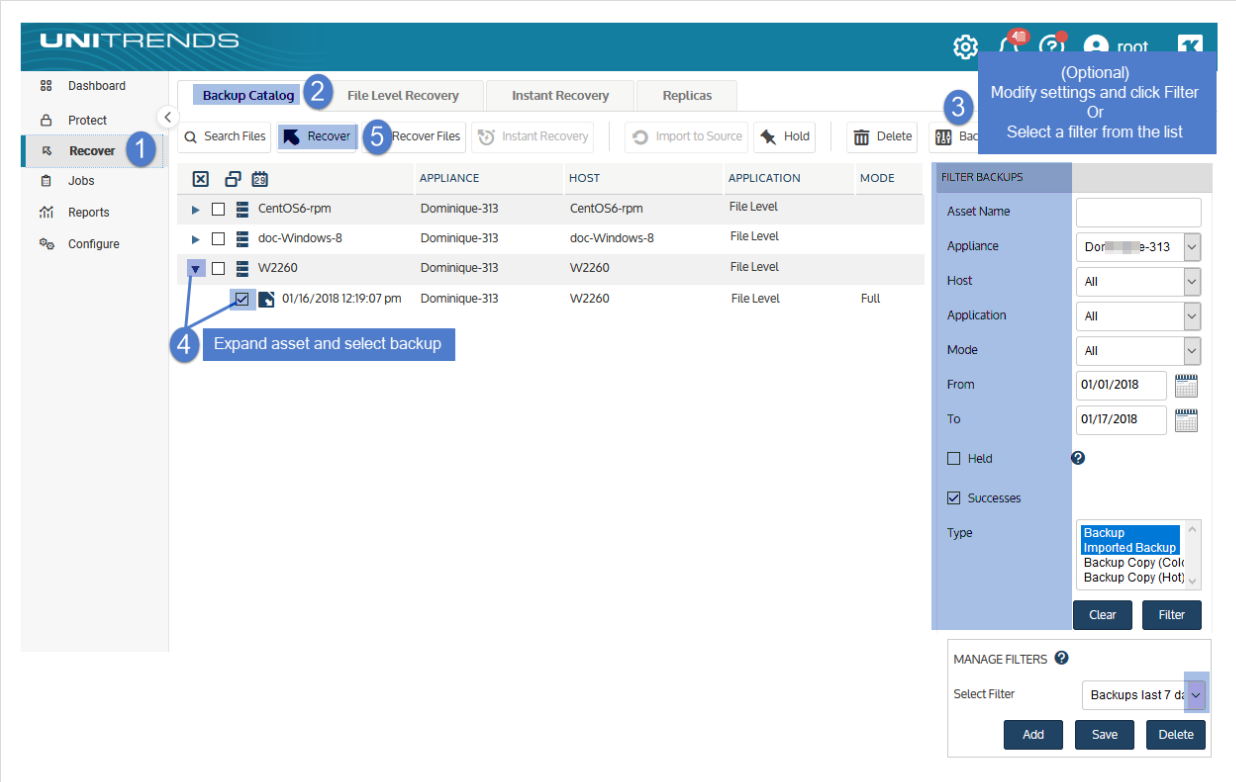

5 Select an **Asset**. Choose from the agent-based assets that have been added to this appliance.

6 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

7 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Commands to run post-restore	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Preserve directory structure	Check this box to preserve the existing file structure within the target directory. **Note:** To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

***Overwrite existing files* and *Restore newer files only* options**

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file. • If the file to recover is older than the one in the Target Directory, does not recover the file.
Overwrite existing files = Yes Restore newer files only = No	Recovers the file and overwrites the existing file.	Recovers the file and overwrites the existing file.
Overwrite existing files = No Restore newer files only = Yes	Does not recover the file.	Does not recover the file.

8 Click **Next**.

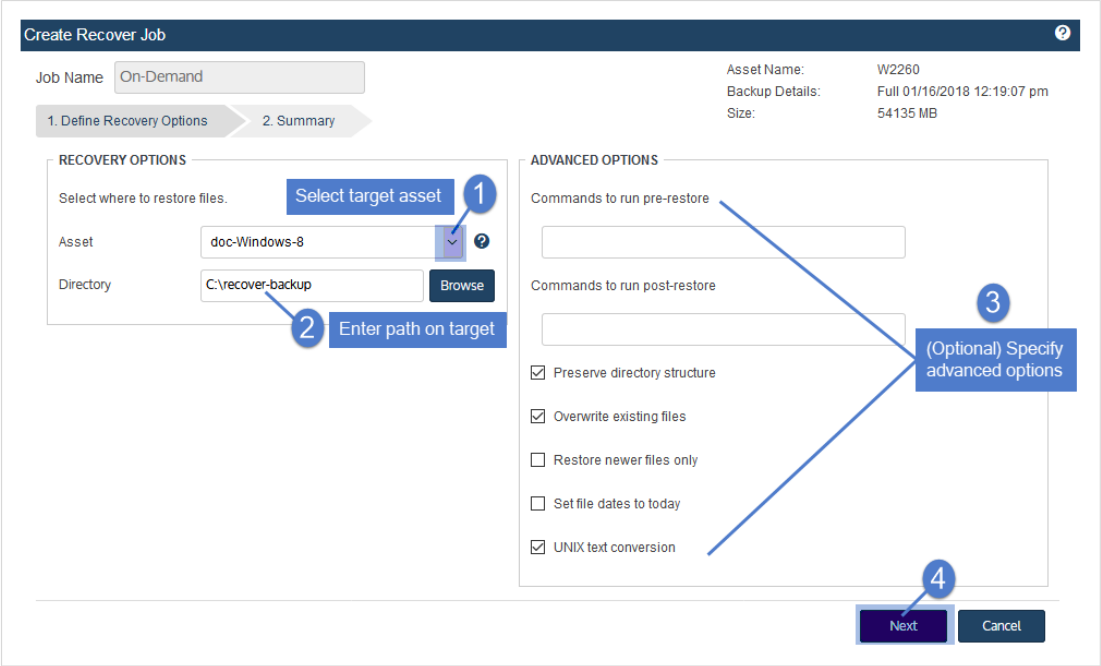

9 Review settings and click **Save**.

10 Click **OK** to close the Information message.

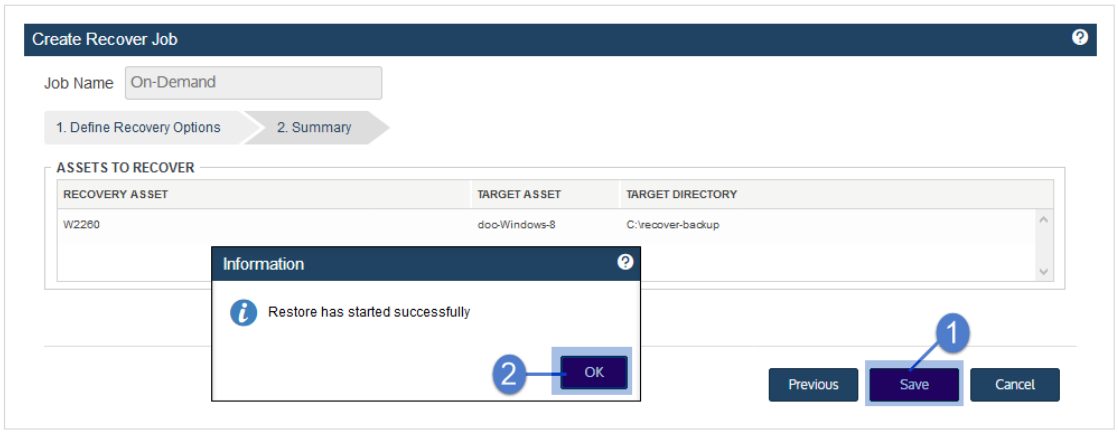

11 To monitor the recovery job:

- Select **Jobs > Active** Jobs.
- Select the job in the list and click **View Details**.

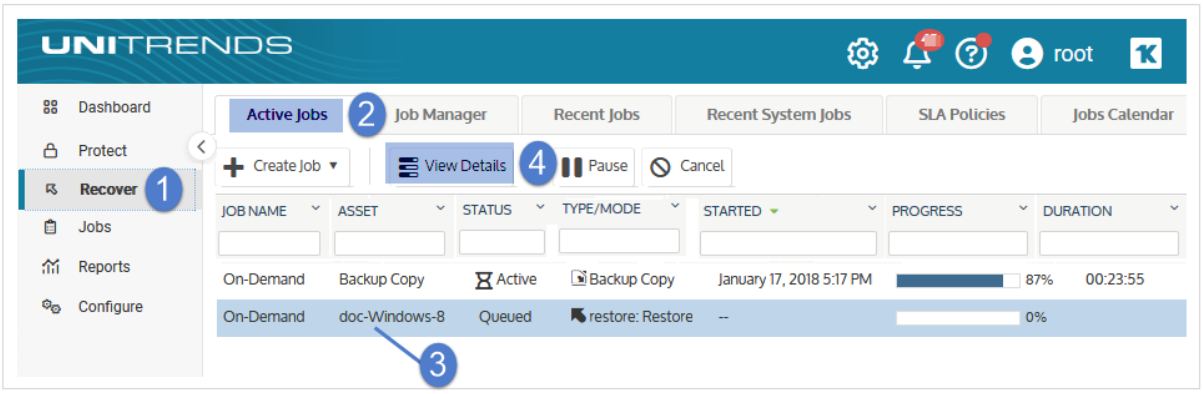

- The recovery is complete when the job's status changes to Success.

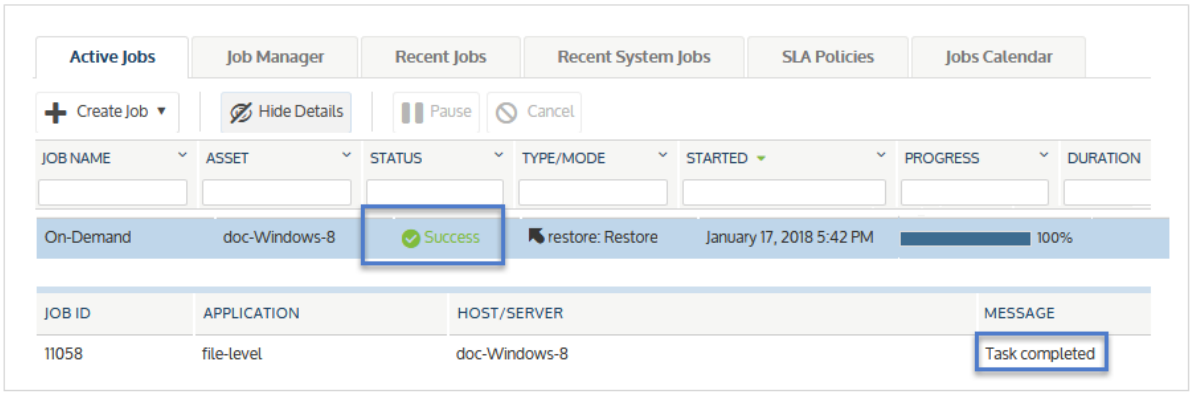

12 Access the recovered files on the recovery target.

Note: Windows assets – Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

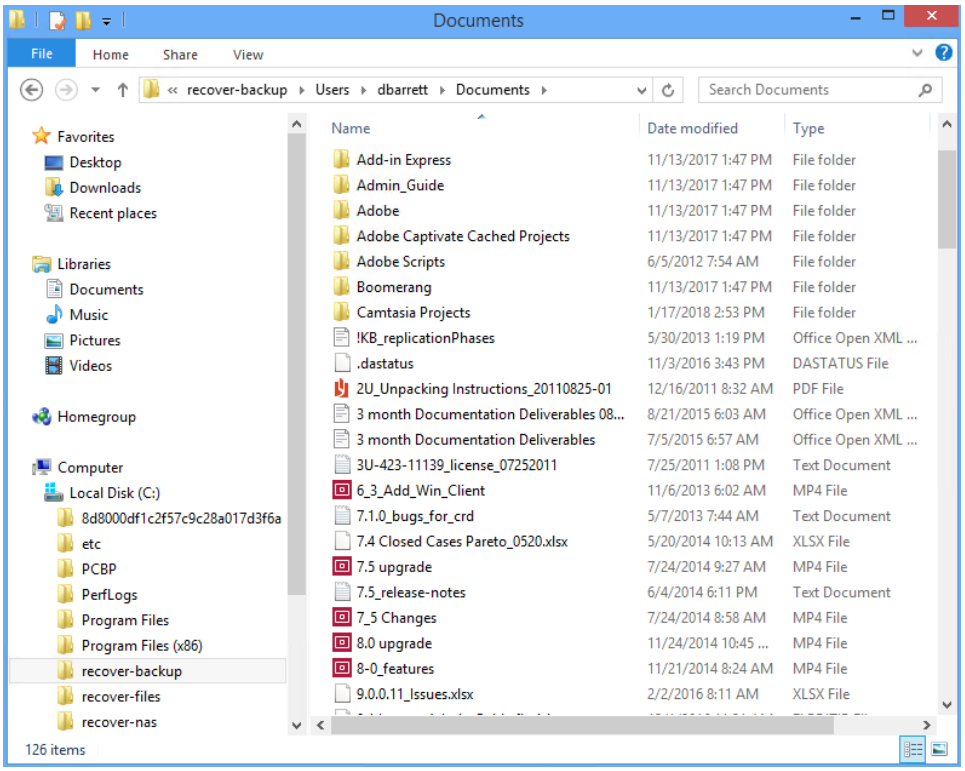

To recover an entire file-level backup by using the Backup Browser

Use this procedure to recover a backup.

Notes:

- You must use the Backup Catalog to recover an imported backup. (See the procedure above.)
- This procedure recovers the backup you select, plus all dependent data in the backup group. For example, recovering an incremental also recovers its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

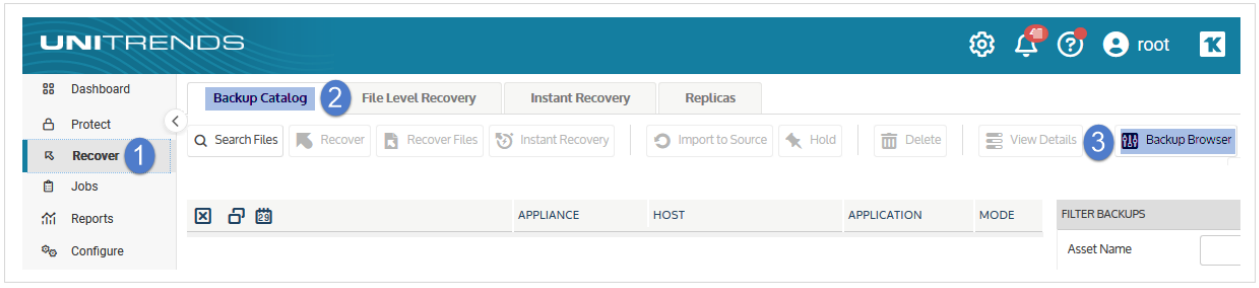

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display:

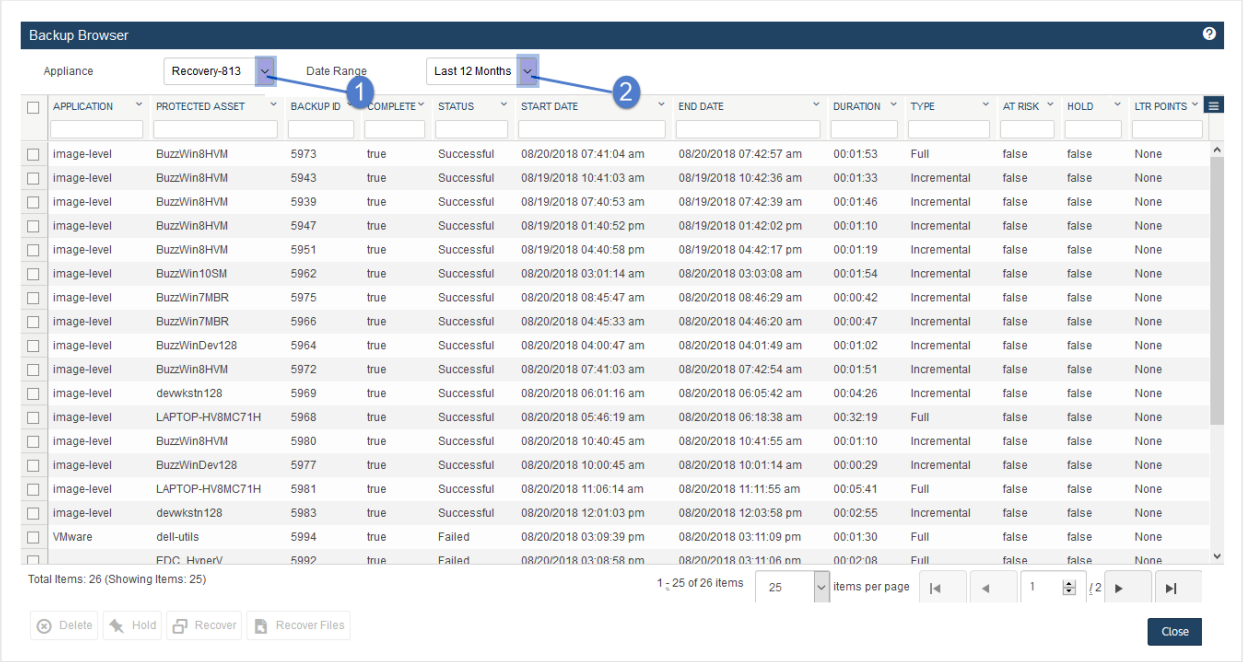

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Select the file-level backup and click **Recover**.

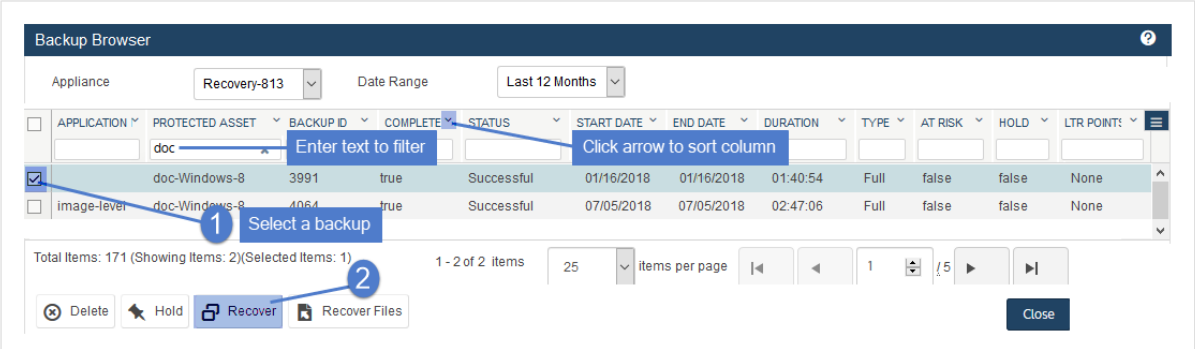

6 Select an **Asset**. Choose from the agent-based assets that have been added to this appliance.

7 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

8 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Commands to run post-restore	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Preserve directory structure	Check this box to preserve the existing file structure within the target directory. **Note:** To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

***Overwrite existing files* and *Restore newer files only* options**

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file. • If the file to recover is older than the one in the Target Directory, does not recover the file.
Overwrite existing files = Yes Restore newer files only = No	Recovers the file and overwrites the existing file.	Recovers the file and overwrites the existing file.
Overwrite existing files = No Restore newer files only = Yes	Does not recover the file.	Does not recover the file.

9 Click **Next**.

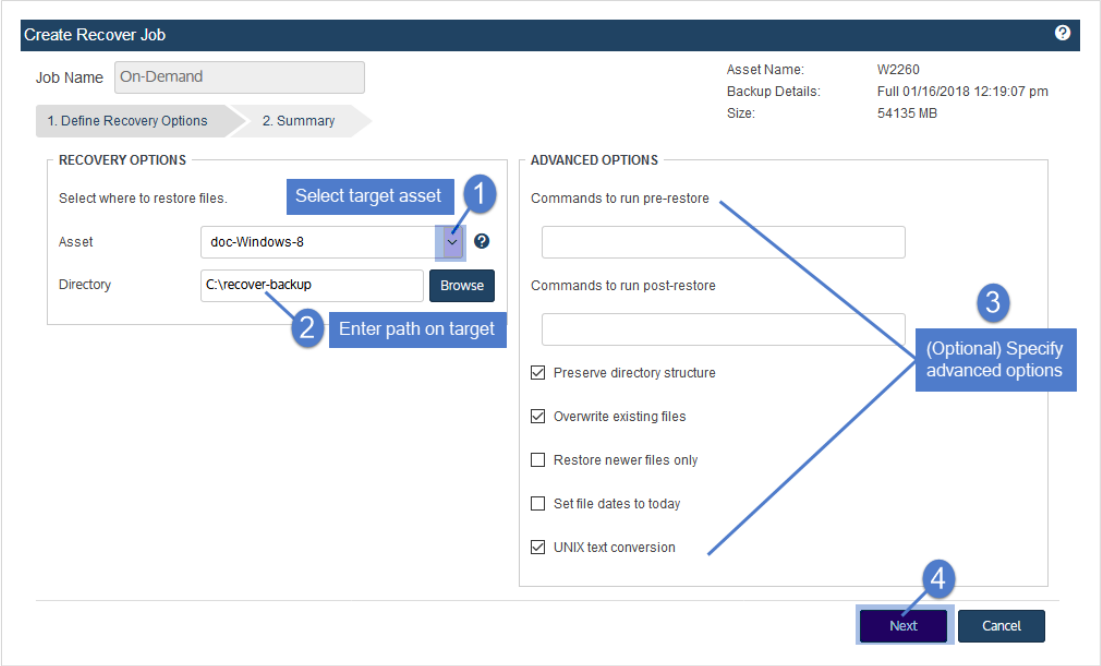

10 Review settings and click **Save**.

11 Click **OK** to close the Information message.

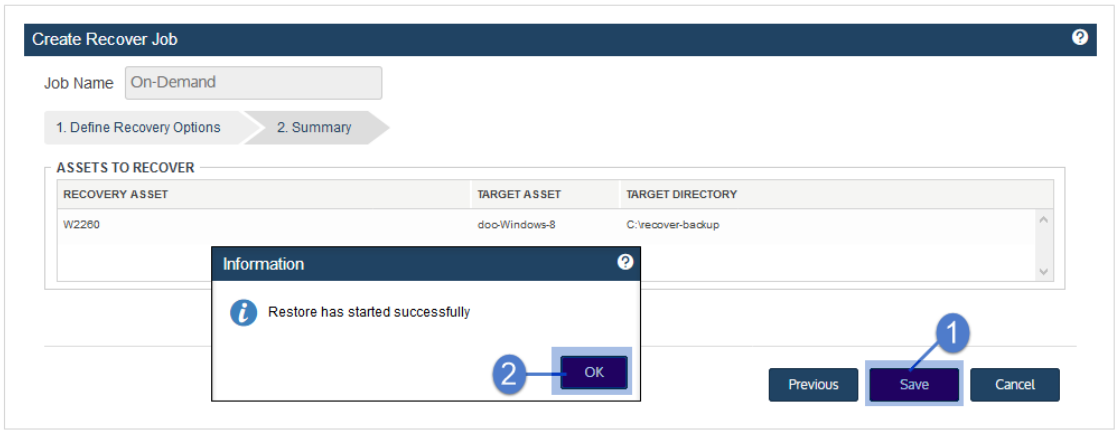

12 To monitor the recovery job:

- Select **Jobs > Active** Jobs.
- Select the job in the list and click **View Details**.

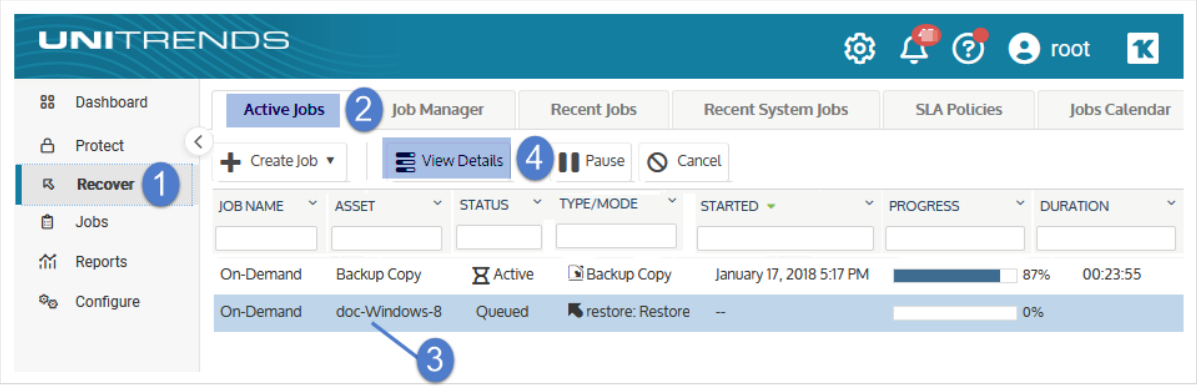

- The recovery is complete when the job's status changes to Success.

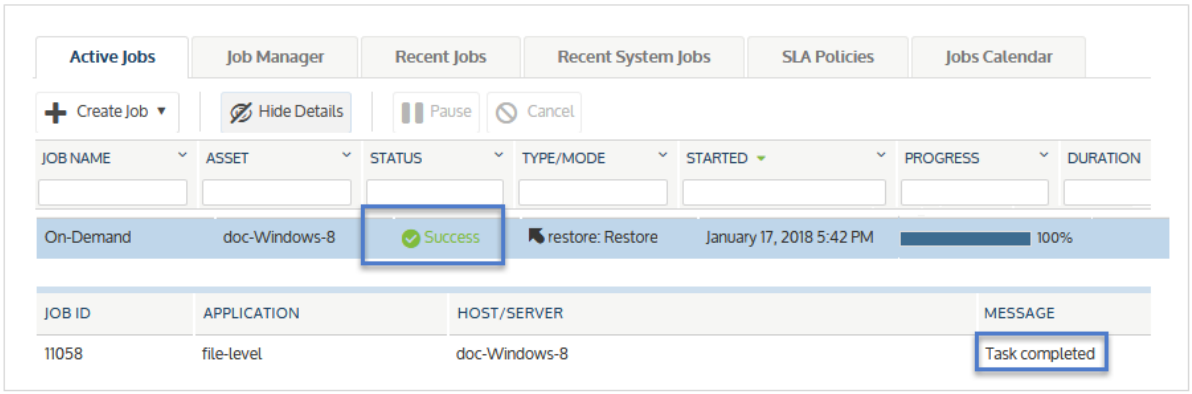

13 Access the recovered files on the recovery target.

Note: Windows assets – Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

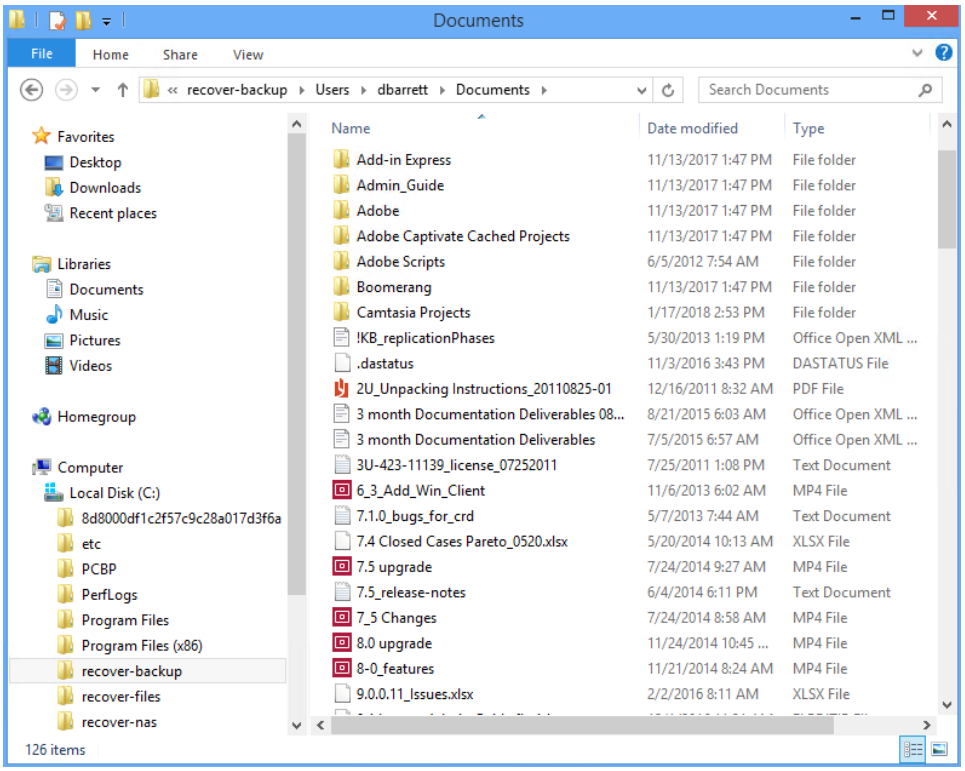

Recover files from cold backup copies

Run these procedures to recover files from cold backup copies:

- "Recover files from a cold backup copy by using Search Files"
- "Recover files from one cold backup copy by using the File Browser" on page 966

Recover files from a cold backup copy by using Search Files

By using Search Files, you can find files in a cold backup copy, import them to the appliance, then recover them to the desired target location. The appliance creates and imports a selective backup containing the files you have picked. Once this backup has been imported, you recover the imported backup.

Search Files searches all cold backup copies that have been created by or imported to this appliance for the selected asset. To import files, the appliance must have access to the associated backup copy. To access the backup copy:

- The backup copy target must be connected to the appliance. To check that the target is connected:
  1. On the **Configure > Appliances** page, select the source backup appliance.
  2. Click the **Backup Copy Targets** tab below.
  3. Click **Scan For Media**. The target displays in Offline status.

- If you are using removable media, the tape(s) or disk(s) where the files are stored must be loaded in the target. Do one of the following:
  - If you know which media contains the files you want to recover, you can load the required media before you run the recovery procedure.
  - If you do NOT know which media contains the files you want to recover, load the required media during the recovery procedure. After you enter file search criteria, the search results show the serial number(s) of the media where the files are stored.

Use these procedures to search an asset's cold backup copies for files that meet specified criteria and recover selected files from the search results.

Find, select, and import files

1 Log in to the backup appliance.

2 Click **Recover > Backup Catalog > Search Files**.

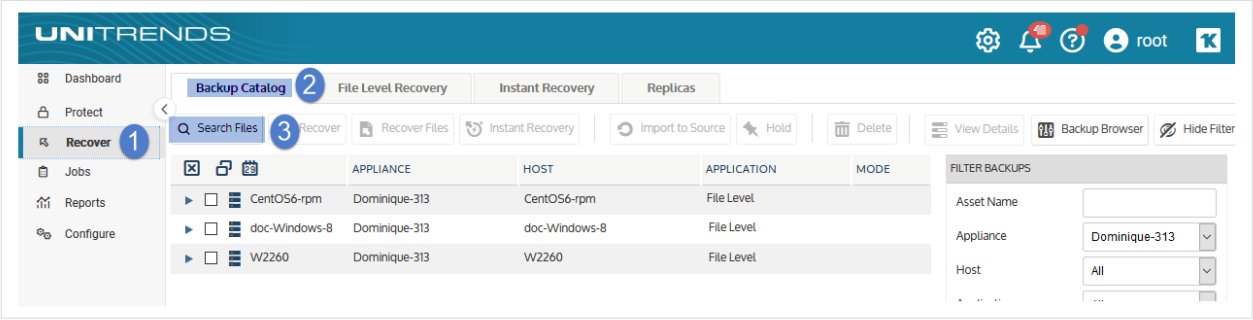

3 In the **Type** list, select **Backup Copy (Cold)**.

4 Select the **Asset** whose backups and imported backup copies will be searched.

5 Enter one or more search options:

Search Options	Description
String	Enter text to search. The wildcard character * is supported.
Match Case	Select to match the letter case of the entered string.
From/To	Use to search for files that were last modified within the specified date range. Results do not include files modified on the From or To date.
Size	Use to search for files that meet this size criteria.

Search Options	Description
Advanced	Click to search using a regular expression.

6 Click **Search**.

All cold backup copies of this asset are searched for matching files.

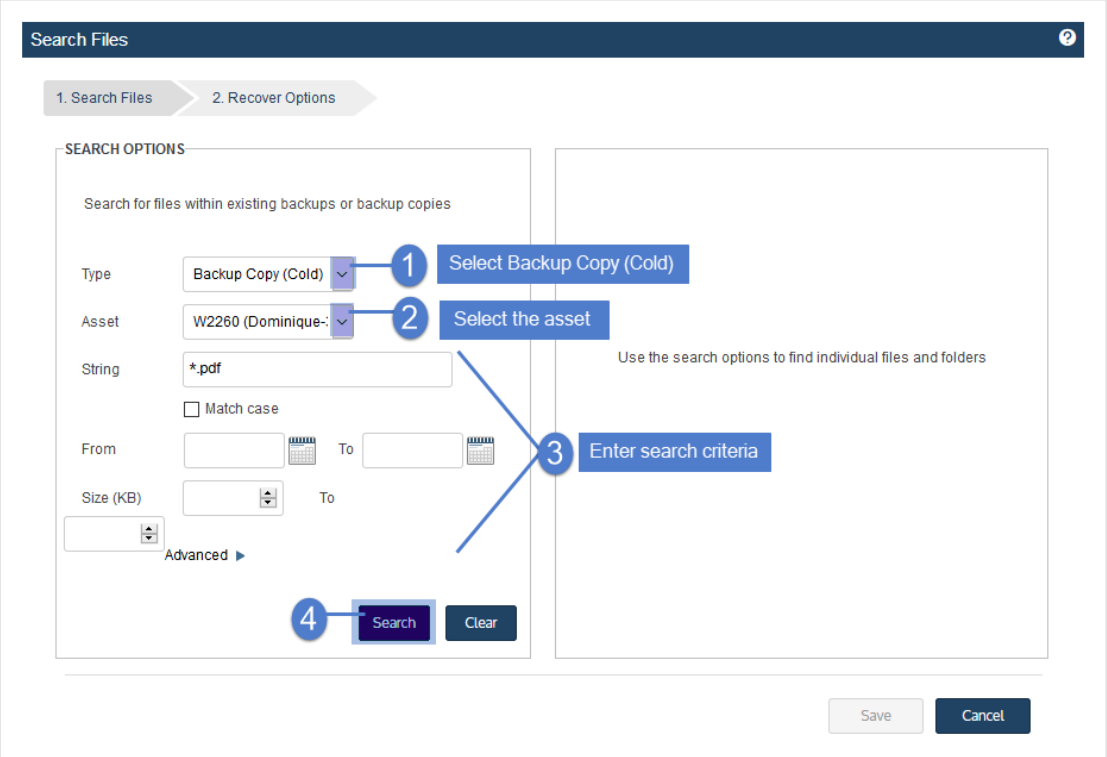

7 The returned files display. For disk or tape targets, the Serial # column shows the serial number of the disk(s) or tape(s) where the copy is stored.

8 (For disk or tape targets) Load the media that contains the files to import.

9 In the results list, click to select files to recover.

10 Click **Save**.

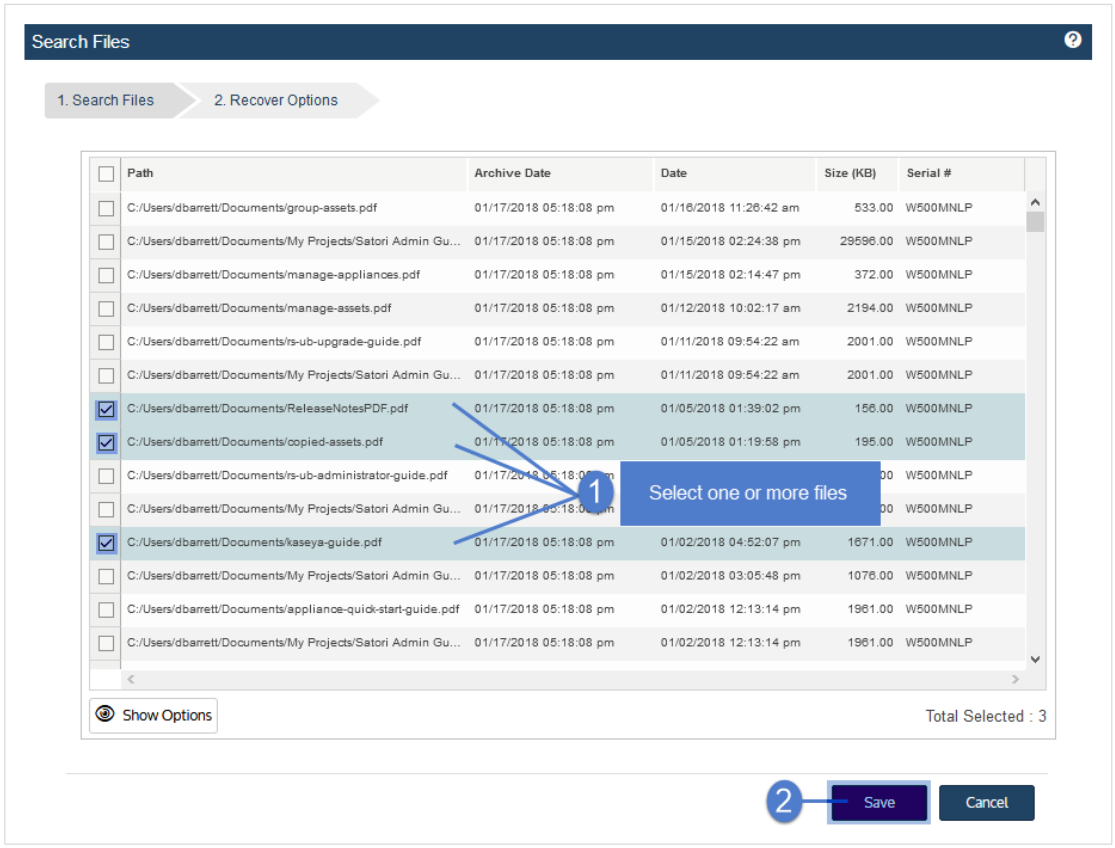

11 The recovery starts. Click **View Jobs**.

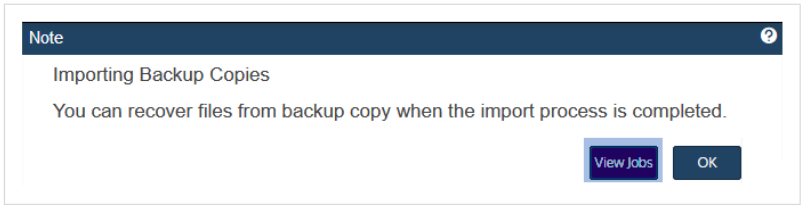

12 The recovery is complete when the job's status changes to Success.

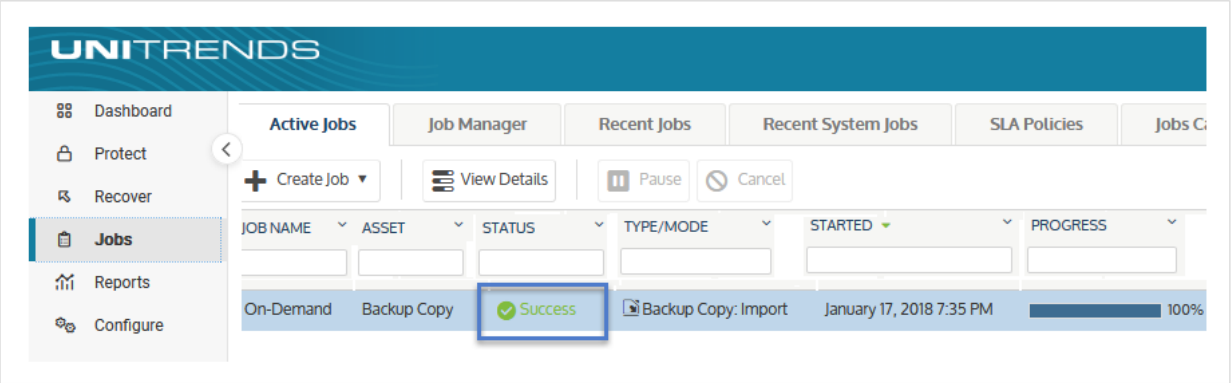

13 Once the import is complete, proceed to "Recover imported files".

Recover imported files

1 Select **Recover** and click the **Backup Catalog** tab.

(Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

2 Expand the asset and select the imported backup copy to use for the recovery.

To locate the files you imported, check the date and time, and look for an imported backup whose Mode is Selective.

3 Click **Recover**.

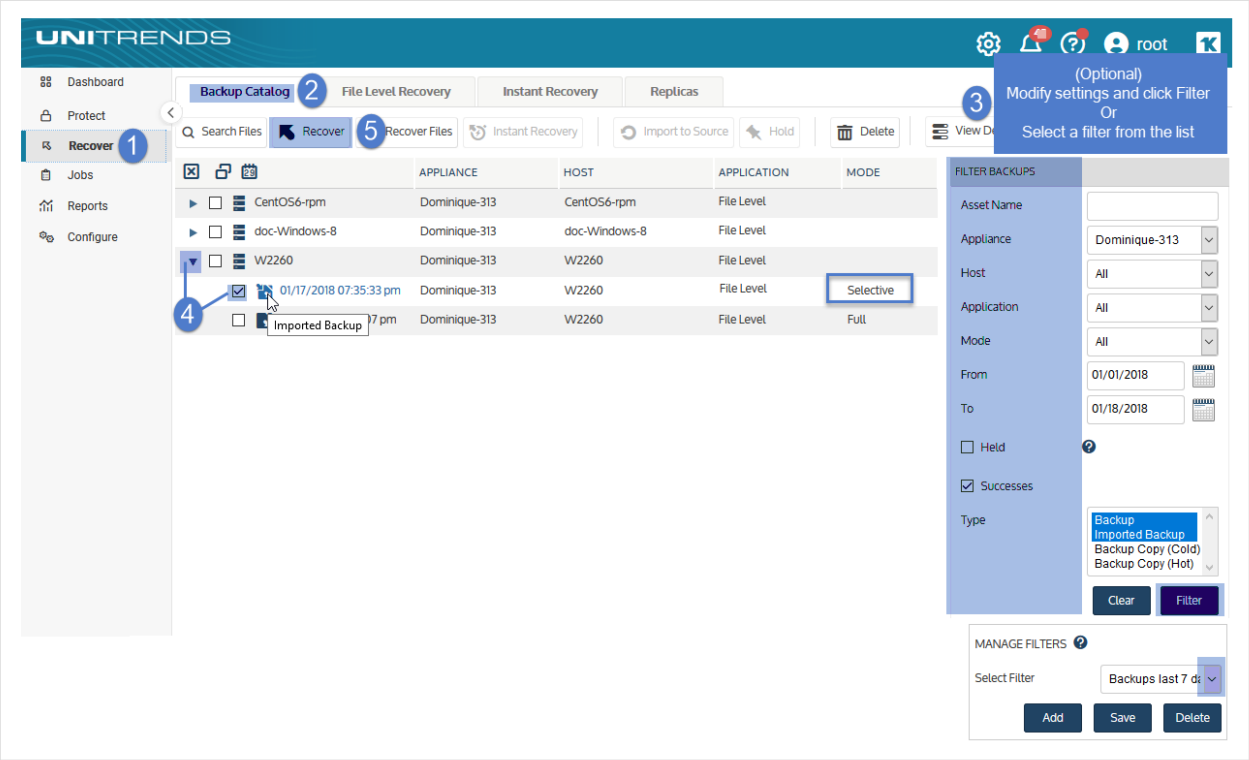

4 Select an **Asset**. Choose from the agent-based assets that have been added to this appliance.

5 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

6 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Commands to run post-restore	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Preserve directory structure	Check this box to preserve the existing file structure within the target directory.

Option	Description
	Note: To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

Overwrite existing files and *Restore newer files only* options

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	**Non-Windows backup**
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file. • If the file to recover is older than the one in the Target Directory, does not recover the file.
Overwrite existing files = Yes	Recovers the file and overwrites the existing file.	Recovers the file and overwrites the existing file.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
Restore newer files only = No		
Overwrite existing files = No Restore newer files only = Yes	Does not recover the file.	Does not recover the file.

7 Click **Next**.

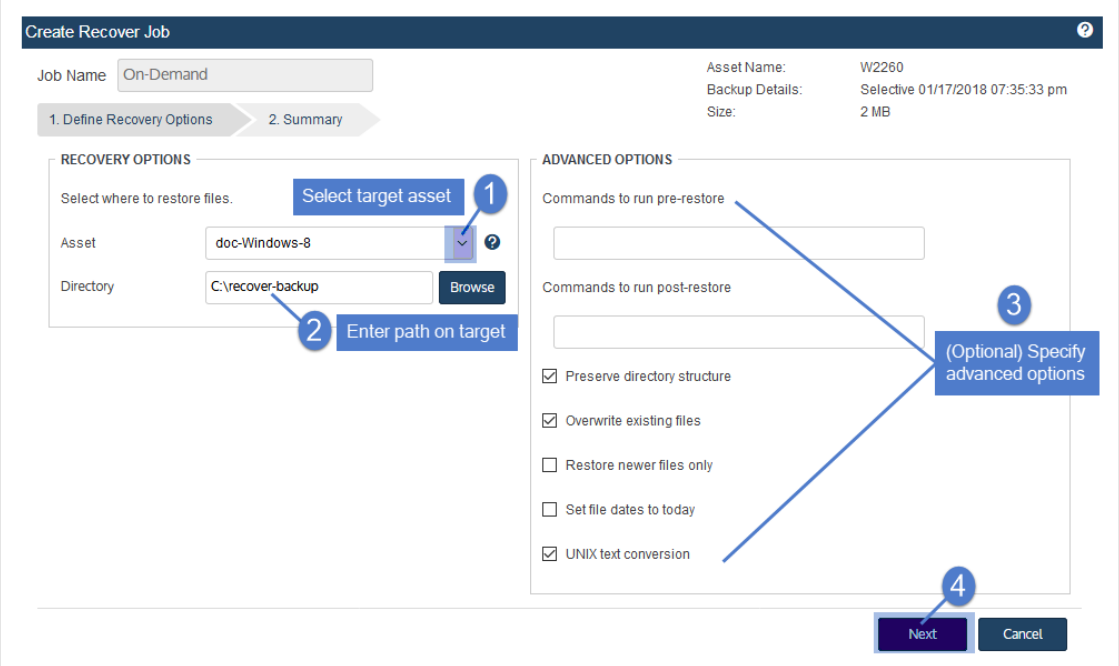

8 Review settings and click **Save**.

9 Click **OK** to close the Information message.

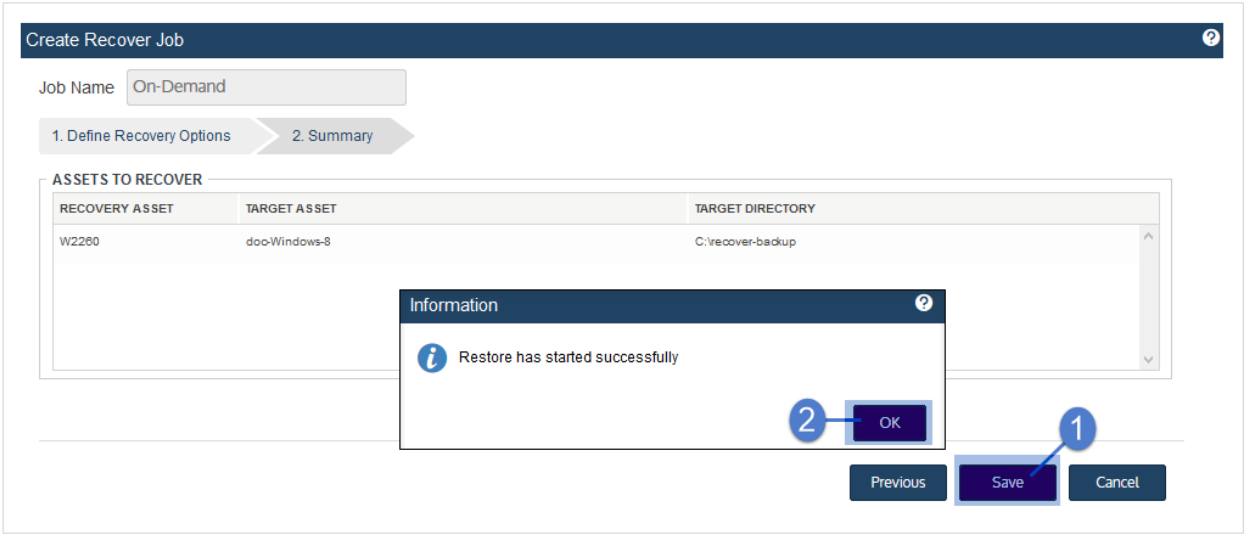

10 To monitor the recovery job:

- Select **Jobs** > **Active** Jobs.
- Select the job in the list and click **View Details**.

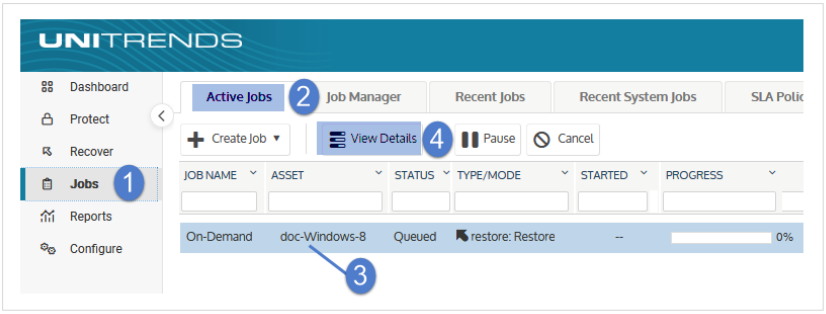

- The recovery is complete when the job's status changes to Success.

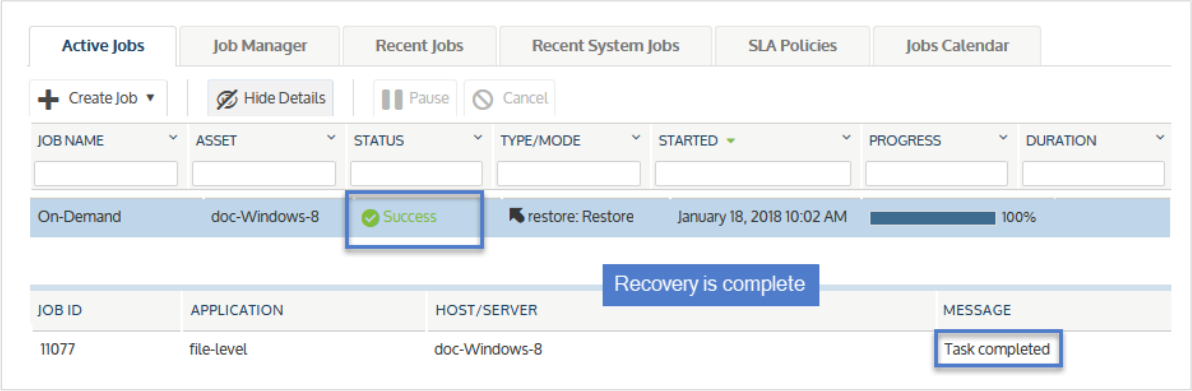

11 Access the recovered files on the recovery target.

Note: Windows assets – Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

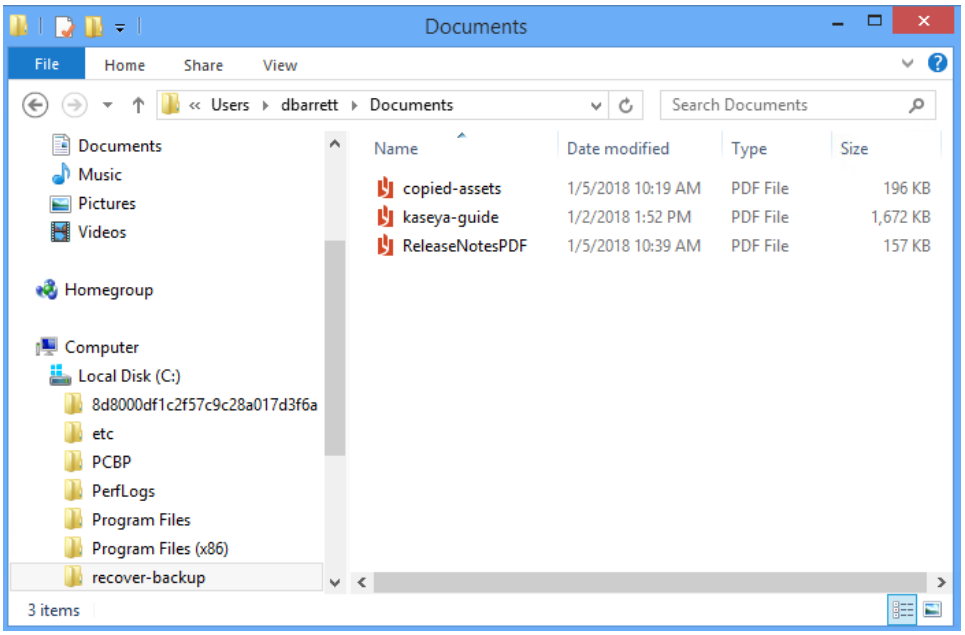

Recover files from one cold backup copy by using the File Browser

By using Recover Files, you can browse a cold backup copy and select specific files to recover. The appliance creates and imports a selective backup containing the files you have picked. Once this backup has been imported, you recover the imported backup.

Recover Files enables you to browse a cold backup copy that currently resides on the cold backup copy target. Verify the following before starting this procedure:

- The backup copy target is connected. To check this:
  - On the **Configure > Appliances** page, select the source backup appliance.
  - Click the **Backup Copy Targets** tab below.
  - (If needed) Click **Scan For Media**. The target displays in Offline status.
- If you are using removable media, the tape(s) or disk(s) where the files are stored must be loaded in the target.

Use these procedures to browse the contents of one cold backup copy and recover selected files.

Note: The file browser contains the backup copy you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

Find, select, and import files

1. Log in to the backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.
3. Use Filter Backups to the right to display cold backup copies:
   - In the Type box, select **Backup Copy (Cold)**.
   - Select other filter options as desired. For details, see "Working with custom filters" on page 67.
   - Click **Filter**.
4. Expand the asset and select a backup copy to use for the recovery.
5. Click **Recover Files**.

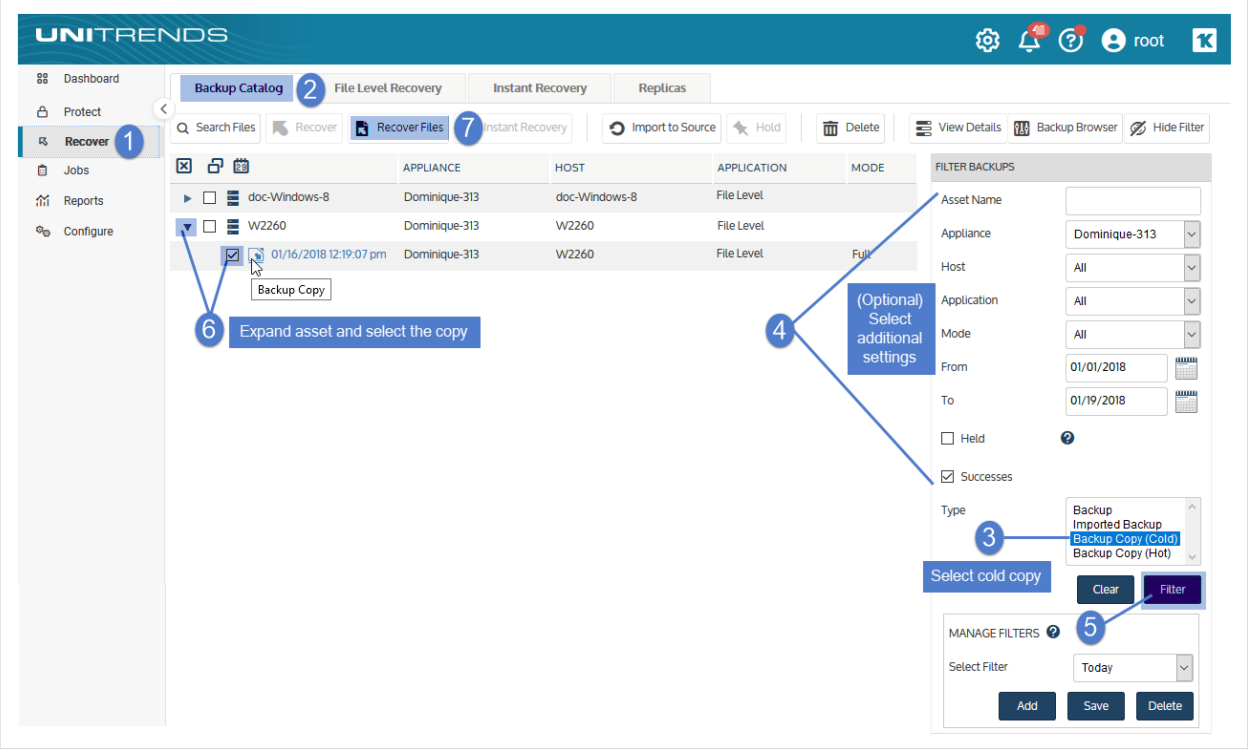

6. Select or drag files and/or directories to recover.
7. Click **Next**.

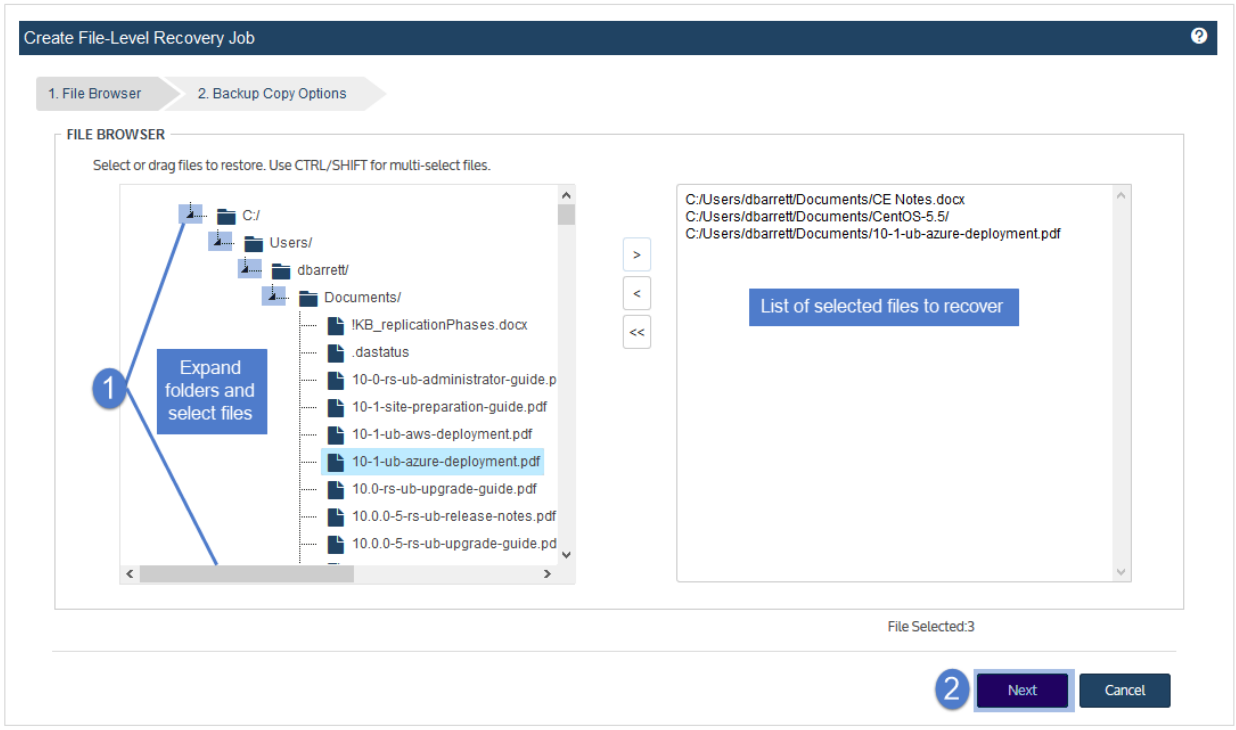

8 You see a message indicating that a selective backup will be created and imported to the appliance. Click **Save** to continue.

9 The recovery starts. Click **OK** to close the Notice message.

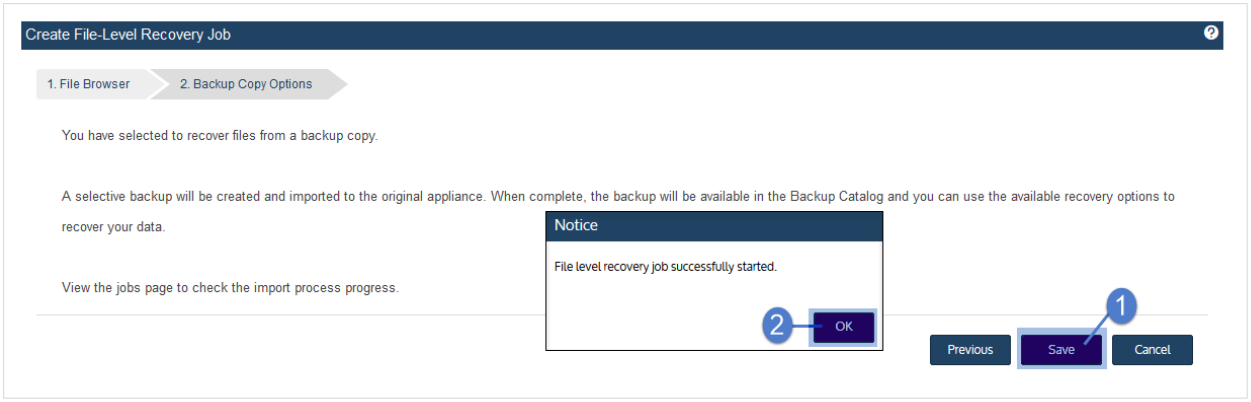

10 To monitor progress of the import:

- Select **Jobs > Active** Jobs.
- Select the job in the list and click **View Details**.

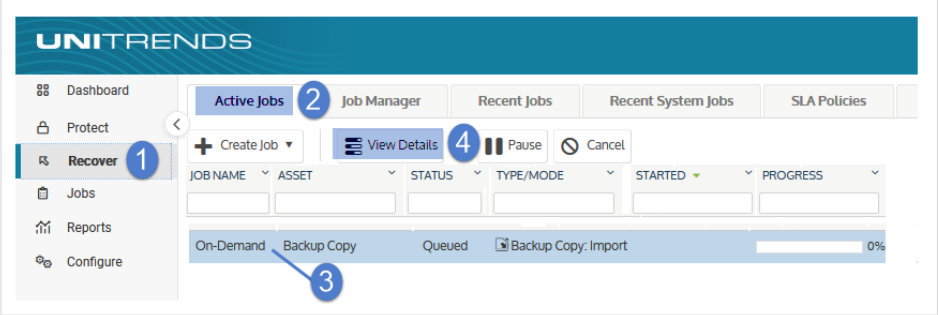

11 The import is complete when the job's status changes to Success.

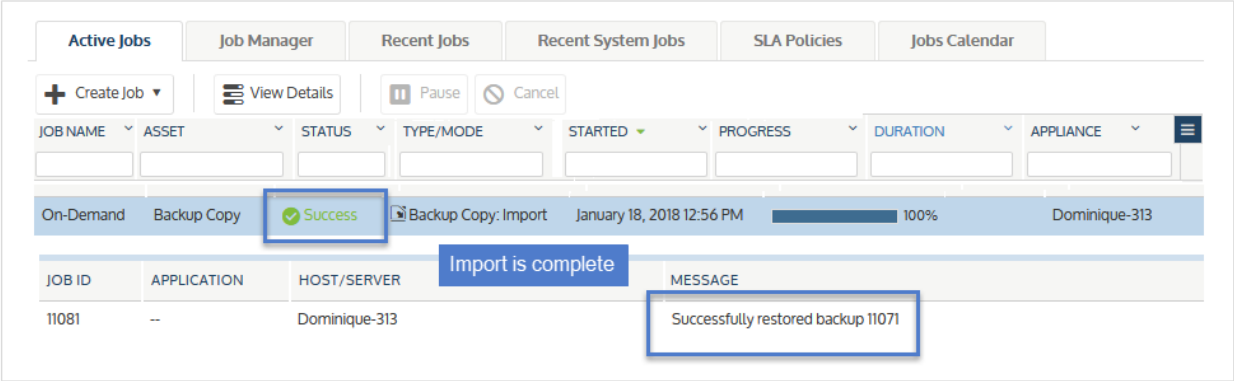

12 Once the import is complete, proceed to the next procedure to recover the imported files.

Recover imported files

1 Select **Recover** and click the **Backup Catalog** tab.

(Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

2 Expand the asset and select the imported backup copy to use for the recovery.

To locate the files you imported, check the date and time, and look for an imported backup whose Mode is Selective.

3 Click **Recover**.

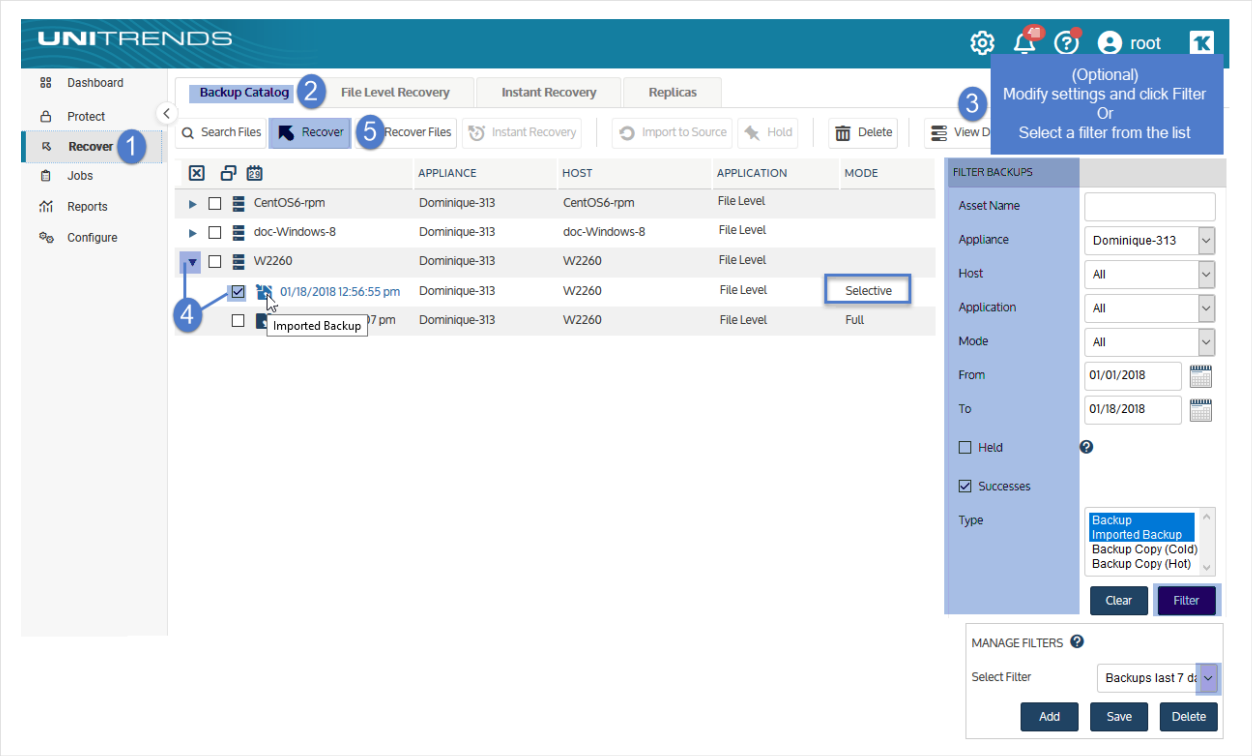

4 Select an **Asset**. Choose from the agent-based assets that have been added to this appliance.

5 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

6 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Commands to run post-restore	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Preserve directory structure	Check this box to preserve the existing file structure within the target directory.

Option	Description
	Note: To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

***Overwrite existing files* and *Restore newer files only* options**

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	**Non-Windows backup**
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file. • If the file to recover is older than the one in the Target Directory, does not recover the file.
Overwrite existing files = Yes	Recovers the file and overwrites the existing file.	Recovers the file and overwrites the existing file.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
Restore newer files only = No		
Overwrite existing files = No Restore newer files only = Yes	Does not recover the file.	Does not recover the file.

7 Click **Next**.

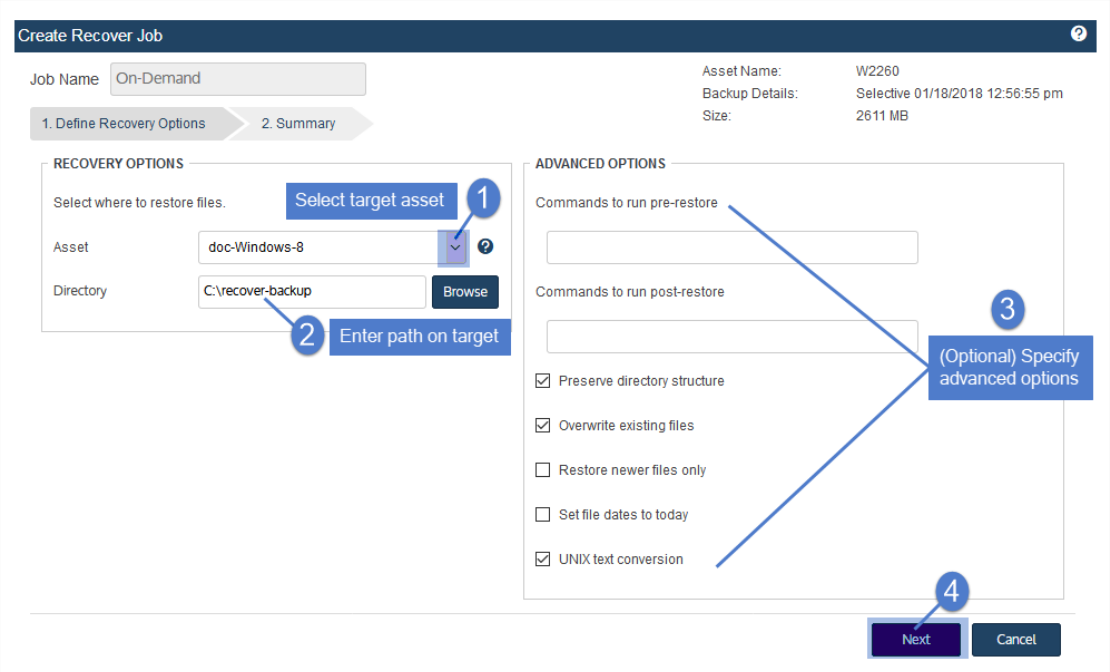

8 Review settings and click **Save**.

9 Click **OK** to close the Information message.

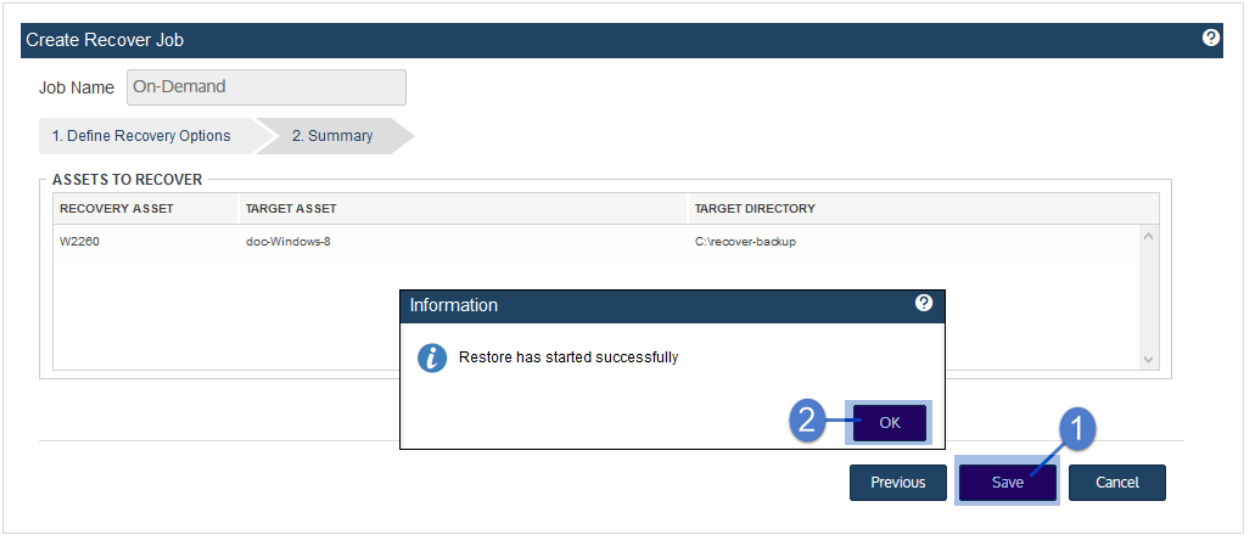

10 To monitor the recovery job:

- Select **Jobs** > **Active** Jobs.
- Select the job in the list and click **View Details**.

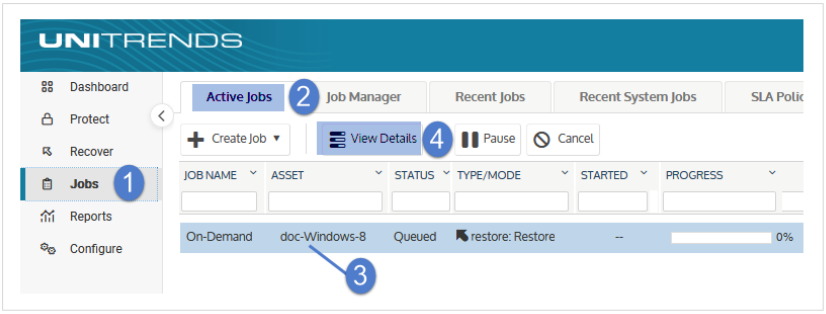

- The recovery is complete when the job's status changes to Success.

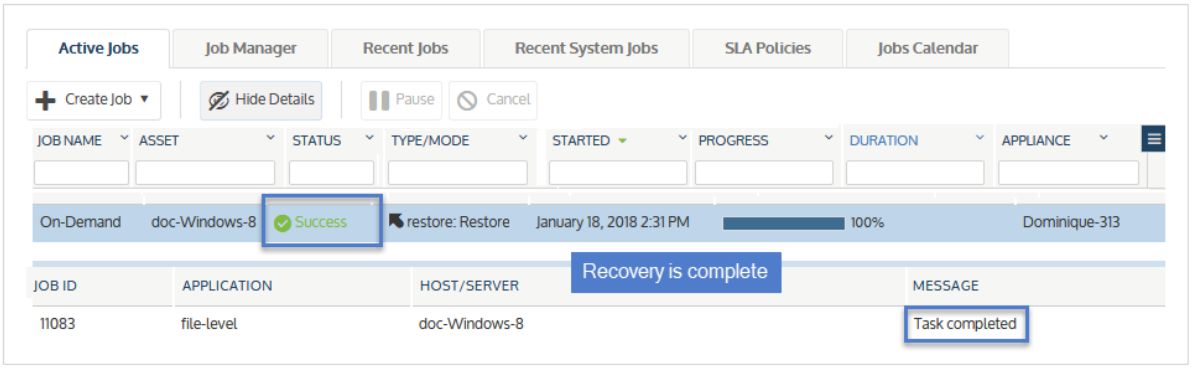

11 Access the recovered files on the recovery target.

Note: Windows assets – Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

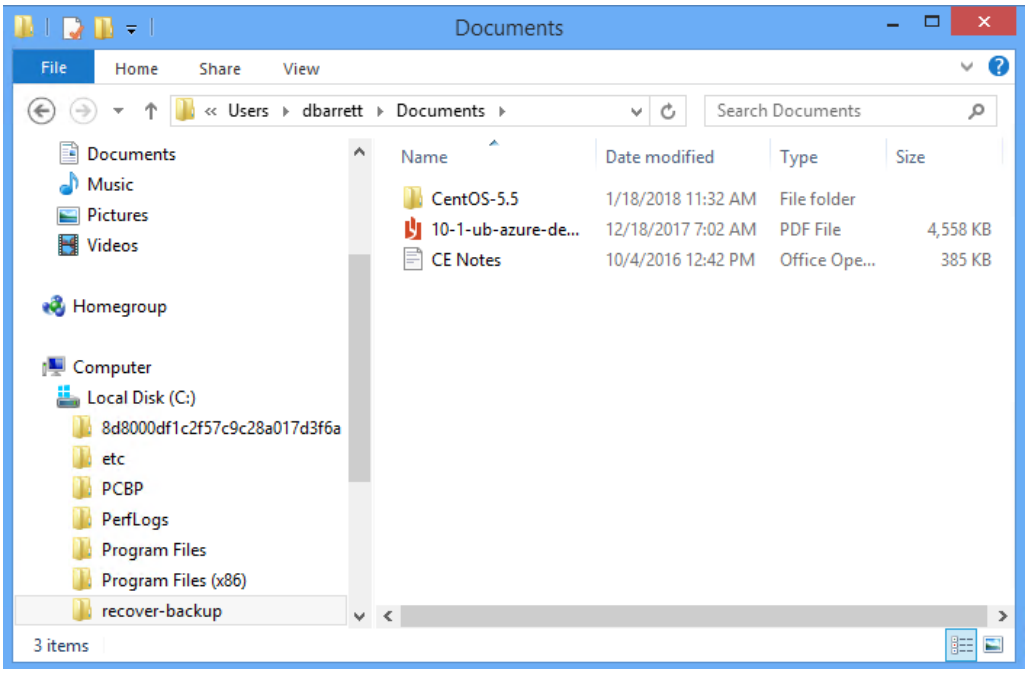

Recover from hot backup copies by running procedures on the target appliance

Run the these procedures from the backup copy target appliance to recover hot backup copies that are stored on that appliance. Before you start, verify that a recovery target is available. You must recover to an agent-based asset that has been added to this target appliance. (The asset must have a Unitrends agent installed and must display on the appliance's Protected Assets tab.) If necessary, add the target asset as described in "To add an agent-based asset" on page 289.

Note: To recover hot backup copies that reside in the Unitrends Cloud, see "Recover from hot backup copies by running procedures on the source appliance" on page 985.

- "Recovering File-level Backups" on page 925
- "To recover files from one file-level hot backup copy by using the File Browser"
- "To recover an entire file-level hot backup copy" on page 980

To recover files from one file-level hot backup copy by using the File Browser

Use this procedure to browse the contents of one backup copy and recover selected files.

Note: The file browser contains the backup copy you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1. Log in to the backup copy target appliance.
2. Select **Recover** and click the **Backup Catalog** tab.
3. Use Filter Backups to display hot backup copies:
   - In the Type box, select **Backup Copy (Hot)**.
   - Select other filter options as desired. For details, see "Working with custom filters" on page 67.
   - Click **Filter**.
4. Expand the asset and select a hot backup copy to use for the recovery.
5. Click **Recover Files**.

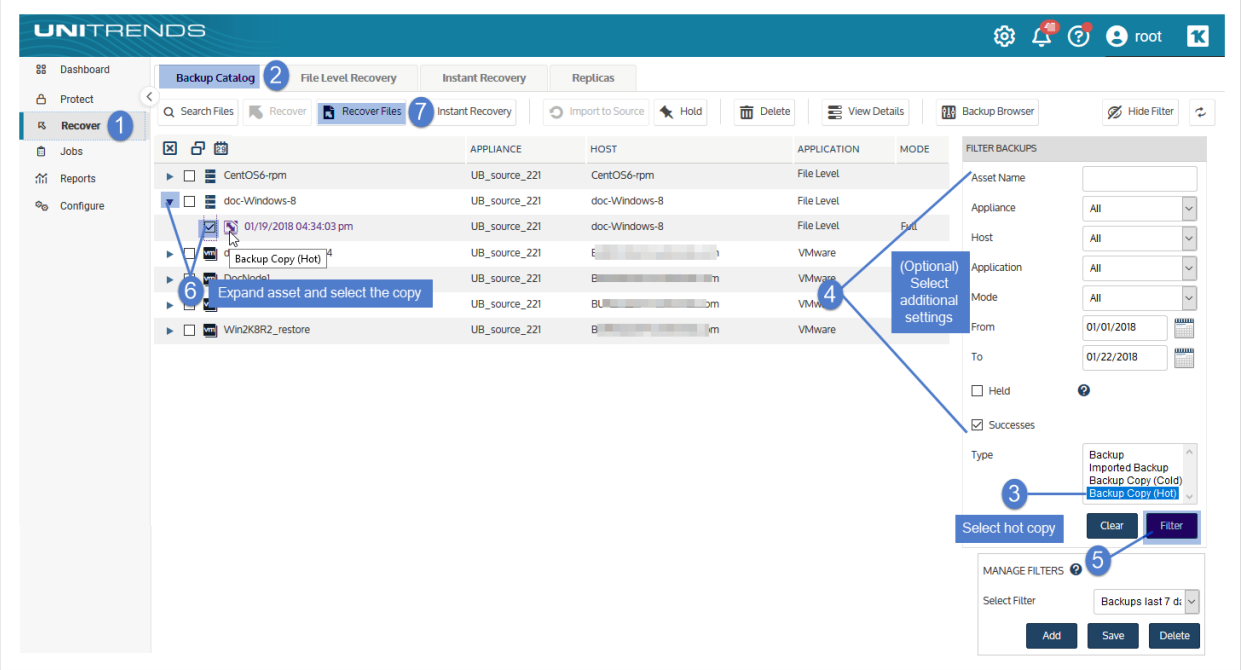

6. In the File Browser, expand folders to view items in the backup copy.
7. Select or drag files and/or directories to recover.
8. Click **Next**.

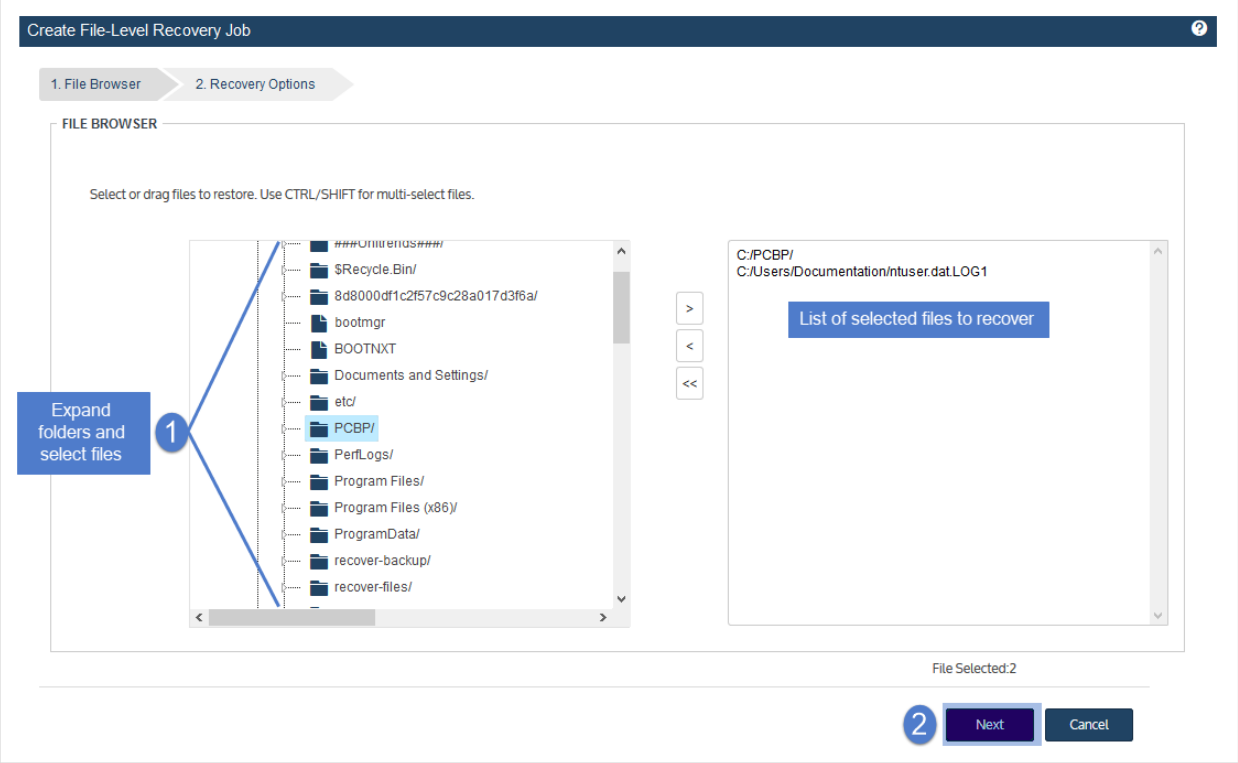

9 Select an **Asset**. Choose from the agent-based assets that have been added to this appliance.

10 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

11 (Optional) Specify Exclusions.

12 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Commands to run post-restore	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Preserve	Check this box to preserve the existing file structure within the target directory.

Option	Description
directory structure	**Note:** To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

Overwrite existing files and *Restore newer files only* options

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	**Non-Windows backup**
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file. • If the file to recover is older than the one in the Target Directory, does not recover the file.
Overwrite existing files = Yes	Recovers the file and overwrites the existing file.	Recovers the file and overwrites the existing file.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
Restore newer files only = No		
Overwrite existing files = No Restore newer files only = Yes	Does not recover the file.	Does not recover the file.

13 Click **Save**.

14 Click **OK** to close the Notice message.

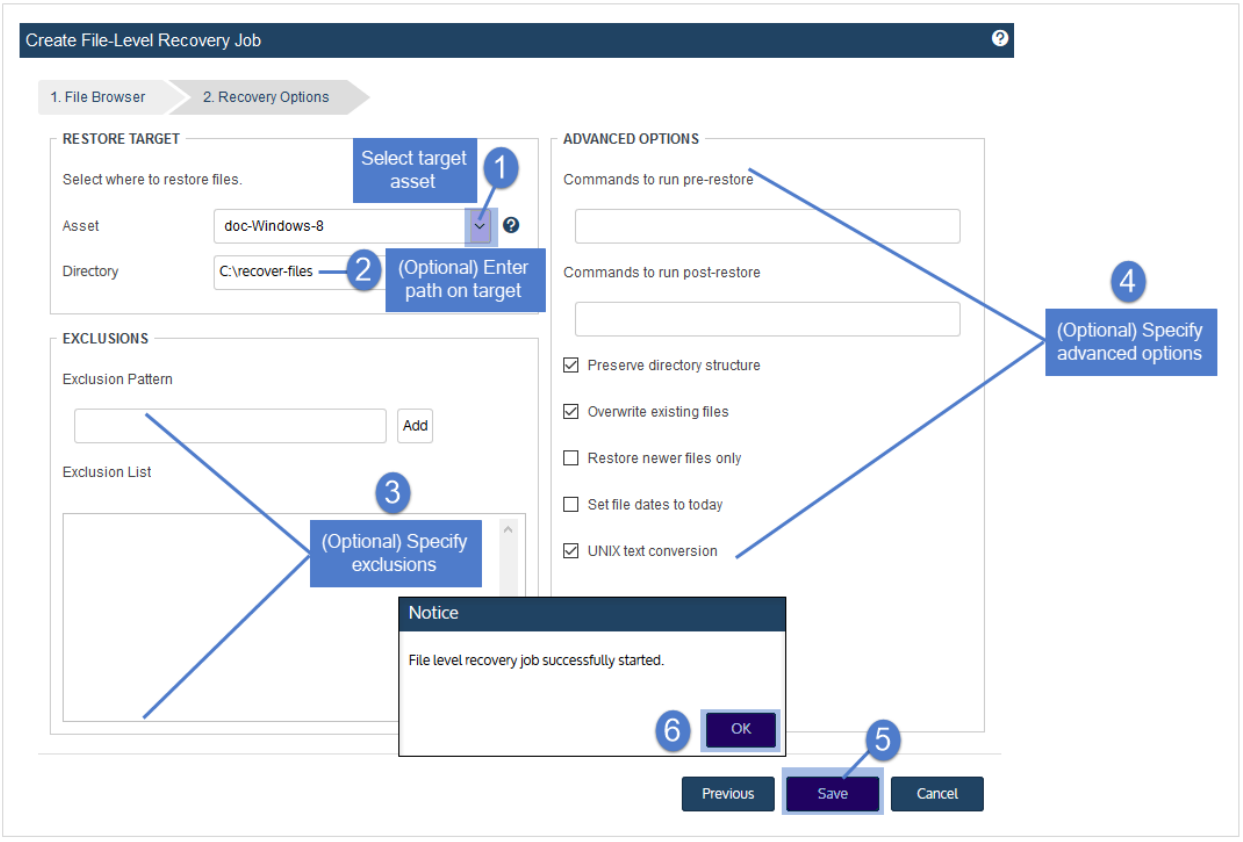

15 To monitor the recovery job:

- Select **Jobs** > **Active** Jobs.
- Select the job in the list and click **View Details**.

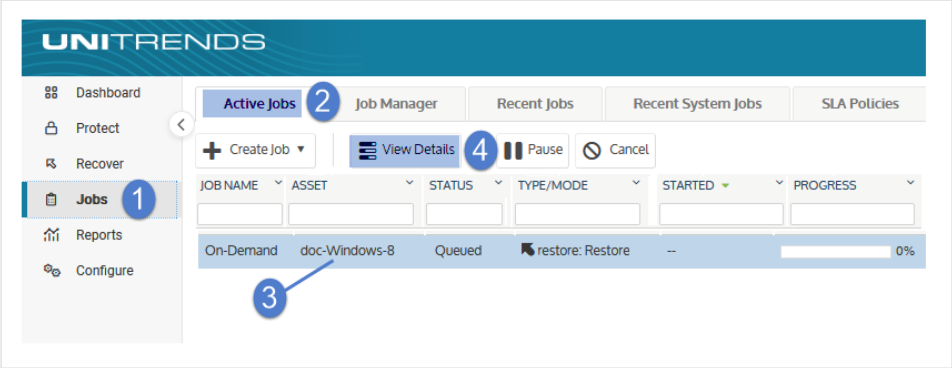

- The recovery is complete when the job's status changes to Success.

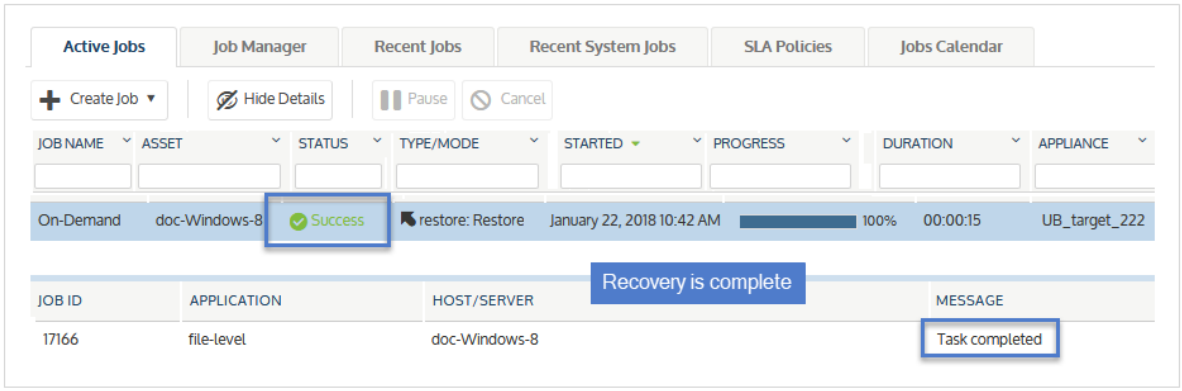

16 Access the recovered files on the recovery target.

Note: Windows assets – Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

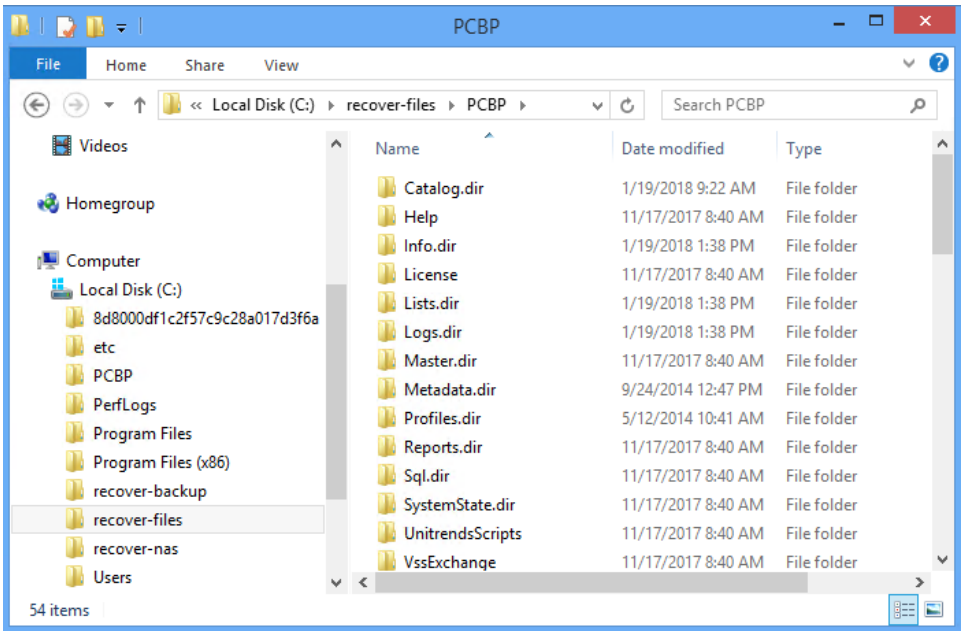

To recover an entire file-level hot backup copy

Note: This procedure recovers the backup copy you select, plus all dependent data in the backup group. For example, recovering an incremental also recovers its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1. Log in to the backup copy target appliance.
2. Select **Recover** and click the **Backup Catalog** tab.
3. Use Filter Backups to the right to display hot backup copies:
   - In the Type box, select **Backup Copy (Hot)**.
   - Select other filter options as desired. For details, see "Working with custom filters" on page 67.
   - Click **Filter**.
4. Expand the asset and select a hot backup copy to use for the recovery.
5. Click **Recover**.

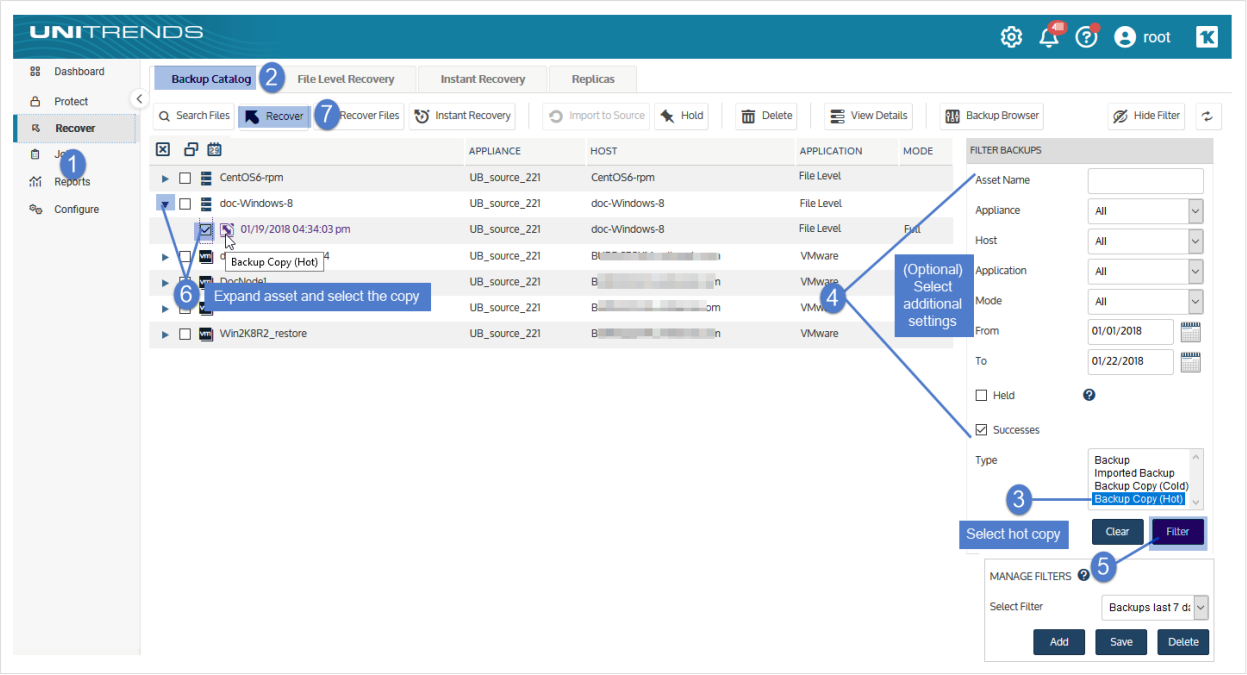

6 Select an **Asset**. Choose from the agent-based assets that have been added to this appliance.

7 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

8 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Commands to run post-restore	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Preserve directory structure	Check this box to preserve the existing file structure within the target directory. **Note:** To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.

Option	Description
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

Overwrite existing files and *Restore newer files only* options

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file. • If the file to recover is older than the one in the Target Directory, does not recover the file.
Overwrite existing files = Yes Restore newer files only = No	Recovers the file and overwrites the existing file.	Recovers the file and overwrites the existing file.
Overwrite	Does not recover the	Does not recover the file.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
existing files = No Restore newer files only = Yes	file.	

9 Click **Next**.

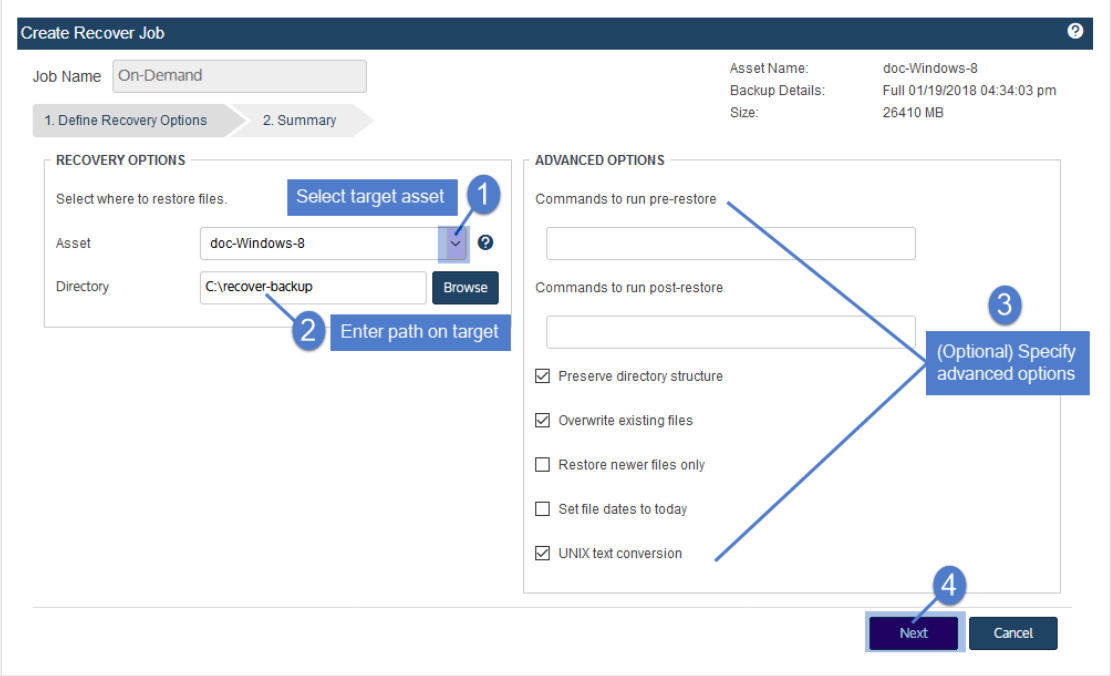

10 Review settings and click **Save**.

11 Click **OK** to close the Information message.

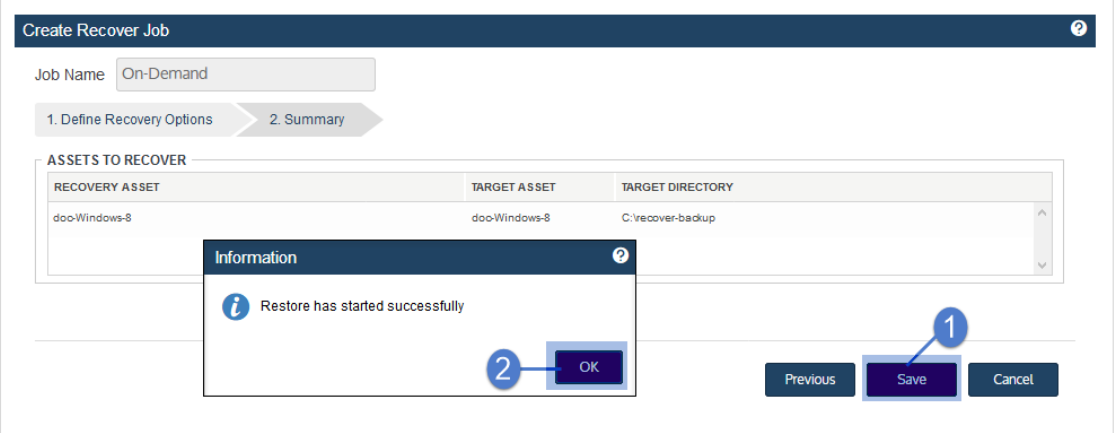

12 To monitor the recovery job:

- Select **Jobs** > **Active** Jobs.
- Select the job in the list and click **View Details**.

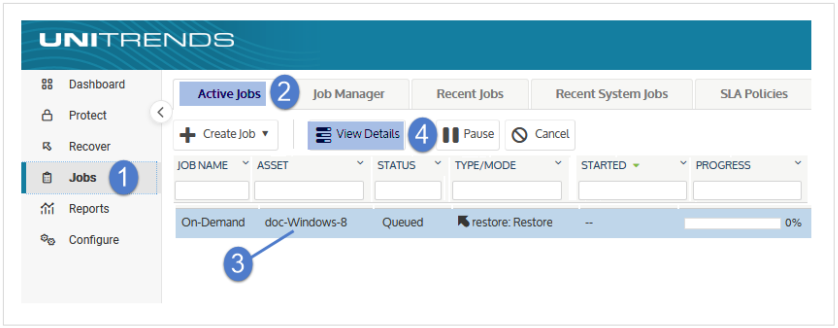

- The recovery is complete when the job's status changes to Success.

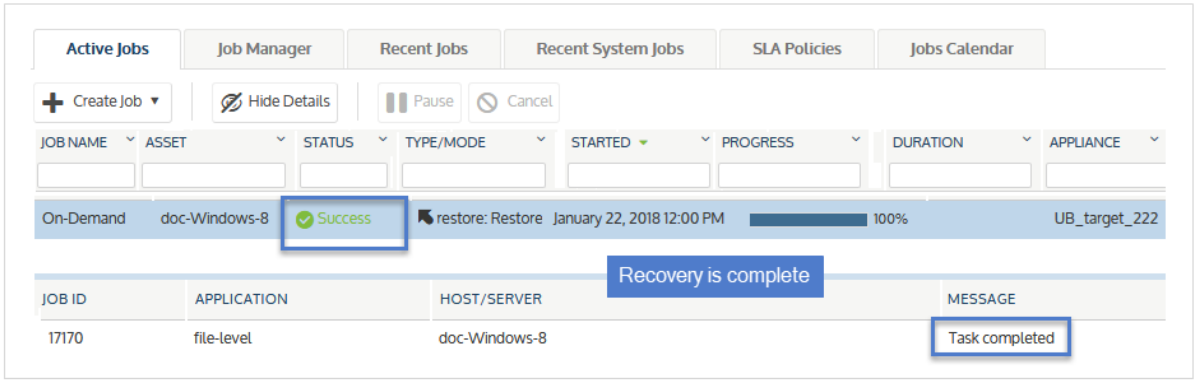

13 Access the recovered files on the recovery target.

Note: Windows assets – Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

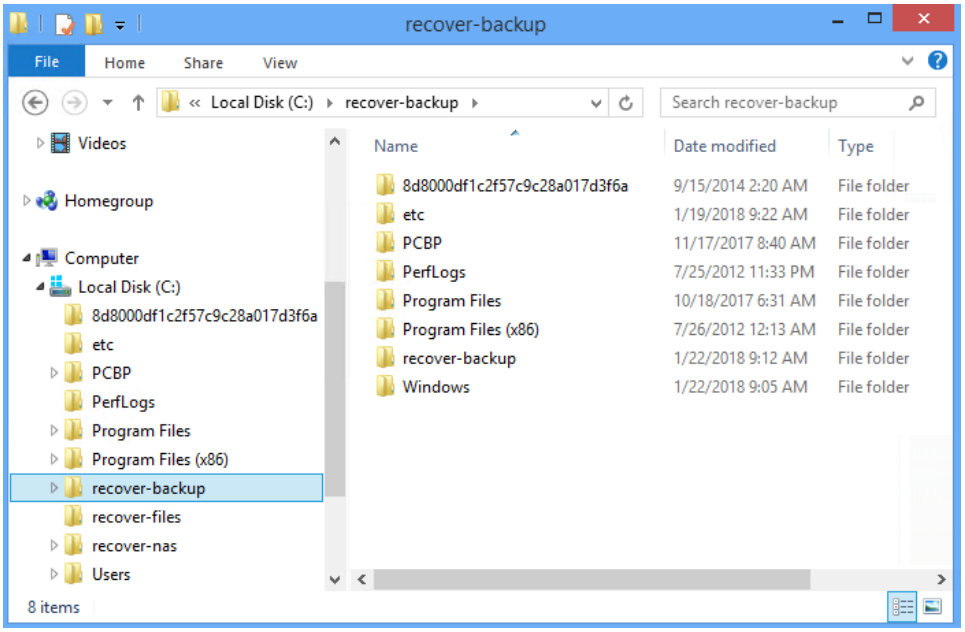

Recover from hot backup copies by running procedures on the source appliance

From the source appliance, you can recover hot backup copies that reside in the Unitrends Cloud or that reside on a target appliance. The recovery procedures either import the backup copy to the source appliance or recover files directly from the backup copy on the target.

To recover an entire backup copy, you must first import the backup copy to the source backup appliance. Once the backup copy has been imported, recover from the imported backup copy by using the "Recover from backups or imported backup copies" on page 926 procedures. For details on importing a backup copy, see "To import a hot backup copy" on page 780.

To recover files from backup copies that reside in the Unitrends Cloud or reside on a target appliance, run these procedures from the source backup appliance:

- "To recover files from a file-level backup copy by using the File Browser" on page 985
- "To recover files from a file-level backup copy by using Search Files" on page 988

To recover files from a file-level backup copy by using the File Browser

Use this procedure to browse a file-level backup copy and recover selected files.

Note: The file browser contains the backup copy you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1. Log in to the source backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.
3. Use Filter Backups to the right to display hot backup copies:
   - In the Type box, select **Backup Copy (Hot)**.
   - Select other filter options as desired. For details, see "Working with custom filters" on page 67.
   - Click **Filter**.
4. Expand the asset and select a *Backup Copy (Hot) on Target* to use for the recovery.

   To be certain you are viewing a backup copy that resides on the target appliance or in the Unitrends Cloud, hover over the backup copy icon and verify the description that displays is *Backup Copy (Hot) on Target*.
5. Click **Recover Files**.

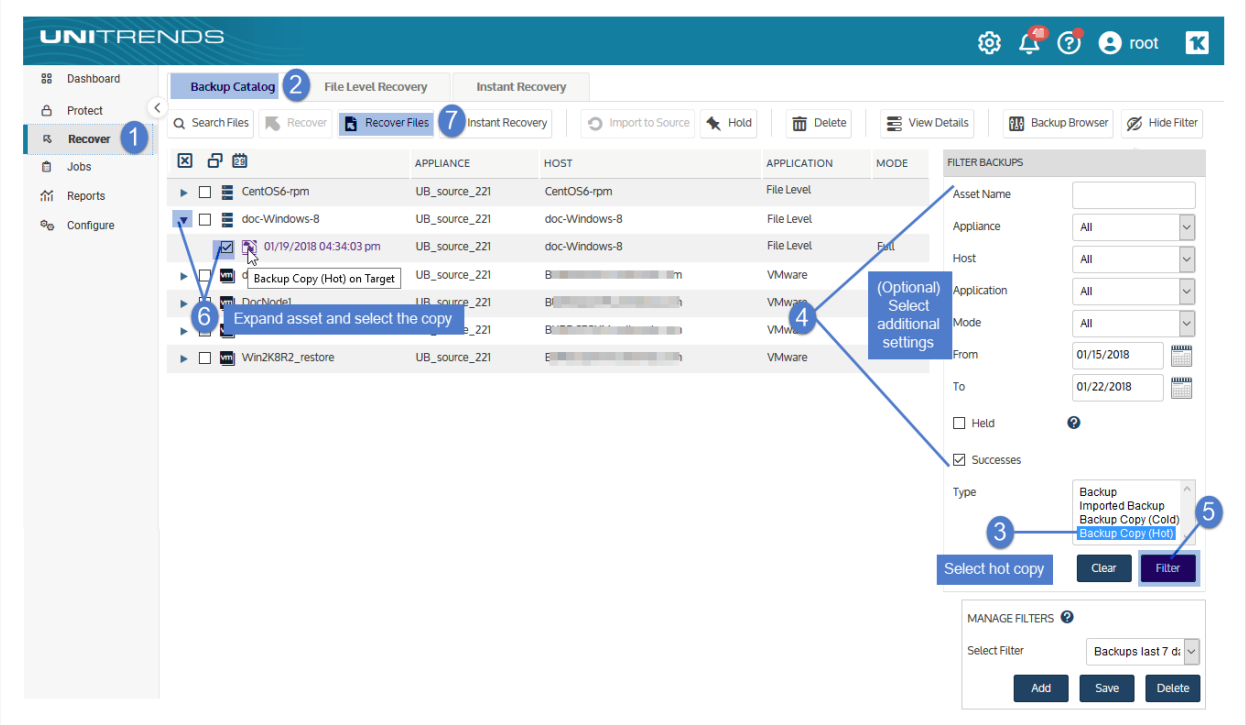

6. In the File Browser, expand folders to view items in the backup copy.
7. Select or drag files and/or directories to recover.

Note: Softlinks (also called *symbolic links*) are excluded from download. If you select a directory that contains files and softlinks, only the files are downloaded.

8 Click **Save**.

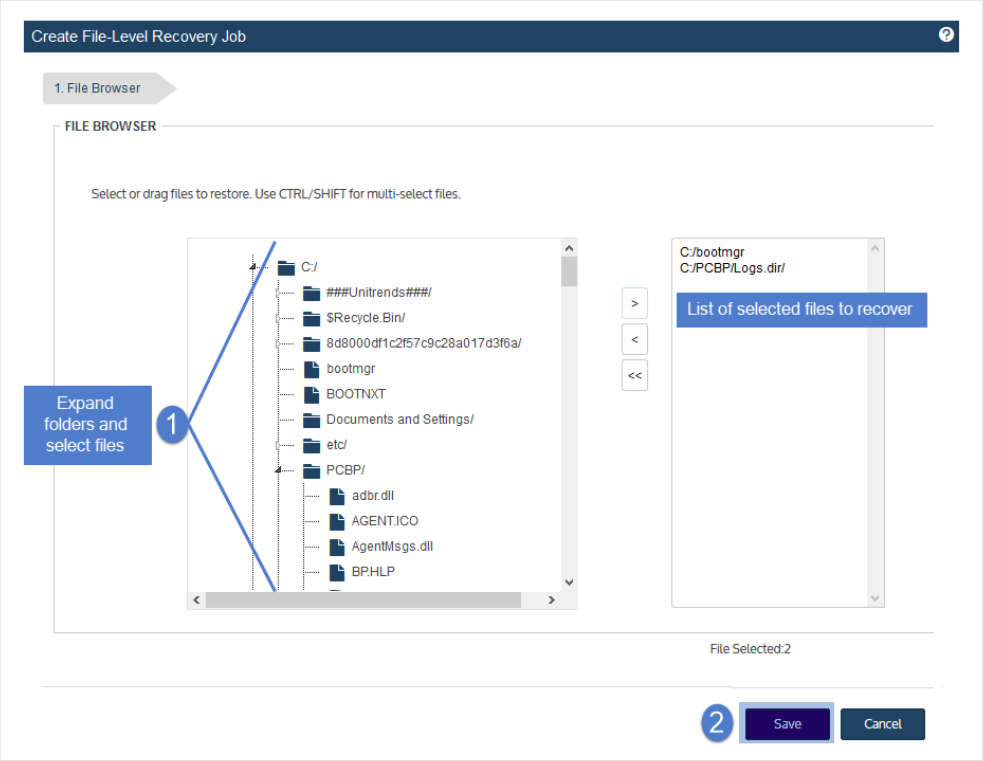

9 The message *Starting File Level Recovery on the Target* displays, indicating that the recovery has started.

10 To verify that the download starts, leave the message dialog open and view the status messages. Job status changes from Queued to Active to Downloading.

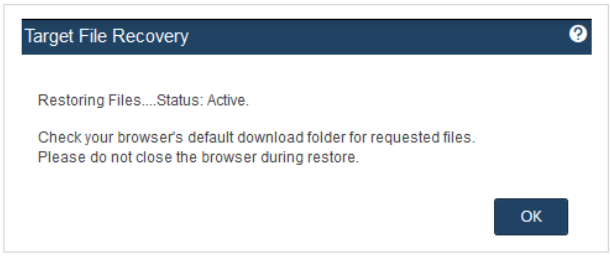

11 A *.zip* file containing the recovered files is placed in the default download location of the browser where the source appliance UI is running.

12 When the download completes, the *Unitrends-Restore.zip* file displays in the browser. Open the *.zip* file to access the recovered files.

Notes:

- The duration of the download is impacted by various factors, such as the size of the files, bandwidth, and download speed.
- Persistent browser and UI sessions are required to create the *.zip* file in the browser's default download location. If you close the browser or UI session during the recovery, you must run a new job.

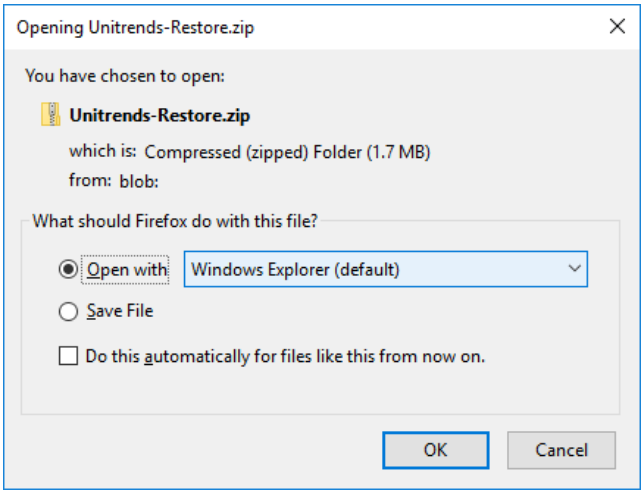

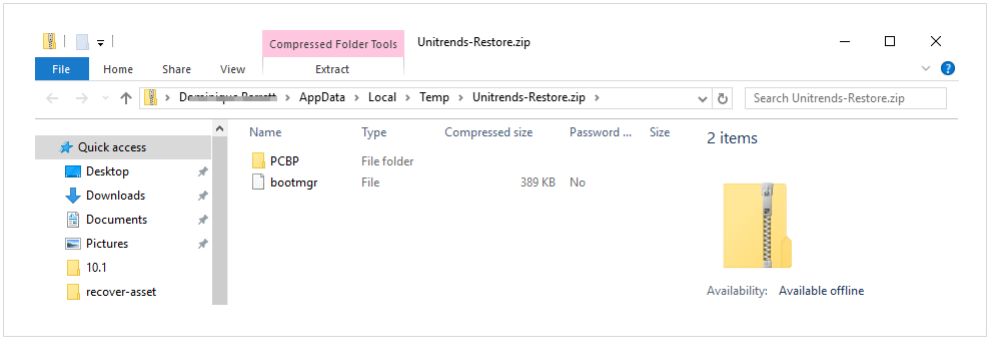

To recover files from a file-level backup copy by using Search Files

Use this procedure to search an asset's backup copies for files that meet specified criteria and recover selected files from the search results.

1. Log in to the source backup appliance.
2. Click **Recover > Backup Catalog > Search Files**.

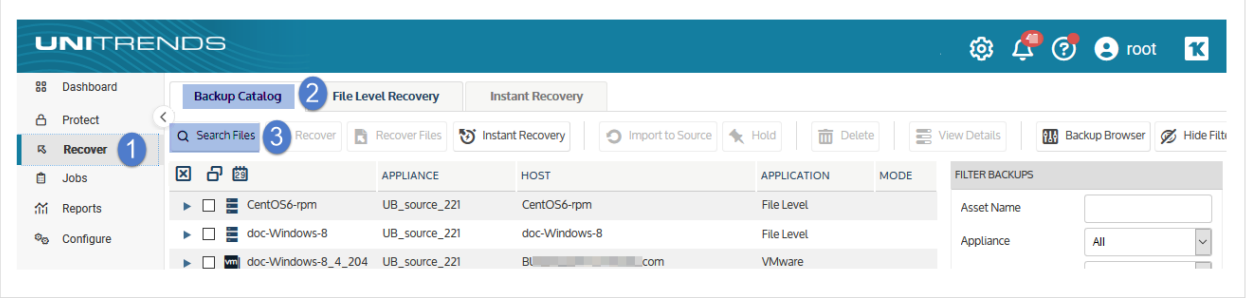

3 Select **Backup Copy (Hot)** from the **Type** list.

4 Select the **Asset** whose hot backup copies will be searched.

Notes:

- The Asset list contains all assets whose backups are being copied to one or more hot backup copy targets. The hot backup copy target displays in parenthesis next to the asset name. If an asset is copying to more than one hot target, the list contains an entry for each. Be sure to select the correct asset and target in the list.
- Multi-target example:

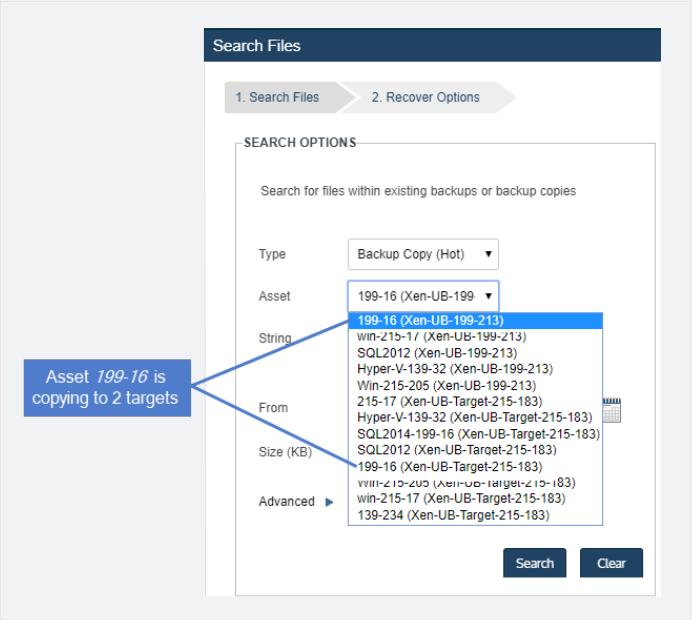

5 Enter one or more search options:

Search Options	Description
String	Enter text to search. The wildcard character * is supported.
Match Case	Select to match the letter case of the entered string.
From/To	Use to search for files that were last modified within the specified date range. Results do not include files modified on the From or To date.
Size	Use to search for files that meet this size criteria.
Advanced	Click to search using a regular expression.

6 Click **Search**.

All hot backup copies of this asset are searched for matching files.

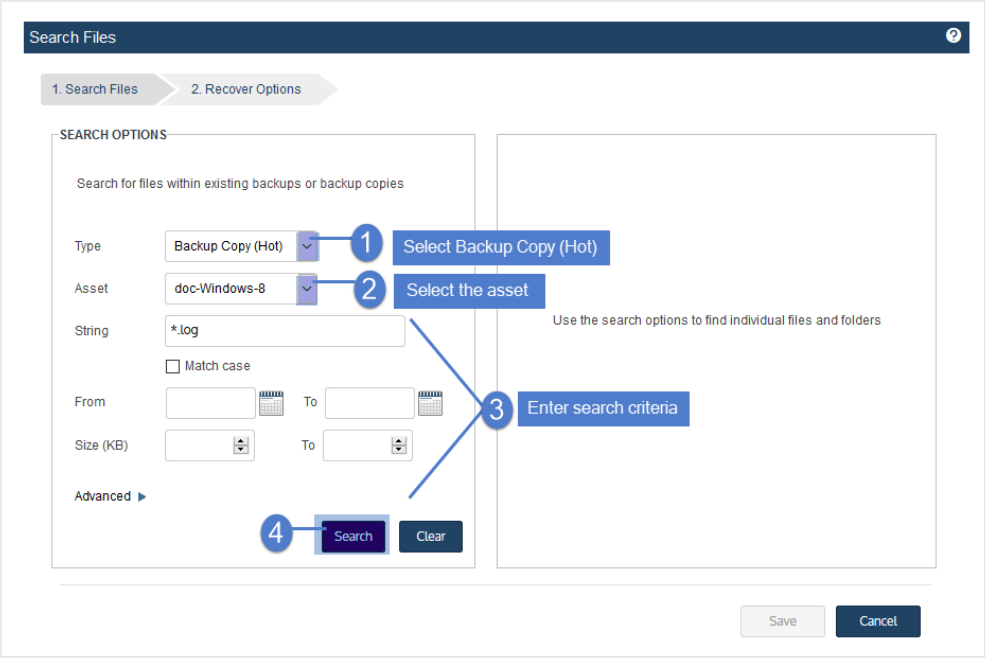

7 In the results list, click to select files to recover.

Notes:

- All files you select must be from a single backup. Check the Backup ID to determine a file's backup. If you select files from multiple backups, the Save button becomes disabled.

- Softlinks cannot be downloaded and are not included in the search results.

8 Click **Save**.

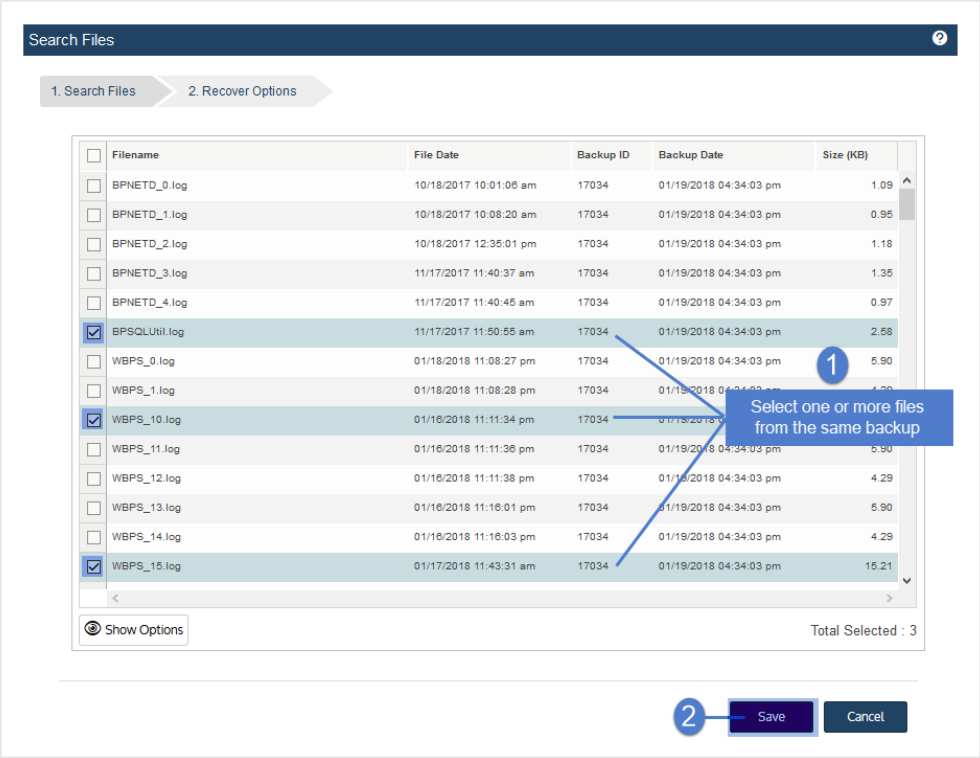

9 The message *Starting File Level Recovery on the Target* displays, indicating that the recovery has started.

10 To verify that the download starts, leave the message dialog open and view the status messages. Job status changes from Queued to Active to Downloading.

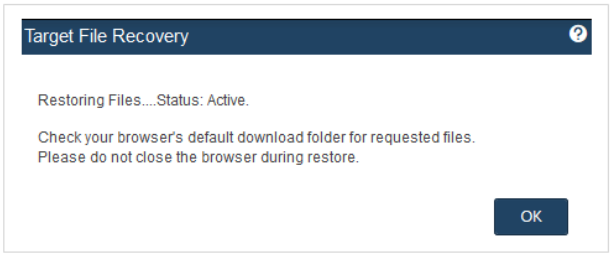

11 A *.zip* file containing the recovered files is placed in the default download location of the browser where the source appliance UI is running.

12 When the download completes, the *Unitrends-Restore.zip* file displays in the browser. Open the *.zip* file to access the recovered files.

Notes:

- The duration of the download is impacted by various factors, such as the size of the files, bandwidth, and download speed.
- Persistent browser and UI sessions are required to create the *.zip* file in the browser's default download location. If you close the browser or UI session during the recovery, you must run a new job.

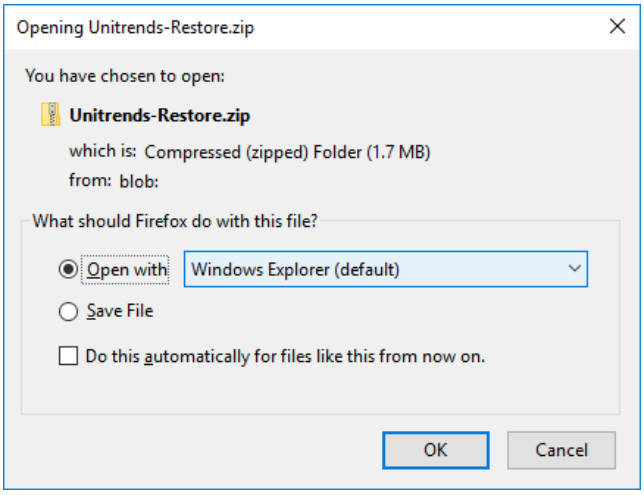

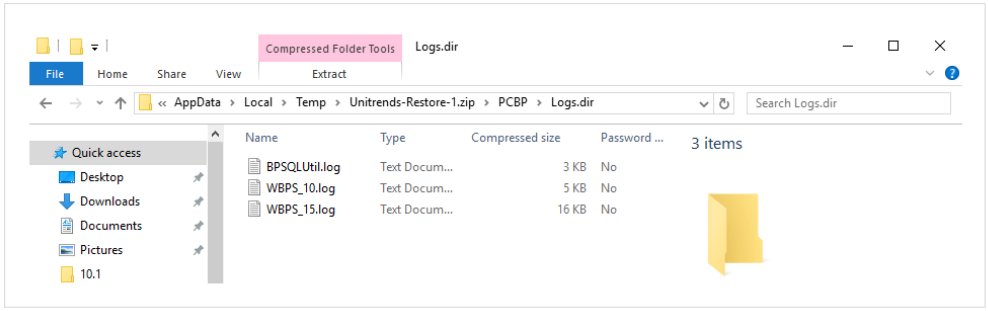

Recover Windows Active Directory information

Active Directory (AD) database and SYSVOL backups are included with the system state in Windows file-level backups. To recover AD information from a file-level backup, special steps are needed. Use this procedure:

1. Boot the AD server into Directory Service Restore Mode.
2. Recover the last incremental or full backup to the AD server as described in "To recover an entire file-level backup by using the Backup Catalog" on page 945.

   The Unitrends agent performs a non-authoritative restore of the AD database.
3. You can proceed with authoritative restore by using the **ntdsutil.exe** command. Follow the best practices around Active Directory Authoritative restores described in the Active Directory Domain Services Operations Guide at https://msdn.microsoft.com/en-us/library/bb727048.aspx.

Note: In certain circumstances, you may wish to recover NTDS or SYSVOL files for a domain controller to an alternate location, and then use Microsoft Utilities to recover from these restored files. You can recover system files to an alternate SystemState.dir location while Active Directory is running without recovering the entire backup. Use the procedure "To browse one file-level backup and recover files by using the Backup Catalog" on page 933.

Recover a Windows cluster database

On Windows servers where the cluster service is active, the cluster database is included with the system state in Windows file-level backups.

A cluster database only needs to be recovered if all nodes in a cluster lose their copy of the database. (If just one node loses the database, the lost database will be restored from another node in the cluster.)

If you need to recover the cluster database, use these steps:

1. Recover the Windows file-level backup as described in "To recover an entire file-level backup by using the Backup Catalog" on page 945.
2. The cluster database (CLUSDB) is recovered to *C:/PCBP/SystemState.dir*. (The exact path may differ if the Unitrends agent was installed to a directory other than the default).
3. From the *PCBP* directory on the Windows server, run the utility **cdrestore.bat** to recover the cluster database. If a cluster database file exists in *C:/PCBP/SystemState.dir*, the utility shows the date that the database was recovered and prompts you to ensure that the Cluster Service (ClusSvc) is stopped on all nodes of the cluster before continuing.
4. Press **y** to continue. The database is recovered. (If the first attempt to restore the database fails, the utility displays the steps required to clear the existing clustering hive from the registry hive so that the database file may be restored.)
5. After the database has been recovered, restart the cluster service.

Windows file-level replicas

The Windows file-level replica feature (formerly known as *Windows instant recovery*) provides a quick way to recover a failed physical Windows asset. It creates a virtual machine replica of the Windows machine, then keeps this replica up-to-date by applying backups of the original asset as they run. In the event of a disaster, you can bring this replica online to immediately assume the role of the failed asset.

To use the feature, simply set up the replica by using the Create Windows Replica dialog. The appliance then creates the replica VM from the most recent backup of the Windows asset, and automatically applies all subsequent backups. Because the replica is continually updated, it is ready for production use at any time.

While creating the replica, you specify the location where the replica VM will reside. The replica can reside on:

- A Recovery Series physical appliance
- A Recovery MAX physical appliance
- An ESXi host

- A Hyper-V server

The replica VM is created as a cold stand-by in the specified location. The replica is powered off and has no network connectivity. Because the replica remains powered off even as backups are applied, it consumes no compute resources.

After the first backup has been applied, replica creation is complete. You can then do the following as needed:

- Audit the replica to verify the integrity of the machine and its data and applications. In audit mode, the replica runs on a private network (inaccessible from the production network). This enables you to check the replica machine while the original Windows server is still operating in production. It is recommended that you periodically audit the replica to ensure it functions as expected.
- Bring the replica online in your production environment to immediately assume the role of the original server. Because the live replica consumes appliance resources, it is intended as a temporary replacement until you can perform a bare metal recovery to restore the failed Windows asset to new hardware. (Or, if the replica resides on an ESXi or Hyper-V server, you can opt to use the replica VM as a permanent replacement.)

See the following topics for details on using the Windows replicas feature:

- "Windows file-level replica requirements" on page 994
- "Setting up a Windows file-level replica" on page 1008
- "Working with Windows file-level replicas" on page 1013

Windows file-level replica requirements

The following topics cover the requirements for Windows file-level replicas:

- "Backup requirements"
- "Replica requirements" on page 995
- "Requirements for protected Windows asset" on page 1002

Notes:

- Only one Windows replica can exist per Windows asset. You cannot run both an image-level replica and a file-level replica of the same asset at the same time. If a replica exists, you must tear it down before creating another for the asset.
- You can opt to run file-level replicas in the Unitrends Cloud. Contact your Account Manager for assistance.

Backup requirements

An agent-based, file-level backup of the physical Windows machine is required to create the replica. (To use a Windows image-level backup, see "Windows image-level replicas" on page 1086.)

- The backup must contain:
  - System files and folders.

  - The system state, which includes the registry, IIS metabase, COM+ certificates, active directory information, and other key components.
  - Disk metadata and layout, file system configurations, and other hardware-related information.
- The backup must be a local backup run by the appliance where you are creating the Windows replica. You cannot create a replica from an imported backup or from a backup copy.

  Notes:

  - By default, file-level backups include all system information needed to create a Windows replica. If you opt to exclude data from backup, use care not to exclude the system state and the boot and critical system (OS) volumes.
  - The replica virtual machine is created based on the backup you select. Volumes that were excluded from backup are not recovered.
  - For SQL, the master, model, and msdb system databases must also be present in the file-level backup of the Windows asset. (These are included by default. If you want the Windows replica to include a hosted SQL application, use care not to exclude these system databases from the file-level backup.)

Back up the Windows asset regularly to keep the replica up to date. See these topics for details:

- To create a job manually, see "To create a file-level backup job" on page 437.
- To create a job by using an SLA policy, see "To create an SLA policy for Windows and Linux file-level assets" on page 537.
- For a comparison of the manual and SLA policy job creation methods, see "About creating backup and backup copy jobs" on page 426.

If the Windows asset is hosting SQL or Exchange, you can configure the replica to include these applications. In this case, SQL or Exchange application backups are also required. See "SQL and Exchange" on page 1004 for details.

Replica requirements

Requirements vary by where the replica resides. See these topics for details:

- "Choosing the replica location"
- "Requirements for running a replica on a Recovery Series or Recovery MAX appliance" on page 996
- "Requirements for running a replica under VMware ESXi" on page 999
- "Requirements for running a replica under Hyper-V" on page 1000
- "Additional requirements for running a replica in a Hyper-V cluster environment" on page 1001

Choosing the replica location

The replica can reside on any of the following: a Recovery Series physical appliance, a Recovery MAX physical appliance, an ESXi server, or a Hyper-V server. Considerations for each are given in the following table.

Replica location	Considerations
Recovery Series or Recovery MAX physical appliance	Running the replica on the backup appliance itself provides a simple, seamless solution. Note that only Recovery Series and Recovery MAX physical appliances have this option. (A replica cannot reside on a Unitrends Backup virtual appliance.) Consider the following to determine whether this is the optimal replica location for your environment: • Requires no additional hardware. Provides near-zero RTO without having to increase CapEx. • Assumes the role of a failed Windows server temporarily, until you can get new hardware and run a bare-metal recovery. • Provides seamless, continuous protection in a failover scenario. The 'live' replica is automatically protected by the existing backup schedule after assuming the identity of the original Windows machine. • Uses the compute resources of the appliance to bring the replica 'live' in a failover scenario. If the appliance is already under high load, there may not be sufficient resources to provide for replicas. • Reduces on-appliance backup retention because a portion of the appliance's storage is reserved for the replica.
ESXi or Hyper-V server	Running the replica on an ESXi or Hyper-V server enables access to the compute and storage in your virtual environment, greatly increasing the pool of resources that can be used for replicas. Consider the following to determine whether this is the optimal replica location for your environment: • Leverages virtual infrastructure for replicas. Disk space and compute resources of the Unitrends appliance are not impacted. • Provides the ability to dynamically scale compute resources in the virtual infrastructure during failover, enabling 'live' replica performance to match that of the original Windows asset. • Provides the option to use the replica VM as a permanent replacement for the failed Windows asset. Can be used to migrate a physical Windows machine to your virtual infrastructure. • Can be used in copy data management jobs for automated failover. For details, see "Recovery Assurance" on page 1263.

Requirements for running a replica on a Recovery Series or Recovery MAX appliance

Ensure that the following requirements have been met before you create the Windows replica.

Requirement	Description
Allocate storage	Appliance storage can be used to store backups, for VM instant recovery, or for Windows replicas. To use the Windows replica and VM instant recovery features, you must allocate a portion of the appliance storage to Instant Recovery in the Edit Storage dialog (as described below). Before allocating storage, determine the percentage to use for backups and the percentage to reserve for Windows replicas and VM instant recovery. Once storage is allocated to instant recovery, it can be used only for Windows replicas and VM instant recovery. The storage is reserved and cannot be used for other purposes, such as backups or deduplication (but you can modify your storage allocation at any time). Because the appliance is designed to retain as many local backups as possible, it is best to reserve instant recovery space soon after initial deployment. If necessary, you can allocate instant recovery space later, but doing so may require the appliance to purge local backups to make room for the newly allocated instant recovery space. **Note:** You do not need to allocate instant recovery storage if the Windows replica will reside on an ESXi or Hyper-V server. This procedure is only required for Windows replicas that reside on the appliance itself. **To allocate storage for Windows replicas** 1 On the **Configure > Appliances** page, select the appliance, then click the **Storage** tab below. 2 Select the **Internal** storage and click **Edit**. 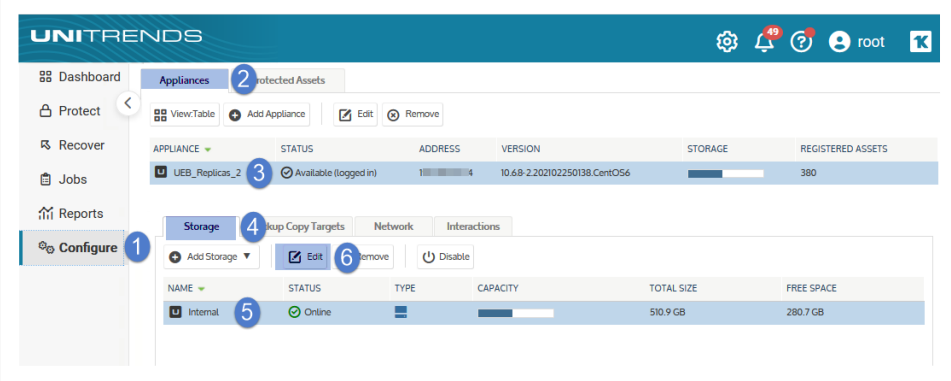 3 Modify the storage percentages allocated to Backups versus Instant Recovery (IR). The minimum IR space needed for a replica is the total amount of space in use on the original asset (the sum of used space on all disks). 4 Click **Save**.

Requirement	Description
	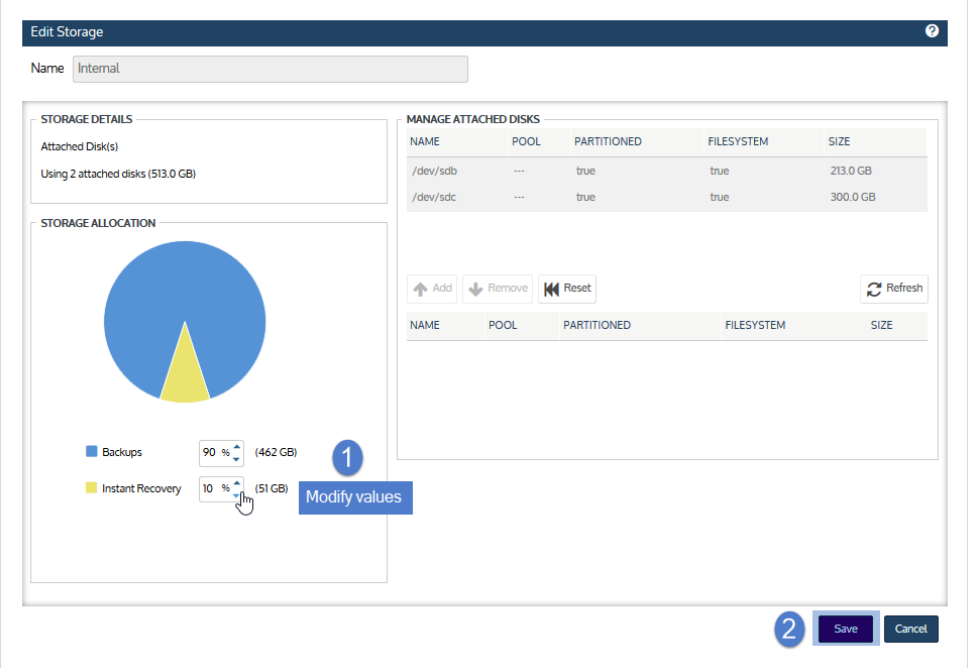
Unitrends appliance resources	The replica uses part of the Unitrends appliance's processors, memory, and storage. This usage may impact the performance of regular system functions (such as backups, backup copies, deduplication, and purging). Monitor the appliance closely and make adjustments as necessary. Minimum replica requirements: • One replica requires a minimum of 2048 MB of memory. • One replica requires a minimum of 1 core processor when in an idle state.
UEFI-based assets	Replicas of UEFI-based assets are supported on the following: • ESXi servers • Hyper-V servers • Recovery Series appliances running the CentOS 7 operating system • Recovery MAX appliances running the CentOS 7 operating system Replicas of UEFI-based assets are NOT supported on physical appliances running the CentOS 6 operating system. To check the appliance OS version, select **? > About** in the appliance UI:

Requirement	Description
	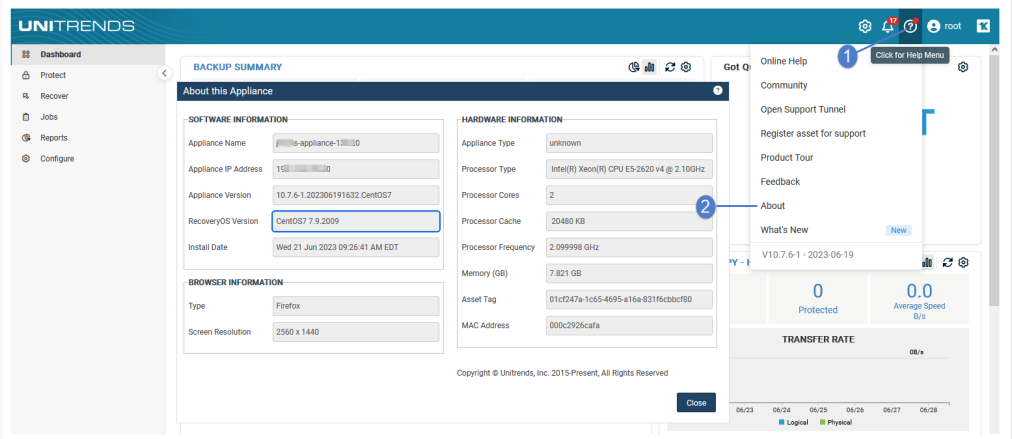
GPT-partitioned assets	Instant recovery is supported for GPT-partitioned volumes.

Requirements for running a replica under VMware ESXi

Ensure that the following requirements have been met before you create the file-level replica.

Requirement	Description
Hypervisor version	The ESXi host must meet these requirements: • Must be running ESXi 5.1 or a higher version listed in the Compatibility and Interoperability Matrix. • Must support the operating system (OS) of the Windows asset. (See the VMware documentation for details.) For example, a replica of a Windows 2016 asset cannot reside on an ESXi 5.1 host.
Virtual host asset	The ESXi server must be added to the appliance as an asset. See "Adding a virtual host" on page 308.
Compute	One replica requires a minimum of 2048 MB of memory. (This number must be a multiple of 4.)
Replica VM changeability	Once you have configured the replica in the Create Windows Replica dialog, do not make any changes to the replica VM. Any alteration to the replica (unless it is in *live* mode) may lead the replica to an inconsistent state.

Requirement	Description
Maximum disk size	The maximum disk size is capped by what the hypervisor supports. The replica's disks will be the same size as those on the original asset. For Windows assets with disks larger than 2 TB, the ESXi server must be running ESXi 5.5 or a higher version listed in the Compatibility and Interoperability Matrix.
Virtual hardware version	The replica VM is configured with the highest hardware version that the hypervisor supports.
Hosted SQL application	If the file-level replica will include a SQL application, the original asset cannot contain a *D* drive letter. If there is a *D* drive on the original asset, SQL restores will fail to process during replica restores. To resolve this issue, use image-level or VMware replicas instead, or change the drive letter of the original asset to no longer use drive *D*.

Requirements for running a replica under Hyper-V

Ensure that the following requirements have been met before you create the file-level replica.

Requirement	Description
Hypervisor version	The hypervisor must be one of the following: • A Windows Server with the Hyper-V role enabled, running 2008 R2 or a higher version listed in the Compatibility and Interoperability Matrix. • A Hyper-V Server running 2008 R2 or a higher version listed in the Compatibility and Interoperability Matrix. The Hyper-V host must support the operating system (OS) of the Windows asset. (See this Microsoft article for details: Should I create a generation 1 or 2 virtual machine in Hyper-V?) For example, a replica of a Windows 2016 asset cannot reside on a Hyper-V 2008 R2 host.
Host agent version	The Hyper-V host must be running Unitrends agent version 9.2 or higher for Windows 2016, 10.4.4 or higher for Windows 2019, and 10.6.2 or higher for Windows 2022. It is best practice to run the latest Unitrends appliance and agent software versions. Older versions do not support all current Unitrends features. For example, to protect SQL Always On availability groups, the appliance and Windows agent must be running release 10.0.0-2 or later.
Virtual host asset	The Hyper-V host must be added to the Unitrends backup appliance as a protected asset. See "Adding a virtual host" on page 308.
Virtual host Samba access	The Hyper-V server must be able to access the appliance's Samba share:

Requirement	Description
	• SMB 2.0 – The SMB 2.0 security option is enabled by default on Unitrends appliances that were originally imaged or deployed with version 10.4.8 or higher. SMB 2.0 must be enabled on the Hyper-V server. • SMB 1.0 – The SMB 1.0 security option is enabled by default on appliances that were originally imaged or deployed with a pre-10.4.8 version. SMB 1.0 must be enabled on the Hyper-V server. **Note:** Upgrading from a pre-10.4.8 version does not change the SMB 1.0 setting. (To configure your appliance to use SMB 2.0, see How Unitrends supports SMBv2.)
Compute	One replica requires a minimum of 2048 MB of memory. (This number must be a multiple of 4).
Replica VM changeability	Once you have configured the replica in the Create Windows Replica dialog, do not make any changes to the replica VM. Any alteration to the replica (unless it is in *live* mode) may lead the replica to an inconsistent state.
Maximum disk size	The maximum disk size is capped by what the hypervisor supports. The replica's disks will be the same size as those on the original asset. For Windows assets with disks larger than 2 TB, the Hyper-V server must be running version 2012 or a higher version listed in the Compatibility and Interoperability Matrix.
Replica VM configuration	The replica VM is created with this configuration: • The asset's firmware interface type determines the generation of the replica VM. BIOS-based assets are created as generation 1 VMs, and UEFI-based assets are created as generation 2 VMs. A replica for a UEFI-based asset cannot run on 2008 R2. • The VM's configuration version is the highest version that the hypervisor supports.
Pass-through disks	Pass-through disks are supported. After you bring the replica online to assume the role of the failed asset, you must refresh and reconnect any existing iSCSI targets.

Additional requirements for running a replica in a Hyper-V cluster environment

Ensure that the following requirements have been met before you create the Windows replica.

Requirement	Description
Cluster asset	To run a replica on a Hyper-V server in a cluster configuration, ensure that these requirements have been met: • The Unitrends Windows agent is installed on each node in the cluster. • Every node in the cluster is running the same agent version. • Each cluster node and the cluster itself has been added to the appliance as an asset. (For details, see "Working with Hyper-V servers" on page 664.)
Storage	The Hyper-V cluster must be configured with Cluster Shared Volumes (CSVs). SMB storage is not supported.
PowerShell FailoverClusters modules	These modules must be installed on every node in the cluster so that the appliance can discover the CSVs.
Selecting the replica location	To create a clustered replica, you must select the cluster itself as the Location in the Create Windows Replica dialog. Do not select an owner node. If you select an individual node in the cluster, the replica will not be clustered.
Network switch selection	For a clustered replica, select the Network Switch common to all nodes in the cluster (in the Create Windows Replica dialog). If you do not select this switch, a 'live' replica that fails over to another node will lose network connectivity.
2008 R2 clusters	To run the replica on 2008 R2 servers in a cluster configuration, enable DCOM and WMI Virtualization access for all nodes in the cluster. For instructions, see Security settings for creating a clustered virtual failover client on Hyper-V server 2008 R2.
Live migration interoperability	During live migration of a clustered replica, the Unitrends appliance cannot apply backups to the replica, verify or audit the replica, or bring the replica online in production to assume the role of the original asset. If the appliance attempts to apply a backup or verify the replica during a live migration, the appliance waits several minutes and then attempts the operation again. If you try to audit the replica or bring it online in production, the appliance notifies you that it cannot run the operation because of the migration and you must try again later.

Requirements for protected Windows asset

The Windows asset must meet the following requirements to use the file-level replica feature:

Requirement	Description
Client Operating Systems	The file-level replica feature is supported for the client operating systems listed below. Additional version limitations apply. See the Compatibility and Interoperability Matrix for details. • Windows XP, 64-bit (SP2 and later) • Windows Vista, 64-bit (SP2) • Windows 7, 64-bit • Windows 8, 64-bit • Windows 8.1, 64-bit • Windows 10, 64-bit • Windows 11, 64-bit **Note:** To run a replica on a virtual host, the Hyper-V or ESXi host must support the guest OS of the replica VM. (See the Microsoft or VMware documentation for details.) For example, a replica running Windows 10 cannot reside on ESXi 5.1 or Hyper-V 2008 R2.
Server Operating Systems	The file-level replica feature is supported for the server operating systems listed below. Additional version limitations apply. See the Compatibility and Interoperability Matrix for details. • Windows 2003, 64-bit (SP2) • Windows 2003 R2, 64-bit • Windows Small Business Server 2003 and later, 64-bit • Windows 2008, 64-bit • Windows 2008 R2, 64-bit • Windows 2012, 64-bit • Windows 2012 R2, 64-bit • Windows 2016, 64-bit • Windows 2019, 64-bit • Windows 2022, 64-bit

Requirement	Description
	Note: To run a replica on a virtual host, the Hyper-V or ESXi host must support the guest OS of the replica VM. (See the Microsoft or VMware documentation for details.) For example, a replica running Windows 2016 cannot reside on ESXi 5.1 or Hyper-V 2008 R2.
SQL and Exchange	If the Windows asset is hosting SQL or Exchange, you can configure the replica to include these applications. The following requirements and limitations apply: • The SQL application must be one of these versions: 2005, 2008, 2012, 2014, 2016, 2017, 2019, or 2022. • The Exchange application must be one of these versions: 2003, 2007, 2010, 2013, or 2016. • Windows server failover clusters (including Exchange DAGs, SQL clustered instances, and SQL availability groups) are *not* supported. You can create replicas to protect the Windows servers that host these applications, but the replicas do not include the applications. • Other applications, such as Hyper-V, Oracle, and SharePoint, are *not* supported. You can create replicas to protect the Windows servers that host these applications, but the replicas do not include the applications. • A successful application backup of each database must reside on the Unitrends appliance. (If there is no backup, the replica includes the application instance, but no database files are available.) • While setting up the replica, you must manually select each database to include by checking boxes on the Exchange or SQL tab in the Create Windows Replica dialog. **Notes:** If you are running both image-level and file-level backups of the Windows asset and have configured image-level backups to use the application aware setting, the following apply to file-level replicas of the asset: • Once an asset is configured with the **Allow application aware** setting, the replica add database option is not supported and any existing SQL and Exchange schedules are disabled. • When selecting a database for the replica, if you receive an error indicating that the asset has been configured for application aware image-level backups, you cannot add a database to the replica. • Delete any existing replica of an application-aware asset if the replica includes SQL or Exchange data. You can then recreate the replica without

Requirement	Description
	adding SQL or Exchange. (If you do not delete the replica, the replica's application data will become stale over time since there will be no new application backups to restore to the replica.) • To include hosted SQL or Exchange data, consider running image-level replicas instead. • A replica's applications are kept up to date by applying application backups. Be sure to run application backups regularly. For details on creating an application backup, see "To create an Exchange backup job" on page 478 or "To create a SQL backup job" on page 483. • For SQL, the master, model, and msdb system databases must be present in the agent-based, file-level backup of the Windows host asset. (These are included by default. Be sure you do not exclude them when creating the backup schedule for the Windows asset.) These system databases must be present to access SQL databases that are available and running when the replica enters *live* mode.
Other applications	Some Windows applications require network access and/or rely on underlying hardware, like network interface MAC addresses, in order to run properly. When booting a Windows replica in audit mode, there is no network interface, so applications requiring network connectivity will not function properly. This is expected behavior. When booting a Windows replica in live or audit mode, applications that rely on unchanging hardware (like MAC addresses) may not function properly or may require re-authentication, re-installation, or other special actions that are application specific in order for them to work properly. You should work with the application vendor to determine what actions are required.
Firmware interface type	The file-level replica feature supports BIOS- and UEFI-based assets. For UEFI-based assets, the replica must reside on one of the following: • An ESXi server running ESXi 5.1 or a higher version listed in the Compatibility and Interoperability Matrix. • A Hyper-V server running 2012 R2 or a higher version listed in the Compatibility and Interoperability Matrix. For a replica running on Hyper-V, the UEFI-based asset's OS must be 64-bit and running Windows 8 or higher. • A Recovery Series or Recovery MAX appliance that is running the CentOS 7 operating system. **Note:** A UEFI-based replica cannot reside on a physical appliance that is running the CentOS 6 operating system.

Requirement	Description
Disk configuration	The file-level replica feature supports Windows machines configured with basic disks and dynamic disks, as long as the boot and system disks are not dynamic. The following types support dynamic volumes configured as data volumes: • RAID 5 • Spanned • Striped • Mirrored • Simple **Notes:** • For Windows 8.1 and Windows 2012 R2, the replica includes the data from all disks, but if created as a Gen 1 VM, only the first four disks are eligible as boot devices. • For Windows 2003, the boot disk must be located on one of the first three disks if the replica resides on an ESXi host.
Disk partition type	The file-level replica feature is supported for assets with GUID Partition Table (GPT) and Master Boot Record (MBR) partitions. For assets with GPT partitions, the replica must reside on one of the following: • An ESXi server • A Hyper-V server • A Recovery Series or Recovery MAX physical appliance that is running the CentOS 7 operating system. To check the appliance OS version, simply click on **? > About**:

Requirement	Description
	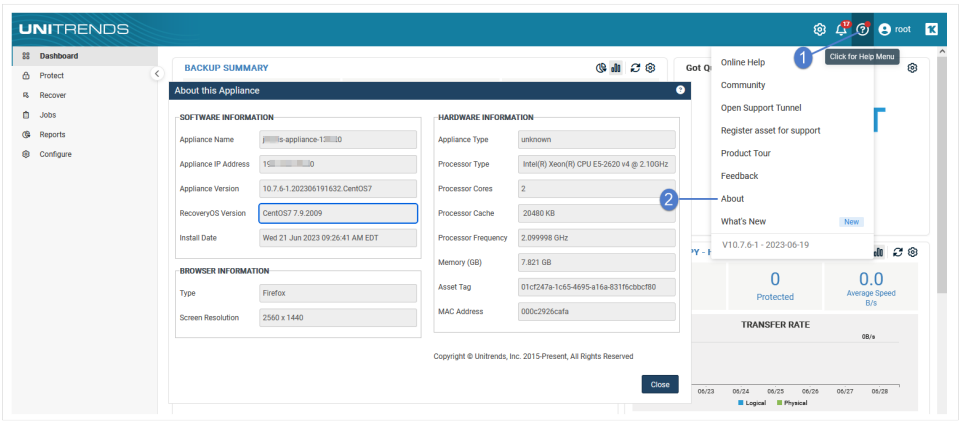
Software RAID volumes	The file-level replica feature is not supported for software RAID configurations.
Deduplicated volumes	Volumes that use Microsoft deduplication are not supported in cases where the size of the data on the volume before it has been deduplicated is greater than the physical capacity of the volume. Because data is applied to the replica in its non-deduplicated form, the volume must have enough capacity to house this non-deduplicated data.
Number of volumes	The Windows asset can have a maximum of 20 volumes, including the System Reserved volume and other unmounted volumes. A replica with more than 20 volumes may fail to boot.
Separate boot and system partitions	For Windows assets with boot and system partitions located on different disks, the system partition must reside on the first disk (Disk 0).
File System Configuration	The file-level replicas feature supports the following file systems: • NTFS • FAT/FAT32 • ReFS (Windows 2012 and later)
Active Directory	The file-level replica feature supports Active Directory database (NTDS) located on the boot volume only. (If it is not on the boot volume, the configuration is not supported and you see an error message when you attempt to create the replica.)

Setting up a Windows file-level replica

Use the following procedure to set up a Windows file-level replica. You can configure the replica to run on:

- A Recovery Series or Recovery MAX physical appliance (backup appliance or backup copy target appliance).
- An external hypervisor (Hyper-V or VMware).

Note: To run the file-level replica in the Unitrends Cloud, contact your Account Manager for assistance.

To set up a Windows file-level replica

1 Select **Recover**, then click the **Replicas** tab.

2 Click **Create Replicas** and select **Windows File-Level**.

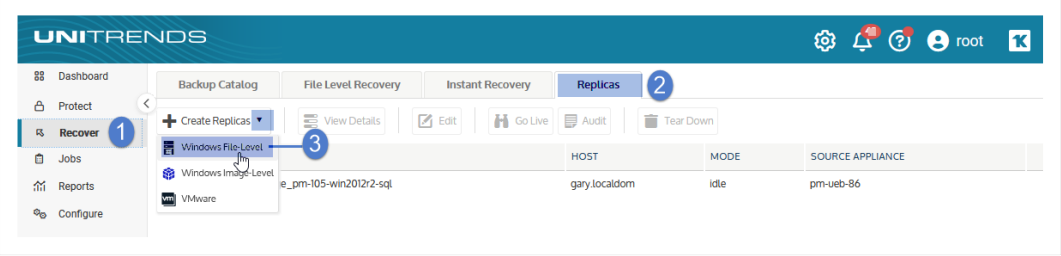

3 Select the Windows asset, then click **Next**.

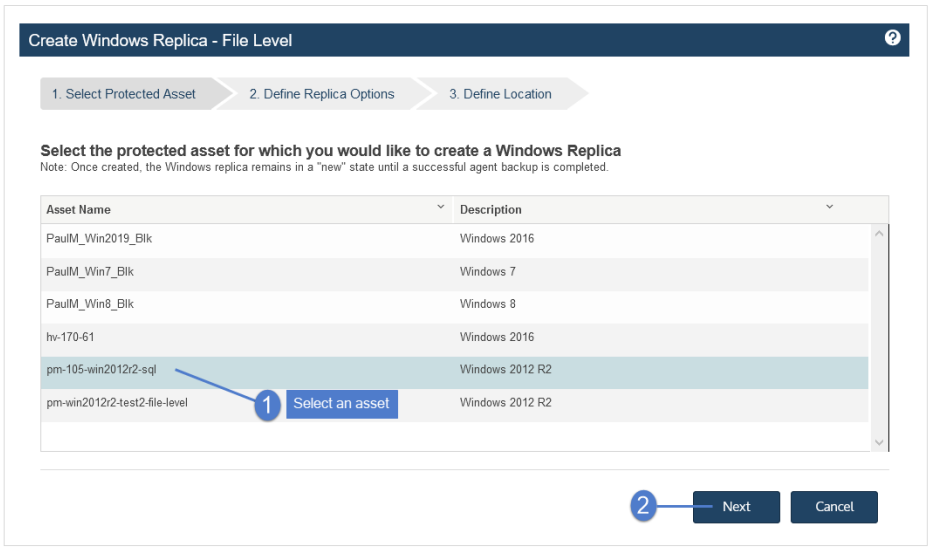

4 Enter replica options. Click **Next**. (See the table below for descriptions of the Define Replica Options fields.)

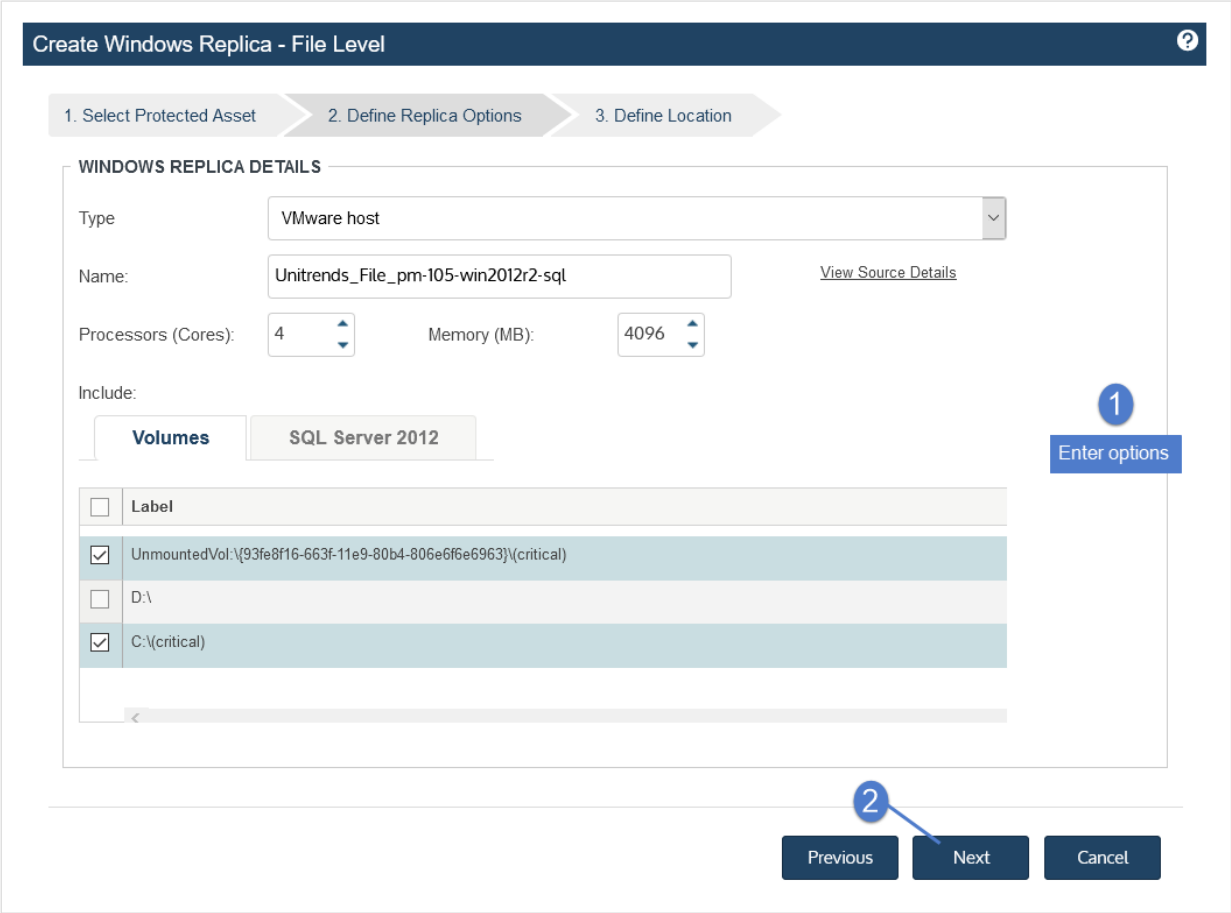

Item	Description
Type	Select a location type from the list (Unitrends Appliance, VMware Host, or Hyper-V Host). The list contains only the types that are available in your environment. For example, the Unitrends Appliance type is not an option for Unitrends Backup virtual appliances. The Hyper-V Host type is not an option if a compatible Hyper-V virtual host asset has not been added to the appliance.
Name	Replica name. By default, the replica is named Unitrends _File<WindowsAssetName>. If you are using aVMware or Hyper-V host, you can opt to edit this name.
View Source Details	Click this link to view details about the original Windows asset. Example:

Item	Description
	
Processors (Cores)	Number of processors connected to the replica. Use care when modifying this value. The compute resources do not have to match the original Windows asset, but you should allocate enough cores for the replica to temporarily replace the original asset.
Memory (MB)	Amount of memory attached to the replica. Use care when modifying this value. The compute resources do not have to match the original Windows asset, but you should allocate enough memory for the replica to temporarily replace the original asset.
Email verification report (Hyper-V or Unitrends Appliance only)	This checkbox displays only for Hyper-V and Unitrends appliance hosts. Check this box to include automated audits of the replica. If selected, the appliance audits the replica and emails a report with a screen shot of the replica running in audit mode. (For details, see "Automated audits for a Windows replica" on page 1015.)
Disks or Volumes tab	Disks or volumes to include when creating the replica. • Check boxes to select disks or volumes to include. • Disks/volumes marked as Critical are required. You can opt to exclude others.
SQL Server tab	This tab displays only if the Windows asset hosts a SQL application. Check boxes to select databases to include when creating the replica. SQL system databases (master, model, and msdb) are required. User databases are optional. Additional SQL requirements apply. See "Windows file-level replica requirements" on page 994 for details. Example:

Item	Description
	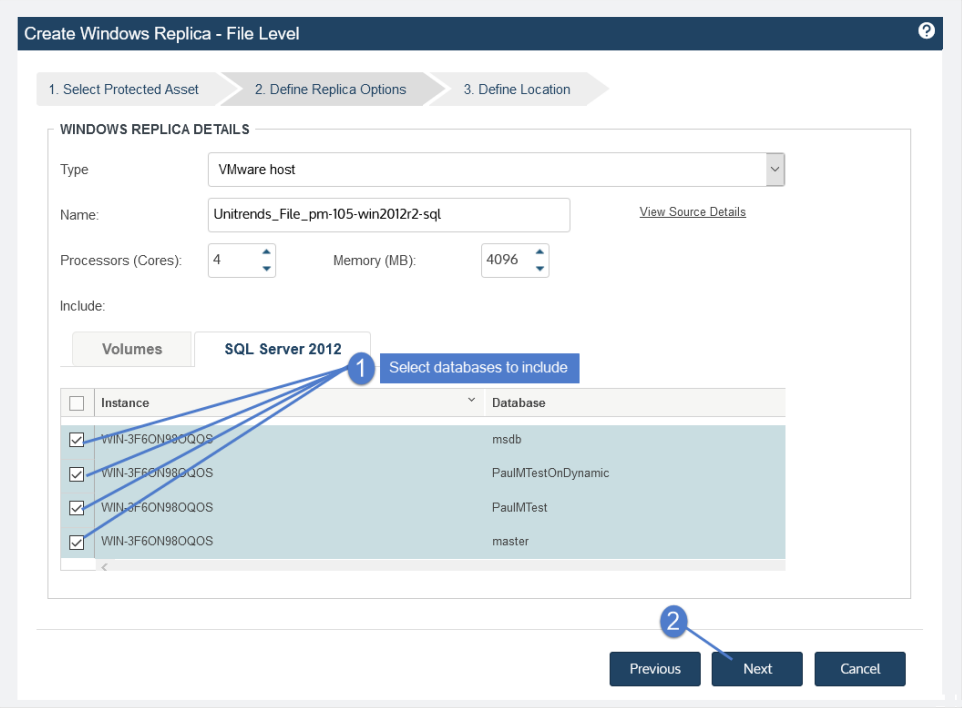
Exchange tab	This tab displays only if the Windows asset hosts an Exchange application. Check boxes to select databases to include when creating the replica. Additional Exchange requirements apply. See "Windows file-level replica requirements" on page 994 for details.

5 Enter details to specify the location where the replica will reside. Click **Save**. (See the table below for descriptions of the Define Location fields.)

The details that display in the Define Location step vary by Type selected:

- If you select **Unitrends Appliance** in the Type list, you do not provide any more location details.
- If you select **VMware Host** or **Hyper-V Host** in the Type list, provide details for the virtual host server.

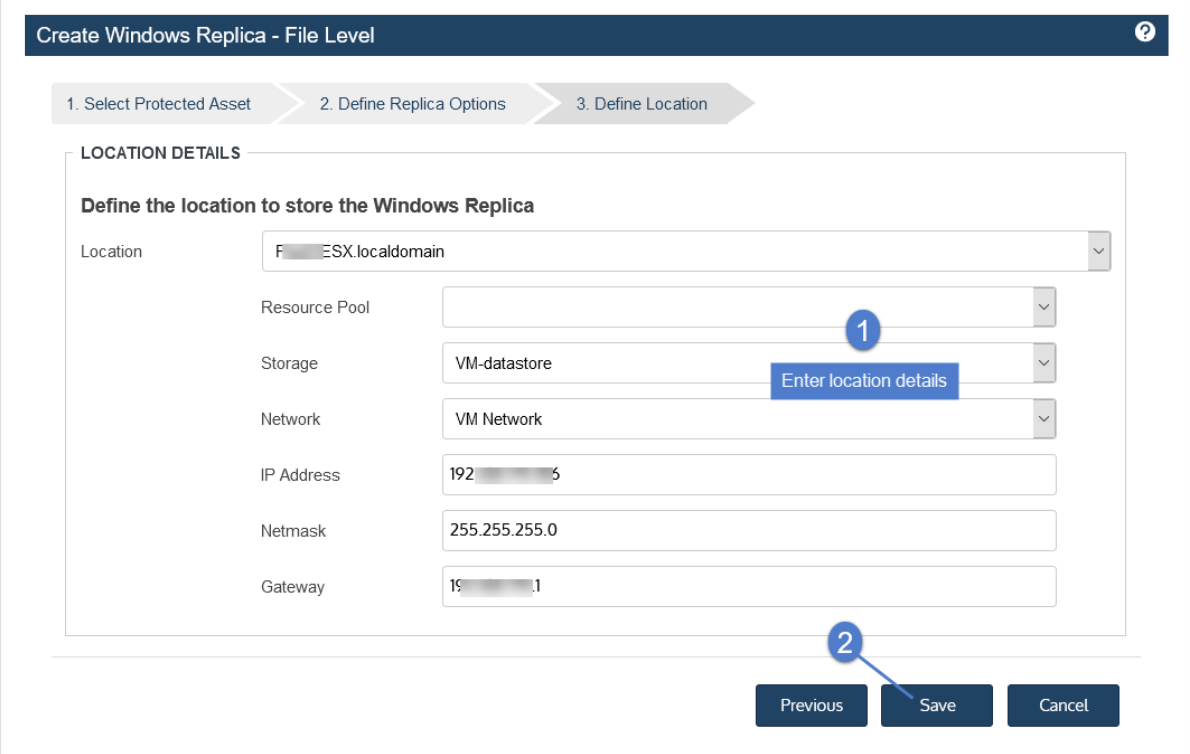

Item	Description
Location	Select a Location from the list (Unitrends Appliance, VMware Host, or Hyper-V Host). The list contains all VMware or Hyper-V virtual host assets that have been added to the appliance and are compatible with the Windows asset. For example, an ESXi 5.1 host does not display if creating a Windows 2016 replica. (For details on adding a virtual host, see "To add a virtual host asset" on page 311.)
Resource Pool (VMware only)	(Optional) If your VMware environment has resource pools, you can opt to select one in the list.
Storage	Select the datastore (VMware) or volume (Hyper-V) that will be used to create the replica VM's disks.
Network	Select a virtual network from the list. The list contains the virtual networks that are discovered and available on the VMware or Hyper-V host.

Item	Description
Network Switch (Hyper-V only)	For Hyper-V hosts only, select a network switch. The list contains the network switches that are discovered and available on the Hyper-V host.
IP Address, Netmask, and Gateway	Use these fields to define the network configuration that the appliance will use to create the replica and to apply backups. **WARNING!** Do NOT enter the IP address of the original Windows asset. Be sure to specify an IP address that is not used by another machine in your environment. When you audit the replica or bring it 'live' in production, the replica assumes the IP address, netmask, and gateway of the original Windows asset and not the configuration specified here.

The appliance creates a replica for the selected Windows asset, then applies the latest backup. The replica is created in *new* mode.

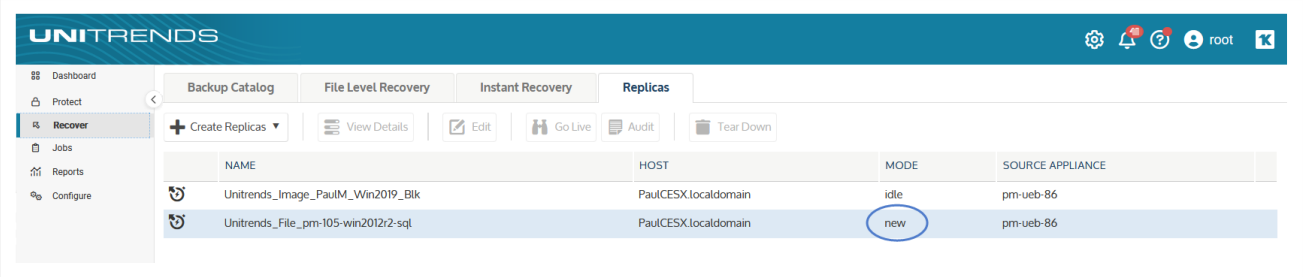

Its mode changes to *restore* while the backup is being applied, then to *idle*. After the replica enters idle mode, you can audit the replica or bring it 'live' as needed. We recommend that you audit the replica soon after it enters Idle mode, to verify its integrity. See "Working with Windows file-level replicas" on page 1013 for details.

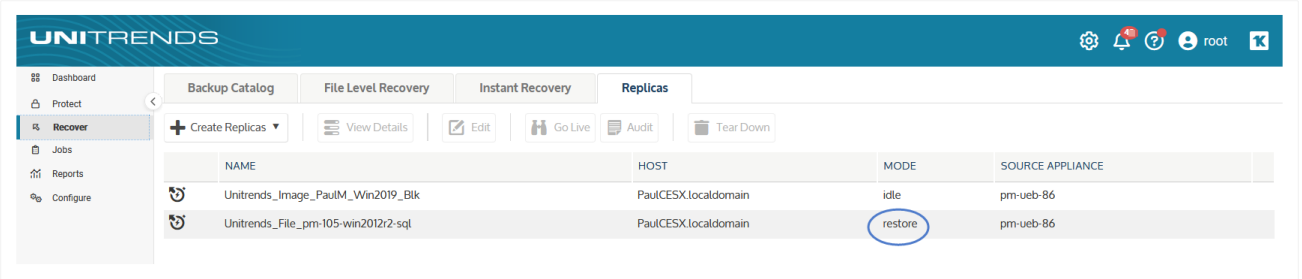

Working with Windows file-level replicas

After setting up a Windows file-level replica, use the following procedures as needed:

Editing a Windows replica

After creating a replica, you can modify the following settings at any time:

- The number of processors connected to the replica.
- The amount of memory attached to the replica.
- Whether to apply new backups to the replica (*Enable virtual restores* checkbox).
- Whether to perform automated audits of the replica (*Email verification report* checkbox, Hyper-V and Unitrends Appliance only).
- Which hosted Exchange databases to include.
- Which hosted SQL databases to include.

To edit a Windows replica

1 Select **Recover**, then click the **Replicas** tab.

2 Select the replica, then click **Edit**.

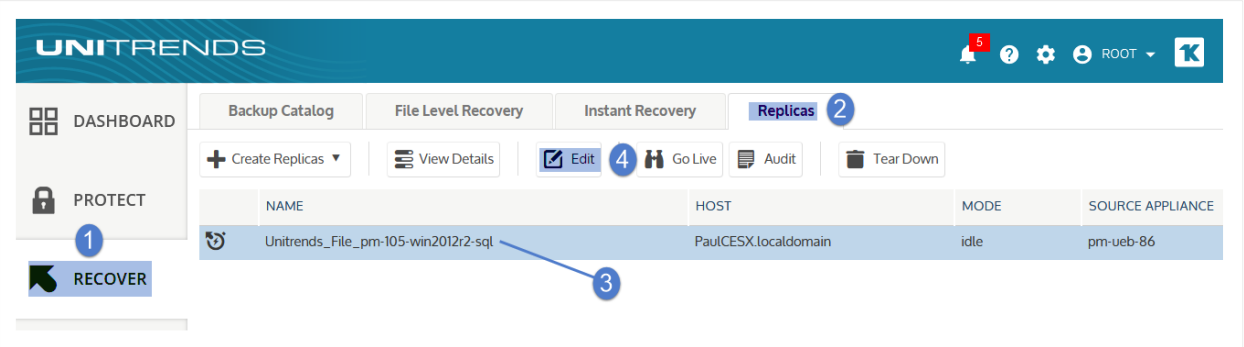

3 Modify settings as desired. Click **Save**. (The settings that display vary by Windows asset. For example, the SQL tab does not display if the Windows asset does not host a SQL application.)

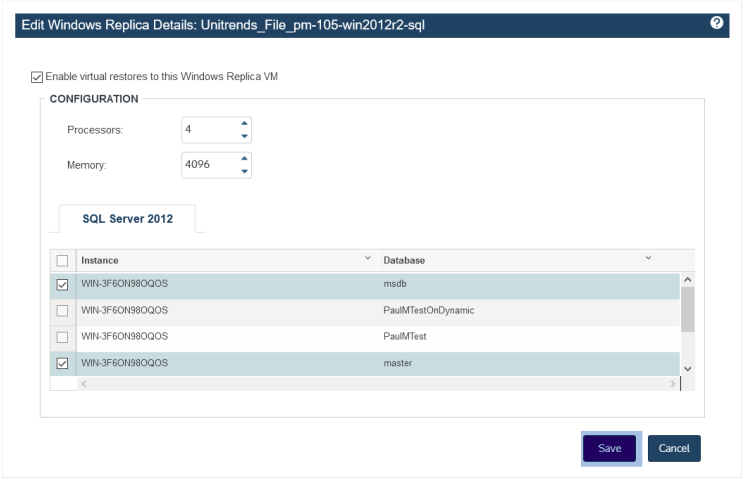

Auditing a Windows replica

Audit mode enables you to run the replica on a private network while the original Windows asset is still operating in production. A replica running in audit mode boots with no network interface. Auditing the replica with the original asset still online does not result in network conflicts or impact the original asset in any way. However, applications on the replica that require network access do not function fully in audit mode.

Note: Some Windows applications require network access and/or rely on underlying hardware, like network interface MAC addresses in order to run properly. When booting a Windows Replica in audit mode, there is no network interface, so applications requiring network connectivity will not function properly. This is expected behavior. When booting a Windows Replica in live or audit mode, applications that rely on unchanging hardware (like MAC addresses) may not function properly or may require re-authentication, reinstallation, or other special actions that are application specific in order for them to work properly. You should work with the application vendor to determine what actions are required.

It is recommended that you audit each newly created replica to ensure it functions as expected, and that you perform additional audits at regular intervals to check subsequent recovery points. You can perform manual audits for all Windows replicas. You can also set up automated audits for replicas that reside on a Unitrends appliance or Hyper-V host.

A newly created replica cannot be audited until at least one backup has been applied. During the audit, no subsequent backups are applied. Upon exiting audit mode, the appliance applies any backups that completed during the audit to bring the replica up to date.

Automated audits for a Windows replica

You can automate the audit process by enabling email verification reports for a replica. The following requirements must be met to use this feature:

- The replica must reside on one of the following: a Recovery Series physical appliance, a Recovery MAX physical appliance, or a Hyper-V server. (Automated audits are not supported for replicas that reside on ESXi servers. Perform manual audits instead.)

- Email reporting must be enabled on the Unitrends appliance where you created the replica. Email must be configured with the System box checked and at least one valid recipient email address. For details, see "Email reporting" on page 117.

To enable verification reports, check the **Email verification report** box while creating or modifying a replica. (See these procedures for details: "Setting up a Windows file-level replica" on page 1008 or "Editing a Windows replica" on page 1014.)

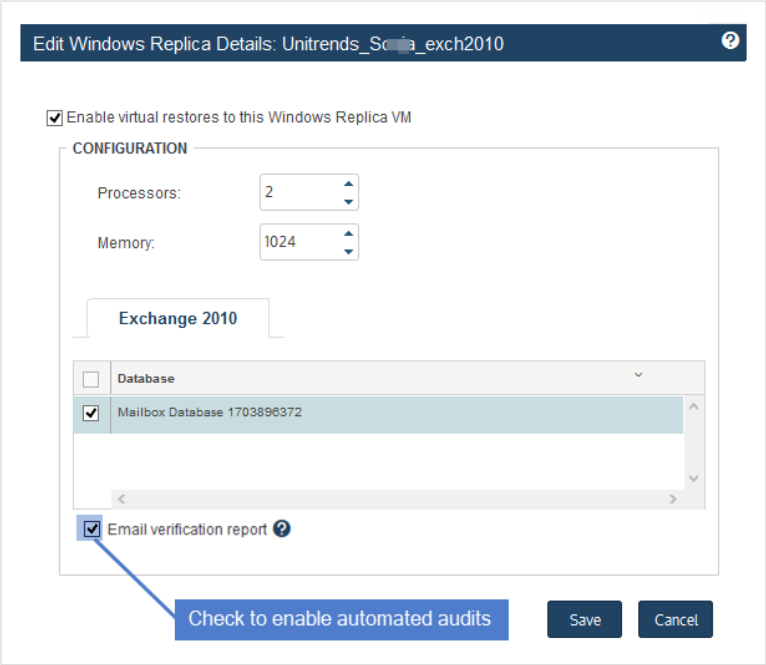

Once email verification is enabled for a replica, the appliance does the following:

- Brings the replica into audit mode after a backup has been applied.
- Takes a screenshot of the Windows login screen (after the replica has had several minutes to boot).
- Sends the screenshot to each email recipient that is configured on the appliance.

The screenshot normally shows the Windows login screen, but it can also show Windows in other boot states, including error conditions.

IMPORTANT! Always view the screenshot to make sure the replica boots correctly.

The report runs once a day, but only after a backup has been applied. If the interval between backups lasts longer than 24 hours, you will not receive a report every day. If the replica cannot boot, you will receive an email report indicating that the replica cannot be verified.

Manually auditing a Windows replica

Manually auditing the replica is a two-part process where you bring the replica into audit mode and then access the replica to verify that it is functioning as expected. During the audit, you should verify the following:

- The replica boots successfully and is operational.
- The replica contains the expected data and applications. (Note that applications requiring network access do not function fully in audit mode.)

After you have finished auditing the replica, you must take it out of audit mode so the appliance can resume applying backups. (Note that any changes made during the audit are lost upon exiting audit mode.)

Audit mode procedures

Use these procedures to manually audit the replica:

- "To bring the replica into audit mode" on page 1017
- "To access a replica on a Recovery Series or Recovery MAX appliance" on page 1018
- "To access a replica on an external hypervisor" on page 1019
- "To exit audit mode" on page 1020

To bring the replica into audit mode

1 Select **Recover**, then click the **Replicas** tab.

2 Select the replica, then click **Audit**. Click **Confirm**.

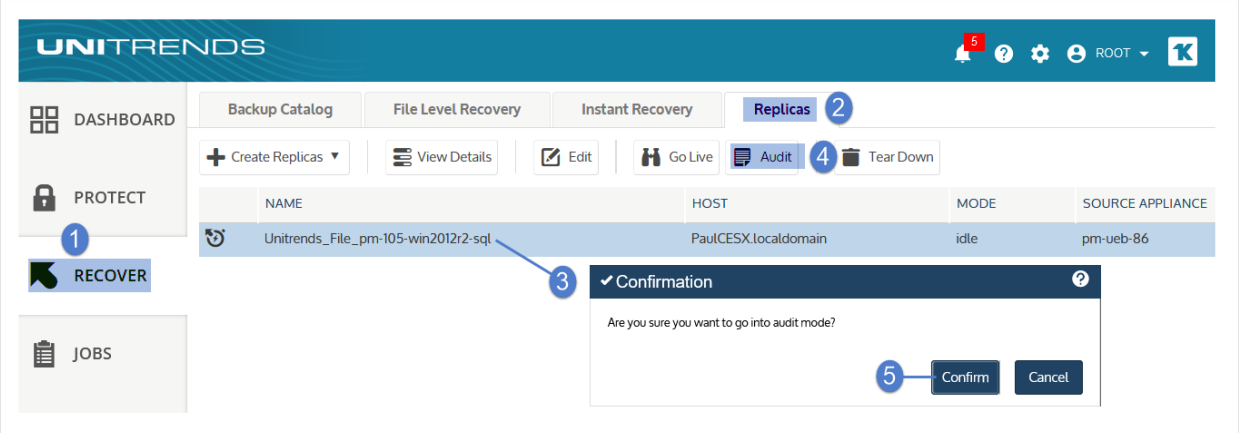

The replica's mode changes to *idle (pending audit)*, then to *audit*.

Note: If a backup is currently being applied, the replica does not enter audit mode until the restore is complete.

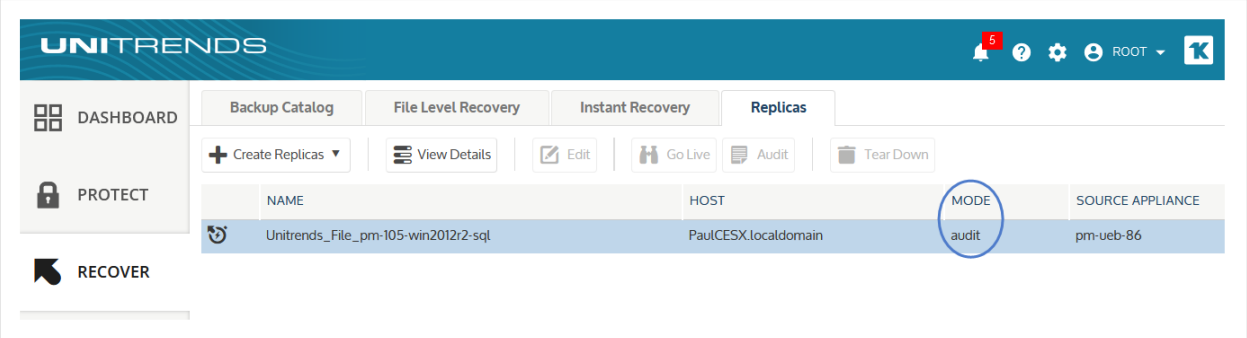

3 After the replica is in audit mode, you can connect to the replica to verify that it is functioning as expected. See the following for details:

- "To access a replica on a Recovery Series or Recovery MAX appliance" on page 1018
- "To access a replica on an external hypervisor" on page 1019

To access a replica on a Recovery Series or Recovery MAX appliance

After a Windows replica has entered audit (or live) mode, use this procedure to access the replica:

Note: You must use a VNC viewer to access the replica in audit or live mode on a Recovery Series or Recovery MAX appliance. If necessary, download one to your workstation before running this procedure.

1 On the Replicas tab, view replica details to obtain the VNC port number:

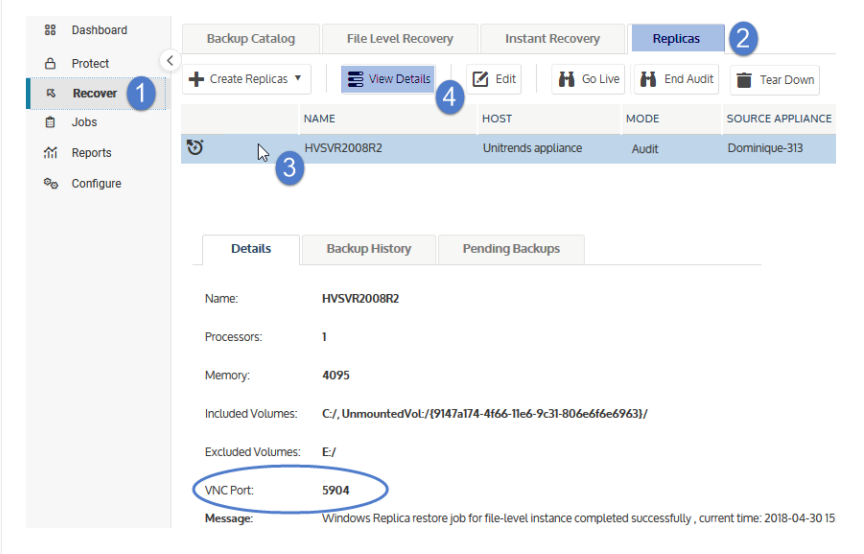

2 Open a VNC viewer. Connect to the replica by entering: <ApplianceIP>::<VNCport>

Exact field names, buttons, and syntax vary by VNC viewer. Typically, one or two colons are required betwe9en the appliance IP address and port number. An example using VNC port 5904 is given here:

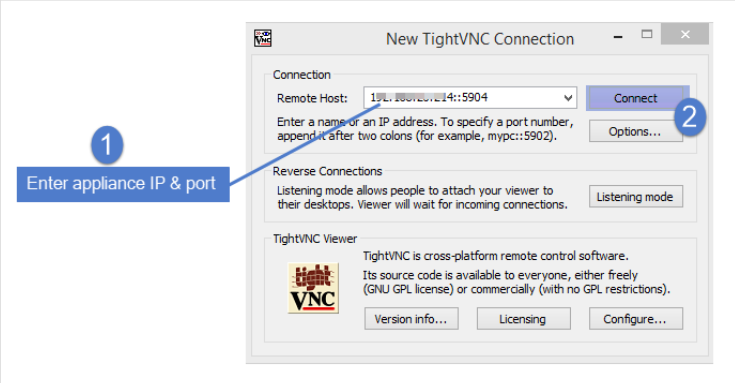

3 The Windows login screen displays, indicating the replica is available.

Note: If you access the replica before it has booted, you may see the first screen of the Unitrends Windows Integrated Bare Metal Recovery Wizard. Do not attempt to complete the steps on the bare metal screen. After several seconds, the login screen for the original Windows asset displays.

4 Log in to the replica by entering the credentials of the original Windows asset.

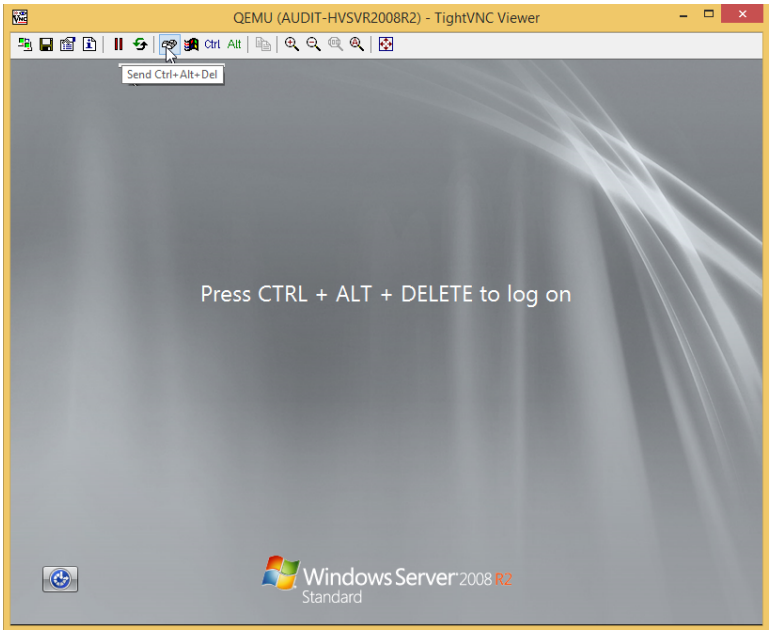

5 After verifying that the replica is running with its recovered data, proceed to "To exit audit mode" on page 1020.

To access a replica on an external hypervisor

After a Windows replica has entered audit (or live) mode, use this procedure to access the replica:

1 Connect to your hypervisor manager.

2 Locate the replica in the list of virtual machines, and access it the same way you access all VMs on the hypervisor.

3 Log in to the replica VM by entering the credentials of the original Windows asset.

> **Note:** If you access the replica before it has booted, you may see the first screen of the Unitrends Windows Integrated Bare Metal Recovery Wizard. Do not attempt to complete the steps on the bare metal screen. After several seconds, the login screen for the original Windows asset displays.

4 After verifying that the replica is running with its recovered data, proceed to "To exit audit mode".

To exit audit mode

1 Select **Recover**, then click the **Replicas** tab.

2 Select the replica, then click **End Audit**.

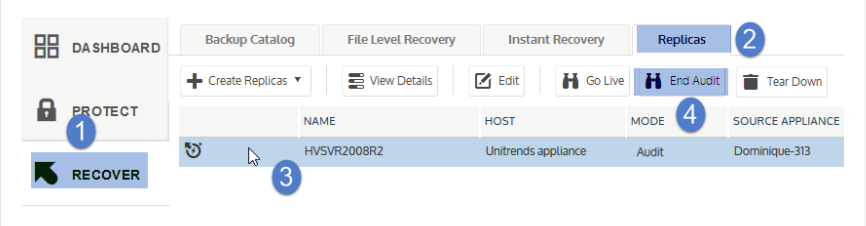

3 The replica exits *audit* mode. Its mode changes to *audit (pending: off)*, then to one of the following:

Restore – One or more backups successfully completed during the audit and the appliance is applying those backups.

Idle - The replica is idle (there are no backups to apply).

Bringing the replica live in production

If disaster strikes and the original asset fails, you can temporarily replace it with the replica by booting into live mode. Because the replica is continually updated with the original asset's data, it can immediately assume the role of the original asset.

The original asset's backup and backup copy schedules protect the replica in live mode, so that any changes made to the replica in live mode are captured under the identity of the original asset. This ensures continuity of recovery points in the asset's backup chain.

See these topics for details:

- "Live mode recommendations"
- "To bring the replica into live mode" on page 1021
- "Using the live replica as a temporary for the original Windows asset" on page 1023
- "Using the live replica as a permanent replacement for the original Windows asset" on page 1024

Live mode recommendations

Review these recommendations before going into live mode:

- Live mode should be used temporarily. The appliance begins sending alerts after a live replica has run for 14 days.
- You can exit live mode after you have recovered to new hardware (supported in all cases) or by retaining the replica as a permanent replacement (supported only for replicas that reside on external hypervisors).
- The backup schedule for the original asset protects the live replica's data.
- A live replica running on a Unitrends appliance uses appliance resources, so it is important that you recover to new hardware as soon as possible by using Unitrends bare metal recovery. (See the "Windows Bare Metal Protection and Recovery" on page 1207 for details.)
- A live replica running on an external hypervisor does not use any appliance resources. Instead, it uses hypervisor resources. The replica can replace the original asset temporarily or be used as a permanent replacement.
- Some Windows applications require network access and/or rely on underlying hardware, like network interface MAC addresses in order to run properly. When booting a Windows Replica in live or audit mode, applications that rely on unchanging hardware (like MAC addresses) may not function properly or may require re-authentication, reinstallation, or other special actions that are application specific in order for them to work properly. You should work with the application vendor to determine what actions are required.

To bring the replica into live mode

This procedure provides instructions for booting a replica in live mode. Be sure to shut down the original asset before running this procedure.

1. Select **Recover**, then click the **Replicas** tab.
2. Select the replica, then click **Go Live**.

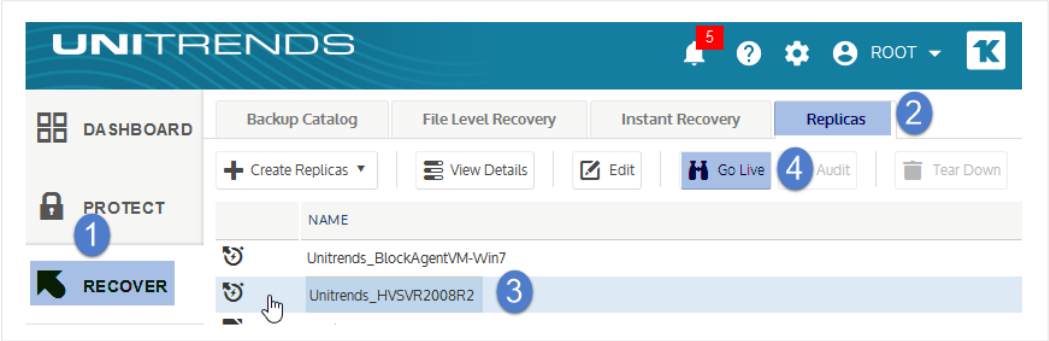

3. Click **Confirm**. The replica's mode changes to *live*.
   - If a backup is currently being applied, the replica does not enter live mode until the restore is complete.
   - Upon entering live mode, the replica assumes the identity of the original Windows asset. The replica is marked *invalid* because the replica role no longer applies.

4 Log in to the replica by using one of these methods:

Note: If you access the replica before it has booted, you may see the first screen of the Unitrends Windows Integrated Bare Metal Recovery Wizard. Do not attempt to complete the steps on the bare metal screen. After several seconds, the login screen for the original Windows asset displays.

Replica location	Description
Recovery Series or Recovery MAX appliance	Connect to the replica by using VNC, then log in using the credentials of the original Windows asset. In the VNC viewer, you will need to enter the appliance IP address and the VNC port (found on the Replicas tab under **View Details**). For detailed steps, see "To access a replica on a Recovery Series or Recovery MAX appliance" on page 1018. (The same steps are used to connect to a replica in audit mode and live mode.)
External hypervisor	Connect to the replica by using the hypervisor manager, then log in using the credentials of the original Windows asset.

5 If you see a message about reactivating Windows, you must activate the operating system by using your product key.

6 Check the disk configuration by using Windows Disk Management. (These steps might be slightly different depending on the Windows version.)

- Press the **Start** button.
- Right-click the **Computer** item.
- Choose **Manage**.
- Choose **Storage > Disk Management**. This application shows a graphical view of all disks and volumes.
- If the disk manager shows any disks in the Offline state, right-click the disk icon and click **Online**.
- If the disk manager shows any dynamic disks as Foreign, right-click the disk icon and click **Import**. All volumes should now display as they did on the original asset.

7 Set the system clock. The asset may be running with the system clock time used by the latest backup. This issue may cause the macine to boot with a past date or time.

8 From the Windows Control Panel, update the network properties for the adapter (the TCP/IPv4 address) by using one of the methods in the following table.

Replica location	Description
Recovery Series or Recovery MAX	Do one of the following: • If the original asset has a static IP address, assign the live replica the same

Replica location	Description
appliance	network settings as the original asset. This ensures that the replica functions as the original asset, and that the original asset's job schedules continue for the live replica. • If you are using DHCP to assign IP addresses and you added the original asset to the backup appliance by using only the asset's name, the appliance detects the live replica after you connect it to your network. The appliance then treats the live replica as if it is the original asset. No additional network configurations are necessary to ensure that scheduled backup and backup copy jobs continue.
External hypervisor	• If the original asset has a static IP address, assign the live replica the same network settings as the original asset. This ensures that the replica functions as the original asset and that the original asset's job schedules continue for the live replica. • If the original asset has a static IP address and the hypervisor does not have a network interface on the same subnet as the original asset, assign the replica a new network setting that uses the same subnet as the hypervisor. You must then modify the settings for the original asset in the Unitrends appliance by entering this new IP address. (For details, see "To edit an agent-based asset" on page 293.) This enables the appliance to treat the live replica as the original asset.

9 Log in to the Unitrends backup appliance and re-save the original Windows asset:

Note: If you recovered by using a backup copy on an appliance backup copy target, perform these steps from the backup appliance where the original asset resides, rather than from the backup copy target appliance.

- Select **Configure > Protected Assets**.
- Select the original Windows asset.
- Click **Edit > Save**.

SQL databases and other applications may require a few minutes to become available.

10 Prepare the replacement machine by doing the steps in "Using the live replica as a temporary for the original Windows asset" on page 1023 or "Using the live replica as a permanent replacement for the original Windows asset" on page 1024. You must do these steps before you tear down the Windows replica.

Using the live replica as a temporary for the original Windows asset

If the replica will replace the original asset only temporarily, do these steps after you have booted the replica in live mode:

1 Recover to new hardware as soon as possible, by using Unitrends bare metal recovery.

   Data from the live replica is protected by the backup schedule of the original asset. Use the latest backup to perform the bare metal recovery. For details, see "Windows Bare Metal Protection and Recovery" on page 1207.

2 If any additional backups completed on the replica (after the backup you used for the bare metal recovery), recover them to the new Windows asset to bring it up to date. For details, see "To recover an entire file-level backup by using the Backup Catalog" on page 945 or "To recover an entire file-level backup by using the Backup Browser" on page 951.

3 After recovering the replica's data to the new Windows asset, delete the replica from the appliance and from the hypervisor (if applicable). For instructions, see "To tear down a Windows replica" on page 1025.

Using the live replica as a permanent replacement for the original Windows asset

If the replica will permanently replace the original asset, do these steps after you have booted the replica in live mode:

1 Determine whether to continue protecting the replica with the backup schedules of the original asset or whether to run virtual machine backups for the VM. For a comparison of each method, see "Protecting Hyper-V virtual machines with file-level backups" on page 661 or "Protecting VMware virtual machines with file-level backups" on page 674.

2 Delete the replica from the appliance as described in "Tearing down a Windows replica" on page 1025 . Be sure to delete the replica from the appliance only, as you have the option to delete it from the hypervisor as well.

3 If you will be switching to VMware or Hyper-V backups, create the VM backup schedule and remove the asset from the original schedule. See "Backup Administration and Procedures" on page 425 for details.

   Notes:

   - The replica VM is not automatically added to any existing VM backup schedule.
   - It can take several minutes for the replica VM to show up in the list of VMs to protect with VMware or Hyper-V backups. To refresh the list of discovered VMs, click the Gear icon in the upper-right of the UI and select **Inventory Sync**.

Do not cancel an active replica restore job

Do not cancel an active replica restore job. Instead, bring the replica into audit mode.

Each successful backup of the original Windows asset is applied to the replica as soon as the backup completes. The appliance applies the backup by running a replica restore job, which displays on the Active Jobs tab as shown here:

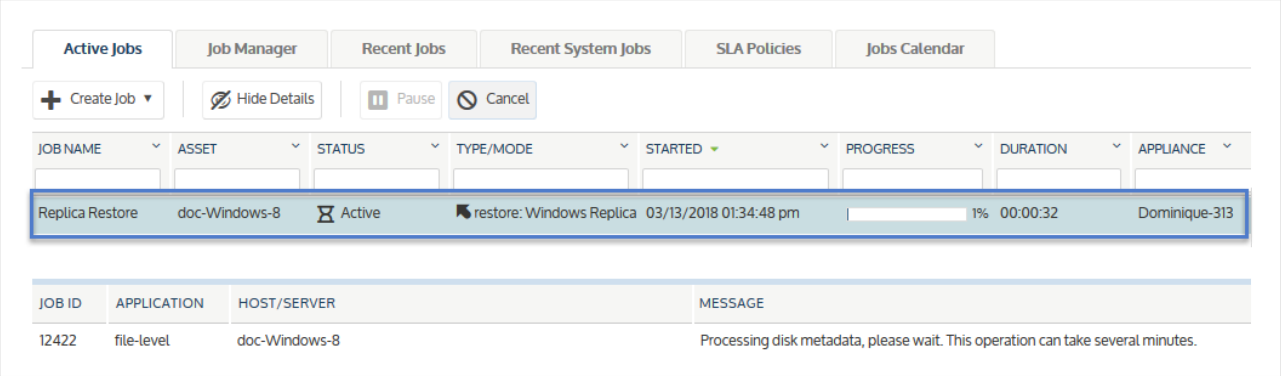

If you cancel a replica restore job by using the Cancel button on the Active Jobs page, the appliance automatically creates a new job to replace the one you canceled. To temporarily stop applying backups, bring the replica into audit mode instead (as described in "Auditing a Windows replica" on page 1015). Use the procedure "To exit audit mode" on page 1020 to start applying backups again. Note that all backups that ran while the replica was in audit mode will be applied to the replica upon exiting audit mode. You cannot skip applying a specific backup to a replica.

Tearing down a Windows replica

This section provides instructions for deleting a Windows replica.

For a replica running on a Unitrends appliance, you should delete the replica as soon as you have recovered the original asset to new physical hardware, to free up appliance resources.

For a replica running on a hypervisor, you have these options:

- Delete the replica from the appliance only – Select this option to use the replica VM as a permanent replacement for the failed Windows asset.
- Delete the replica from the appliance and delete the replica VM from the hypervisor itself – Select this option if you have recovered the original asset to new physical hardware and will not be using the replica VM as a permanent replacement.

To tear down a Windows replica

IMPORTANT! If you are running a live replica as a temporary replacement for an asset that has failed, do not tear down the replica until you have recovered the original asset to new physical hardware. (For details, see "Using the live replica as a temporary for the original Windows asset" on page 1023.)

1. Select **Recover**, then click the **Replicas** tab.
2. Select the replica, then click **Tear Down**.

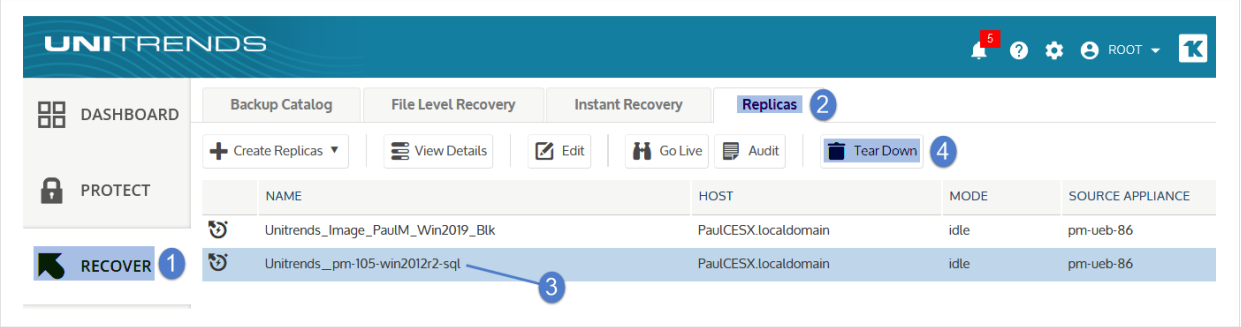

3 Do one of the following (options differ by replica location):

- Replica residing on a Recovery Series or Recovery MAX appliance – Click **Delete** to delete the replica.
- Replica residing on an external hypervisor – A box displays with options to delete the replica from the appliance only or from both the appliance and the hypervisor. Select the desired option and click **Tear Down**.

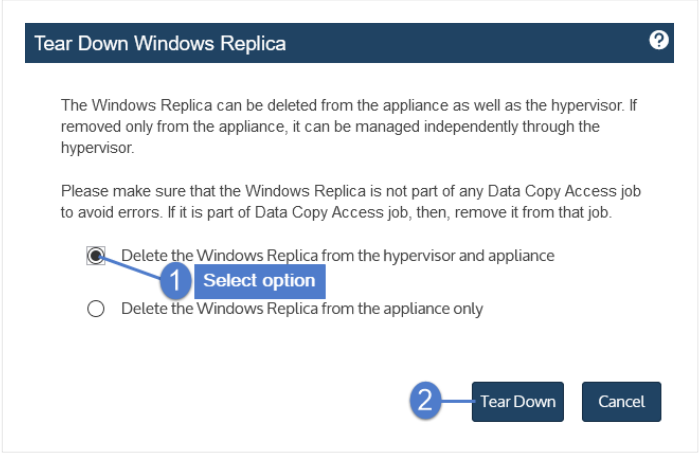

It can take several minutes for the appliance to purge all information about the replica. If you need to create a new replica for the original asset, you must wait for this information to purge. If it has not yet purged, the original asset does not display in the list of assets for which you can create a replica.

Monitoring Windows replicas

Use these procedures to check the status and details of existing Windows replicas:

- "To view all Windows replicas"
- "To view Windows replica details" on page 1027
- "Windows replica modes" on page 1029

To view all Windows replicas

1 Select **Recover**, then click the **Replicas** tab.

2 All Windows replicas and VM replicas display in a list on the Replicas tab.

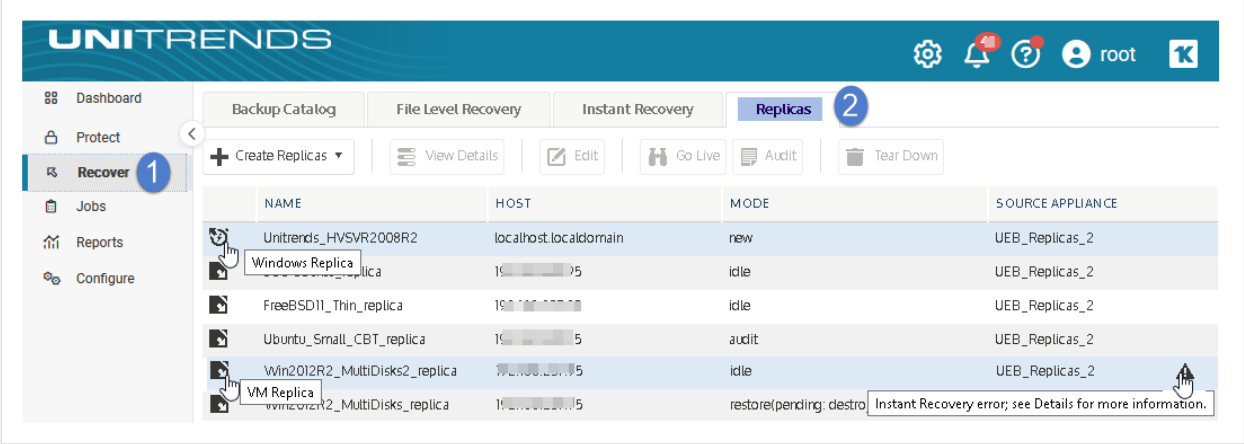

The following information is given for each replica:

Column	Description
Replica type icon	Indicates the replica type: Windows or VM.
Name	Replica name. By default, the replica is named Unitrends _<WindowsAssetName>.
Host	Host where the replica resides: • For ESXi servers, *localhost.localdomain* or server IP address. • For Hyper-V servers, hostname of the Hyper-V server asset. • For Unitrends appliances, *Unitrends appliance*.
Mode	Replica mode. Examples: new, audit, restore, or idle. See "Windows replica modes" on page 1029 for additional details.
Source Appliance	Appliance where the replica was created.
Alert icon	Indicates that an alert has been generated for the replica. Hover over the icon for details.

To view Windows replica details

1 Select **Recover**, then click the **Replicas** tab.

2 Select the replica, then click **View Details**.

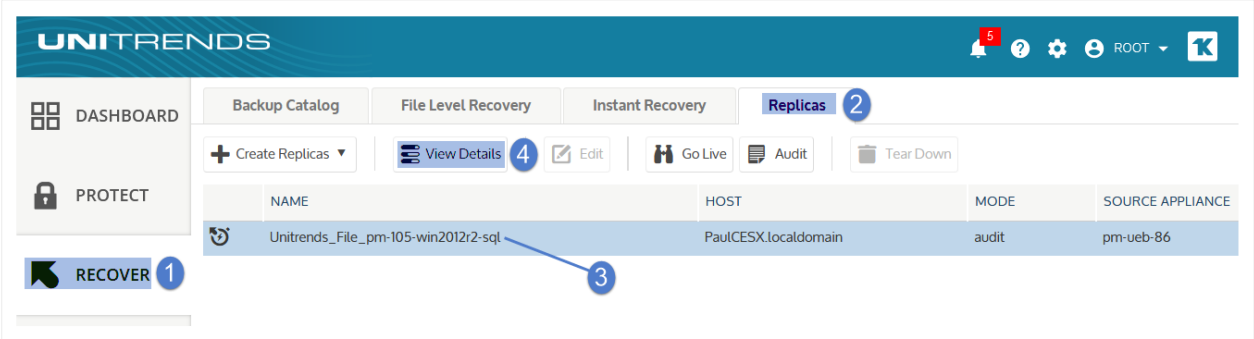

3 Tabs containing replica details display below:

Details tab	Description
Details	Shows the replica name, IP address (if applicable), number of processors, amount of memory, included volumes, excluded volumes (if any), VNC port (if applicable), and message describing the last replica operation. Example: 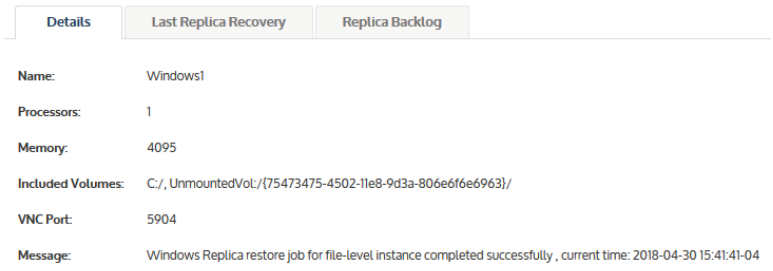
Last Replica Recovery	Shows the most recent backup that has been applied to the replica. If the replica includes Exchange or SQL databases, the list shows the most recent file-level backup and the most recent application backup for each protected database. Example: 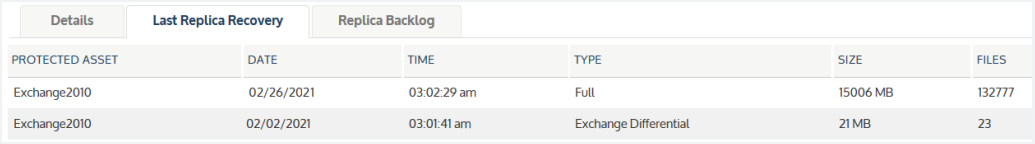
Replica Backlog	List of backups that have not yet been applied to the replica. No backups are listed if the replica is up to date. Example:

Details tab	Description
	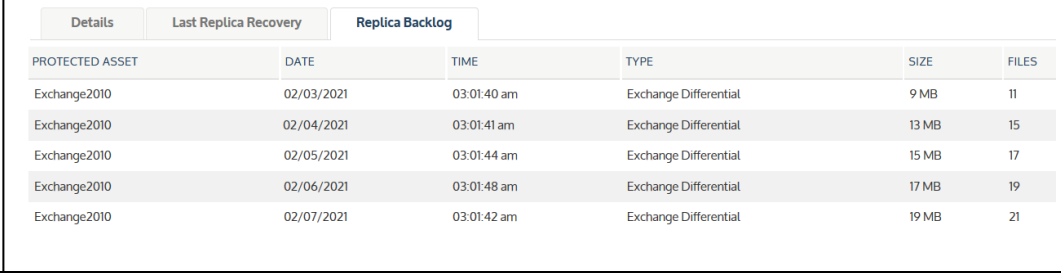

4 (Optional) Click **Hide Details** to stop displaying details for the selected replica.

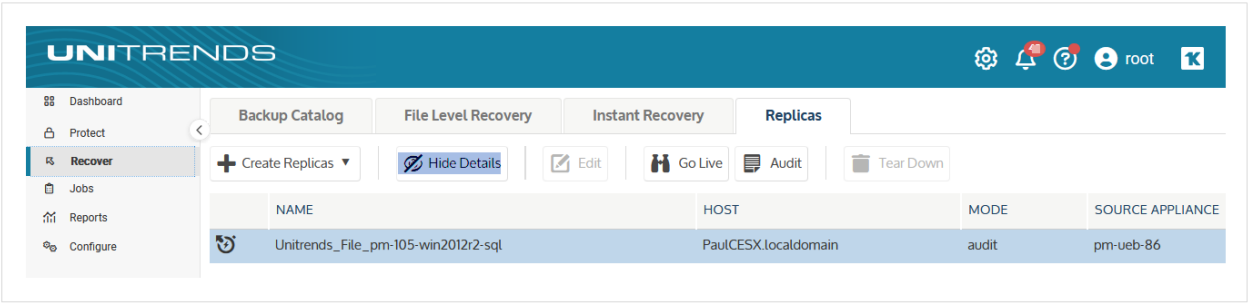

Windows replica modes

You can monitor a Windows replica by checking its mode on the Replicas tab. The mode indicates what is currently happening with the replica (for example, whether it is newly created, whether a backup is being applied, or whether it is in audit mode.)

Windows replica modes are described in the following table:

Mode	Description
New	The replica is new and no backup has been applied. A replica in *new* mode cannot be audited or booted into live mode.
Restore	A backup has completed, and the appliance is applying it to the replica. The replica remains in *restore* mode until the restore completes.
Idle	At least one backup has been applied to the replica, but currently no action is occurring.
Halted	A backup has completed, and the appliance has requested a restore. The replica goes into a *halted* state if the restore cannot be performed. The following can occur when a replica is in this mode: • If the restore could not be performed because the appliance could not reach the replica,

Mode	Description
	it tries again after several minutes, and the mode changes from *halted* to *idle*. After three failed attempts, the replica becomes invalid, and it remains in *halted* mode until a user deletes it. • If the restore could not be performed because a configuration change was made to the original asset, the replica becomes invalid, and it remains in *halted* mode until a user deletes it.
Audit	A user is performing a manual audit and the replica has booted in *audit* mode. For details, see "Auditing a Windows replica" on page 1015.
Verify	A user has enabled verification reports (automated audits). The appliance has booted the replica in *audit* mode to take a screenshot of the replica's login screen. For details about verification reports, see "Automated audits for a Windows replica" on page 1015.
Live	A user has booted the replica in *live* mode to replace the original asset. For details, see "Bringing the replica live in production" on page 1020. Once the replica is live, the only other mode it can enter is *off*.
Off	A user has taken the replica out of *live* mode. Once the replica is *off*, the only other mode it can enter is *live*.

Chapter 16: Recovering Windows Image-level Backups

Unitrends provides a variety of methods for recovering image-level backups of Windows assets. You can recover the entire asset or selected files from backup. For quick recovery of critical assets, you can create image-level replicas or use the instant recovery feature. See the following table for descriptions of each recovery method:

Recovery method	Description
File recovery	Use these procedures to recover files from an image-level backup: • "Recovering from an indexed image-level backup by using Search Files" – Use to search an asset's indexed image-level backups for files/folders that meet specified criteria and recover selected items from the search results. This procedure can be used only for backups that were run with the Edit Asset > Index Image-Level Backups option. For details on configuring this option, see "To edit an agent-based asset" on page 293. • "Recovering files by browsing a Windows image-level backup" on page 1040 – Use to recover files by browsing the image-level backup.
Windows image-level replicas	Creates a stand-by virtual replica of the Windows asset that you can bring online in minutes. As backups of the original asset run, they are applied to the replica to keep it up-to-date. Take the replica 'live' to assume the role of a failed Windows asset. To meet near-zero RTOs, set up the replica before the asset fails. For detailed requirements and procedures, see "Windows image-level replicas" on page 1086.
Instant recovery of Windows image-level backups	Recovers a failed Windows asset in minutes. The recovered asset can immediately assume the role of the original, failed machine. Unitrends recommends that you plan for instant recovery (IR) before an asset fails, by reviewing requirements, allocating IR storage, and using audit mode to test IR of your critical assets. For detailed requirements and procedures, see "Instant recovery of Windows image-level backups" on page 1055.
"Windows unified bare metal recovery" on page 1209	Use for disaster recovery (DR) of a failed Windows asset. Image-level backups capture the disk metadata needed for DR. You perform DR using a standard 32-bit or 64-bit ISO image provided on the Unitrends backup appliance. The target for the recovery can be a physical or virtual machine. For details, see "Windows unified bare metal recovery" on page 1209.

Selecting a backup to recover

To perform the recovery, you will start by selecting a backup or backup copy. For backups, you can do this in the Backup Catalog or in the Backup Browser. For backup copies and imported backups, you must use the Backup

Catalog.

In the Backup Catalog, backups and backup copies display under the protected asset. You can modify the display by entering filter criteria. Expand an asset to view its backups and backup copies:

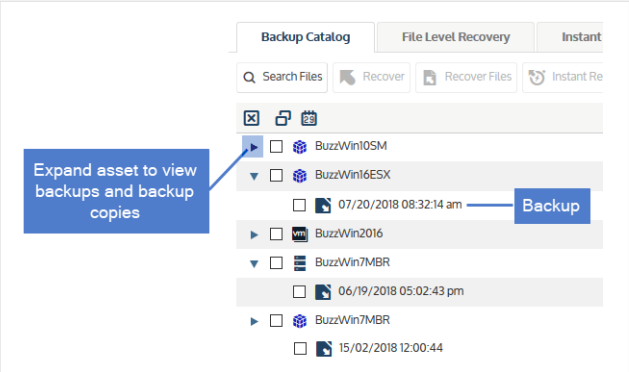

The Backup Browser provides advanced search and filter options. Backups are not grouped under the protected asset. Search for backups by selecting an appliance and date range. Filter the display by entering text in the column fields. Click an arrow to sort by column:

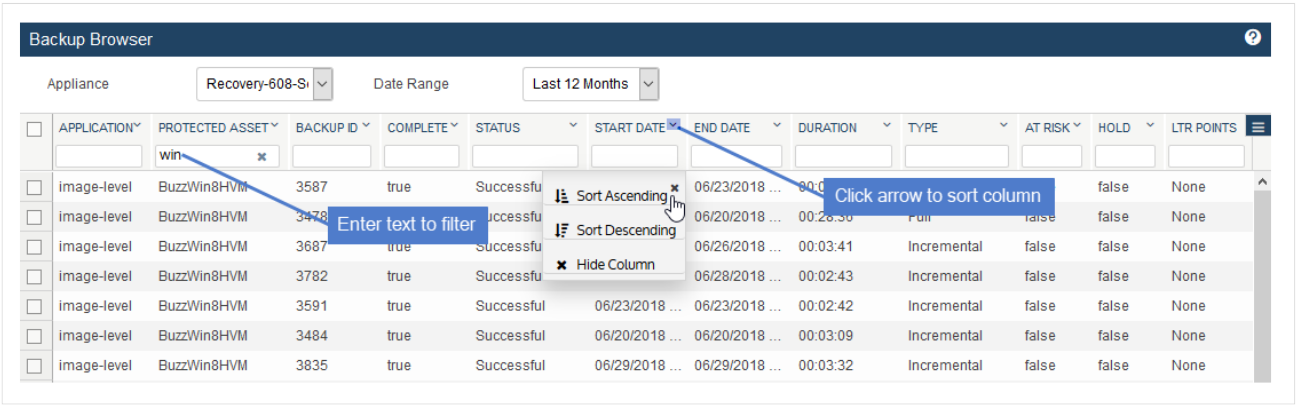

Considerations for recovering SQL clusters, SQL availability groups, and Exchange DAGs

Before recovering from the image-level backup, review the following considerations:

- SQL clusters – The following apply when recovering from an image-level backup of a clustered SQL instance:
  - You are able to recover the cluster node and clustered SQL instance from an image-level backup. Depending on the configuration and cluster dynamics, after a restore and reboot, the cluster may accept the restored system back into the cluster. If not, it may have to be removed from the cluster and re-added.
  - By definition, database files for a clustered instance have to be on shared storage, so they are not included in the image-level backup of the clustered node.

  - Image-level backups of nodes containing Cluster Shared Volumes (CSVs) are not supported, hence clustered SQL instance nodes using CSVs are also not supported.

- SQL availability groups – The following apply when recovering from an image-level backup of SQL instances that contain SQL availability groups:
  - You are able to recover the cluster node and availability groups from an image-level backup. Depending on the configuration and cluster dynamics, after a restore and reboot, the cluster may accept the restored system back into the cluster. If not, it may have to be removed from the cluster and re-added.
  - Availability group databases with a secondary role on the node will have to be deleted and re-added.
  - Availability group databases with a primary role on the node may have to be re-synced to secondary mirrors. In some cases, the availability group may have to be deleted entirely and re-added on the node. Consult Microsoft's documented procedures for dealing with SQL availability groups and databases.
- Exchange DAGs – You are able to recover a DAG node from an image-level backup, but the integrity of Exchange on the recovered node is not guaranteed.

Recovering files from Windows image-level backups

Use the procedures in this section to recover files from an image-level backup.

Note: You must recover to an agent-based asset that has been added to the backup appliance. (The asset must have a Unitrends agent installed and must display on the appliance's Protected Assets tab.) If necessary, add the asset (as described in "To add an agent-based asset" on page 289) before you start the recovery.

See these procedures to recover files:

- "Recovering from an indexed image-level backup by using Search Files" – Run this procedure from the backup appliance to recover from an indexed backup. (This procedure is not supported for imported backup copies).
- "Recovering files by browsing a Windows image-level backup" on page 1040 – Run this procedures from the backup appliance to recover from any of the following: a backup, an imported backup copy, or a hot backup copy that resides on a target appliance or in the Unitrends Cloud. Run this procedure from the target appliance to recover from a hot backup copy that resides on the that target appliance.

Recovering from an indexed image-level backup by using Search Files

Use this procedure to search an asset's indexed image-level backups for files/folders that meet specified criteria and recover selected items from the search results.

Notes:

- This procedure can only be used for local backups that were run with the Edit Asset > Index Image-Level Backups option. For details on configuring this option, see "To edit an agent-based asset" on page 293.
- This procedure is not supported for recovery of ReFS filesystems. Recover by browsing the backup instead (see "Recovering files by browsing a Windows image-level backup" on page 1040).

- File search of imported backup copies is not supported. Recover by browsing the backup instead (see "Recovering files by browsing a Windows image-level backup" on page 1040).

1 Log in to the backup appliance.

2 Click **Recover > Backup Catalog > Search Files**.

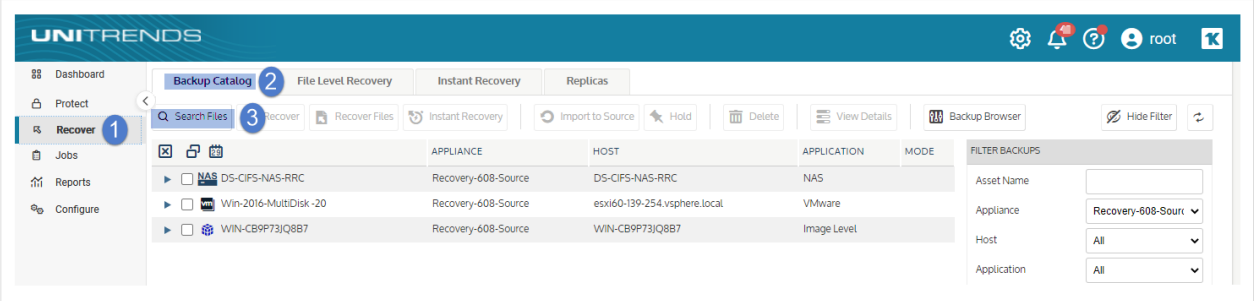

3 In the **Type** list, select **Backup**.

4 Select the **Asset** whose backups will be searched.

5 Enter one or more search options:

Search Options	Description
String	Enter text to search. The wildcard character * is supported.
Match Case	Select to match the letter case of the entered string.
From/To	Use to search for files that were last modified within the specified date range. Results do not include files modified on the From or To date.
Size	Use to search for files that meet this size criteria.
Advanced	Click to search using a regular expression.

6 Click **Search**.

All backups of this asset stored on the appliance are searched for matching files.

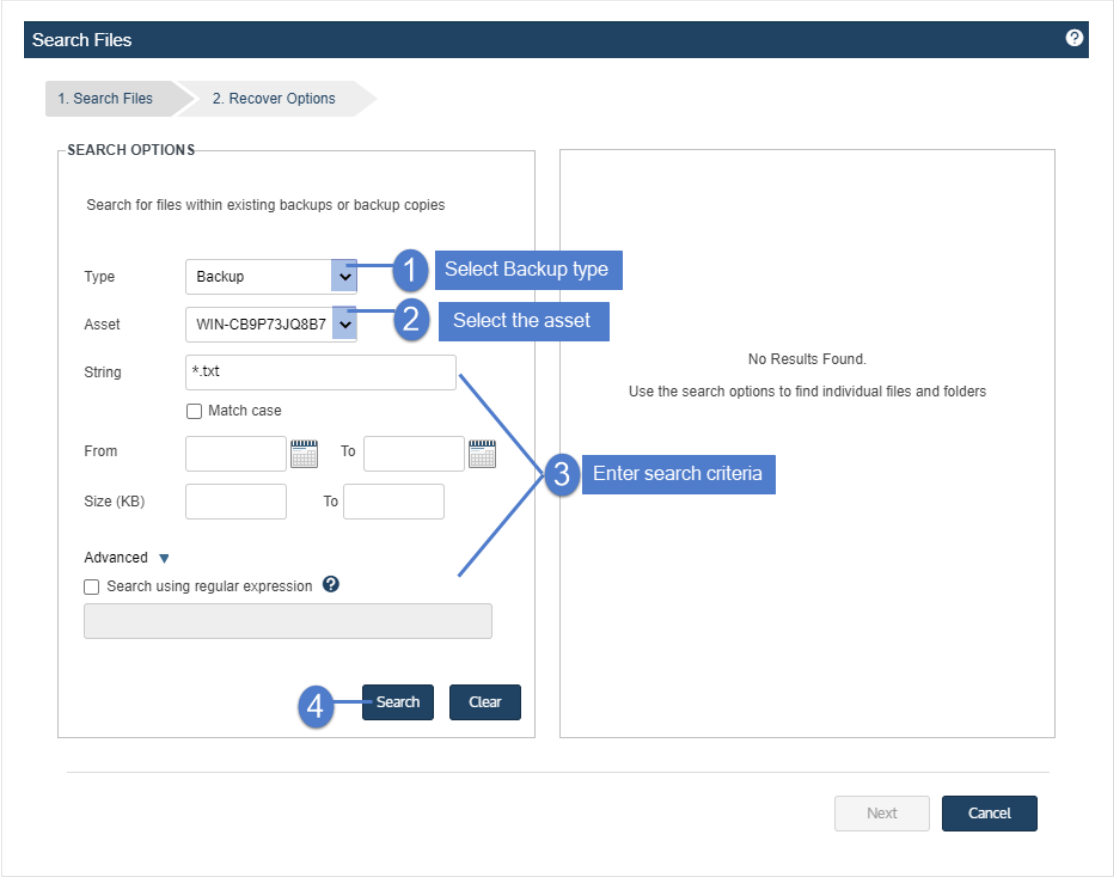

7 In the results list, click to select files to recover.

Notes:

- All files you select must be from a single backup. Check the Backup ID to determine a file's backup. If you select files from multiple backups, the Next button becomes disabled.
- Softlinks cannot be downloaded and are not included in the search results.

8 Click **Next**.

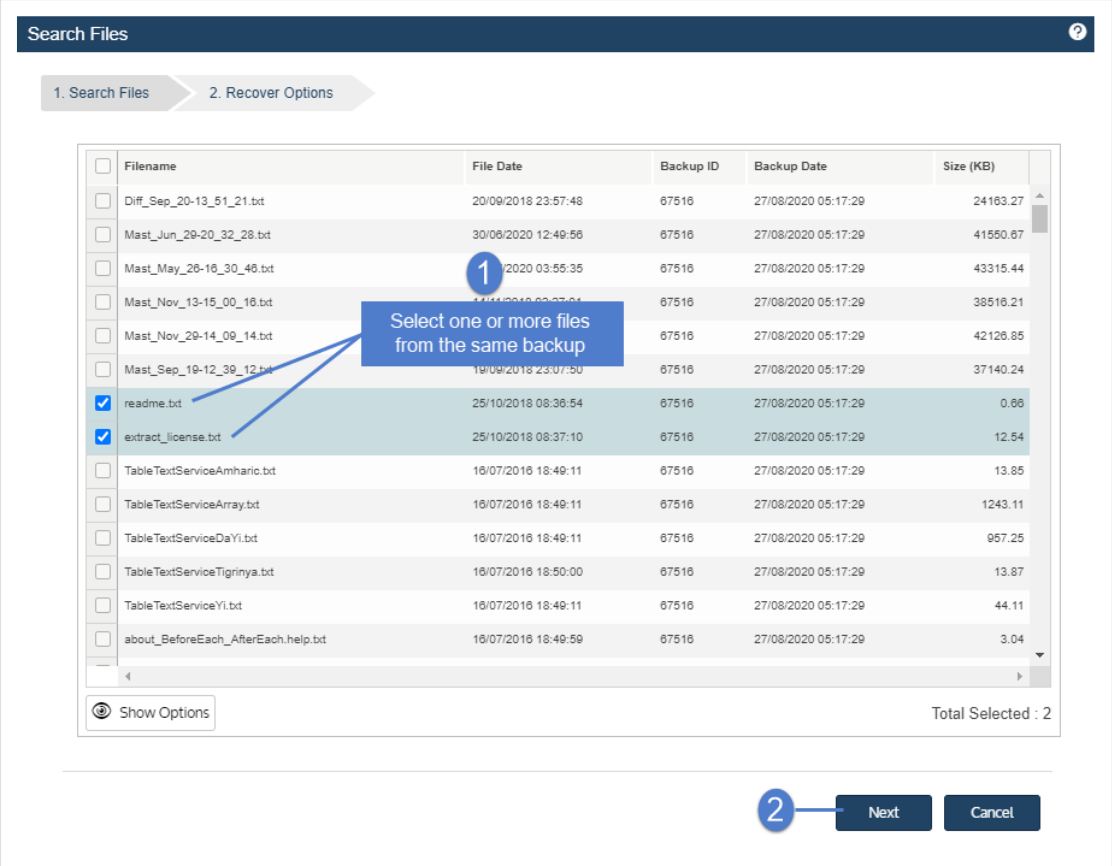

9 Select the **Asset** where the files will be recovered.

10 (Optional) Enter a Directory path or click **Browse** and select a Directory path from the drop-down list.

- If the directory does not exist, the job creates it on the target asset.
- Leave the Directory field empty to recover files to their original location.

11 (Optional) Specify Exclusions.

12 (Optional) Specify Advanced Options. See the following table for descriptions of these options.

Option	Description
Commands to run pre-restore	To run a command or script on the asset before the recovery, enter the full path to the command or script in the **Commands to run pre-restore** field. For example, C:\Data\script.bat or /usr/jsmith/script.sh.
Commands to run post-	To run a command or script on the asset after the recovery, enter the full path to the command or script in the **Commands to run post-restore** field. For example,

Option	Description
restore	C:\Data\script.bat or /usr/jsmith/script.sh.
Preserve directory structure	Check this box to preserve the existing file structure within the target directory. **Note:** To recover the file(s) to the original location, Preserve Directory Structure must be selected. If you attempt to recover to the original location and uncheck this box, the recovery fails.
Overwrite existing files	If this box is checked, files in the Target Directory may be overwritten. (See "Overwrite existing files and Restore newer files only options" for details.) This is useful if you are recovering an updated version of a document and only want the most up to date version.
Restore newer files only	Check this box to recover a file only if its date is newer than the existing version in the Target Directory. (See "Overwrite existing files and Restore newer files only options" for details.) If the file does not exist in the Target Directory, the file is recovered.
Set file dates to today	Check this box to update the recovered files with the recover date and time. If not checked, file dates are not updated during the recovery.
UNIX text conversion	When recovering UNIX Text files to MS-DOS systems, checking this option prevents new lines from being converted to CR-LF.

Overwrite existing files and *Restore newer files only* options

This table describes how *Overwrite existing files* and *Restore newer files only* work if the file exists in the Target Directory.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
Overwrite existing files = Yes Restore newer files only = Yes	Recovers the file and overwrites the existing file.	• If the file to recover is newer than the one in the Target Directory, recovers the file and overwrites the existing file. • If the file to recover is older than the one in the Target Directory, does not recover the file.

Option selected?	Recovery behavior	
	Windows backup	Non-Windows backup
Overwrite existing files = Yes Restore newer files only = No	Recovers the file and overwrites the existing file.	Recovers the file and overwrites the existing file.
Overwrite existing files = No Restore newer files only = Yes	Does not recover the file.	Does not recover the file.

13 Click **Save**.

14 Click **OK** to close the Notice message.

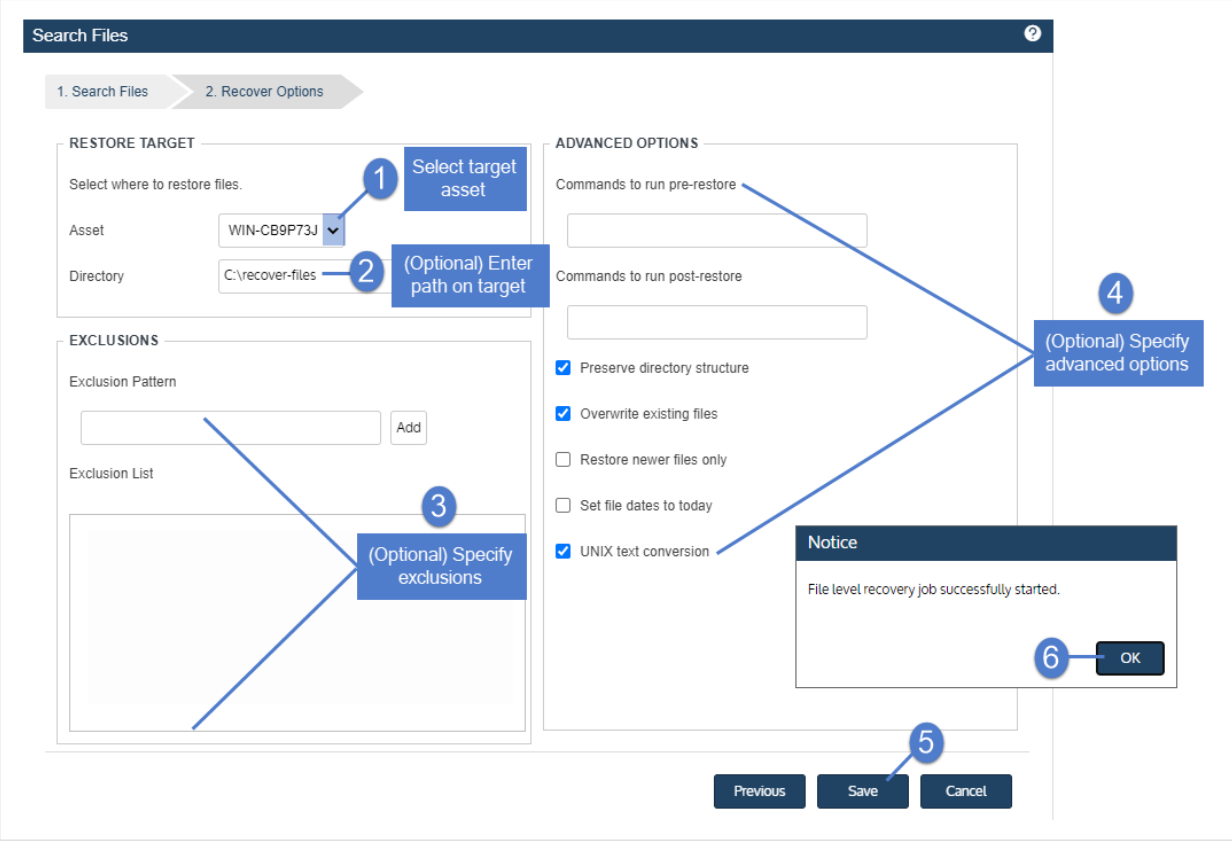

15 To monitor the recovery job:

- Select **Jobs** > **Active** Jobs.
- Select the job in the list and click **View Details**.

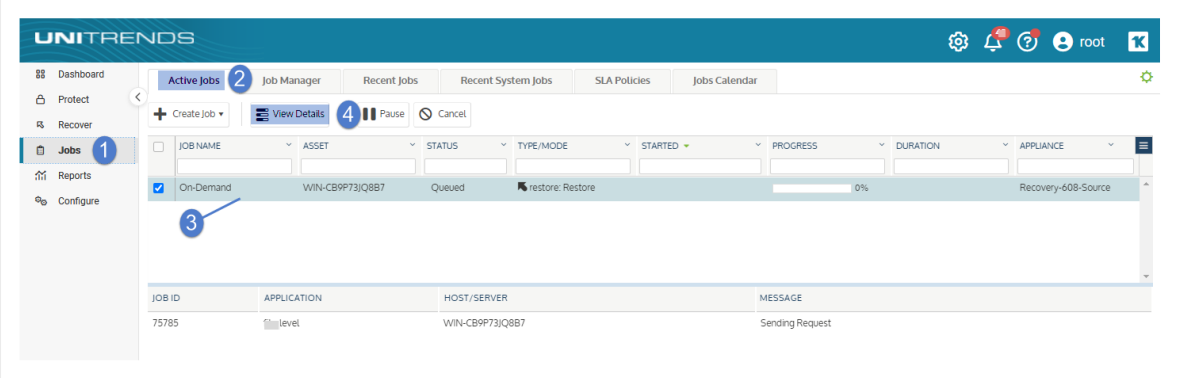

- The recovery is complete when the job's status changes to Success.

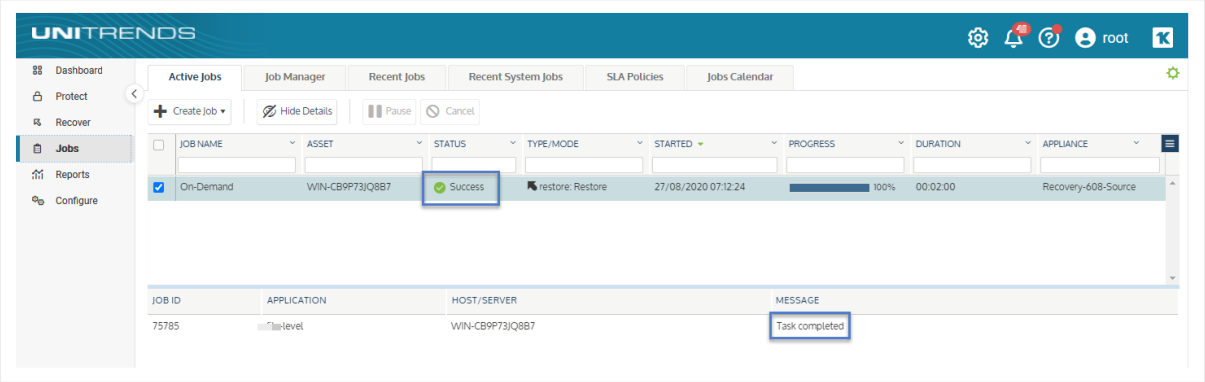

16 Access the recovered files on the recovery target.

Note: Files that are locked because they are in use in the location where you are performing the restore (system files, database files, etc.) cannot be restored while the system is live. During recovery, these files are moved to a temporary location and an entry is added into the registry telling Windows to restore the files to their actual final location on the next boot.

Recovering files by browsing a Windows image-level backup

To recover files by browsing the backup, ensure that "Prerequisites and considerations" have been met, then proceed to "File recovery procedures" on page 1041.

Prerequisites and considerations

The following requirements and considerations apply to recovering files by browsing a Windows image-level backup:

Prerequisite or consideration	Description
Backup	You can use an indexed or non-indexed Windows image-level backup for the recovery.
Supported recovery methods	To recover from a backup or imported backup copy, you must log in to the backup appliance directly. Logging in to an appliance that is managing the backup appliance is not supported. To recover from a hot backup copy that resides in the Unitrends Cloud or on another Unitrends appliance, you can either run the procedure from the source backup appliance (without first importing the backup copy) or run the procedure from the target appliance. During the recovery procedure, the appliance creates a recovery object that contains the backup's files. For some Windows assets, this object is also exposed as a CIFS (Samba) share and/or an iSCSI LUN on the backup appliance. After you create the recovery object, you will view it on the File Level Recovery tab to see whether the CIFS and iSCSI options are available. You can recover files from this object in several ways. Options include:

Prerequisite or consideration	Description
	• Browse the recovery object and download selected files to a *.zip* file. This is the simplest method. • Mount the CIFS share on a recovery target machine. From the target machine, select files to recover. • Mount the iSCSI LUN on a recovery target machine. From the target machine, select files to recover. (You must use an iSCSI LUN in some cases. For details, see "When to use an iSCSI LUN" below.)
Recovery target requirements	The target can be configured with basic, GUID Partition Table (GPT), or dynamic disks. All configured disks must have unique names.
When to use an iSCSI LUN	You must recover by mounting the iSCSI LUN to perform the following tasks: • Recover access control information on files and folders. • Recover New Technology File System (NTFS) encrypted files. • Recover Resilient File System (ReFS) files. **Note:** ReFS limitation - ReFS file system versions are not compatible with all Windows operating system versions. To avoid compatibility issues, recover ReFS files by mounting the iSCSI LUN on a machine whose operating system version is the same or later than that of the machine where the backup was taken. • Recover exFAT files. **Note:** For the recovery, iSCSI disks are writable and a 1 GB write limit is enforced. Errors display on the recovery target machine if more than 1 GB is required. In this case, you cannot recover by using iSCSI. Recover files by downloading to a *.zip* file or by mounting the CIFS share, or recover the entire asset (by using "Instant recovery of Windows image-level backups" on page 1055 or "Windows unified bare metal recovery" on page 1209).

File recovery procedures

Use the following procedures to recover files from a Windows image-level backup. Before you start, be sure all "Prerequisites and considerations" have been met.

- "Step 1: Create the recovery object"
- "Step 2: Recover files" on page 1047

- "Step 3: Remove the recovery object from the appliance" on page 1054

Step 1: Create the recovery object

Use one of these procedures to create the recovery object:

Note: If a previously-created recovery object is still mounted for the Windows asset, you must remove it before creating a new one.

- "To create the recovery object and recover from a backup or imported backup copy" on page 1042 – Run on the backup appliance to recover from a backup or imported backup copy.
- "To create the recovery object and recover from a hot backup copy by using the source backup appliance" on page 1044 – Run on the backup appliance to recover from a backup copy that resides on a target appliance or in the Unitrends Cloud.
- "To create the recovery object and recover from a hot backup copy by using the target appliance" on page 1045 – Run on the target appliance to recover from a backup copy that resides on that target appliance.

To create the recovery object and recover from a backup or imported backup copy

1. Log in to the backup appliance.

   Note: You must log in to the backup appliance directly. Logging in to an appliance that is managing the backup appliance is not supported.

2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. (For details, see "Working with custom filters" on page 67.)

3. Expand the Windows asset and select the image-level backup from which you want to recover files.
4. Click **Recover Files**.

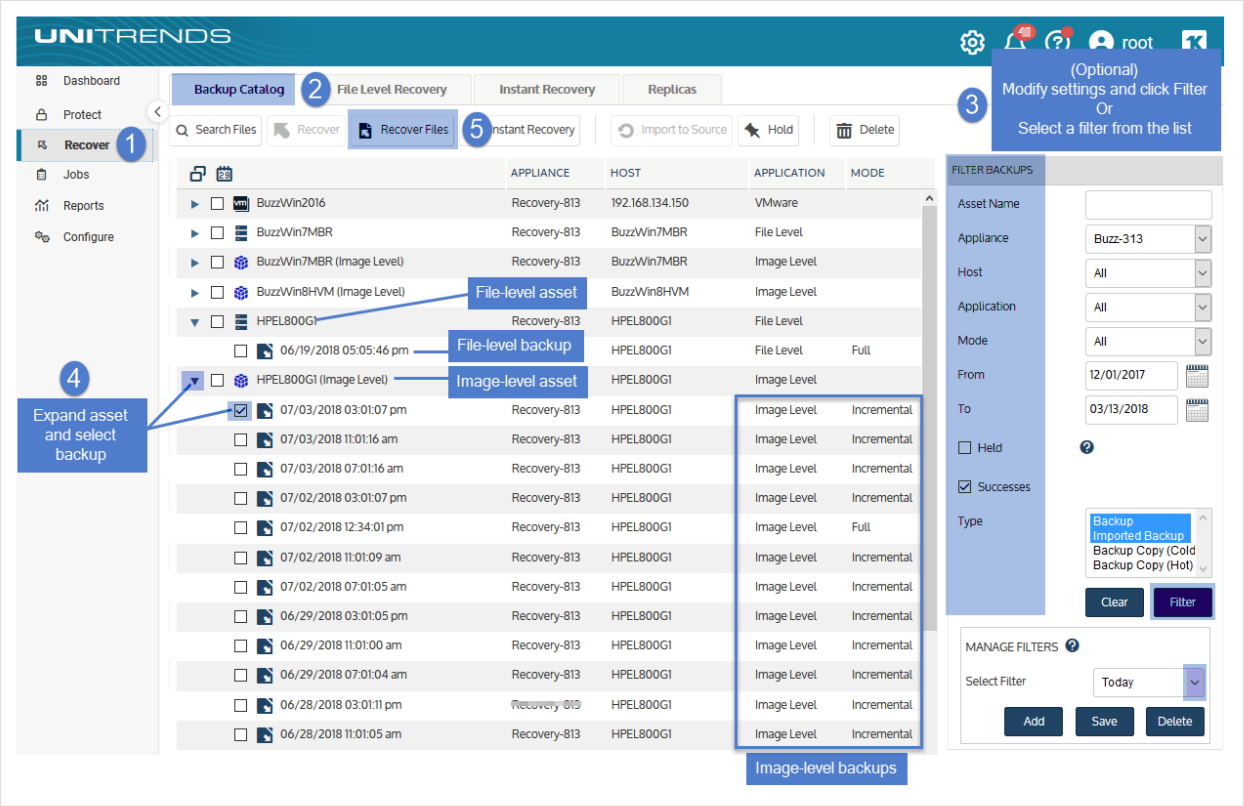

5 Click **Confirm** to continue. The appliance creates the recovery object.

Note: If you receive an error on a Unitrends Backup appliance while creating the recovery object, increase the memory allocation for the Unitrends Backup VM by using the host that manages it.

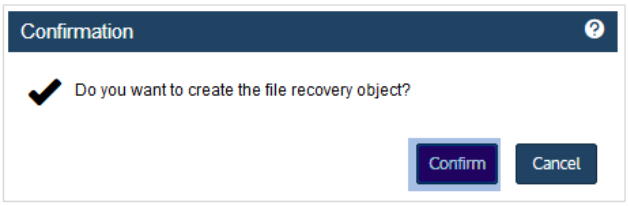

6 Click **View FLR**.

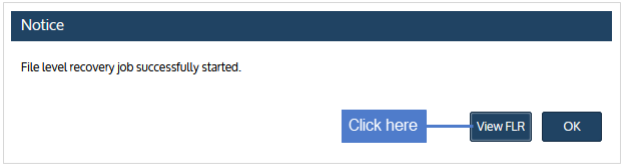

7 Proceed to "Step 2: Recover files" on page 1047.

To create the recovery object and recover from a hot backup copy by using the source backup appliance

Run this procedure on the backup appliance to recover from a backup copy that resides on a target appliance or in the Unitrends Cloud.

> **Note:** You must log in to the backup appliance directly. Logging in to an appliance that is managing the backup appliance is not supported.

1. Log in to the source backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.
3. Use Filter Backups to the right to display hot backup copies:
   - In the Type box, select **Backup Copy (Hot)**.

     (Optional) Use Filter Backups to the right to customize the backups that display. (For details, see "Working with custom filters" on page 67.)
   - Click **Filter**.
4. Expand the asset and select the hot backup copy from which you want to recover files.
5. Click **Recover Files**.

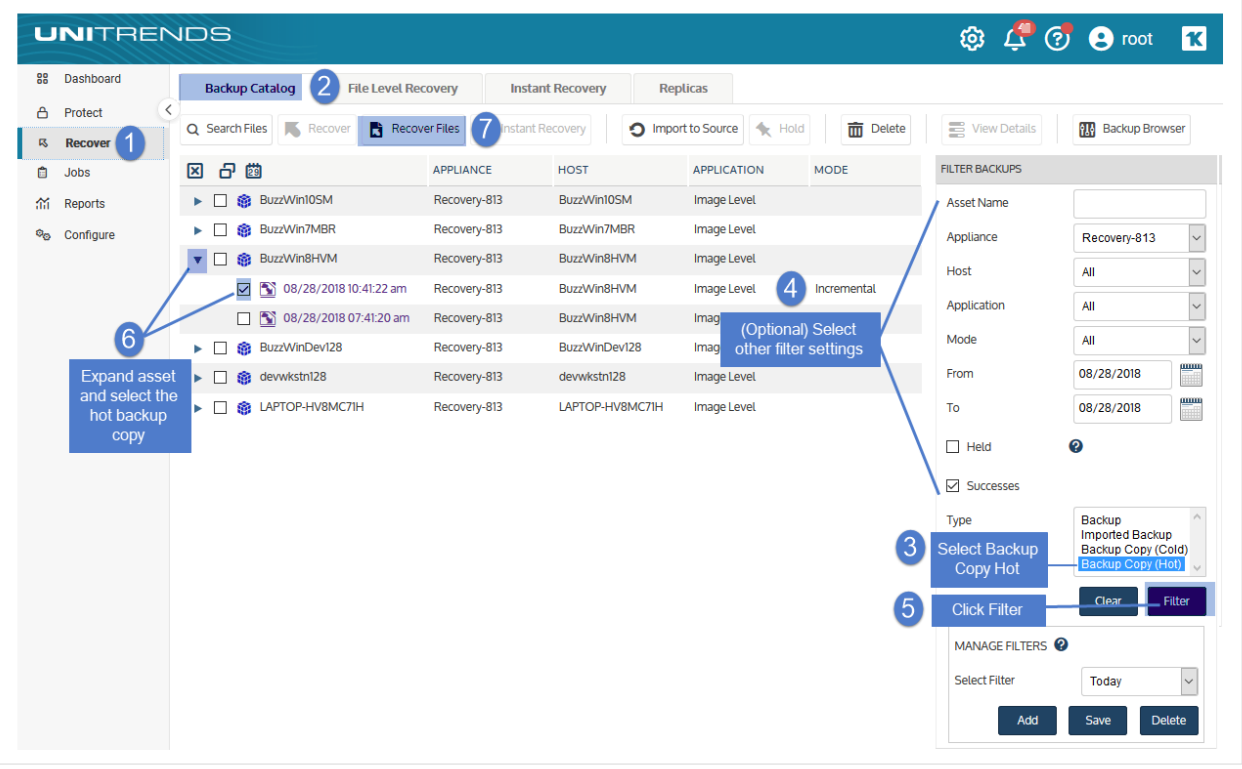

6. Click **Confirm** to continue. The appliance creates the recovery object in the Cloud or on the target appliance.

Notes:

- If your appliance is a Unitrends Backup virtual appliance and you receive an error while creating the recovery object, increase the memory allocation for the Unitrends Backup VM using the host that manages it.
- Recovery objects created in the Unitrends Cloud are automatically removed after 96 hours.

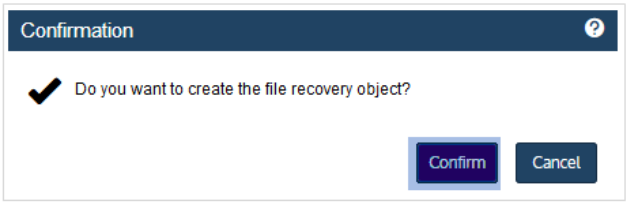

7 Click **View FLR** to view the recovery object on the File Level Recovery tab. The recovery object Name is *AssetName (on the target)*.

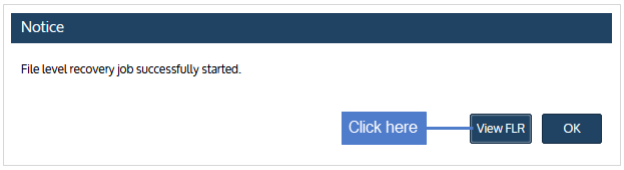

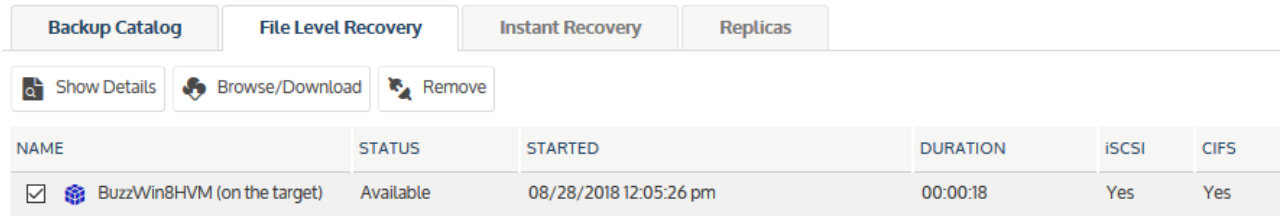

8 Proceed to "Step 2: Recover files" on page 1047.

To create the recovery object and recover from a hot backup copy by using the target appliance

Run this procedure on the target appliance to recover from a backup copy that resides on that target appliance.

1 Log in to the backup copy target appliance.

2 Select **Recover** and click the **Backup Catalog** tab.

3 Use Filter Backups to the right to display hot backup copies:

- In the Type box, select **Backup Copy (Hot)**.

  (Optional) Use Filter Backups to the right to customize the backups that display. (For details, see "Working with custom filters" on page 67.)

- Click **Filter**.

4 Expand the asset and select the hot backup copy from which you want to recover files.

5 Click **Recover Files**.

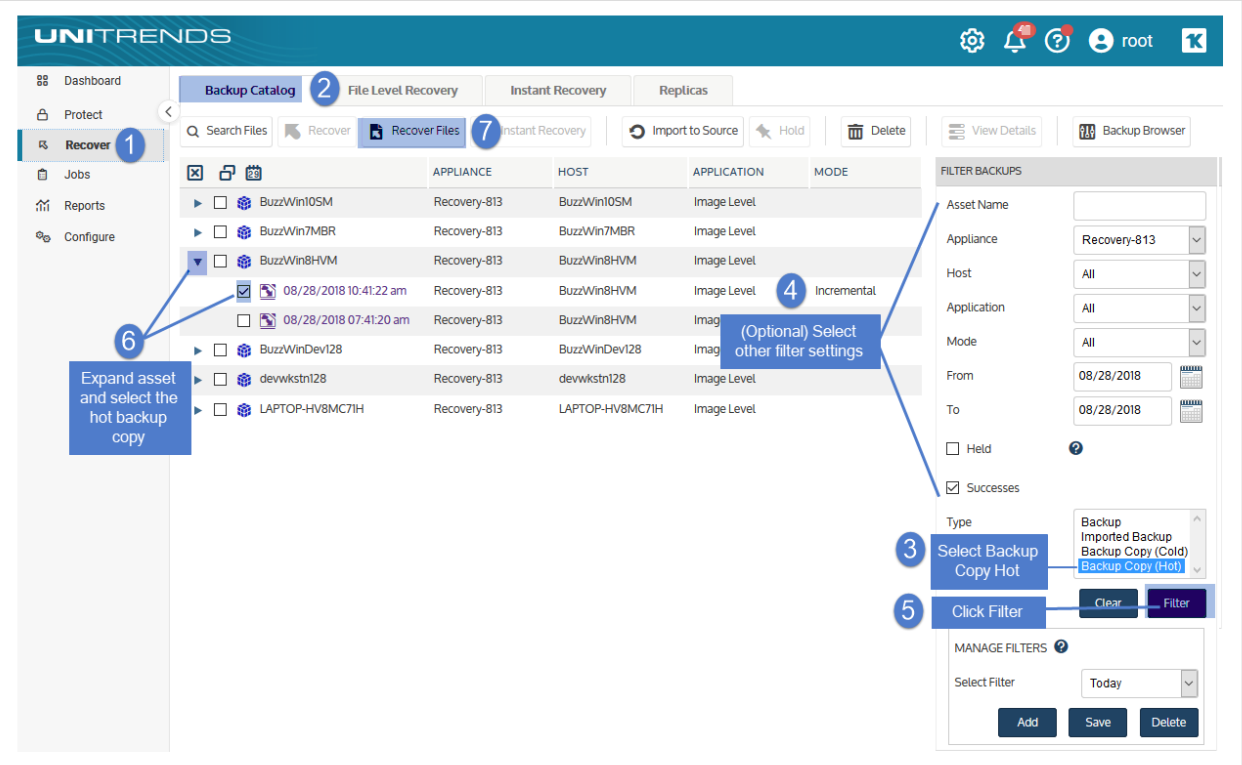

6 Click **Confirm** to continue. The appliance creates the recovery object on the target appliance.

Note: If your appliance is a Unitrends Backup virtual appliance and you receive an error while creating the recovery object, increase the memory allocation for the Unitrends Backup VM using the host that manages it.

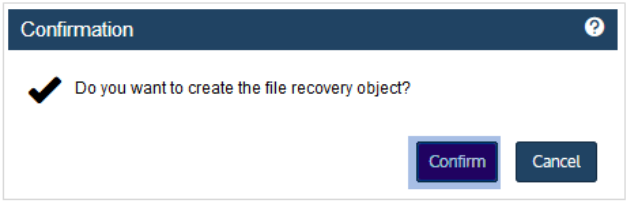

7 Click **View FLR** to view the recovery object on the File Level Recovery tab.

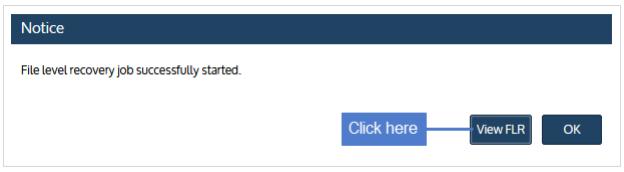

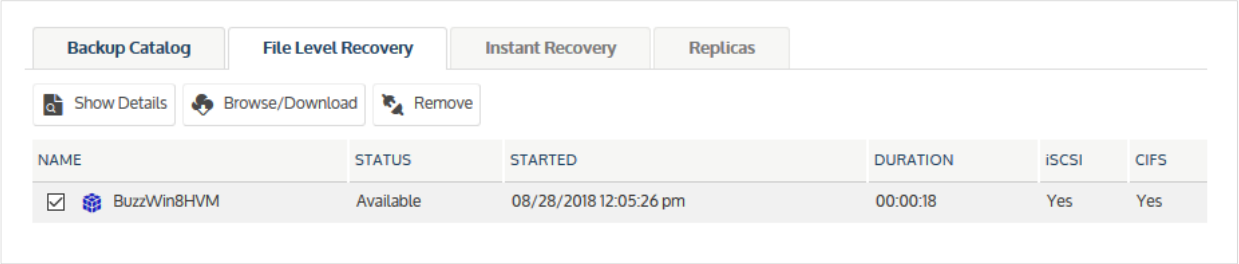

8 Proceed to "Step 2: Recover files" on page 1047.

Step 2: Recover files

View the recovery object on the File Level Recovery tab to see which recovery options are supported for the asset you selected. Use one of the following procedures to recover files:

- "To recover files by browsing and downloading to a .zip file"
- "To recover files by mounting the CIFS share" on page 1049
- "To recover files by mounting the iSCSI LUN" on page 1051

To recover files by browsing and downloading to a *.zip* file

1 On the **File Level Recovery** tab, locate the recovery object.

Recovery objects display on the tab with the following details: the name of the Windows asset, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

2 Select the recovery object and click **Browse/Download**.

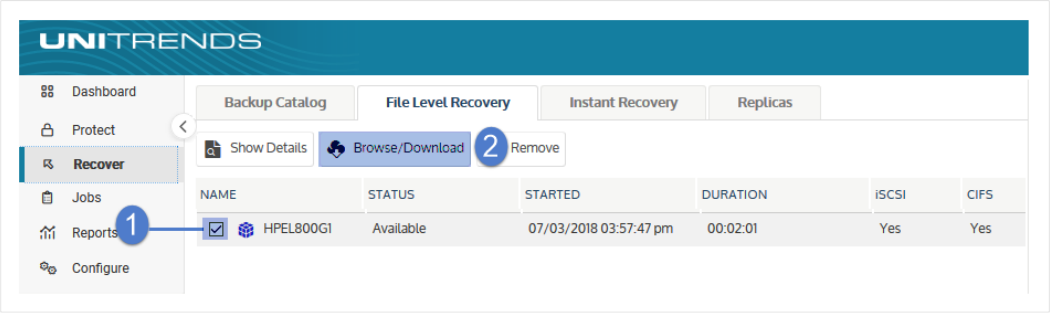

3 In the File Browser, select or drag files and/or directories to recover.

> **Note:** Softlinks (also called *symbolic links*) are excluded from download. If you select a directory that contains files and softlinks, only the files are downloaded.

4 Click **Download**.

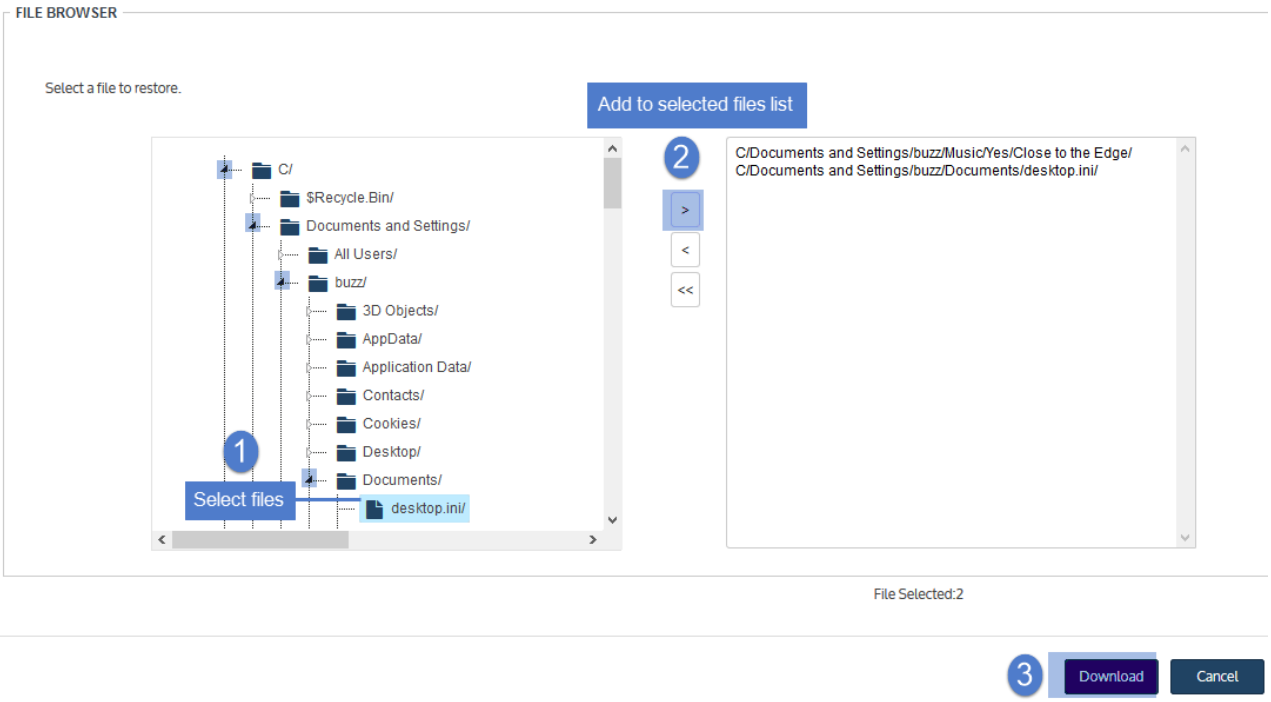

5 Click **Confirm** to download the selected files to a *.zip* file. The *.zip* file is downloaded to your browser's default location.

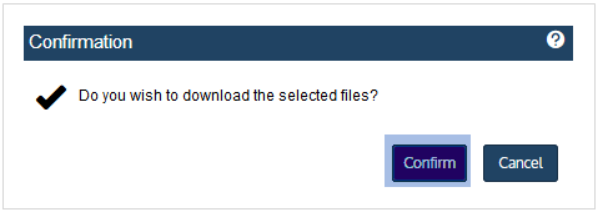

Notes:

- Volumes are assigned letters during recovery that may not match the letters from the original disks.
- The duration of the download is impacted by various factors, such as the size of the files, bandwidth, and download speed.
- Persistent browser and UI sessions are required to create the *.zip* file in the browser's default download location. Do not close the browser or UI session during the recovery.

6 When the download completes, the *Unitrends-Restore.zip* file displays in the browser. Select whether to open or save the file.

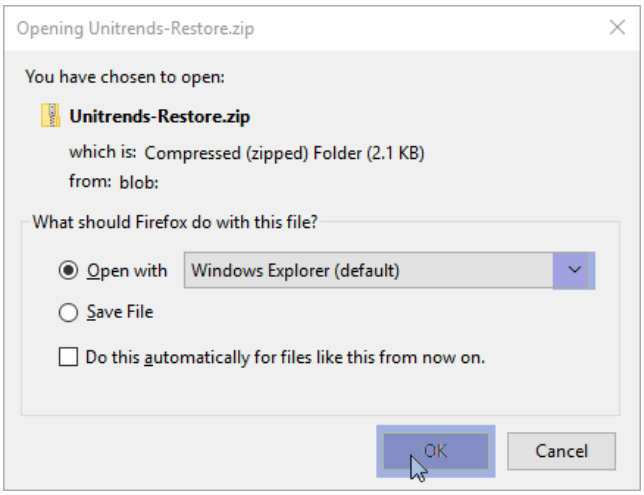

7 Access the recovered files in the download location and move them to another location on the local machine.

Note: The Windows file explorer contains a setting to hide protected/system files. Turn off this setting to access all files.

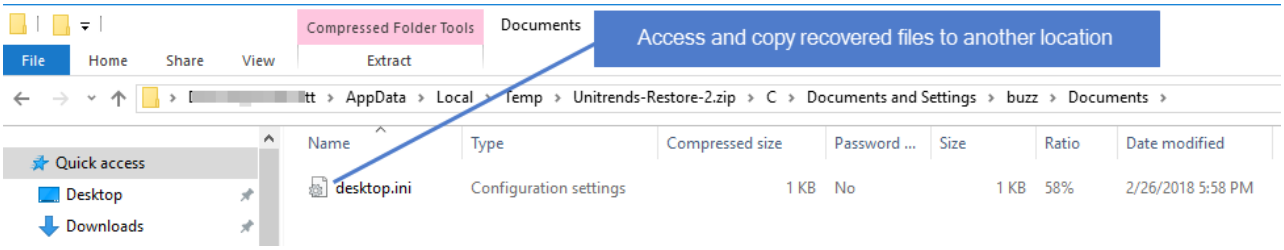

8 Proceed to "Step 3: Remove the recovery object from the appliance" on page 1054.

To recover files by mounting the CIFS share

1 Select **Recover** and click the **File Level Recovery** tab.

Recovery objects display with the following details: the name of the Windows asset, the status of the object, the date and time it was created, the length of time it has existed on the appliance, and whether it can be accessed through iSCSI or CIFS.

2 Select the recovery object and click **Show Details**.

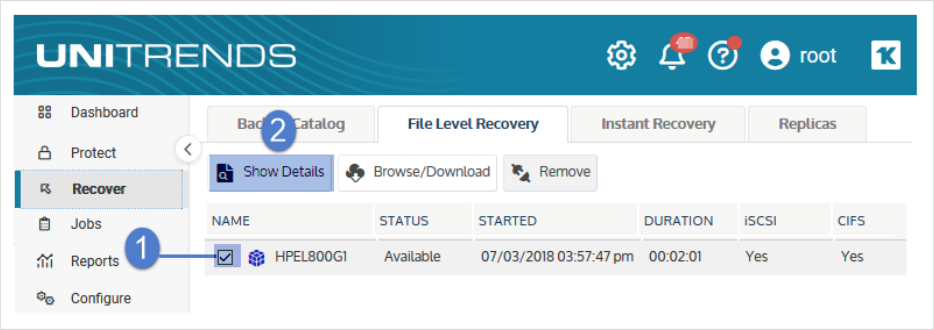

3 Note the CIFS path that displays in the File Level Recovery Details window. You will need this path to mount the CIFS share on the target machine.

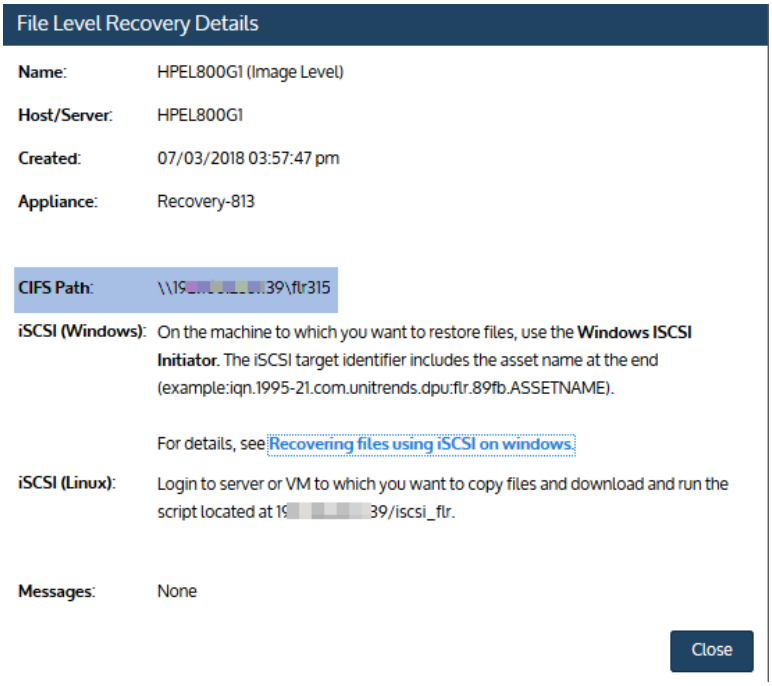

4 Log in to the recovery target workstation.

5 Enter the CIFS path into a file browser on the recovery target.

6 Browse the share to locate the files you want to recover.

Notes:

- Volumes are assigned letters during recovery that may not match the letters from the original disks.

- The Windows file explorer contains a setting to hide protected/system files. Turn off this setting to access all files.

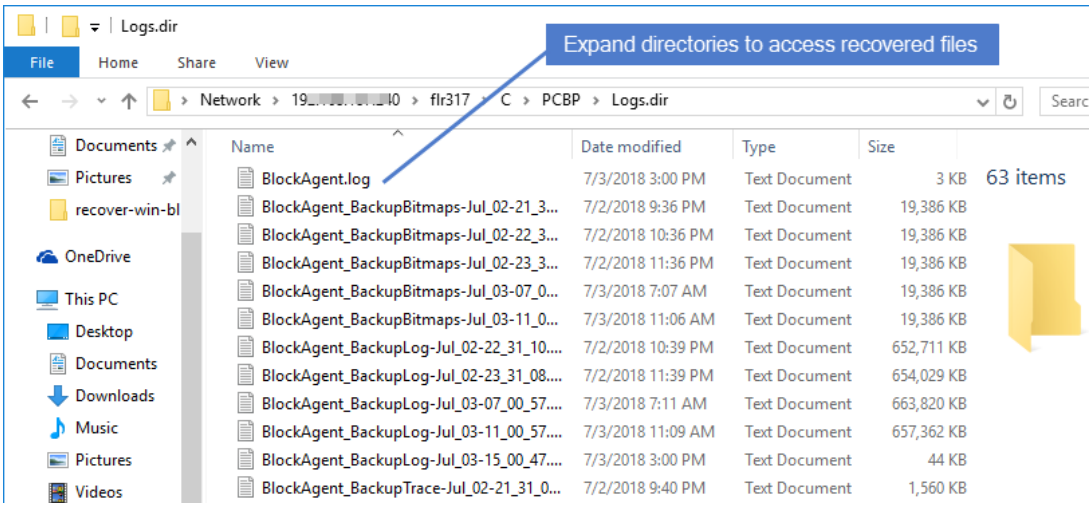

7 Move selected files to another location on the local machine.

8 Disconnect the network share by right-clicking the share and selecting **Disconnect**.

9 Proceed to "Step 3: Remove the recovery object from the appliance" on page 1054.

To recover files by mounting the iSCSI LUN

1 Log in to the recovery target.

2 Launch the iSCSI Initiator from **Administrative Tools** in the **Control Panel**.

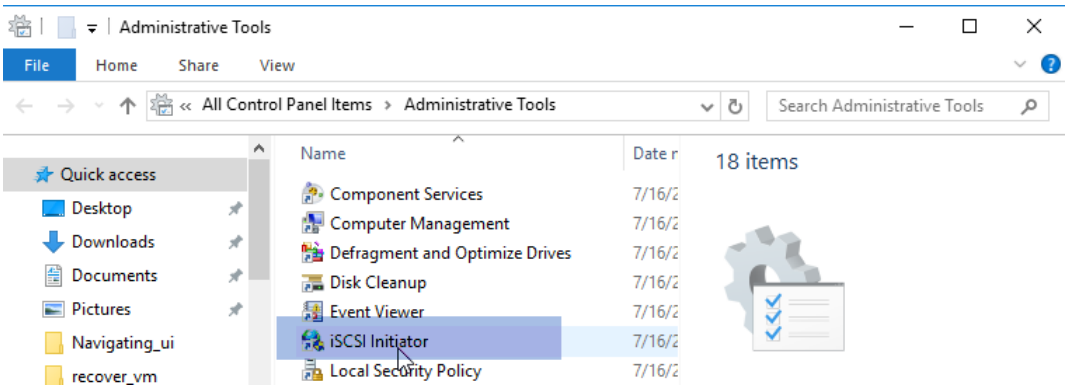

3 In the **Target** field, enter the appliance IP address and click **Quick Connect...**.

The **Discovered targets** field populates with a list of iSCSI LUN targets.

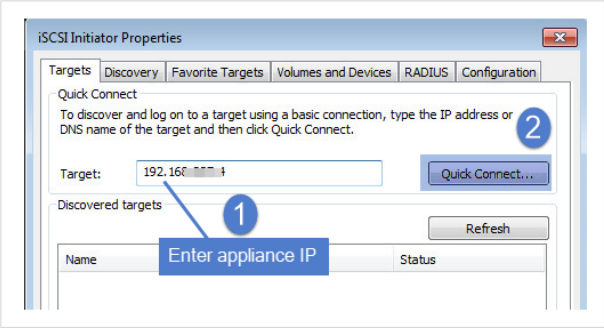

4 Select the iSCSI target from the list.

5 The iSCSI target is discovered and connected to the local machine. Click **Done**.

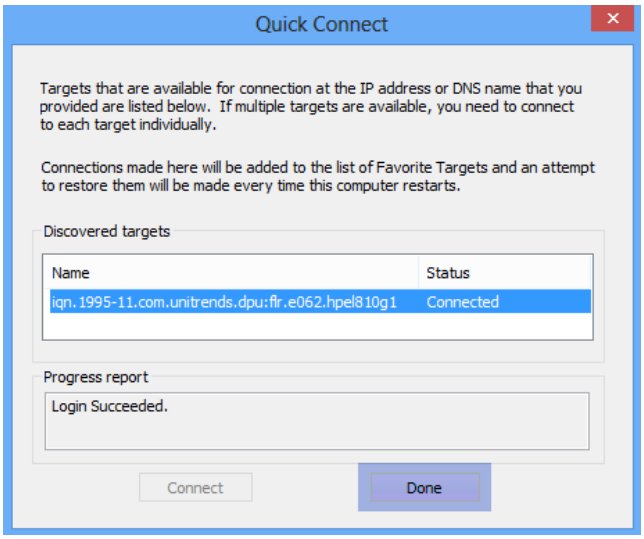

6 Use Disk Manager or diskpart to verify that the mounted iSCSI disk is online. If not, bring the drive online.

7 Return to the iSCSI Initiator. On the Volumes and Devices tab, click **Auto Configure** to map drives from the iSCSI target to the local machine (or map them manually if you prefer).

> **Note:** Volumes are assigned letters during recovery that may not match the letters from the original disks.

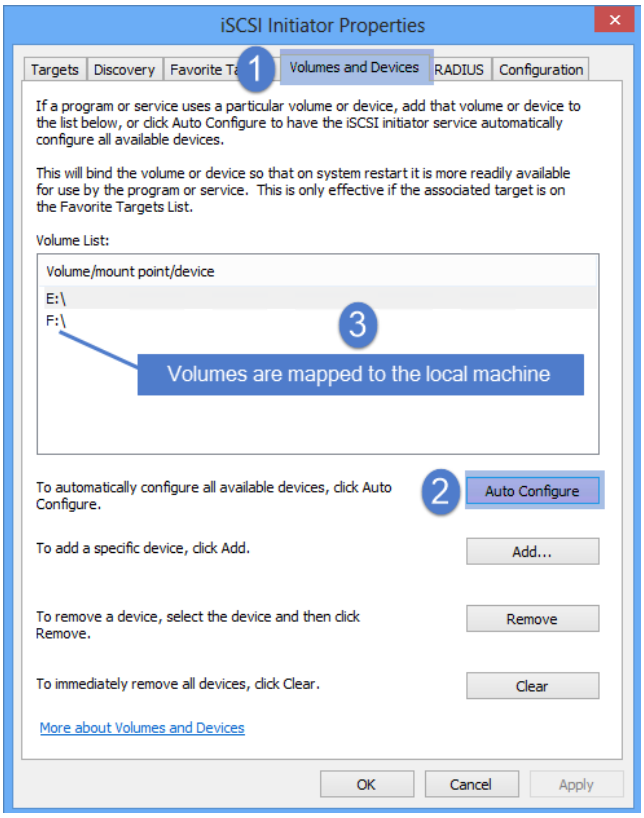

8 Access the files under the mapped drives and move them to another location on the local machine.

Note: The Windows file explorer contains a setting to hide protected/system files. Turn off this setting to access all files.

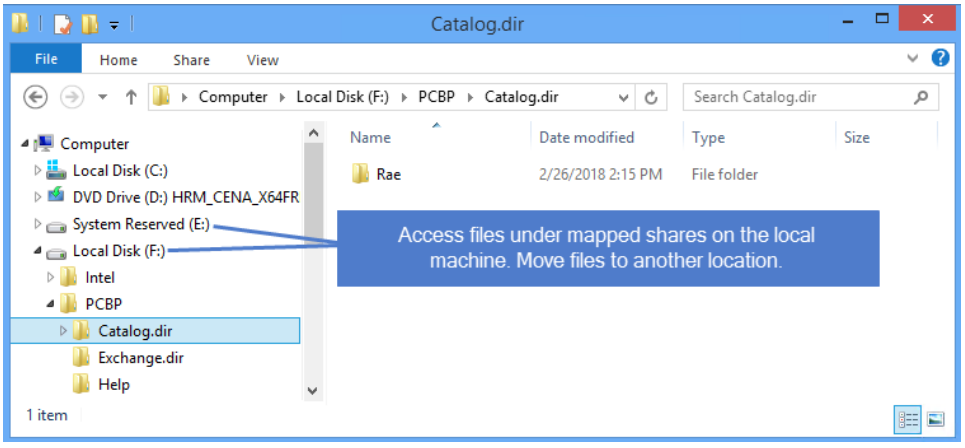

9 Use the iSCSI Initiator to disconnect from the LUN.

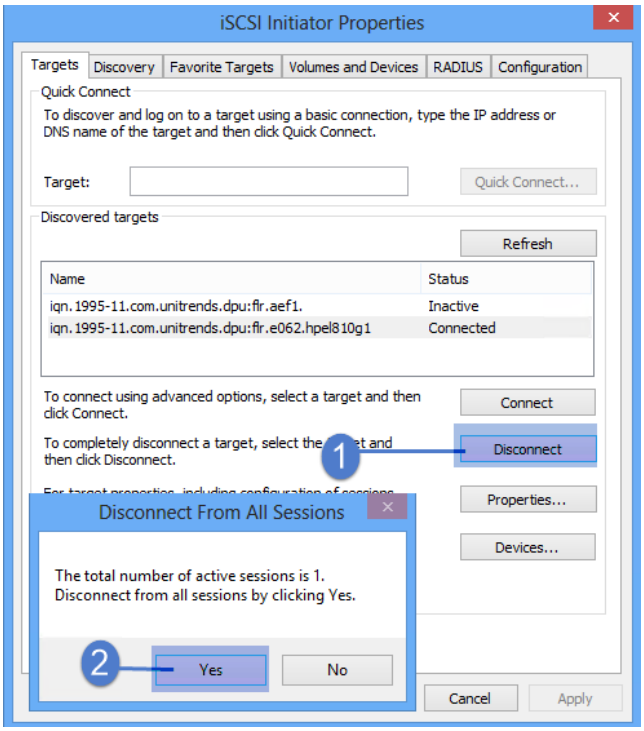

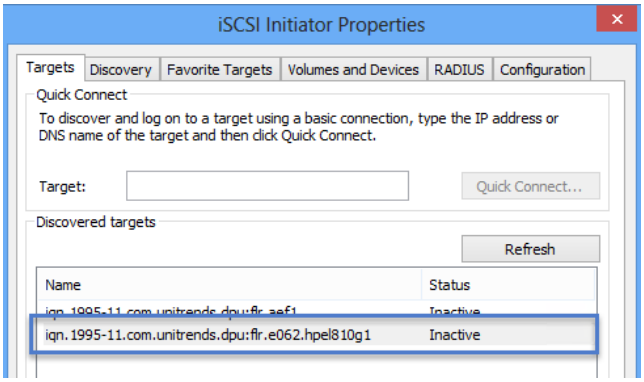

10 Proceed to "Step 3: Remove the recovery object from the appliance".

Step 3: Remove the recovery object from the appliance

To ensure optimal performance, remove the recovery object from the appliance.

WARNING! If you mounted the CIFS share or iSCSI LUN, be sure to unmount it from the target before you remove the recovery object. Removing the recovery object while the target is still connected causes undesired results and errors on the target machine.

To remove a file-level recovery object

1 Select **Recover** and click the **File Level Recovery** tab.

2 Select the recovery object.

3 Click **Remove**.

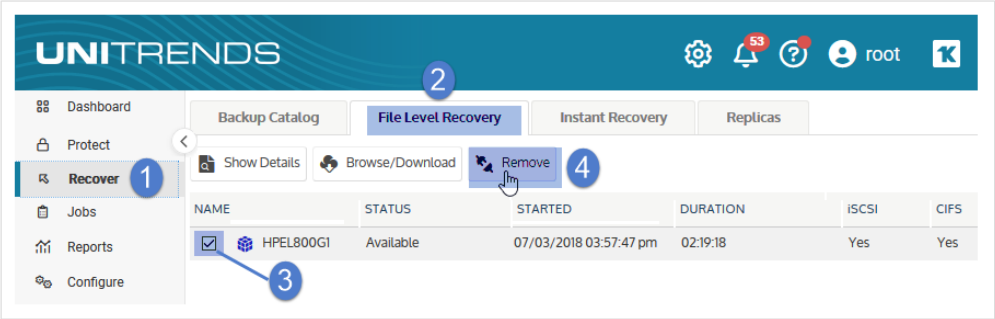

4 Click **Confirm** to continue. The object is removed and no longer displays on the File Level Recovery tab.

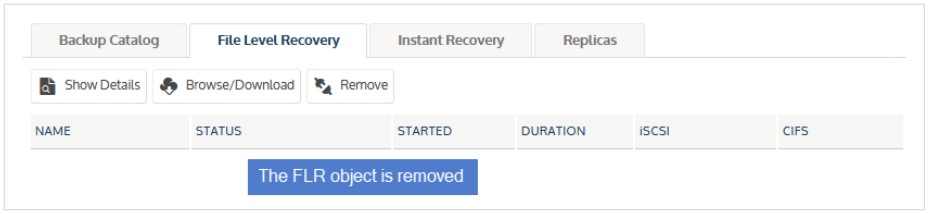

Instant recovery of Windows image-level backups

Instant recovery (IR) enables you to recover a failed or corrupted Windows machine and access it in minutes.

To perform instant recovery, you specify a recovery point (by selecting an image-level backup or backup copy) and a target location where the recovered asset will reside. The recovered asset can reside on:

- A Recovery Series physical appliance
- A Recovery MAX physical appliance
- An ESXi host
- A Hyper-V server

Instant recovery then creates a disk image *recovery object* on the Unitrends appliance and a virtual machine (VM) on the backup appliance or target virtual host. This takes just a few minutes. The IR VM is created by using the backup you

selected and has the same network settings as the original Windows asset, so it can immediately assume production operations.

Upon creating the IR VM, instant recovery migrates data from the on-appliance recovery object to the new VM. The IR VM remains fully operational during the migration.

After the data has been migrated, instant recovery is complete. If you have recovered to a virtual host, the recovery object is no longer needed and you can tear down the instant recovery session. If you have recovered to a Recovery Series or Recovery MAX appliance, the instant recovery session is needed as long as you are using the IR VM.

Tearing down the instant recovery session deletes the recovery object, freeing the appliance resources and reserved space. For appliance performance, it is best to tear down the instant recovery session as soon as possible.

See these topics for details on using the instant recovery feature:

- "Instant recovery modes"
- "Preparing for instant recovery" on page 1057
- "Prerequisites for Windows image-level instant recovery" on page 1058
- "Allocate storage for instant recovery" on page 1068
- "Perform instant recovery in audit mode" on page 1069
- "Perform instant recovery for a failed asset" on page 1075
- "Working with the instant recovery session" on page 1081

Instant recovery modes

You can choose to perform the recovery in *audit* mode or *instant recovery* mode. Use audit mode to verify recovery points for Windows assets that are still running in production. Use instant recovery mode to replace a failed or corrupted Windows asset. Descriptions of each mode are given in the following table:

Mode	Description
Audit	Enables you to verify that a VM can be created from an image-level backup or backup copy. The appliance uses data from the selected backup or backup copy to create a disk image on the appliance and a new VM on the appliance or on the target host. If recovering to a virtual host, the VM resides on the host but it runs from the disk image on the appliance. All other resources, such as the processors and memory, reside on the host. A VM in audit mode is not intended for production use. It does not have network connectivity, and changes made to the VM in audit mode are not backed up on the Unitrends appliance. Applications on the IR VM that require network access are not fully functional in audit mode. Recovering in audit mode has no impact on the original Windows asset. It is not necessary to shut down the original Windows asset during the audit. After verifying that the VM has booted and its data is accessible, you tear down the recovery session. This deletes the IR VM from the appliance or target host and deletes the on-appliance recovery object, freeing the appliance resources.

Mode	Description
Instant recovery	Enables you to replace a failed or corrupted Windows asset. The appliance uses data from the selected backup or backup copy to create a disk image recovery object on the appliance and a VM on the appliance or on the target host. The IR VM is available for use immediately. • IR VM on a Recovery Series or Recovery MAX appliance – The instant recovery session is needed for the IR VM to run. When you are no longer using the IR VM, tear down the IR session to remove both the recovery object and the IR VM. • IR VM on a virtual host – The Unitrends appliance uses Storage vMotion (VMware) or Storage Live Migration (Hyper-V) to copy the data from the disk image to the target Hyper-V or ESXi host. Once data migration is complete, you tear down the recovery session to free appliance resources. The IR VM remains on the virtual host and continues to operate.

Preparing for instant recovery

Unitrends recommends planning for instant recovery before a Windows asset fails. Following is a summary of the steps needed to set up instant recovery for your Windows machines. Steps include links to detailed instructions for each procedure.

Step 1: Ensure that all requirements have been met. For details, see "Prerequisites for Windows image-level instant recovery" on page 1058.

Step 2: Run image-level backups of the Windows assets.

Back up the Windows asset regularly to ensure recent recovery points are available:

- To create a job manually, see "To create an image-level backup job" on page 449.
- To create a job by using an SLA policy, see "To create an SLA policy for Windows image-level assets" on page 545.
- For a comparison of the manual and SLA policy job creation methods, see "About creating backup and backup copy jobs" on page 426.

Step 3: Reserve space on the appliance for instant recovery. For details, see "Allocate storage for instant recovery" on page 1068.

Step 4: (Hypervisor target only) Add target virtual hosts to the Unitrends appliance. (Skip this step if you will be running IR on the backup appliance.)

While running the IR procedure, you select a Hyper-V or ESXi host where the IR VM will be created. If needed, add virtual hosts to the appliance to make them available for IR. For details, see "Adding a virtual host" on page 308.

Note: For VMware, a vCenter is required. Add the ESXi hosts and the vCenter managing the hosts.

Step 5: Run IR in audit mode to check the IR VMs. For details, see "Perform instant recovery in audit mode" on page 1069.

Prerequisites for Windows image-level instant recovery

See the following topics for instant recovery requirements:

- "Backup requirements for IR"
- "Target requirements for IR" on page 1058
- "Windows asset requirements for IR" on page 1065

Backup requirements for IR

An image-level backup of the physical Windows machine is required for instant recovery. The backup must contain all critical system volumes.

Notes:

- By default, image-level backups include all system information needed for instant recovery. If you opt to exclude volumes from backup, use care not to exclude the boot and critical system (OS) volumes.
- The instant recovery virtual machine is created based on the backup you select. Volumes that were excluded from backup are not recovered.
- For SQL, the master, model, and msdb system databases must also be present in the image-level backup of the Windows asset. (These are included by default. If you want the recovered asset to include a hosted SQL application, use care not to exclude these system databases from the image-level backup.)

Back up the Windows asset regularly to ensure recent recovery points are available:

- To create a job manually, see "To create an image-level backup job" on page 449.
- To create a job by using an SLA policy, see "To create an SLA policy for Windows image-level assets" on page 545.
- For a comparison of the manual and SLA policy job creation methods, see "About creating backup and backup copy jobs" on page 426.

Requirements vary by where the IR VM resides. See these topics for details:

- "Choosing the target location"
- "Requirements for recovering to a Recovery Series or Recovery MAX appliance" on page 1059
- "Requirements for recovering to a VMware ESXi host" on page 1060
- "Requirements for recovering to a Hyper-V host" on page 1062

Choosing the target location

Instant recovery creates a virtual machine based on the image-level backup you select. The IR VM can reside on any of the following: a Recovery Series physical appliance, a Recovery MAX physical appliance, an ESXi server, or a Hyper-V server. Considerations for each are give in the following table.

Target location	Considerations
Recovery Series or Recovery MAX physical appliance	Running the IR VM on the backup appliance itself provides a simple, seamless solution. Note that only physical appliances have this option. (An IR VM cannot reside on a Unitrends Backup virtual appliance.) Consider the following to determine whether this is the optimal location for your environment: • Requires no additional hardware. Provides near-zero RTO without having to increase CapEx. • Assumes the role of a failed Windows server temporarily, until you can get new hardware and run a bare metal recovery. • Provides seamless, continuous protection in a failover scenario. The 'live' IR VM is automatically protected by the existing backup schedule after assuming the identity of the original Windows machine. • Uses the compute resources of the appliance to bring the IR VM 'live' in a failover scenario. If the appliance is already under high load, there may not be sufficient resources to provide for instant recovery. • Reduces on-appliance backup retention because a portion of the appliance's storage is reserved for the IR VM.
ESXi or Hyper-V server	Running the IR VM on an ESXi or Hyper-V server enables access to the compute and storage in your virtual environment, greatly increasing the pool of resources that can be used for instant recovery. Consider the following to determine whether this is the optimal location for your environment: • Leverages virtual infrastructure for IR VMs. Disk space and compute resources of the Unitrends appliance are minimally impacted. • Provides the option to use the IR VM as a permanent replacement for the failed Windows asset. Can be used to migrate a physical Windows machine to your virtual infrastructure.

Requirements for recovering to a Recovery Series or Recovery MAX appliance

Ensure that the following requirements have been met before you set up instant recovery.

Requirement	Description
Allocate IR storage	Appliance storage can be used to store backups, for instant recovery, or for Windows replicas. To recover to the Unitrends appliance, you must allocate a portion of the appliance storage to Instant Recovery in the Edit Storage dialog (as described in "Allocate storage for instant recovery" on page 1068). The amount of available instant recovery storage space must be at least 20% of the space used on the original Windows asset.

Requirement	Description
Unitrends appliance resources	Instant recovery uses part of the Unitrends appliance's processors, memory, and storage. This usage may impact the performance of regular system functions (such as backups, backup copies, deduplication, and purging). Monitor the appliance closely and make adjustments as necessary. The IR VM is assigned 2 CPUs and 4GB of RAM by default. If needed, you can modify these settings while creating your IR VM.
IR VM changeability	Consider the following when modifying the IR VM: • In audit mode - The IR VM has no network connectivity. Changes made in audit mode are not captured by Unitrends backups and are lost after you tear down the IR session. • In instant recovery mode - Changes to the IR VM are captured by Unitrends backups of the original Windows asset.
Assets with more than 4 disks	On-appliance instant recovery is supported for assets with up to 4 disks. For assets with more than 4 disks, recover to an ESXi or Hyper-V server instead.
GPT partitioned and UEFI-based assets	Instant recovery of GPT partitioned and UEFI-based assets is supported on Recovery Series and Recovery MAX appliances that are running the CentOS 7 operating system. To check the appliance OS version, select **? > About** in the appliance UI: 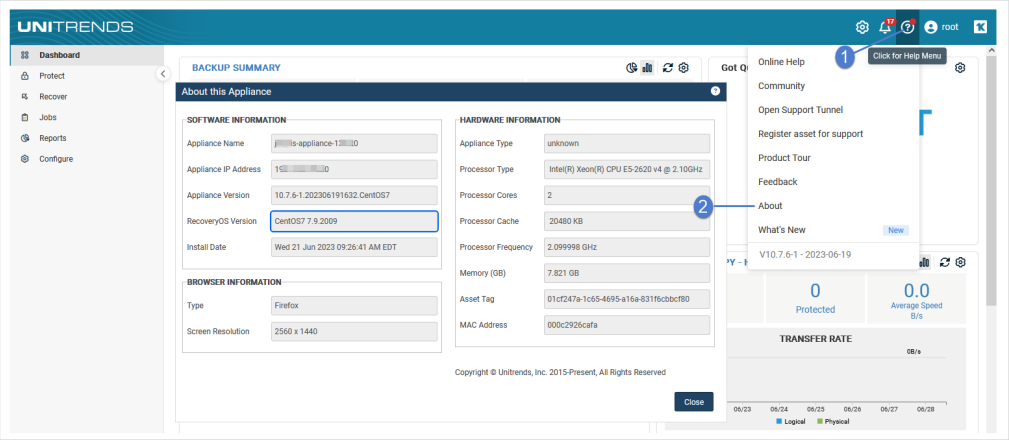

Requirements for recovering to a VMware ESXi host

Ensure that the following requirements have been met before you set up instant recovery.

Requirement	Description
vCenter version and license	To perform instant recovery, the ESXi server used as the instant recovery target must be managed by a vCenter that meets the following requirements: • Must be running one of the following: vCSA 5, vCenter version 5, or a higher vCenter version listed in the Compatibility and Interoperability Matrix. • Must have a license that supports Storage vMotion. • Must be a added to the Unitrends appliance from which you are performing the instant recovery. (For details, see "Adding a virtual host" on page 308.) **Note:** You can perform the audit process using a stand-alone ESXi server, but instant recovery mode is not supported.
ESXi server	The ESXi server used as the instant recovery target must meet the following requirements: • Must be managed by a vCenter that meets the version and license criteria above. • Must be running ESXi version 5 or a higher version listed in the Compatibility and Interoperability Matrix. • Must be added to the Unitrends appliance from which you are performing the instant recovery. (For details, see "Adding a virtual host" on page 308.) • Must support the operating system (OS) of the Windows asset you are recovering. (See the VMware documentation for details.) For example, you cannot recover a Windows 2016 asset to ESXi 5.1. • Must have sufficient space and compute resources for the IR VM.
IR VM resources	The IR VM is assigned resources based on the original asset. CPU is assigned as follows: • If the original asset had more than 6 CPU, the IR VM is assigned 4 CPU. • If the original asset had more than 2 CPU, the IR VM is assigned 2 CPU. • If the original asset had less than 2 CPU, the IR VM is assigned 1 CPU. Memory is assigned as follows: • If the original asset had more than 8GB of RAM, the IR VM is assigned 8GB. • If the original asset had more than 4GB of RAM, the IR VM is assigned 4GB. • If the original asset had more than 2GB of RAM, the IR VM is assigned 2GB. • If the original asset had less than 2GB of RAM, the IR VM is assigned 1GB.

Requirement	Description
IR VM changeability	Consider the following when modifying the IR VM: • In audit mode - The IR VM has no network connectivity. Changes made in audit mode are not captured by Unitrends backups and are lost after you tear down the IR session. **Note:** Performing operations through the hypervisor on an IR VM that is running in audit mode on an ESXi server is not recommended and may yield undesirable results. • In instant recovery mode - Changes to the IR VM are captured by Unitrends backups of the original Windows asset and are retained after you tear down the IR session.
Maximum disk size	The maximum disk size is capped by what the hypervisor supports. The IR VM's disks will be the same size as those on the original asset. For Windows assets with disks larger than 2 TB, the ESXi server must be running ESXi 5.5 or a higher version listed in the Compatibility and Interoperability Matrix.
IR VM configuration	The following configuration settings apply to the IR VM: • The Mac address and network settings of the IR VM match those of the source backup used for the recovery. • The IR VM is configured with the latest hardware version that is supported by the target hypervisor. For example, if recovering to an ESXi 5.5 server, the IR VM is hardware version 10. • The IR VM's disks are provisioned as Thick Eager Zeroed.

Requirements for recovering to a Hyper-V host

Ensure that the following requirements have been met before you set up instant recovery.

Requirement	Description
Hyper-V server	The Hyper-V server used as the instant recovery target must meet the following requirements. • Must be one of the following: – A Windows Server with the Hyper-V role enabled, running Windows 2012 or a higher version listed in the Compatibility and Interoperability Matrix. – A Hyper-V Server running Windows 2012 or a higher version listed in the

Requirement	Description
	Compatibility and Interoperability Matrix. • Must support the operating system (OS) of the Windows asset. (See this Microsoft article for details: Should I create a generation 1 or 2 virtual machine in Hyper-V?) • Must be added to the Unitrends appliance from which you are performing the instant recovery. (For details, see "Adding a virtual host" on page 308.) • Must have sufficient space and compute resources for the IR VM. • Must be running Unitrends agent version 10.3 or higher with the CBT volume driver installed. To enable incremental image-level backups, you must reboot the Hyper-V server after installing the image driver. (For details, see "Installing the Windows agent" on page 362.) **Notes:** • It is best practice to run the latest Unitrends appliance and agent software versions. Older versions do not support all current Unitrends features. • Unitrends recommends reloading the list of VMs on the appliance after installing the agent on the host Hyper-V servers. • To reload the list of VMs, click the **Gear** icon in the upper-right of the UI and select **Inventory Sync**:  • Must be configured for adequate simultaneous storage migrations. The host must be able to run simultaneous storage migrations for all sessions that are currently running in instant recovery mode. (Storage migration is not needed to run IR in audit mode). Inadequate simultaneous storage migrations leads to undesirable results. To modify the number of simultaneous storage migrations, open Hyper-V manager, right-click the host server and select **Properties**, select **Storage Migrations** and modify the **Simultaneous storage migrations** setting:

Requirement	Description
	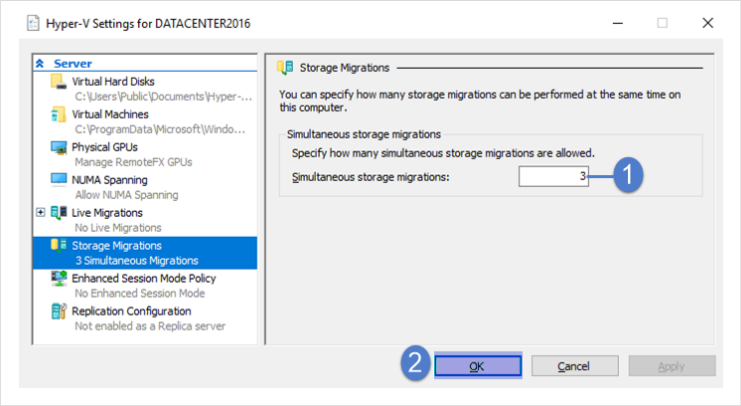
IR VM resources	The IR VM is assigned resources based on the original asset. CPU is assigned as follows: • If the original asset had more than 6 CPU, the IR VM is assigned 4 CPU. • If the original asset had more than 2 CPU, the IR VM is assigned 2 CPU. • If the original asset had less than 2 CPU, the IR VM is assigned 1 CPU. Dynamic memory is used for the IR VM: • If the original asset had more than 2048MB of RAM, the IR VM uses 2GB at startup. Maximum RAM used equals the amount of RAM assigned to the original asset. • If the original asset had less than 2048MB of RAM, the IR VM uses the amount of RAM assigned to the original asset at startup. Maximum RAM used equals twice the amount of RAM assigned to the original asset.
IR VM changeability	Consider the following when modifying the IR VM: • In audit mode - The IR VM has no network connectivity. Changes made in audit mode are not captured by Unitrends backups and are lost after you tear down the IR session. **Note:** Performing operations through the hypervisor on an IR VM that is running in audit mode on a Hyper-V server is not recommended and may yield undesirable results. For example, creating a Hyper-V checkpoint of an IR VM in audit mode causes an error on the hypervisor. • In instant recovery mode - Changes to the IR VM are captured by Unitrends

Requirement	Description
	backups of the original Windows asset and are retained after you tear down the IR session.
Maximum disk size	The maximum disk size is capped by what the hypervisor supports. The IR VM's disks will be the same size as those on the original asset.
IR VM configuration	The IR VM is created with this configuration: • Disk names match those of the source backup used for the recovery. • All disks created on the IR VM must have unique names. • The network settings of the IR VM match those of the source backup used for the recovery. • The IR VM is configured with the latest hardware generation version that is supported by the target Hyper-V server. • The IR VM's configuration version is the highest version that the hypervisor supports.
Hyper-V clusters	If recovering to a Hyper-V host that is part of a cluster, instant recovery treats the host as a stand-alone Hyper-V server. The IR VM is not part of the cluster and does not fail over to other cluster nodes.

Windows asset requirements for IR

The Windows asset must meet the following requirements to use the instant recovery feature:

Requirement	Description
Client Operating Systems	Instant recovery is supported for the client operating systems listed below. Additional version limitations apply. See the Compatibility and Interoperability Matrix for details. • Windows 7, 64-bit only • Windows 8, 64-bit only • Windows 8.1, 64-bit only • Windows 10, 64-bit only • Windows 11, 64-bit only

Requirement	Description
	Note: To run instant recovery on a virtual host, the Hyper-V or ESXi host must support the guest OS of the IR VM. (See the Microsoft or VMware documentation for details.) For example, a Windows 10 IR VM cannot reside on ESXi 5.1 or Hyper-V 2008 R2.
Server Operating Systems	Instant recovery is supported for the server operating systems listed below. Additional version limitations apply. See the Compatibility and Interoperability Matrix for details. • Windows 2008 R2, 64-bit only • Windows 2012, 64-bit only • Windows 2012 R2, 64-bit only • Windows 2016, 64-bit only • Windows 2019, 64-bit only • Windows 2022, 64-bit only **Note:** To run instant recovery on a virtual host, the Hyper-V or ESXi host must support the guest OS of the IR VM. (See the Microsoft or VMware documentation for details.) For example, a Windows 2016 IR VM cannot reside on ESXi 5.1 or Hyper-V 2008 R2.
Firmware interface type	The instant recovery feature is supported for BIOS- and UEFI-based assets. For BIOS-based assets, the boot and system disks must be MBR. For UEFI-based assets, these requirements apply: • Boot and system disks must be GPT. • The IR VM must reside on one of the following: – An ESXi server running ESXi 5.1 or a higher version listed in the Compatibility and Interoperability Matrix. – A Hyper-V server running 2012 R2 or a higher version listed in the Compatibility and Interoperability Matrix. – A Recovery Series or Recovery MAX appliance that is running the CentOS 7 operating system. **Note:** A UEFI-based asset cannot run on a physical appliance that is running CentOS 6. For details on checking the appliance OS, see "GPT partitioned and UEFI-based assets" on page 1060.

Requirement	Description
Disk partition type	Instant recovery is supported for Master Boot Record (MBR) and GUID Partition Table (GPT) partition types. For GPT partitioned assets, the IR VM must reside on one of the following: • An ESXi server • A Hyper-V server • A Recovery Series or Recovery MAX appliance that is running the CentOS 7 operating system. **Note:** A GPT partitioned asset cannot run on a physical appliance that is running CentOS 6. For details on checking the appliance OS, see "GPT partitioned and UEFI-based assets" on page 1060.
Deduplicated volumes	Data is applied to the IR VM in its non-deduplicated form. • If recovering to a Recovery Series or Recovery MAX appliance, ensure that enough IR storage space has been allocated to house this non-deduplicated data. • If recovering to a virtual host, ensure that the datastore (VMware) or volume (Hyper-V) that will be used to create the VM disks has enough capacity to house this non-deduplicated data.
Number of volumes	The Windows asset can have a maximum of 20 volumes, including the System Reserved volume and other unmounted volumes. An IR VM with more than 20 volumes may fail to boot.
Separate boot and system partitions	For Windows assets with boot and system partitions located on different disks, the system partition must reside on the first disk (Disk 0).
File System Configuration	Instant recovery is supported for the following file systems: • NTFS • FAT/FAT32/exFAT • ReFS (Windows 2012 and later)
Active Directory	Instant recovery supports Active Directory database (NTDS) located on the boot volume only. (If it is not on the boot volume, the configuration is not supported and you see an error message when you attempt to create the IR VM.)

Allocate storage for instant recovery

Instant recovery can be performed at any time, as long as there are image-level backups or backup copies of your Windows assets and sufficient backup storage is allocated as instant recovery space. Unitrends strongly recommends reserving space for instant recovery soon after initial deployment and before an asset fails. If necessary, you can allocate instant recovery space later, but doing so may require the appliance to purge local backups to make room for the newly allocated instant recovery space.

Because the disks for a recovered asset reside on the appliance until storage migration completes, you must allocate a portion of your backup storage for instant recovery. Allocate at least 20% of the space used on the original Windows asset. Additionally, ensure that there is enough extra storage to account for anticipated growth during the recovery process, especially for long-running recoveries. Once disks have been migrated to another location, or BMR is complete and the IR is torn down, the appliance's instant recovery storage is no longer needed.

Notes:

- Appliance storage that is allocated to instant recovery can also be used for the Windows replica feature.
- Storage allocated for instant recovery cannot be used for backups.
- Physical appliances come with a set amount of backup storage. Backup storage allocation can be modified to increase the amount used for instant recovery. Backup storage cannot be added to the appliance.
- Unitrends Backup virtual appliances are deployed as virtual machines. After initial deployment, you can add more backup storage as desired. For details, see "About adding backup storage to a Unitrends Backup appliance" on page 199.

To allocate storage for instant recovery

1 On the **Configure > Appliances** page, select the appliance.

2 On the **Storage** tab, select the **Internal** storage and click **Edit**.

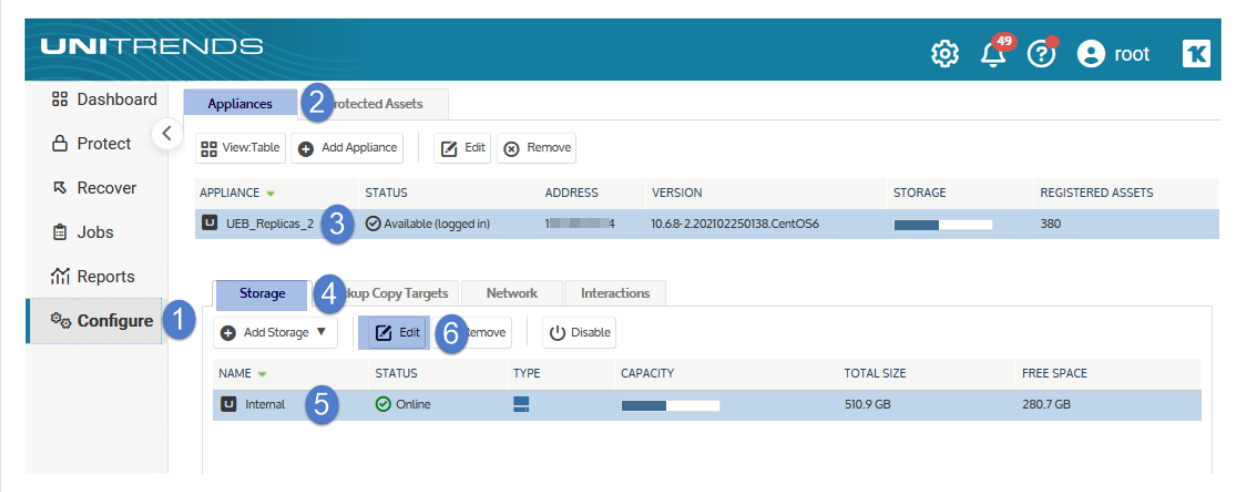

3 Modify the percentages used for backups versus instant recovery, and click **Save**.

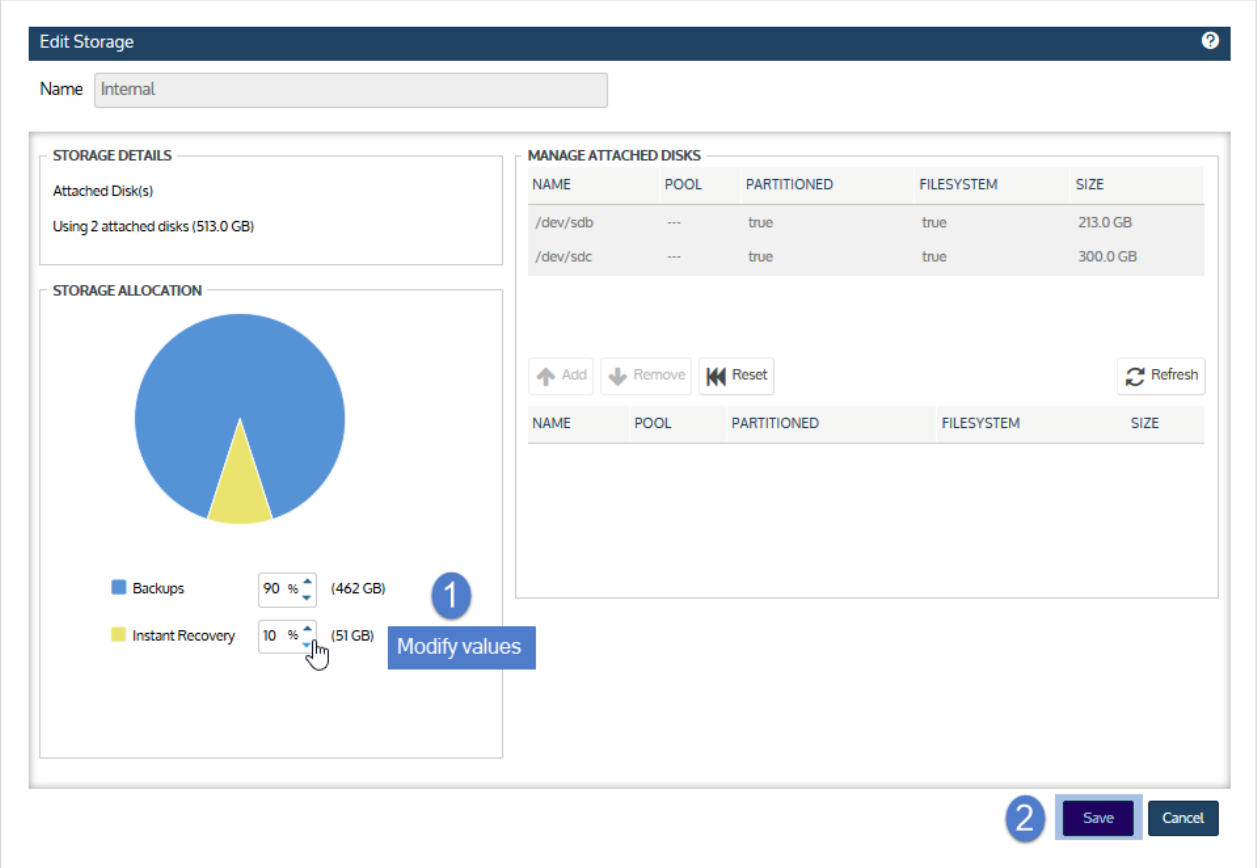

Perform instant recovery in audit mode

Perform instant recovery in audit mode to verify that backups and backup copies can be used to recover the Windows asset in the event of a disaster. Repeat this procedure as needed to test new backups.

To perform instant recovery in audit mode

Run this procedure from the backup appliance to recover from a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover from a hot backup copy.

1. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. (For details, see "Working with custom filters" on page 67.)
3. Expand the Windows asset and select one of the following to use for the recovery:
   - An image-level backup.

- An imported image-level backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
- A hot image-level backup copy (supported if performing IR on the target appliance where the hot copy resides).

Note: If you are running both file-level and image-level backups for the Windows asset, two asset instances display in the catalog: one for file-level backups and one for image-level backups. Be sure to select an image-level backup (where the application is *image level*).

4 Click **Instant Recovery**.

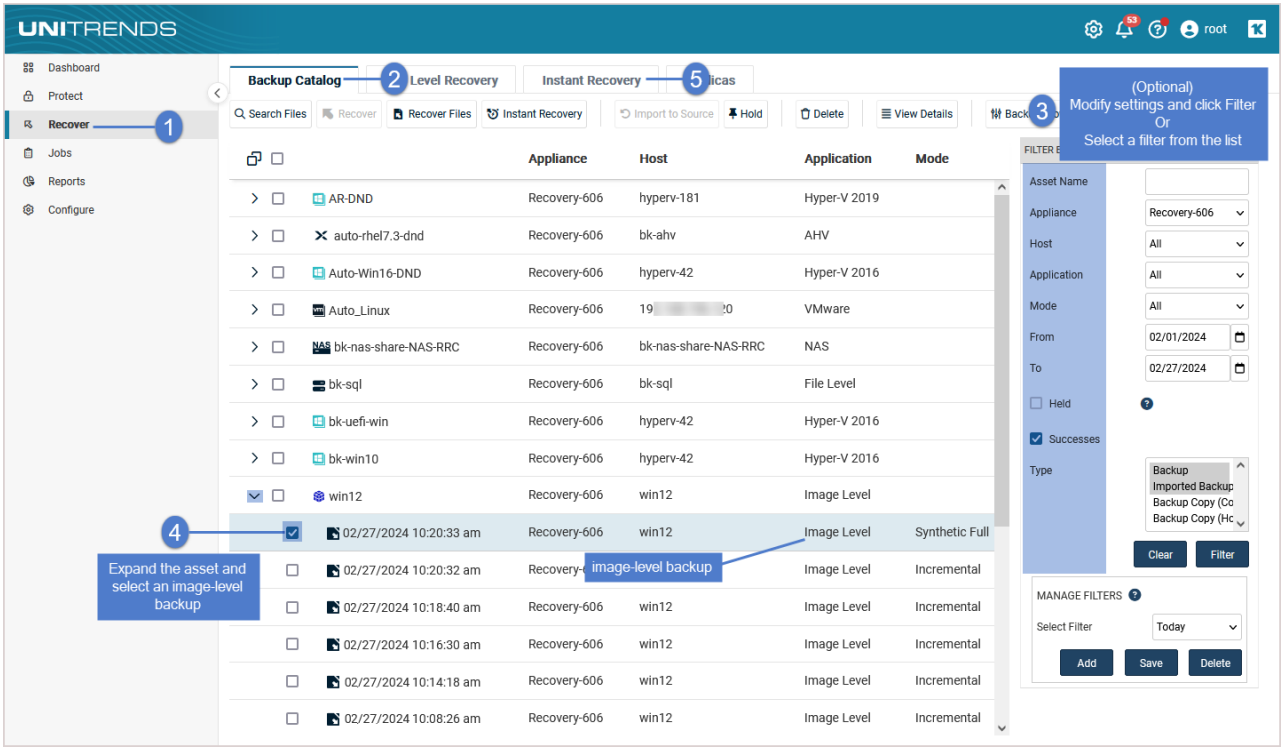

5 Check the **Recover this VM in Audit Mode** box.

6 Enter Location Details to specify the location where the IR VM will reside.

The details that display vary by Target Hypervisor Type. The list contains only the types that are available in your environment. For example, the Unitrends Appliance type is not an option for Unitrends Backup virtual appliances. The Hyper-V Host type is not an option if a compatible Hyper-V virtual host asset has not been added to the appliance.

- If you select **Unitrends Appliance**, the IR VM is assigned 2 CPUs and 4GB of RAM by default. Modify these settings as needed. Then continue to step 7.

Note: Use care when modifying these settings to ensure that the appliance has adequate CPU and RAM for regular operations. For example, Unitrends recommends reserving at least 50% of the available RAM.

- If you select **VMware Host** or **Hyper-V Host**, provide details for the virtual host server. See the table below for descriptions of the Location Details fields.

Item	Description
Target Hypervisor Type	Select a location type from the list (Unitrends Appliance, VMware Host, or Hyper-V Host).
Target Location	Select a VMware or Hyper-V host from the list. The list contains all VMware and Hyper-V virtual host assets that have been added to the appliance and are compatible with the Windows asset. For example, an ESXi 5.1 host does not display for a Windows 2016 asset. (For details, see "Adding a virtual host" on page 308.)
Target Resource Pool (VMware only)	(Optional) If your VMware environment has resource pools, you can opt to select one in the list.
Target Storage	Select the datastore (VMware) or volume (Hyper-V) that will be used to create the VM's disks.
Target Appliance Network	Select a virtual network from the list. The list contains the virtual networks that are discovered and available on the VMware or Hyper-V host.
Network Switch (Hyper-V only)	If the Hyper-V host has multiple switches, select one from the list. The list contains all network switches that are discovered and available on the Hyper-V host.

7 Click **Next**.

Virtual host example:

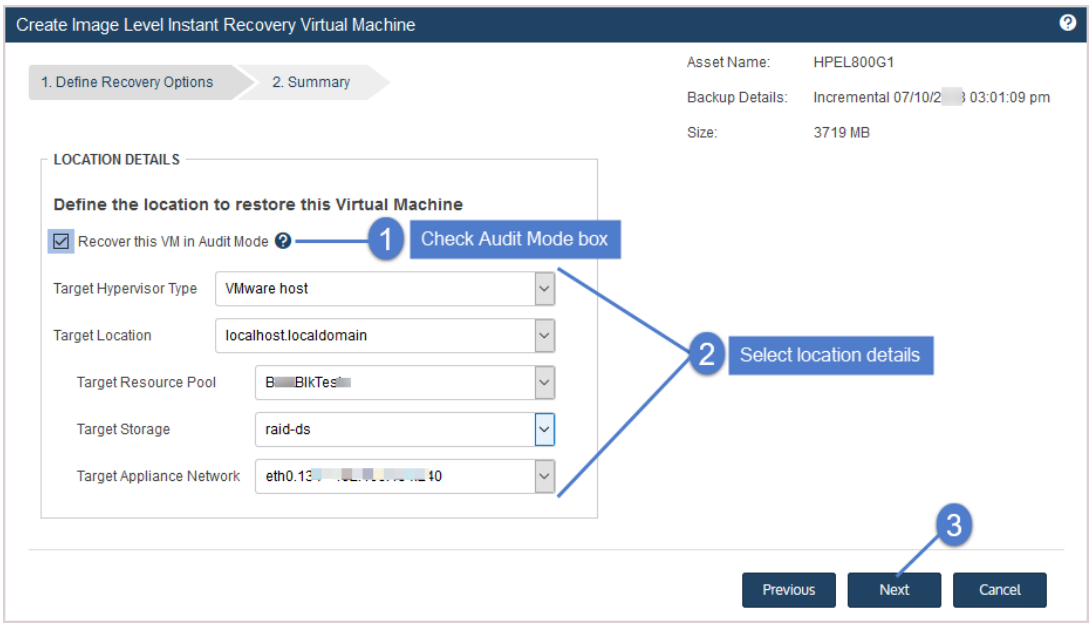

Unitrends appliance example:

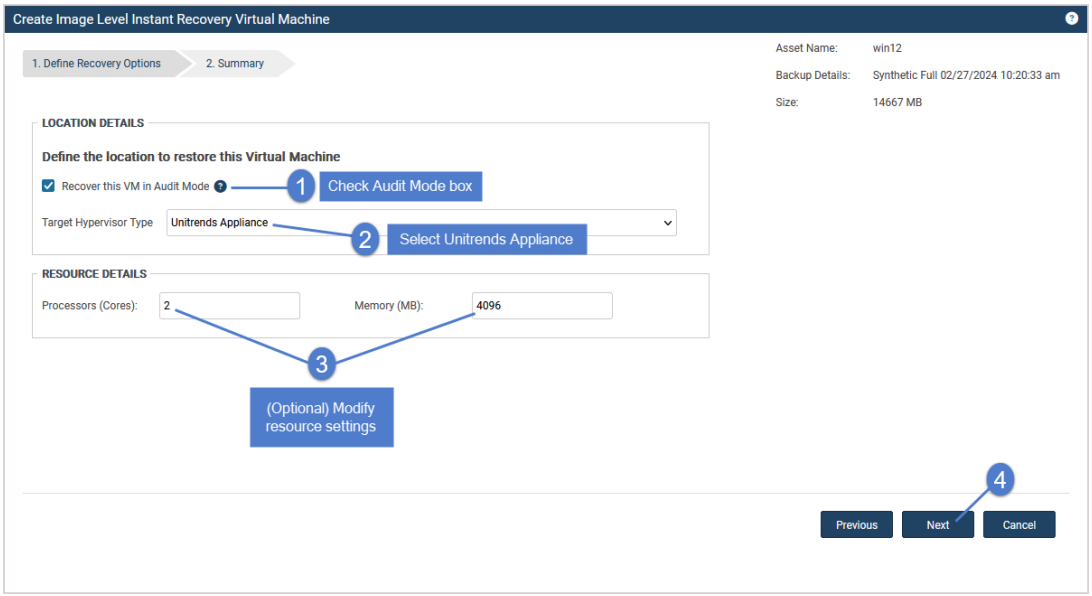

8 Review the recovery settings.

The VM is created with the following default name: <*Windows_asset_name*>*_restore*. To change this name, click the pencil icon.

9 Click **Save** to start the recovery.

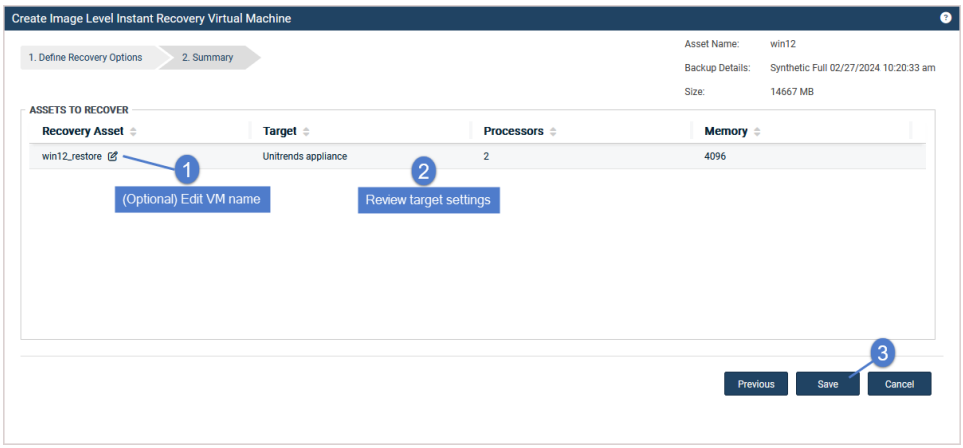

10 (If needed) Address any error or warning messages as shown below.

- Error – If you see an error, click **Back** to resolve the issue.

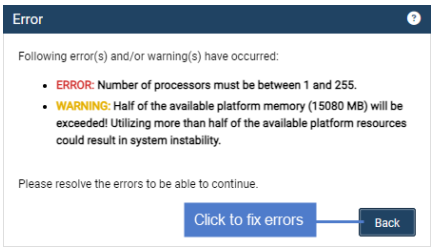

- Warning – If you see a warning message, click **Back** to address the warning or **Continue** to create the IR VM without addressing the warning condition.

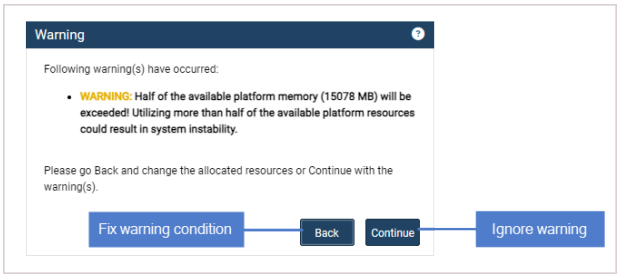

11 Click **OK** to close the Information message.

The VM is created on the appliance or target virtual host. The VM has no network settings.

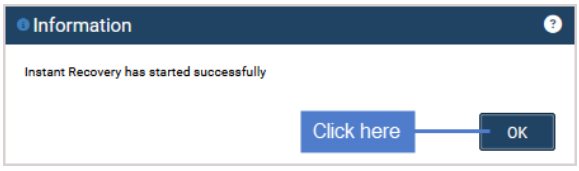

12 Check IR progress by viewing details on the Instant Recovery tab.

Note: If you are recovering an incremental that has a long chain of dependent incrementals, it can take extra time to create the IR VM.

- Click the **Instant Recovery** tab. Active IR sessions display by VM name.
- Click to select the IR session you created.
- Click **View Details** above.
- Status messages display in the Details tab below. IR moves through several phases. The VM can be accessed when you see this message: *VM is available for use*.

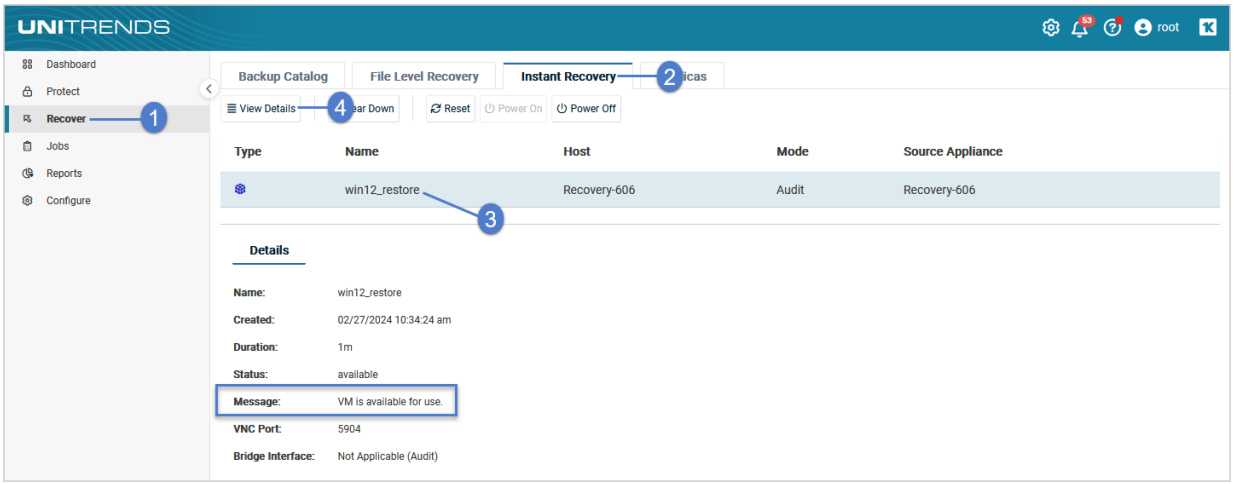

13 Access the IR VM to verify that it is functioning as expected. See one of the following for details:

- "To access an IR VM on a Recovery Series or Recovery MAX appliance" on page 1081

OR

- "To access an IR VM on an external hypervisor" on page 1083

Notes:

- Any features or applications that require network access do not have full functionality in audit mode.
- Any changes made to the IR VM in audit mode are lost when you tear down the IR session. These changes are not applied to the production Windows asset that you are auditing.
- In rare instances, after IR is performed for a Windows server, a disk may be inaccessible because it has been placed in an offline state. To bring disks into an online state, login to the VM, go to Disk Management, right-click on the offline disk, and select **Online** from the drop-down menu.
- While running the IR VM, you can power on, power off, or reset the IR VM right from the Instant Recovery tab. For details, see "Working with the instant recovery session" on page 1081.

14 Tear down the IR session:

- On the Instant Recovery tab, click to select the IR session.

- Click **Tear Down**.

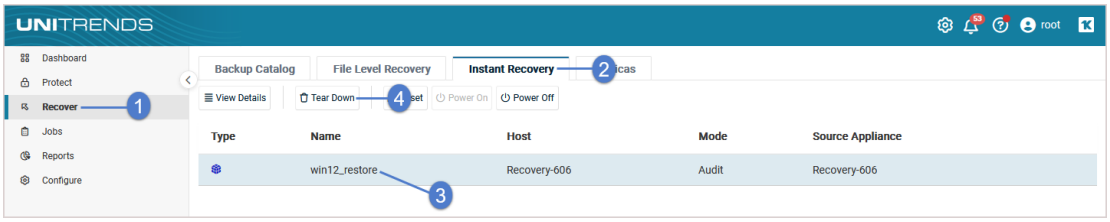

- Click **Confirm**.

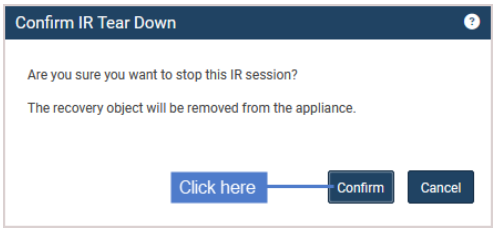

The recovery object is removed from the appliance and the IR VM is removed from the appliance or target virtual host.

Perform instant recovery for a failed asset

Perform instant recovery after the Windows asset fails. (If the Windows asset has not failed, you can use audit mode to verify that the IR VM can be created from a backup. See "Perform instant recovery in audit mode" on page 1069 for details.)

Recommendations for instant recovery mode

Review these recommendations before going into instant recovery mode:

- Instant recovery mode should be used temporarily. The appliance begins sending alerts after a live IR VM has run for 14 days.
- You can tear down the IR session after you have recovered to new hardware (supported in all cases) or by retaining the IR VM as a permanent replacement (supported only for IR VMs that reside on external hypervisors).
- A live IR VM running on an ESXi or Hyper-V server uses hypervisor resources. Disk space and compute resources of the Unitrends appliance are minimally impacted. Tear down the IR session as soon as VM data has been migrated to the hypervisor. This removes the IR session only. The IR VM remains fully functional on the target hypervisor.

To perform instant recovery

IMPORTANT! Be sure to shut down the original Windows asset before running this procedure.

Run this procedure from the backup appliance to recover from a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover from a hot backup copy.

1 Log in to the backup appliance or target appliance (if recovering from a hot backup copy).

2 Select **Recover** and click the **Backup Catalog** tab.

(Optional) Use Filter Backups to the right to customize the backups that display. (For details, see "Working with custom filters" on page 67.)

3 Expand the Windows asset and select one of the following to use for the recovery:

- An image-level backup.
- An imported image-level backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
- A hot image-level backup copy (supported if performing IR on the target appliance where the hot copy resides).

If you are running both file-level and image-level backups for the Windows asset, two asset instances display in the catalog: one for file-level backups and one for image-level backups. Be sure to select an image-level backup (where the application is *image level*).

4 Click **Instant Recovery**.

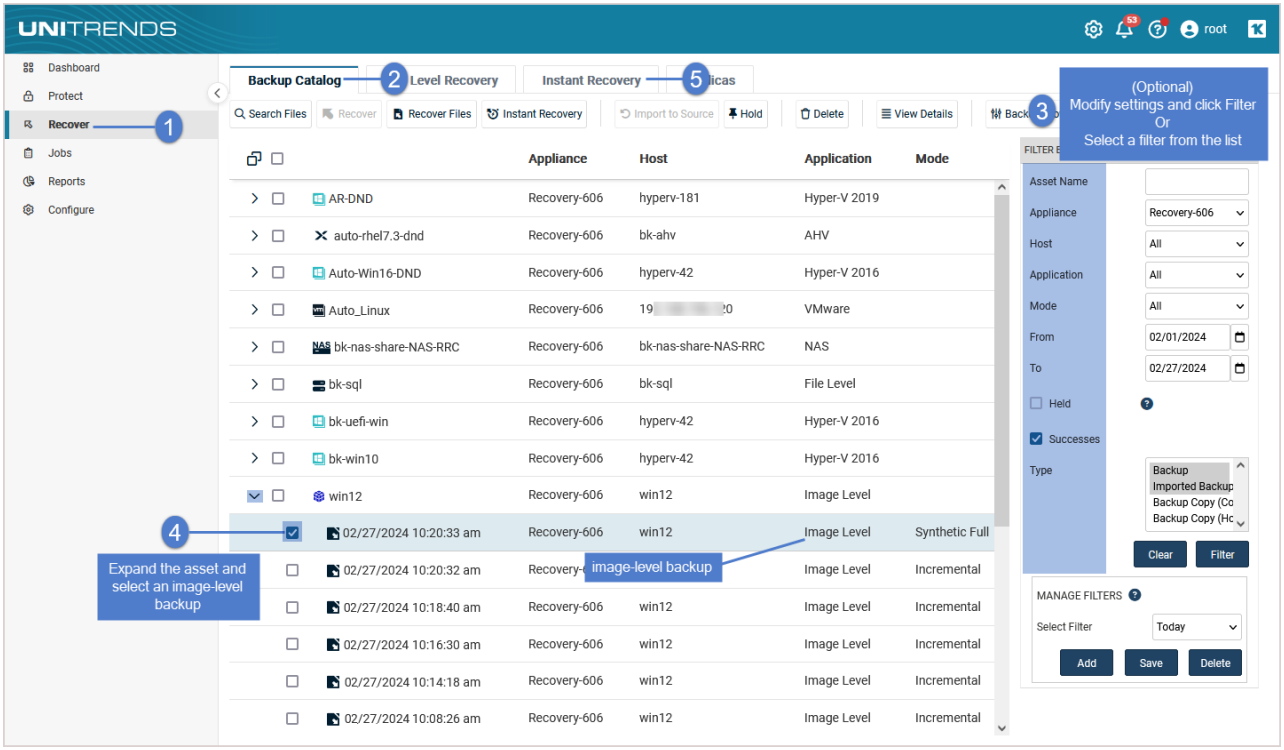

5 Do *not* check the **Recover this VM in Audit Mode** box.

This box must be unchecked to recover the failed Windows asset with its original network settings.

Note: If the Windows asset has not failed, you can "Perform instant recovery in audit mode" by checking this box.

6 Enter Location Details to specify the location where the IR VM will reside.

The details that display vary by Target Hypervisor Type. The list contains only the types that are available in your environment. For example, the Unitrends Appliance type is not an option for Unitrends Backup virtual appliances. The Hyper-V Host type is not an option if a compatible Hyper-V virtual host asset has not been added to the appliance.

- If you select **Unitrends Appliance**, the IR VM is assigned 2 CPUs and 4GB of RAM by default. Modify these settings as needed. Then continue to step 7.

  Note: Use care when modifying these settings to ensure that the appliance has adequate CPU and RAM for regular operations. For example, Unitrends recommends reserving at least 50% of the available RAM.

- If you select **VMware Host** or **Hyper-V Host**, provide details for the virtual host server. See the table below for descriptions of the Location Details fields.

Item	Description
Target Hypervisor Type	Select a location type from the list (Unitrends Appliance, VMware Host, or Hyper-V Host).
Target Location	Select a VMware or Hyper-V host from the list. The list contains all VMware and Hyper-V virtual host assets that have been added to the appliance and are compatible with the Windows asset. For example, an ESXi 5.1 host does not display for a Windows 2016 asset. (For details, see "Adding a virtual host" on page 308.)
Target Resource Pool (VMware only)	(Optional) If your VMware environment has resource pools, you can opt to select one in the list.
Target Storage	Select the datastore (VMware) or volume (Hyper-V) that will be used to create the VM's disks.
Target Appliance Network	Select a virtual network from the list. The list contains the virtual networks that are discovered and available on the VMware or Hyper-V host.
Network Switch (Hyper-V only)	If the Hyper-V host has multiple switches, select one from the list. The list contains all network switches that are discovered and available on the Hyper-V host.

7 Click **Next**.

Virtual host example:

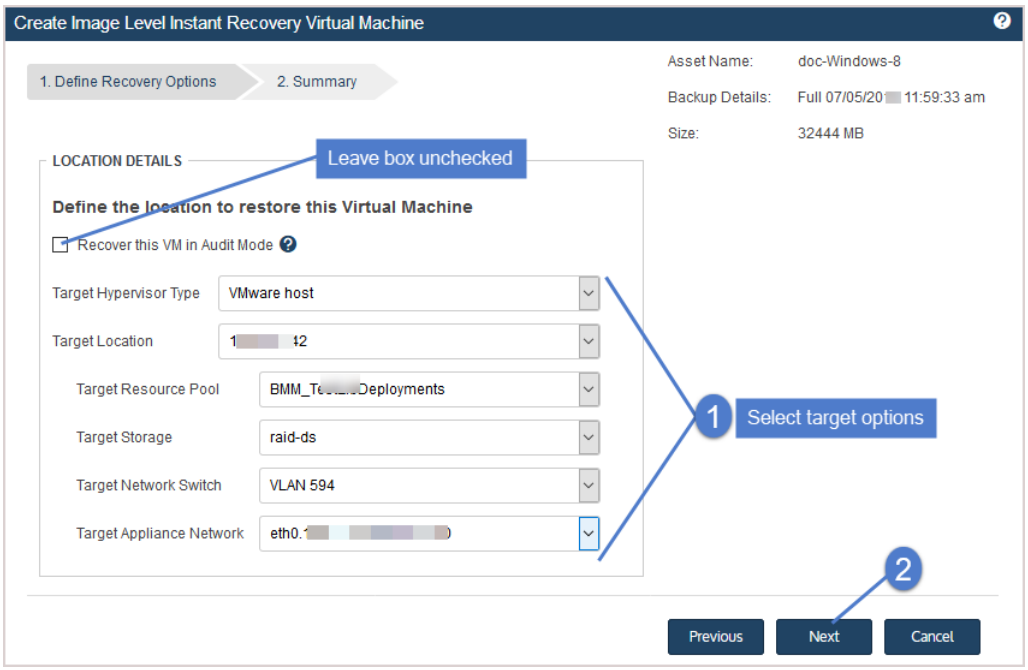

Unitrends appliance example:

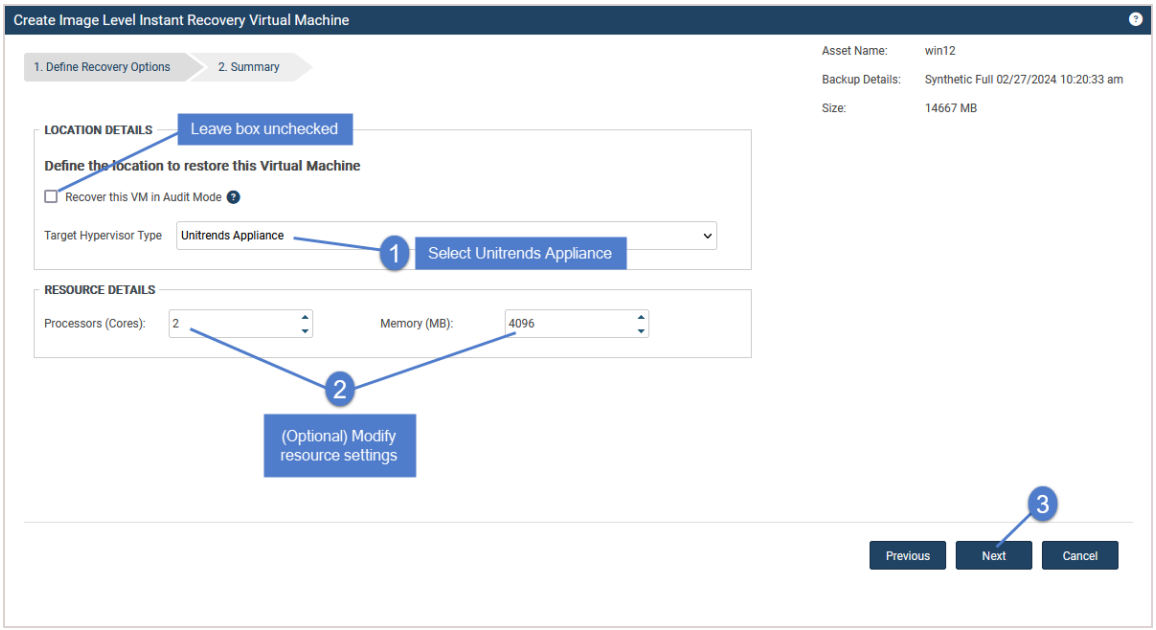

8 Review the recovery settings.

The VM is created with the following default name: <*Windows_asset_name*>*_restore*. To change this name, click the pencil icon.

9 Click **Save** to start the recovery.

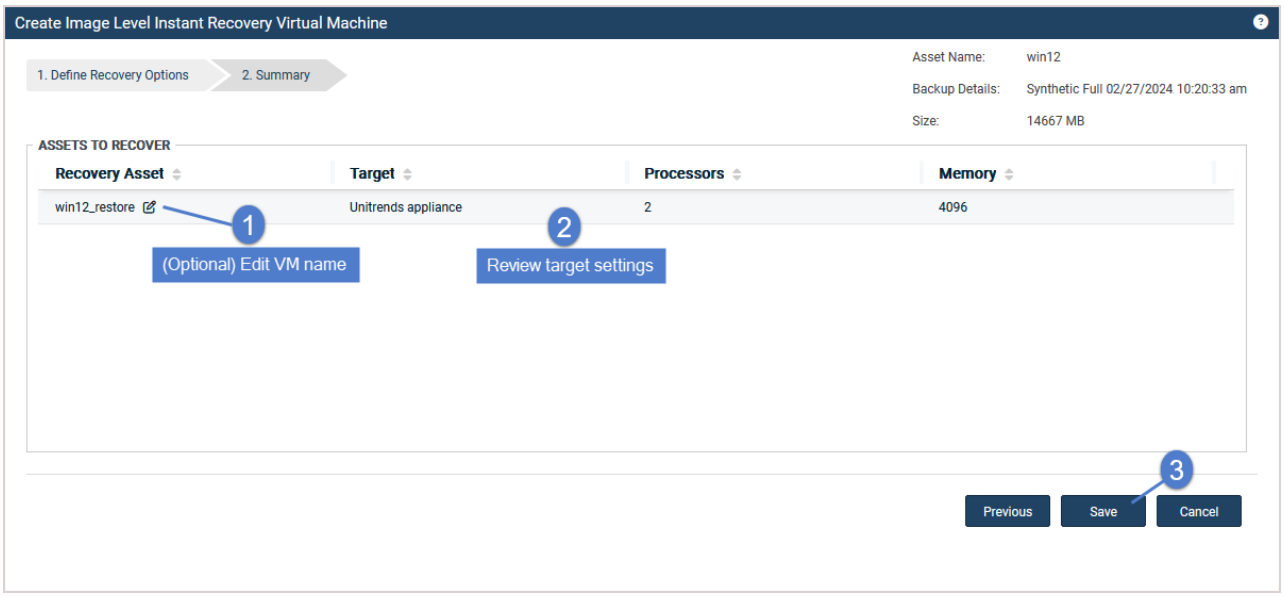

10 (If needed) Address any error or warning messages as shown below.

- Error – If you see an error, click **Back** to resolve the issue.

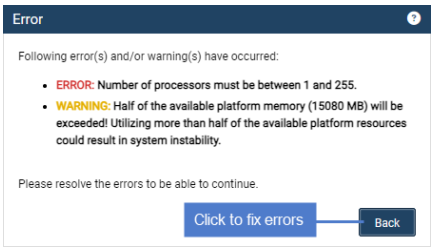

- Warning – If you see a warning message, click **Back** to address the warning or **Continue** to create the IR VM without addressing the warning condition.

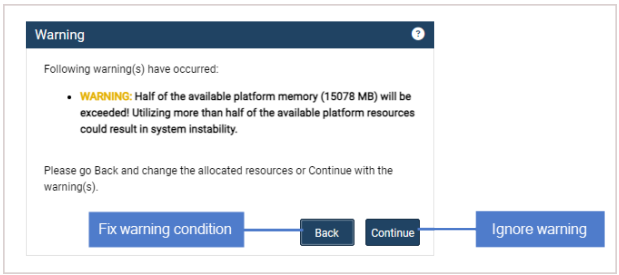

11 Click **OK** to close the Information message.

The recovery object and IR VM are created, and data migration begins. During this time, the VM is fully operational.

IMPORTANT! Do not tear down the recovery session or modify the VM in any way until data migration is complete.

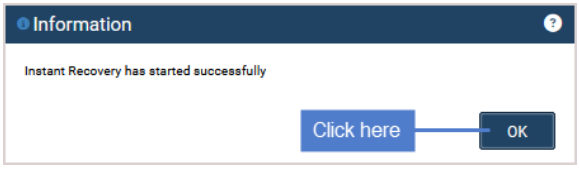

12 Check IR progress on the Instant Recovery tab:

- Click the **Instant Recovery** tab. Active IR sessions display by asset name.
- Click to select the IR session you created.
- Click **View Details**.
- Status messages display in the Details tab below. IR moves through several phases. Data migration is complete when you see this message: *VM is available for use*.

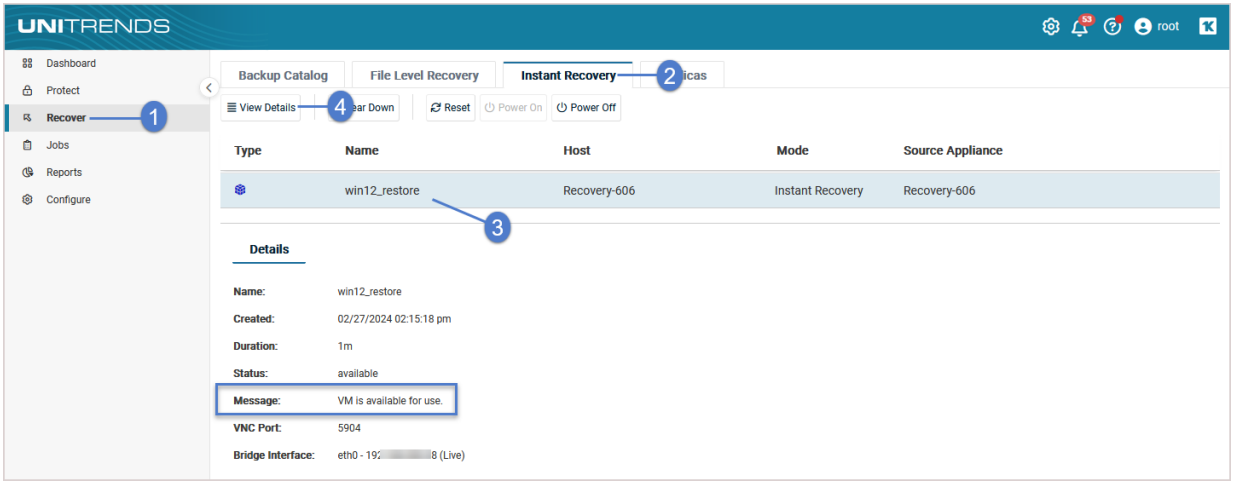

13 When you see the *VM is available for use* message, all data has been migrated and IR is complete.

- The IR VM has the same network settings and username/password credentials as the original Windows asset.
- Access the VM and verify that it is functioning as expected in production. See these topics for details: "To access an IR VM on a Recovery Series or Recovery MAX appliance" on page 1081 or "To access an IR VM on an external hypervisor" on page 1083.

Note: In rare instances, after a restore is performed for a Windows server, a disk may be inaccessible because it has been placed in an offline state. To bring disks into an online state, log in to the VM, go to Disk Management, right-click the offline disk, and select **Online** from the drop-down menu.

14 Modify VM settings and backup schedules as needed to begin protecting the IR VM. The next backup of the IR VM is promoted to a full.

15 Next steps vary by where the IR VM resides. Do one of the following:

- IR VM running on a Recovery Series or Recovery MAX appliance – Tearing down the IR session removes the IR VM from the appliance. After you have recovered the original asset to new physical hardware, tear down the IR session as described in "Tearing down the instant recovery session" on page 1085.
- IR VM running on a VMware or Hyper-V host – The IR session is not needed once data has migrated to the IR VM. Tear down the IR session as described in "Tearing down the instant recovery session" on page 1085.

Note: While running the IR VM, you can power on, power off, or reset the virtual machine right from the Instant Recovery tab. For details, see "Working with the instant recovery session" on page 1081.

Working with the instant recovery session

While your IR is running in audit or live mode, use these procedures as needed:

- "To access an IR VM on a Recovery Series or Recovery MAX appliance"
- "To access an IR VM on an external hypervisor" on page 1083
- "To power off the IR VM" on page 1083
- "To power on the IR VM" on page 1083
- "To reset the IR VM" on page 1084
- "Tearing down the instant recovery session" on page 1085

To access an IR VM on a Recovery Series or Recovery MAX appliance

Use this procedure to access an IR VM that is running on a physical appliance in audit or live mode:

Note: You must use a VNC viewer to access the IR VM in audit or live mode on a physical appliance. If necessary, download one to your workstation before running this procedure.

1 On the Instant Recovery tab, view details to obtain the VNC port number:

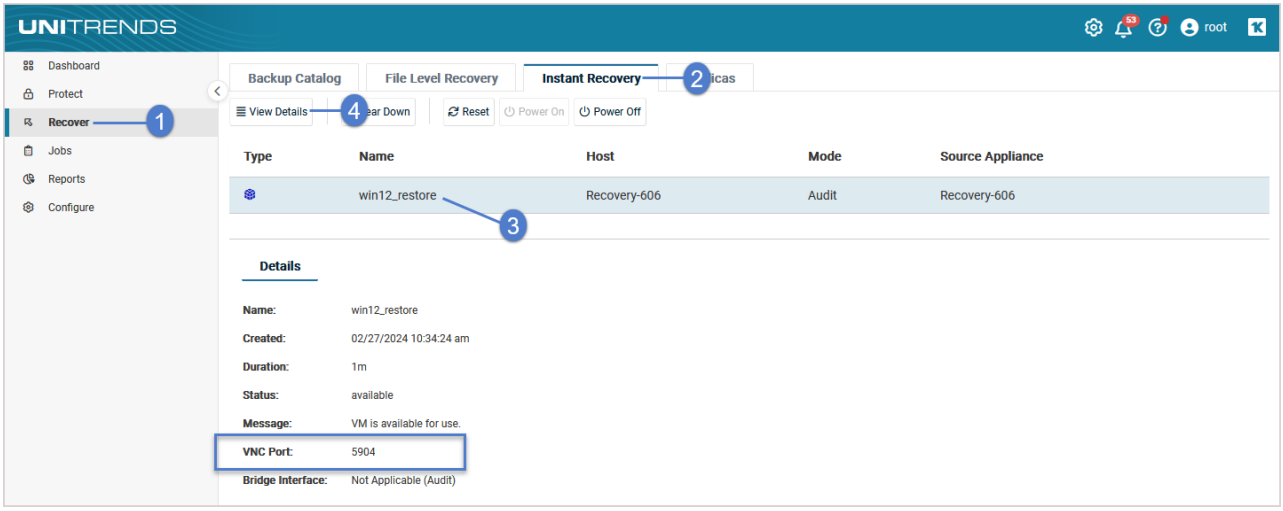

2 Open a VNC viewer. Connect to the VM by entering: **<ApplianceIP>::<VNCport>**

> **Note:** If you are unable to connect, check the appliance port security settings and temporarily switch to *Low* if the setting is *Medium* or *High*. For details, see "To view or edit port security settings" on page 113.

Exact field names, buttons, and syntax vary by VNC viewer. Typically, one or two colons are required between the appliance IP address and port number. An example using VNC port 5904 is given here:

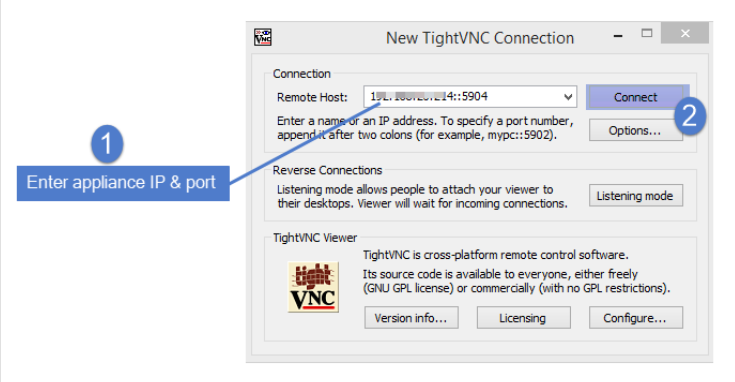

3 The Windows login screen displays, indicating the VM is available.

> **Note:** If you access the VM before it has booted, you may see the first screen of the Unitrends Windows Unified Bare Metal Recovery Wizard. Do not attempt to complete the steps on the bare metal screen. After several seconds, the login screen for the original Windows asset displays.

4 Log in to the VM by entering the credentials of the original Windows asset.

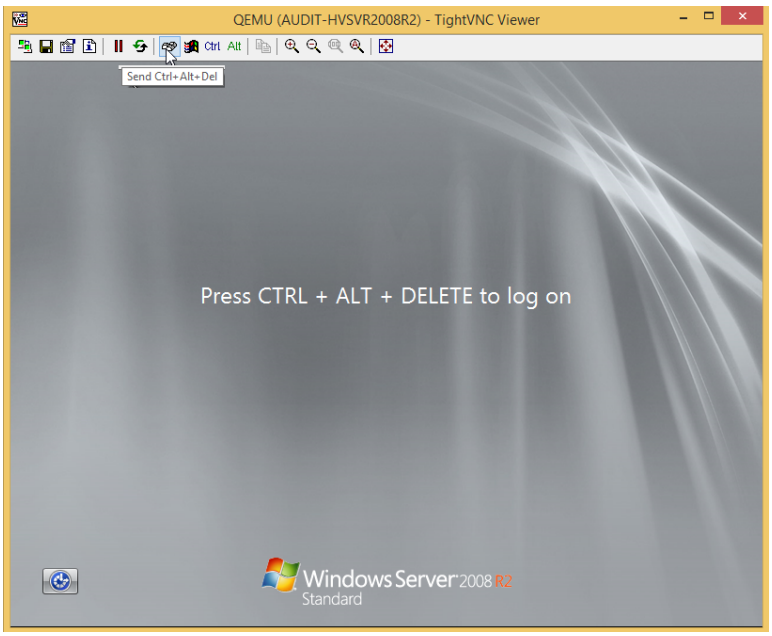

To access an IR VM on an external hypervisor

Use this procedure to access an IR VM that is running on a VMware or Hyper-V host in audit or live mode:

1. Connect to your hypervisor manager.
2. Locate the VM in the list of virtual machines and launch the VM console.
3. Log in to the VM by entering the credentials of the original Windows asset.

Notes:

- If you access the VM before it has booted, you may see the first screen of the Unitrends Windows Unified Bare Metal Recovery Wizard. Do not attempt to complete the steps on the bare metal screen. After several seconds, the login screen for the original Windows asset displays.
- Performing operations through the hypervisor on an IR VM that is running in audit mode on a Hyper-V or ESXi server is not recommended and may yield undesirable results. For example, creating a Hyper-V checkpoint of an IR VM in audit mode causes an error on the hypervisor.

To power off the IR VM

1. On the appliance used for the recovery, select **Recover** and click the **Instant Recovery** tab.
2. Click to select the instant recovery session.
3. Click **Power Off**.

Note: If the Power Off button is disabled, either the VM is already powered off or you do not have the permissions needed to run this procedure.

4. Click **Confirm**.

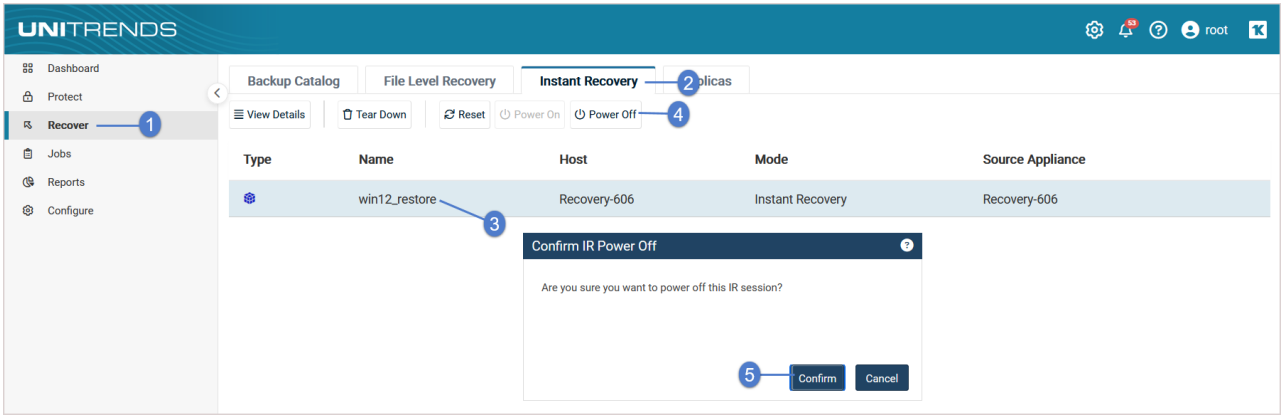

To power on the IR VM

1. On the appliance used for the recovery, select **Recover** and click the **Instant Recovery** tab.
2. Click to select the instant recovery session.
3. Click **Power On**.

Note: If the Power On button is disabled, either the VM is already powered on or you do not have the permissions needed to run this procedure.

4 Click **Confirm**.

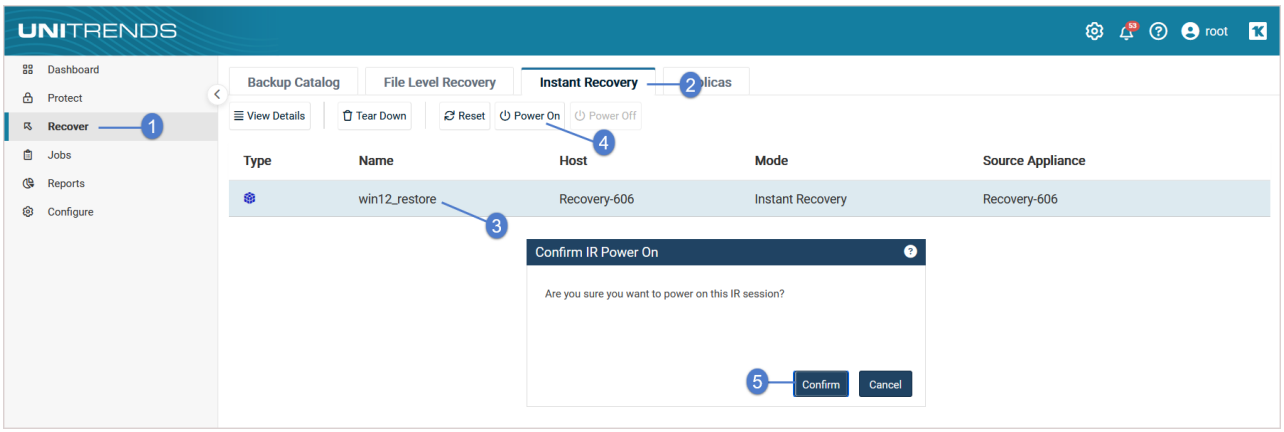

To reset the IR VM

IMPORTANT! This procedure performs a hard reset of the IR VM. If possible, use the power off/power on procedures instead.

1 On the appliance used for the recovery, select **Recover** and click the **Instant Recovery** tab.

2 Click to select the instant recovery session.

3 Click **Reset**.

Note: If the Reset button is disabled, either the VM is not powered on or you do not have the permissions needed to run this procedure.

4 Click **Confirm**.

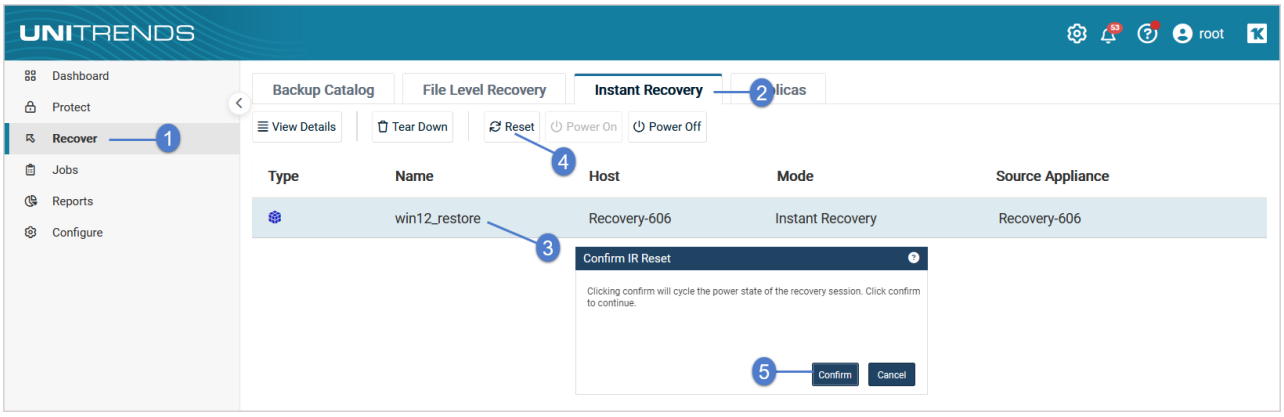

Tearing down the instant recovery session

Tearing down the instant recovery session deletes the recovery object, freeing the appliance resources and reserved space. For appliance performance, it is best to tear down the instant recovery session as soon as possible.

Tearing down the IR session deletes the IR VM in these cases:

- The IR VM resides on a Recovery Series or Recovery MAX appliance – The instant recovery session is needed for the IR VM to run. Tearing down the IR session removes both the recovery object and the IR VM from the appliance. Do not tear down the session until you are no longer using the IR VM.
- The IR VM is running in audit mode on a virtual host – The instant recovery session is needed for the IR VM to run. Tearing down the IR session removes the recovery object from the appliance and the IR VM from the Hyper-V or ESXi host. Do not tear down the session until you are no longer using the IR VM.
- The IR VM is running in instant recovery (live) mode on a virtual host– Tearing down the IR session removes the recovery object from the appliance. The IR VM remains on the virtual host server. The Unitrends appliance uses Storage vMotion (VMware) or Storage Live Migration (Hyper-V) to copy the data from the disk image to the target Hyper-V or ESXi host. Do not tear down the IR session until all data has been migrated and the VM is ready for use. Tearing down too early invalidates the IR VM (and you must perform IR again to create a new one).

  To verify that data has been migrated, go to the Instant Recovery tab, select the session and click **Details**. Data migration is complete when you see the message *VM is available for use*:

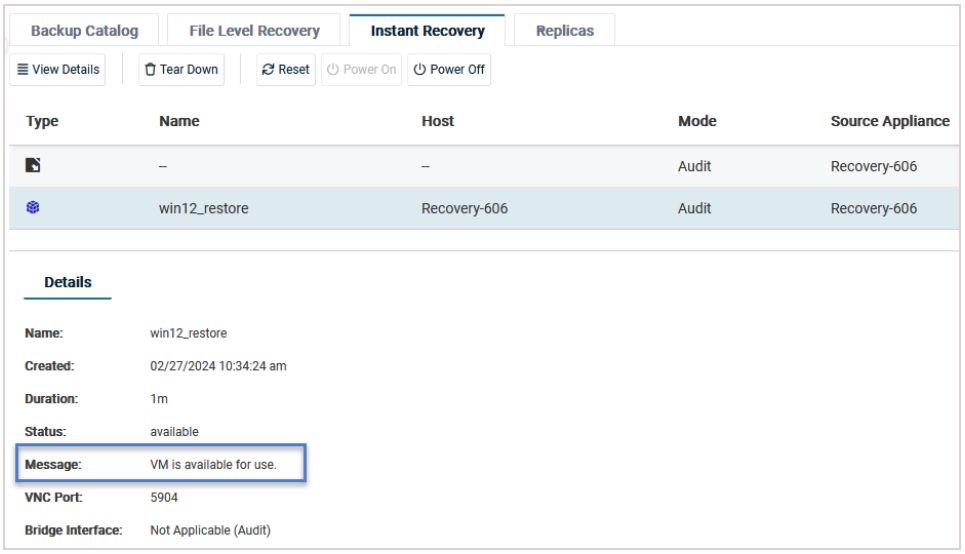

Use these steps to tear down the instant recovery session:

Note: Be sure to review the information above before running this procedure.

1. On the appliance used for the recovery, select **Recover** and click the **Instant Recovery** tab.

   Active IR sessions display by asset name.
2. Click to select the instant recovery session.
3. Click **Tear Down**.

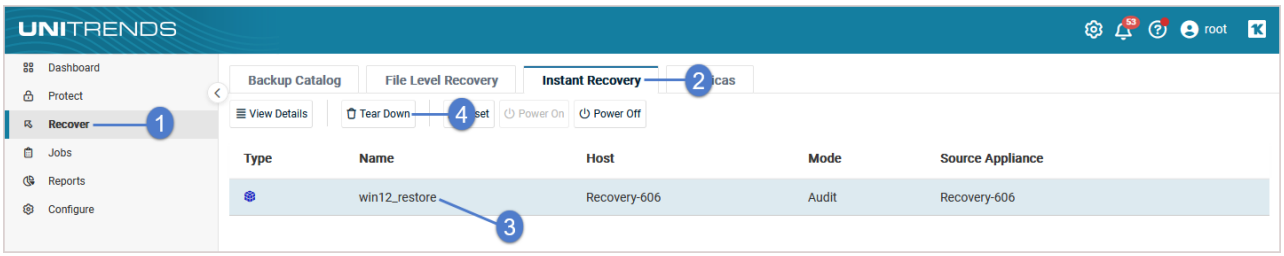

4 Click **Confirm**.

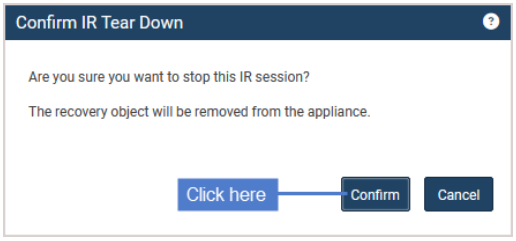

The session is removed from the appliance and no longer displays on the Instant Recovery tab:

- If the IR VM was running on a Recovery Series or Recovery MAX appliance, the IR VM is removed from the appliance.
- If the IR VM was running in audit mode on a VMware or Hyper-V host, it is removed from the virtual host.
- If the IR VM was running in instant recovery (live) mode on a VMware or Hyper-V host, it is retained on the virtual host.

Windows image-level replicas

The Windows image-level replica feature provides a quick way to recover a failed physical Windows asset. It creates a virtual machine replica of the Windows machine, then keeps this replica up-to-date by applying backups of the original asset as they run. In the event of a disaster, you can bring this replica online to immediately assume the role of the failed asset.

To use the feature, simply set up the replica by using the Create Windows Replica dialog. The appliance then creates the replica VM from the most recent backup of the Windows asset, and automatically applies all subsequent backups. Because the replica is continually updated, it is ready for production use at any time.

While creating the replica, you specify the location where the replica VM will reside. The replica can reside on:

- A Recovery Series physical appliance
- A Recovery MAX physical appliance
- An ESXi host
- A Hyper-V server

The replica VM is created as a cold stand-by in the specified location. The replica is powered off and has no network connectivity. Because the replica remains powered off even as backups are applied, it consumes no compute resources.

After the first backup has been applied, replica creation is complete. You can then do the following as needed:

- Audit the replica to verify the integrity of the machine and its data and applications. In audit mode, the replica runs on a private network (inaccessible from the production network). This enables you to check the replica machine while the original Windows server is still operating in production. It is recommended that you periodically audit the replica to ensure it functions as expected.
- Bring the replica online in your production environment to immediately assume the role of the original server. The live replica is intended as a temporary replacement until you can perform a bare metal recovery to restore the failed Windows asset to new hardware. Or you can opt to use the replica VM as a permanent replacement.
- Add the replica to copy data management jobs to run scheduled tests, and perform failover to your recovery network or instant lab. For details, see "Recovery Assurance" on page 1263.

See the following topics for details on using the Windows replicas feature:

- "Image-level replica requirements"
- "Setting up an image-level replica" on page 1098
- "Working with image-level replicas" on page 1102

Image-level replica requirements

The following topics cover the requirements for image-level replicas:

- "Backup requirements"
- "Replica requirements" on page 1088
- "Requirements for protected Windows asset" on page 1095

Notes:

- Only one Windows replica can exist per Windows asset. You cannot run both an image-level replica and a file-level replica of the same asset at the same time. If a replica exists, you must tear it down before creating another for the asset.
- You can opt to run image-level replicas in the Unitrends Cloud. Contact your Account Manager for assistance.

Backup requirements

An image-level backup of the physical Windows machine is required to create the replica. (To use a Windows file-level backup, see "Windows file-level replicas" on page 993.)

- The backup must contain all boot and critical system volumes.
- The backup must be a local backup run by the appliance where you are creating the Windows replica or a hot backup copy run on the backup copy target appliance. You cannot create a replica from an imported backup.

Notes:

- By default, image-level backups include all system information needed for replicas. If you opt to exclude volumes from backup, use care not to exclude the boot and critical system (OS) volumes.
- The replica virtual machine is created based on the backup you select. Volumes that were excluded from backup are not recovered.
- If the image-level backup includes SQL or Exchange, the replica automatically includes these applications. For details on how these applications are handled by image-level backups, see "Hosted applications" on page 712.
- For SQL, the master, model, and msdb system databases must also be present in the image-level backup of the Windows asset. (These are included by default. If you want the replica to include a hosted SQL application, use care not to exclude these system databases from the image-level backup.)

Back up the Windows asset regularly to keep the replica up to date. See these topics for details:

- To create a job manually, see "To create an image-level backup job" on page 449.
- To create a job by using an SLA policy, see "To create an SLA policy for Windows image-level assets" on page 545.
- For a comparison of the manual and SLA policy job creation methods, see "About creating backup and backup copy jobs" on page 426.

Replica requirements

Requirements vary by where the replica resides. See these topics for details:

- "Choosing the replica location"
- "Requirements for running an image-level replica on a Recovery Series or Recovery MAX appliance" on page 1089
- "Requirements for running an image-level replica under VMware ESXi" on page 1092
- "Requirements for running an image-level replica under Hyper-V" on page 1093
- "Additional requirements for running an image-level replica in a Hyper-V cluster environment" on page 1094

Choosing the replica location

The replica can reside on any of the following: a Recovery Series physical appliance, a Recovery MAX physical appliance, an ESXi server, or a Hyper-V server. Considerations for each are given in the following table.

Replica location	Considerations
Recovery Series or Recovery MAX physical appliance	Running the replica on the backup appliance itself provides a simple, seamless solution. Note that only Recovery Series and Recovery MAX physical appliances have this option. (A replica cannot reside on a Unitrends Backup virtual appliance.) Consider the following to determine whether this is the optimal replica location for your environment: • Requires no additional hardware. Provides near-zero RTO without having to increase

Replica location	Considerations
	CapEx. • Assumes the role of a failed Windows server temporarily, until you can get new hardware and run a bare-metal recovery. • Provides seamless, continuous protection in a failover scenario. The 'live' replica is automatically protected by the existing backup schedule after assuming the identity of the original Windows machine. • Uses the compute resources of the appliance to bring the replica 'live' in a failover scenario. If the appliance is already under high load, there may not be sufficient resources to provide for replicas. • Reduces on-appliance backup retention because a portion of the appliance's storage is reserved for the replica.
ESXi or Hyper-V server	Running the replica on an ESXi or Hyper-V server enables access to the compute and storage in your virtual environment, greatly increasing the pool of resources that can be used for replicas. Consider the following to determine whether this is the optimal replica location for your environment: • Leverages virtual infrastructure for replicas. Disk space and compute resources of the Unitrends appliance are not impacted. • Provides the ability to dynamically scale compute resources in the virtual infrastructure during failover, enabling 'live' replica performance to match that of the original Windows asset. • Provides the option to use the replica VM as a permanent replacement for the failed Windows asset. Can be used to migrate a physical Windows machine to your virtual infrastructure.

Requirements for running an image-level replica on a Recovery Series or Recovery MAX appliance

Ensure that the following requirements have been met before you create the image-level replica.

Requirement	Description
Appliance and agent version	The Unitrends appliance where the replica will reside must be running version 10.5.3 or later. The Windows asset must be running agent version 10.5.3 or later.
Allocate storage	Appliance storage can be used to store backups, for VM instant recovery, or for

Requirement	Description
	Windows replicas. To use the Windows replica and VM instant recovery features, you must allocate a portion of the appliance storage to Instant Recovery in the Edit Storage dialog (as described below). Before allocating storage, determine the percentage to use for backups and the percentage to reserve for Windows replicas and VM instant recovery. Once storage is allocated to instant recovery, it can be used only for Windows replicas and VM instant recovery. The storage is reserved and cannot be used for other purposes, such as backups or deduplication (but you can modify your storage allocation at any time). Because the appliance is designed to retain as many local backups as possible, it is best to reserve instant recovery space soon after initial deployment. If necessary, you can allocate instant recovery space later, but doing so may require the appliance to purge local backups to make room for the newly allocated instant recovery space. **Note:** You do not need to allocate instant recovery storage if the Windows replica will reside on an ESXi or Hyper-V server. This procedure is only required for Windows replicas that reside on the appliance itself. **To allocate storage for Windows replicas** 1 On the **Configure > Appliances** page, select the appliance, then click the **Storage** tab below. 2 Select the **Internal** storage and click **Edit**. 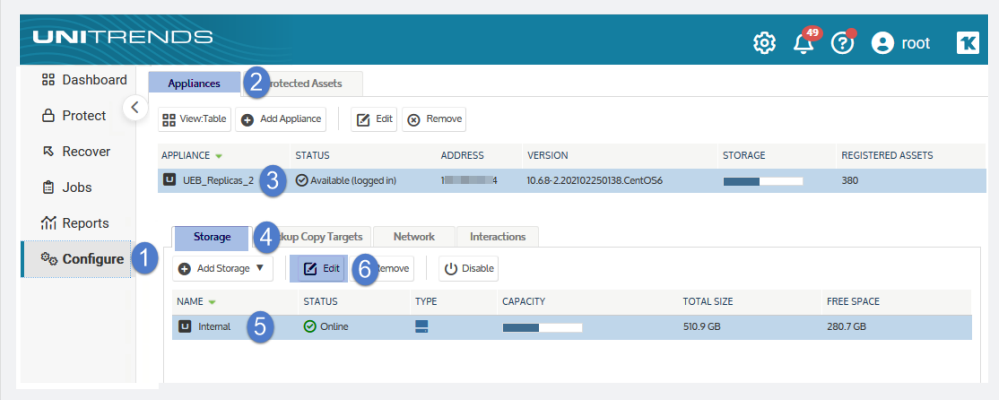 3 Modify the storage percentages allocated to Backups versus Instant Recovery (IR). The minimum IR space needed for a replica is the total amount of space in use on the original asset (the sum of used space on all disks). 4 Click **Save**.

Requirement	Description
	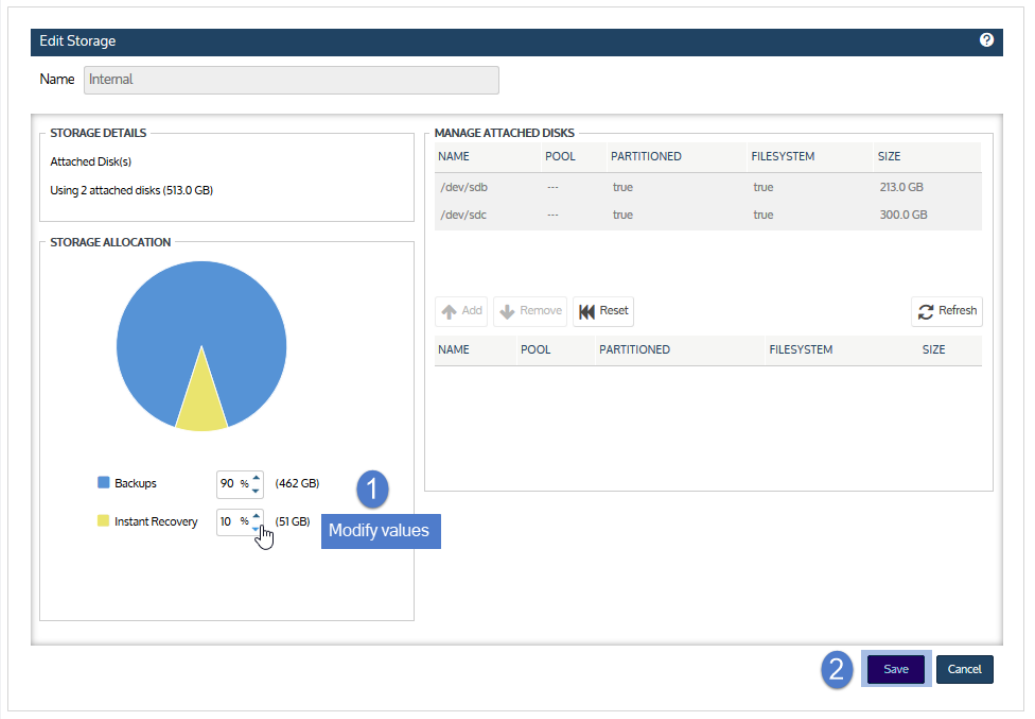
Unitrends appliance resources	The replica uses part of the Unitrends appliance's processors, memory, and storage. This usage may impact the performance of regular system functions (such as backups, backup copies, deduplication, and purging). Monitor the appliance closely and make adjustments as necessary. Minimum replica requirements: • One replica requires a minimum of 2048 MB of memory. • One replica requires a minimum of 1 core processor when in an idle state.
UEFI-based assets	Windows replicas of UEFI-based assets are supported on the following: • ESXi servers • Hyper-V servers • Recovery Series appliances running the CentOS 7 operating system • Recovery MAX appliances running the CentOS 7 operating system Windows replicas of UEFI-based assets are NOT supported on physical appliances running the CentOS 6 operating system. To check the appliance OS version, select **? > About** in the appliance UI:

Requirement	Description
	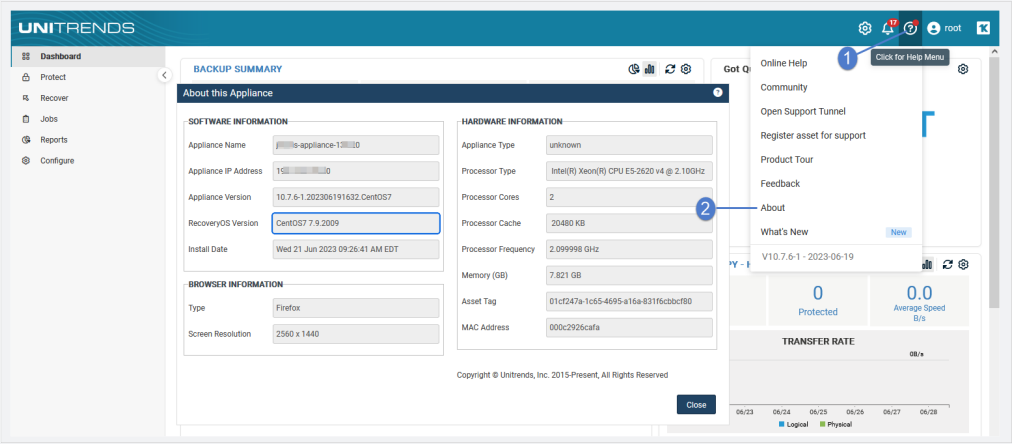
GPT-partitioned assets	Instant recovery is supported for GPT-partitioned volumes.

Requirements for running an image-level replica under VMware ESXi

Ensure that the following requirements have been met before you create the image-level replica.

Requirement	Description
Hypervisor version	The ESXi host must meet these requirements: • Must be running ESXi 5.1 or a higher version listed in the Compatibility and Interoperability Matrix. • Must support the operating system (OS) of the Windows asset. (See the VMware documentation for details.) For example, a replica of a Windows 2016 asset cannot reside on an ESXi 5.1 host.
Virtual host asset	The ESXi server must be added to the appliance as an asset. See "Adding a virtual host" on page 308.
Compute	One replica requires a minimum of 2048 MB of memory. (This number must be a multiple of 4.)
Replica VM changeability	Once you have configured the replica in the Create Windows Replica dialog, do not make any changes to the replica VM. Any alteration to the replica (unless it is in *live* mode) may lead the replica to an inconsistent state.

Requirement	Description
Maximum disk size	The maximum disk size is capped by what the hypervisor supports. The replica's disks will be the same size as those on the original asset. For Windows assets with disks larger than 2 TB, the ESXi server must be running ESXi 5.5 or a higher version listed in the Compatibility and Interoperability Matrix.
Virtual hardware version	The replica VM is configured with the highest hardware version that the hypervisor supports.

Requirements for running an image-level replica under Hyper-V

Ensure that the following requirements have been met before you create the image-level replica.

Requirement	Description
Hypervisor version	The hypervisor must be one of the following: • A Windows Server with the Hyper-V role enabled, running 2008 R2 or a higher version listed in the Compatibility and Interoperability Matrix. • A Hyper-V Server running 2008 R2 or a higher version listed in the Compatibility and Interoperability Matrix. The Hyper-V host must support the operating system (OS) of the Windows asset. (See this Microsoft article for details: Should I create a generation 1 or 2 virtual machine in Hyper-V?) For example, a replica of a Windows 2016 asset cannot reside on a Hyper-V 2008 R2 host.
Host agent version	The Hyper-V host must be running Unitrends agent version 10.5.1 or higher. It is best practice to run the latest Unitrends appliance and agent software versions. Older versions do not support all current Unitrends features.
Virtual host asset	The Hyper-V host must be added to the Unitrends backup appliance as a protected asset. See "Adding a virtual host" on page 308.
Virtual host Samba access	The Hyper-V server must be able to access the appliance's Samba share: • SMB 2.0 – The SMB 2.0 security option is enabled by default on Unitrends appliances that were originally imaged or deployed with version 10.4.8 or higher. SMB 2.0 must be enabled on the Hyper-V server. • SMB 1.0 – The SMB 1.0 security option is enabled by default on appliances that were originally imaged or deployed with a pre-10.4.8 version. SMB 1.0 must be enabled on the Hyper-V server.

Requirement	Description
	Note: Upgrading from a pre-10.4.8 version does not change the SMB 1.0 setting. (To configure your appliance to use SMB 2.0, see How Unitrends supports SMBv2.)
Compute	One replica requires a minimum of 2048 MB of memory. (This number must be a multiple of 4).
Replica VM changeability	Once you have configured the replica in the Create Windows Replica dialog, do not make any changes to the replica VM. Any alteration to the replica (unless it is in *live* mode) may lead the replica to an inconsistent state.
Maximum disk size	The maximum disk size is capped by what the hypervisor supports. The replica's disks will be the same size as those on the original asset. For Windows assets with disks larger than 2 TB, the Hyper-V server must be running version 2012 or a higher version listed in the Compatibility and Interoperability Matrix.
Replica VM configuration	The replica VM is created with this configuration: • The asset's firmware interface type determines the generation of the replica VM. BIOS-based assets are created as generation 1 VMs, and UEFI-based assets are created as generation 2 VMs. A replica for a UEFI-based asset cannot run on 2008 R2. • The VM's configuration version is the highest version that the hypervisor supports.
Pass-through disks	Pass-through disks are supported. After you bring the replica online to assume the role of the failed asset, you must refresh and reconnect any existing iSCSI targets.

Additional requirements for running an image-level replica in a Hyper-V cluster environment

Ensure that the following requirements have been met before you create the image-level replica.

Requirement	Description
Cluster asset	To run a replica on a Hyper-V server in a cluster configuration, ensure that these requirements have been met: • The Unitrends Windows agent is installed on each node in the cluster. • Every node in the cluster is running the same agent version.

Requirement	Description
	• Each cluster node and the cluster itself has been added to the appliance as an asset. (For details, see "Working with Hyper-V servers" on page 664.)
Storage	The Hyper-V cluster must be configured with Cluster Shared Volumes (CSVs). SMB storage is not supported.
PowerShell FailoverClusters modules	These modules must be installed on every node in the cluster so that the appliance can discover the CSVs.
Selecting the replica location	To create a clustered replica, you must select the cluster itself as the Location in the Create Windows Replica dialog. Do not select an owner node. If you select an individual node in the cluster, the replica will not be clustered.
Network switch selection	For a clustered replica, select the Network Switch common to all nodes in the cluster (in the Create Windows Replica dialog). If you do not select this switch, a 'live' replica that fails over to another node will lose network connectivity.
2008 R2 clusters	To run the replica on 2008 R2 servers in a cluster configuration, enable DCOM and WMI Virtualization access for all nodes in the cluster. For instructions, see Security settings for creating a clustered virtual failover client on Hyper-V server 2008 R2.
Live migration interoperability	During live migration of a clustered replica, the Unitrends appliance cannot apply backups to the replica, verify or audit the replica, or bring the replica online in production to assume the role of the original asset. If the appliance attempts to apply a backup or verify the replica during a live migration, the appliance waits several minutes and then attempts the operation again. If you try to audit the replica or bring it online in production, the appliance notifies you that it cannot run the operation because of the migration and you must try again later.

Requirements for protected Windows asset

The Windows asset must meet the following requirements to use the image-level replica feature:

Requirement	Description
Operating System	The Windows OS must meet these requirements: • If running the replica on a virtual host, the Hyper-V or ESXi host must support the guest OS of the replica VM. (See the Microsoft or VMware documentation for details.) For example, a replica running Windows 10 cannot reside on ESXi 5.1 or Hyper-V 2008 R2. • For Windows 8.1 and Windows 2012 R2, the replica includes the data from all

Requirement	Description
	disks, but if created as a Gen 1 VM, only the first four disks are eligible as boot devices. • For Windows 2008 R2 SP1, these additional requirements apply: – These Windows security updates must be installed: Update for Windows Server 2008 R2 x64 Edition (KB2533623) and Security Update for Windows 7 for x64-based Systems (KB3033929). (If these updates have not been installed, you are prompted to install them during agent installation.) – The Unitrends Volume CBT driver (used to run image-level incremental backups) cannot be installed along with the Unitrends Windows agent. You must install it manually. During agent installation the Volume CBT installer is placed here: *C:\PCBP\Installers\uvcbt.msi*. To install the driver, simply run *uvcbt.msi*. After installing the driver, you must enable it by rebooting the Windows asset.
Applications	Some Windows applications require network access and/or rely on underlying hardware, like network interface MAC addresses, in order to run properly. When booting a replica in audit mode, there is no network interface, so applications requiring network connectivity will not function properly. This is expected behavior. When booting a replica in live or audit mode, applications that rely on unchanging hardware (like MAC addresses) may not function properly or may require re-authentication, re-installation, or other special actions that are application specific in order for them to work properly. You should work with the application vendor to determine what actions are required.
Firmware interface type	The replica feature supports BIOS- and UEFI-based assets. For UEFI-based assets, the replica must reside on one of the following: • An ESXi server running ESXi 5.1 or a higher version listed in the Compatibility and Interoperability Matrix. • A Hyper-V server running 2012 R2 or a higher version listed in the Compatibility and Interoperability Matrix. For a replica running on Hyper-V, the UEFI-based asset's OS must be 64-bit and running Windows 8 or higher. • A Recovery Series or Recovery MAX appliance that is running the CentOS 7 operating system. **Note:** A UEFI-based replica cannot reside on a physical appliance that is running the CentOS 6 operating system.
Disk partition type	The replica feature is supported for assets with GUID Partition Table (GPT) and Master Boot Record (MBR) partitions. For assets with GPT partitions, the replica must reside on one of the following:

Requirement	Description
	• An ESXi server • A Hyper-V server • A Recovery Series or Recovery MAX physical appliance that is running the CentOS 7 operating system. To check the appliance OS version, simply click on **? > About**: 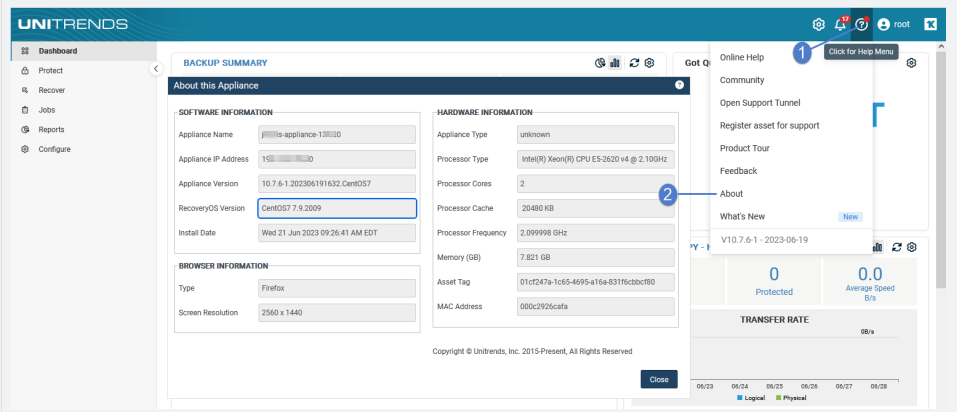
Software RAID volumes	The Windows image-level replica feature is not supported for software RAID configurations.
Deduplicated volumes	Volumes that use Microsoft deduplication are not supported in cases where the size of the data on the volume before it has been deduplicated is greater than the physical capacity of the volume. Because data is applied to the replica in its non-deduplicated form, the volume must have enough capacity to house this non-deduplicated data.
Number of volumes	The Windows asset can have a maximum of 20 volumes, including the System Reserved volume and other unmounted volumes. A replica with more than 20 volumes may fail to boot.
Separate boot and system partitions	For Windows assets with boot and system partitions located on different disks, the system partition must reside on the first disk (Disk 0).
File System Configuration	The replicas feature supports the following file systems: • NTFS • FAT/FAT32 • ReFS (Windows 2012 and later)

Requirement	Description
Active Directory	The replica feature supports Active Directory database (NTDS) located on the boot volume only. (If it is not on the boot volume, the configuration is not supported and you see an error message when you attempt to create the replica.)

Setting up an image-level replica

Use the following procedure to set up an image-level replica. You can configure the replica to run on:

- A Recovery Series or Recovery MAX physical appliance (backup appliance or backup copy target appliance).
- An external hypervisor (Hyper-V or VMware).

Note: To run the image-level replica in the Unitrends Cloud, contact your Account Manager for assistance.

To set up a Windows image-level replica

1 Select **Recover**, then click the **Replicas** tab.

2 Click **Create Replicas** and select **Windows Image Level**.

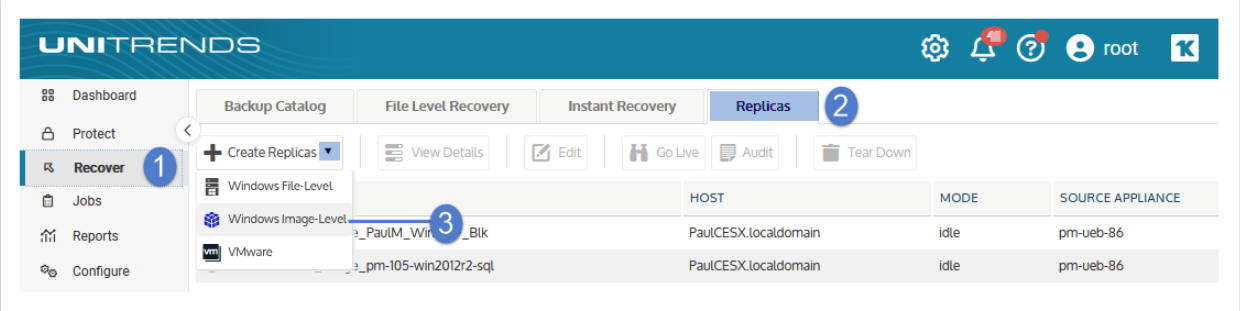

3 Select the Windows asset. Click **Next**.

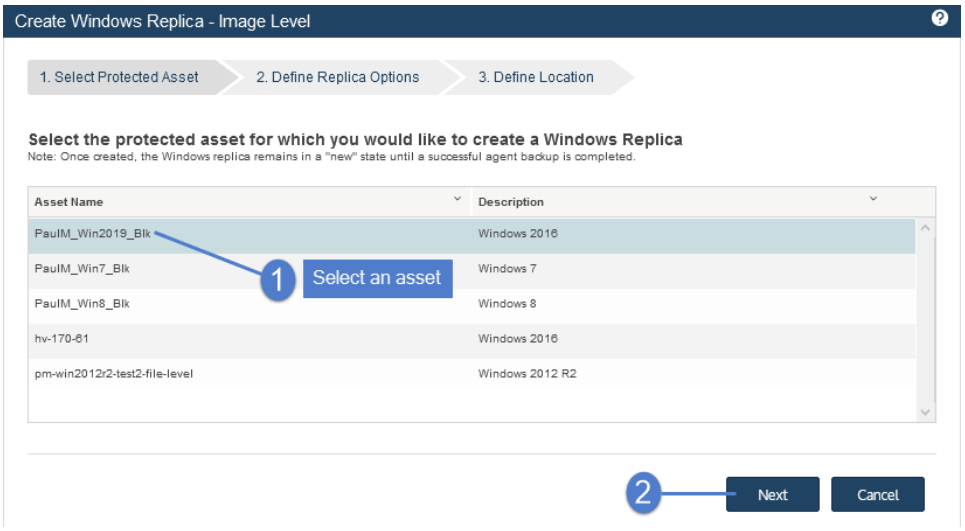

4 Enter replica options. Click **Next**. (See the table below for descriptions of the Define Replica Options fields.)

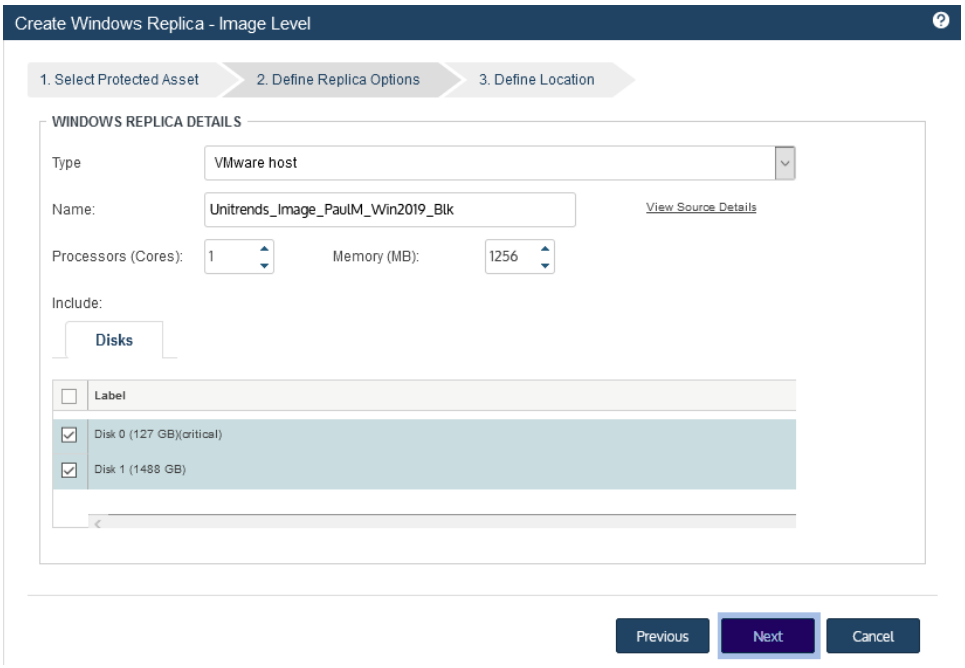

Item	Description
Type	Select a location type from the list (Unitrends Appliance, VMware Host, or Hyper-V Host). The list contains only the types that are available in your environment. For example, the Unitrends Appliance type is not an option for Unitrends Backup virtual appliances. The Hyper-V Host type is not an option if a compatible Hyper-V virtual host asset has not been added to the appliance.

Item	Description
Name	Replica name. By default, the replica is named Unitrends_Image_<WindowsAssetName>. If you are running the replica on a VMware or Hyper-V host, you can opt to edit this name.
View Source Details	Click this link to view details about the original Windows asset. Example: 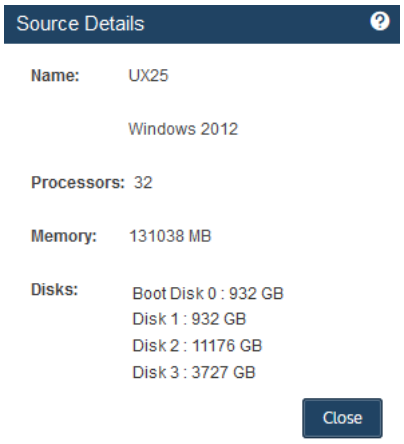
Processors (Cores)	Number of processors connected to the replica. Use care when modifying this value. The compute resources do not have to match the original Windows asset, but you should allocate enough cores for the replica to temporarily replace the original asset.
Memory (MB)	Amount of memory attached to the replica. Use care when modifying this value. The compute resources do not have to match the original Windows asset, but you should allocate enough memory for the replica to temporarily replace the original asset.
Email verification report (Hyper-V or Unitrends Appliance only)	This checkbox displays only for Hyper-V and Unitrends appliance hosts. Check this box to include automated audits of the replica. If selected, the appliance audits the replica and emails a report with a screen shot of the replica running in audit mode. (For details, see "Automated audits for a Windows image-level replica" on page 1104.)
Disks or Volumes tab	Disks or volumes to include when creating the replica. • Check boxes to select disks or volumes to include. • Disks/volumes marked as Critical are required. You can opt to exclude others.

5 Enter details to specify the location where the replica will reside. Click **Save**.

The details that display in the Define Location step vary by Type selected in the Define Replica Options step:

- If you selected **Unitrends Appliance** in the Type list, you do not provide any more location details.
- If you selected **VMware Host** or **Hyper-V Host** in the Type list, provide details for the virtual host server. See the table below for descriptions of the Define Location fields.

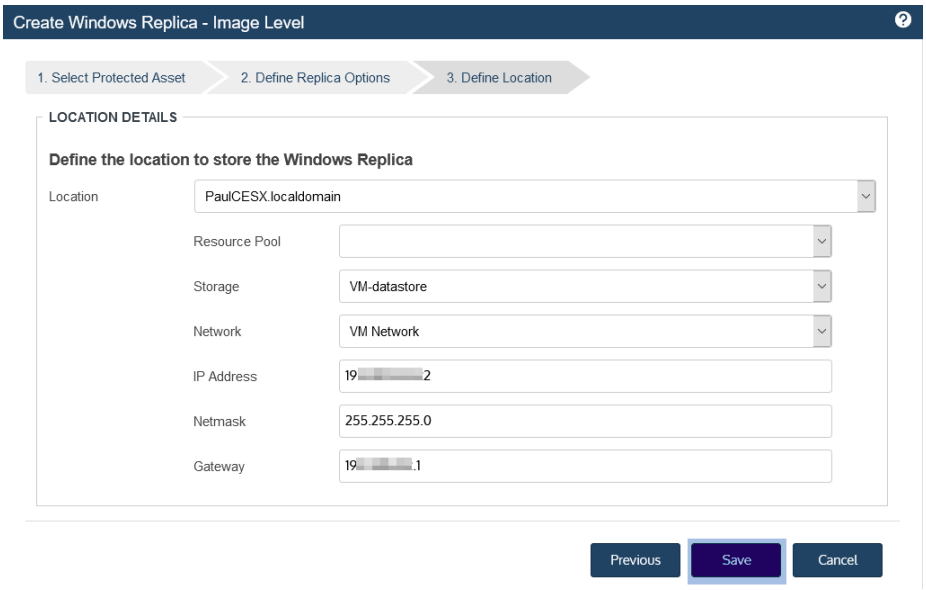

Item	Description
Location	Select a location type from the list (Unitrends Appliance, VMware Host, or Hyper-V Host). The list contains only the types that are available in your environment and are compatible with the Windows asset. For example, the Unitrends Appliance type is not an option for Unitrends Backup virtual appliances. The Hyper-V Host type is not an option if a compatible Hyper-V virtual host asset has not been added to the appliance.(For details on adding a virtual host, see "To add a virtual host asset" on page 311.)
Resource Pool (VMware only)	(Optional) If your VMware environment has resource pools, you can opt to select one in the list.
Storage	Select the datastore (VMware) or volume (Hyper-V) that will be used to create the replica VM's disks.
Network	Select a virtual network from the list. The list contains the virtual networks that are discovered and available on the VMware or Hyper-V host.

Item	Description
Network Switch (Hyper-V only)	For Hyper-V hosts only, select a network switch. The list contains the network switches that are discovered and available on the Hyper-V host.
IP Address, Netmask, and Gateway	Use these fields to define the network configuration that the appliance will use to create the replica and to apply backups. **WARNING!** Do NOT enter the IP address of the original Windows asset. Be sure to specify an IP address that is not used by another machine in your environment. When you audit the replica or bring it 'live' in production, the replica assumes the IP address, netmask, and gateway of the original Windows asset and not the configuration specified here.

The appliance creates a replica for the selected Windows asset, then applies the latest image-level backup. The replica is created in *new* mode.

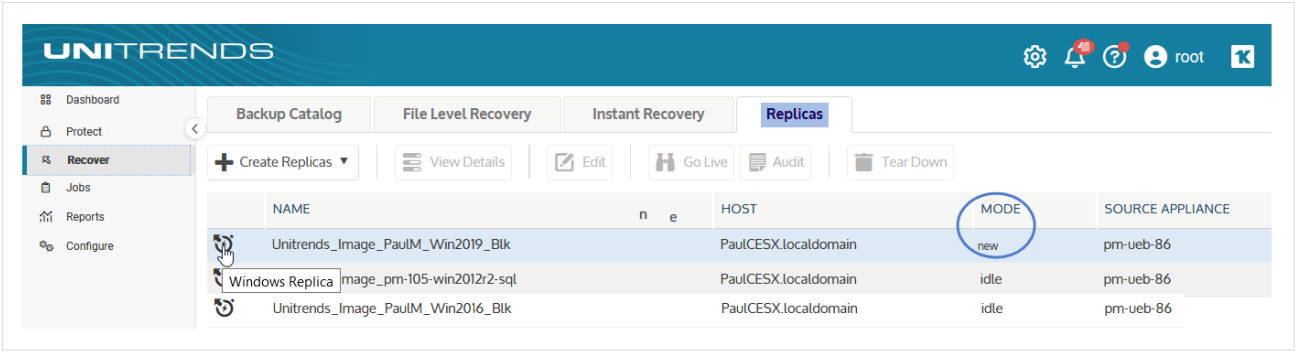

Its mode changes to *restore* while the backup is being applied, then to *idle*. After the replica enters idle mode, you can audit the replica or bring it 'live' as needed. We recommend that you audit the replica soon after it enters Idle mode, to verify its integrity. See "Working with image-level replicas" on page 1102 for details.

Working with image-level replicas

After setting up Windows image-level replicas, use the following procedures as needed:

- "Editing a Windows image-level replica"
- "Auditing a Windows image-level replica"
- "Bringing the replica live in production" on page 1109
- "Do not cancel an active replica restore job" on page 1113

- "Tearing down a Windows replica" on page 1114
- "Monitoring Windows image-level replicas" on page 1115

Editing a Windows image-level replica

After creating a replica, you can modify the following settings at any time:

- The number of processors connected to the replica.
- The amount of memory attached to the replica.
- Whether to apply new backups to the replica (*Enable virtual restores* checkbox).
- Whether to perform automated audits of the replica (*Email verification report* checkbox, Hyper-V and Unitrends Appliance only).

To edit a Windows image-level replica

1. Select **Recover**, then click the **Replicas** tab.
2. Select the replica, then click **Edit**.

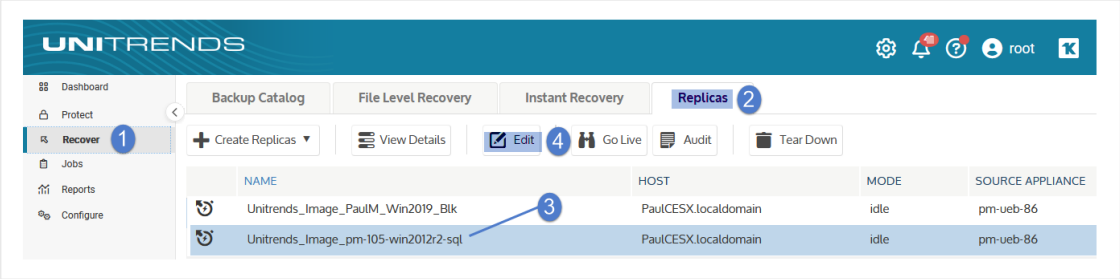

3. Modify settings as desired.
4. Click **Save**.

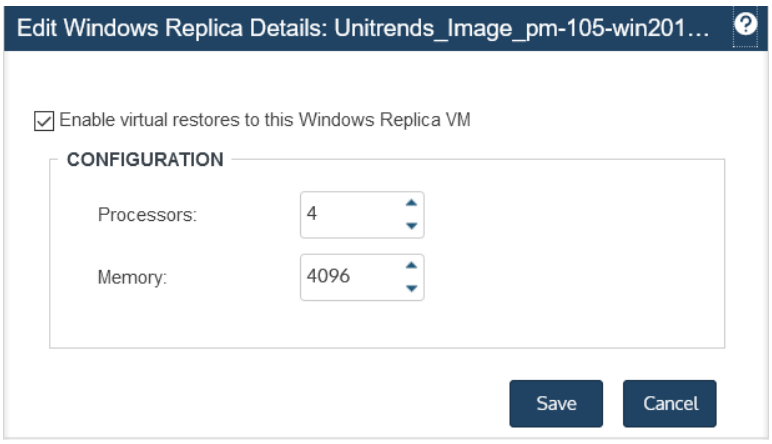

Auditing a Windows image-level replica

Audit mode enables you to run the replica on a private network while the original Windows asset is still operating in production. A replica running in audit mode boots with no network interface. Auditing the replica with the original asset still online does not result in network conflicts or impact the original asset in any way. However, applications on the replica that require network access do not function fully in audit mode.

Note: Some Windows applications require network access and/or rely on underlying hardware, like network interface MAC addresses in order to run properly. When booting a replica in audit mode, there is no network interface, so applications requiring network connectivity will not function properly. This is expected behavior. When booting a replica in live or audit mode, applications that rely on unchanging hardware (like MAC addresses) may not function properly or may require re-authentication, re-installation, or other special actions that are application specific in order for them to work properly. You should work with the application vendor to determine what actions are required.

It is recommended that you audit each newly created replica to ensure it functions as expected, and that you perform additional audits at regular intervals to check subsequent recovery points. You can perform manual audits for all image-level replicas. You can also set up automated audits for replicas that reside on a Unitrends appliance or Hyper-V host.

A newly created replica cannot be audited until at least one backup has been applied. During the audit, no subsequent backups are applied. Upon exiting audit mode, the appliance applies any backups that completed during the audit to bring the replica up to date.

Automated audits for a Windows image-level replica

You can automate the audit process by enabling email verification reports for a replica. The following requirements must be met to use this feature:

- The replica must reside on one of the following: a Recovery Series physical appliance, a Recovery MAX physical appliance, or a Hyper-V server. (Automated audits are not supported for replicas that reside on ESXi servers. Perform manual audits instead.)
- Email reporting must be enabled on the Unitrends appliance where you created the replica. Email must be configured with the System box checked and at least one valid recipient email address. For details, see "Email reporting" on page 117.

To enable verification reports, check the **Email verification report** box while creating or modifying a replica. (See these procedures for details: "Setting up an image-level replica" on page 1098 or "Editing a Windows image-level replica" on page 1103.)

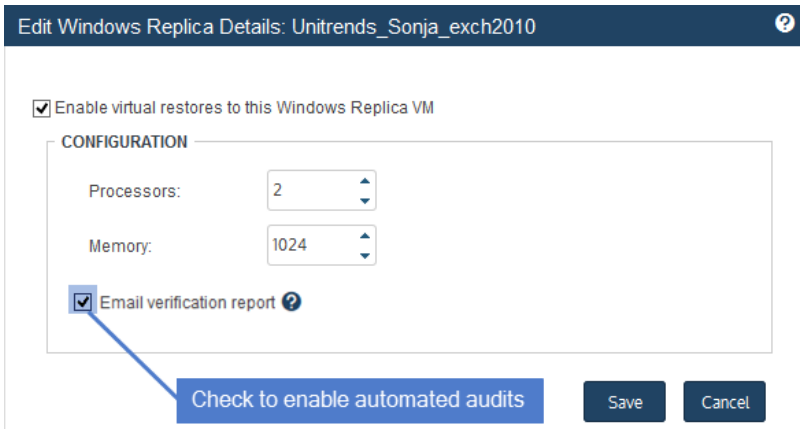

Once email verification is enabled for a replica, the appliance does the following:

- Brings the replica into audit mode after a backup has been applied.
- Takes a screenshot of the Windows login screen (after the replica has had several minutes to boot).
- Sends the screenshot to each System email recipient that is configured on the appliance.

The screenshot normally shows the Windows login screen, but it can also show Windows in other boot states, including error conditions.

IMPORTANT! Always view the screenshot to make sure the replica boots correctly.

The report runs once a day, but only after a backup has been applied. If the interval between backups lasts longer than 24 hours, you will not receive a report every day. If the replica cannot boot, you will receive an email report indicating that the replica cannot be verified.

Manually auditing a Windows replica

Manually auditing the replica is a two-part process where you bring the replica into audit mode and then access the replica to verify that it is functioning as expected. During the audit, you should verify the following:

- The replica boots successfully and is operational.
- The replica contains the expected data and applications. (Note that applications requiring network access do not function fully in audit mode.)

After you have finished auditing the replica, you must take it out of audit mode so the appliance can resume applying backups. (Note that any changes made during the audit are lost upon exiting audit mode.)

Audit mode procedures

Use these procedures to manually audit the replica:

- "To bring the replica into audit mode"
- "To access a replica on a Recovery Series or Recovery MAX appliance" on page 1106
- "To access the replica on an ESXi or Hyper-V server" on page 1108

- "To exit audit mode" on page 1108

To bring the replica into audit mode

1 Select **Recover**, then click the **Replicas** tab.

2 Select the replica, then click **Audit**. Click **Confirm**.

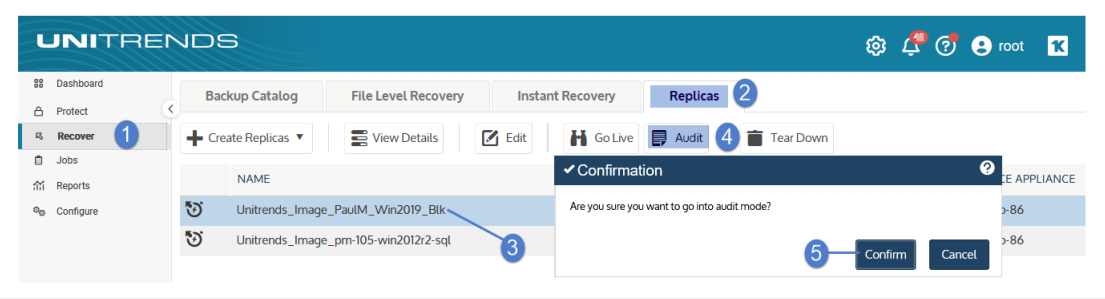

The replica's mode changes to *idle (pending audit)*, then to *audit*.

Note: If a backup is currently being applied, the replica does not enter audit mode until the restore is complete.

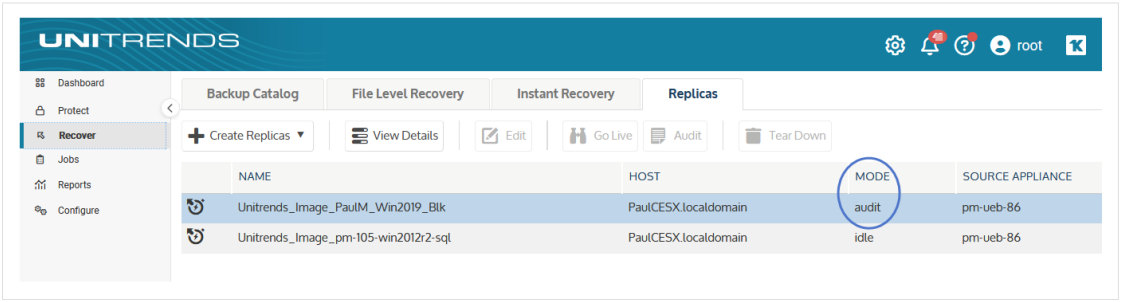

3 After the replica is in audit mode, you can connect to the replica to verify that it is functioning as expected. See the following for details:

- "To access a replica on a Recovery Series or Recovery MAX appliance"
- "To access the replica on an ESXi or Hyper-V server" on page 1108

To access a replica on a Recovery Series or Recovery MAX appliance

After a Windows replica has entered audit (or live) mode, use this procedure to access the replica:

Note: You must use a VNC viewer to access the replica in audit or live mode on a Recovery Series or Recovery MAX appliance. If necessary, download one to your workstation before running this procedure.

1 On the Replicas tab, view replica details to obtain the VNC port number:

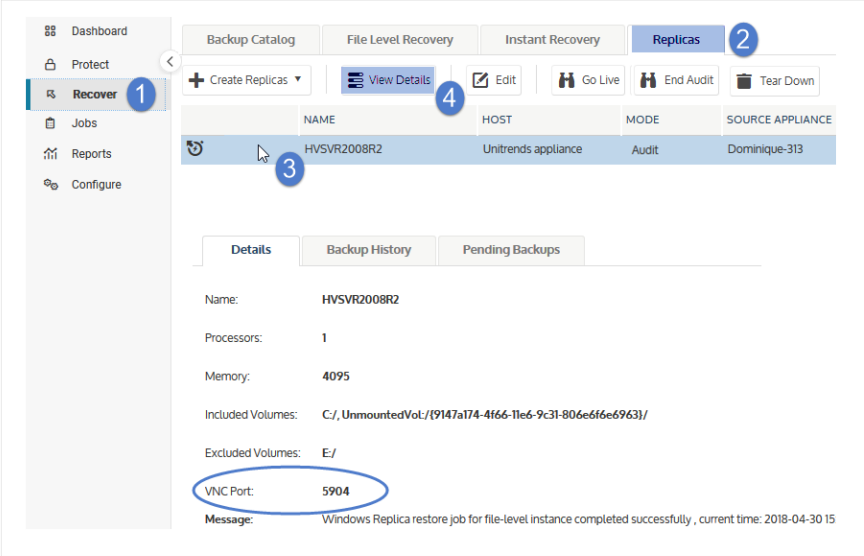

2 Open a VNC viewer. Connect to the replica by entering: <ApplianceIP>::<VNCport>

Exact field names, buttons, and syntax vary by VNC viewer. Typically, one or two colons are required between the appliance IP address and port number. An example using VNC port 5904 is given here:

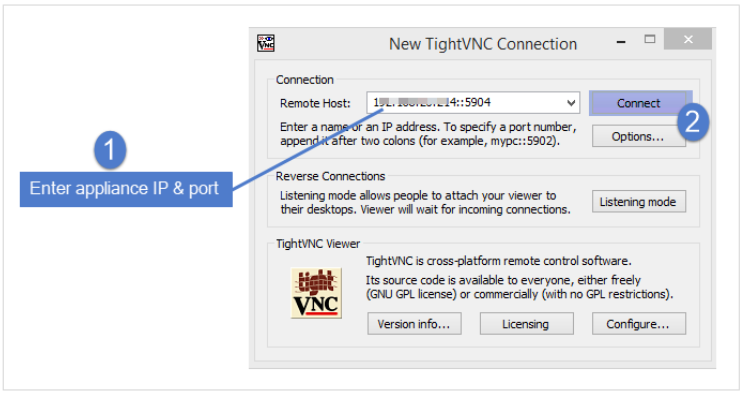

3 The Windows login screen displays, indicating the replica is available.

Note: If you access the replica before it has booted, you may see the first screen of the Unitrends Windows Integrated Bare Metal Recovery Wizard. Do not attempt to complete the steps on the bare metal screen. After several seconds, the login screen for the original Windows asset displays.

4 Log in to the replica by entering the credentials of the original Windows asset.

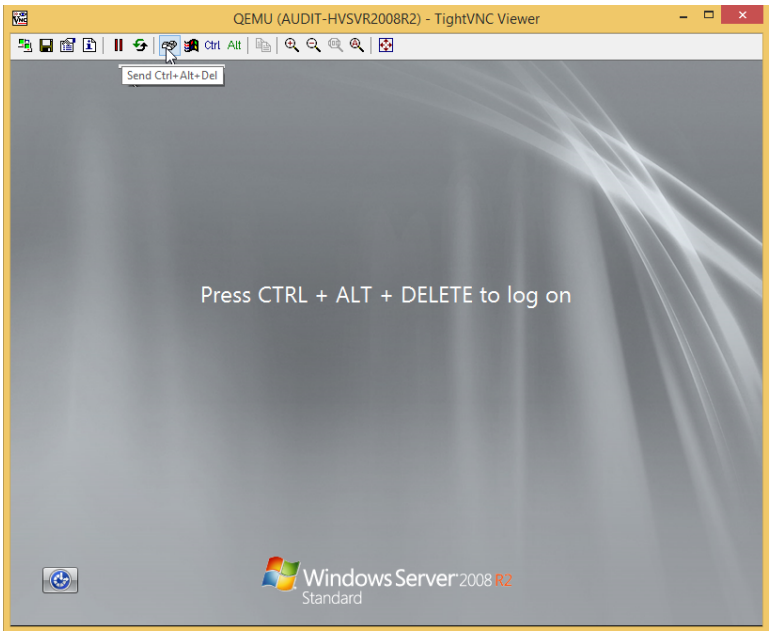

5 After verifying that the replica is running with its recovered data, proceed to "To exit audit mode" on page 1108.

To access the replica on an ESXi or Hyper-V server

After a Windows replica has entered audit (or live) mode, use this procedure to access the replica:

1 Connect to your hypervisor manager.

2 Locate the replica in the list of virtual machines, and access it the same way you access all VMs on the hypervisor.

3 Log in to the replica VM by entering the credentials of the original Windows asset.

> **Note:** If you access the replica before it has booted, you may see the first screen of the Unitrends Windows Integrated Bare Metal Recovery Wizard. Do NOT attempt to complete the steps on the bare metal screen. After several seconds, the login screen for the original Windows asset displays.

4 After verifying that the replica is running with its recovered data, proceed to "To exit audit mode".

To exit audit mode

1 Select **Recover**, then click the **Replicas** tab.

2 Select the replica, then click **End Audit**.

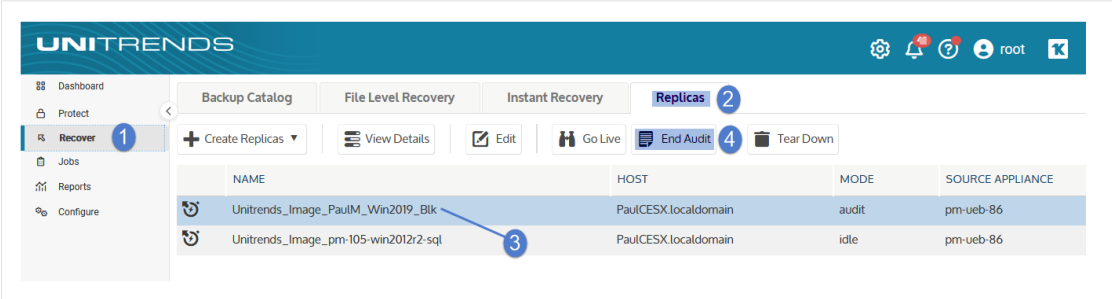

3 The replica exits *audit* mode. Its mode changes to *audit (pending: off)*, then to one of the following:

Restore – One or more backups successfully completed during the audit and the appliance is applying those backups.

Idle - The replica is idle (there are no backups to apply).

Bringing the replica live in production

If disaster strikes and the original asset fails, you can temporarily replace it with the replica by booting into live mode. Because the replica is continually updated with the original asset's data, it can immediately assume the role of the original asset.

The original asset's backup and backup copy schedules protect the replica in live mode, so that any changes made to the replica in live mode are captured under the identity of the original asset. This ensures continuity of recovery points in the asset's backup chain.

See these topics for details:

- "Live mode recommendations"
- "To bring the replica into live mode" on page 1110
- "Using the live replica as a temporary for the original Windows asset" on page 1112
- "Using the live replica as a permanent replacement for the original Windows asset" on page 1113

Live mode recommendations

Review these recommendations before going into live mode:

- Live mode should be used temporarily. The appliance begins sending alerts after a live replica has run for 14 days.
- You can exit live mode after you have recovered to new hardware (supported in all cases) or by retaining the replica as a permanent replacement (supported only for replicas that reside on external hypervisors).
- A live replica running on a Unitrends appliance uses appliance resources, so it is important that you recover to new hardware as soon as possible.
- A live replica running on an external hypervisor does not use any appliance resources. Instead, it uses hypervisor resources. The replica can replace the original asset temporarily or be used as a permanent replacement.

- Some Windows applications require network access and/or rely on underlying hardware, like network interface MAC addresses in order to run properly. When booting a Windows Replica in live or audit mode, applications that rely on unchanging hardware (like MAC addresses) may not function properly or may require re-authentication, re-installation, or other special actions that are application specific in order for them to work properly. You should work with the application vendor to determine what actions are required.

To bring the replica into live mode

This procedure provides instructions for booting a replica in live mode. Be sure to shut down the original asset before running this procedure.

1 Select **Recover**, then click the **Replicas** tab.

2 Select the replica, then click **Go Live**.

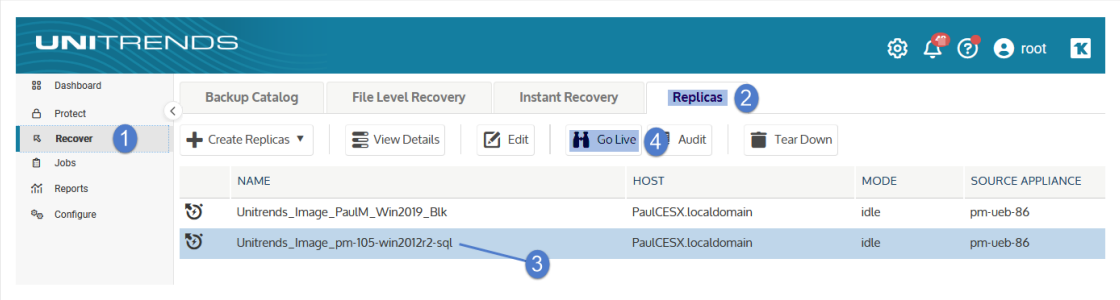

3 Click **Confirm**. The replica's mode changes to *live*.

- If a backup is currently being applied, the replica does not enter live mode until the restore is complete.
- Upon entering live mode, the replica assumes the identity of the original Windows asset. The replica is marked *invalid* because the replica role no longer applies.

4 Log in to the replica by using one of these methods:

Note: If you access the replica before it has booted, you may see the first screen of the Unitrends Windows Integrated Bare Metal Recovery Wizard. Do NOT attempt to complete the steps on the bare metal screen. After several seconds, the login screen for the original Windows asset displays.

Replica location	Description
Recovery Series or Recovery MAX appliance	Connect to the replica by using VNC, then log in using the credentials of the original Windows asset. In the VNC viewer, you will need to enter the appliance IP address and the VNC port (found on the Replicas tab under **View Details**). For detailed steps, see "To access a replica on a Recovery Series or Recovery MAX appliance" on page 1106. (The same steps are used to connect to a replica in audit mode and live mode.)

Replica location	Description
External hypervisor	Connect to the replica by using the hypervisor manager, then log in using the credentials of the original Windows asset.

5 If you see a message about reactivating Windows, you must activate the operating system by using your product key.

6 Check the disk configuration by using Windows Disk Management. (These steps might be slightly different depending on the Windows version.)

- Press the **Start** button.
- Right-click the **Computer** item.
- Choose **Manage**.
- Choose **Storage > Disk Management**. This application shows a graphical view of all disks and volumes.
- If the disk manager shows any disks in the Offline state, right-click the disk icon and click **Online**.
- If the disk manager shows any dynamic disks as Foreign, right-click the disk icon and click **Import**. All volumes should now display as they did on the original asset.

7 Set the system clock. The asset may be running with the system clock time used by the latest backup. This issue may cause the macine to boot with a past date or time.

8 From the Windows Control Panel, update the network properties for the adapter (the TCP/IPv4 address) by using one of the methods in the following table.

Replica location	Description
Recovery Series or Recovery MAX appliance	Do one of the following: • If the original asset has a static IP address and is on the same subnet as the replica, assign the live replica the same network settings as the original asset. This ensures that the replica functions as the original asset, and that the original asset's job schedules continue for the live replica. • If the original asset has a static IP address and the replica is NOT on the same subnet as the original asset, assign the replica a new network setting that uses the same subnet as the appliance. You must then modify the settings for the original asset in the Unitrends appliance by entering this new IP address. (For details, see "To edit an agent-based asset" on page 293.) This enables the appliance to treat the live replica as the original asset. • If you are using DHCP to assign IP addresses and you added the original asset to

Replica location	Description
	the backup appliance by using only the asset's name, the appliance detects the live replica after you connect it to your network. The appliance then treats the live replica as if it is the original asset. No additional network configurations are necessary to ensure that scheduled backup and backup copy jobs continue.
External hypervisor	Do one of the following: • If the original asset has a static IP address and is in the same subnet as the replica, assign the live replica the same network settings as the original asset. This ensures that the replica functions as the original asset and that the original asset's job schedules continue for the live replica. • If the original asset has a static IP address and the hypervisor does not have a network interface on the same subnet as the original asset, assign the replica a new network setting that uses the same subnet as the hypervisor. You must then modify the settings for the original asset in the Unitrends appliance by entering this new IP address. (For details, see "To edit an agent-based asset" on page 293.) This enables the appliance to treat the live replica as the original asset.

9 Log in to the Unitrends backup appliance and re-save the original Windows asset:

Note: If you recovered by using a backup copy on an appliance backup copy target, perform these steps from the backup appliance where the original asset resides, rather than from the backup copy target appliance.

- Select **Configure > Protected Assets**.
- Select the original Windows asset.
- Click **Edit > Save**.

SQL databases and other applications may require a few minutes to become available.

10 Prepare the replacement machine by doing the steps in "Using the live replica as a temporary for the original Windows asset" or "Using the live replica as a permanent replacement for the original Windows asset" on page 1113. You must do these steps before you tear down the Windows replica.

Using the live replica as a temporary for the original Windows asset

If the replica will replace the original asset only temporarily, do these steps after you have booted the replica in live mode:

1 Recover to new hardware as soon as possible, by using Unitrends bare metal recovery.

Data from the live replica is protected by the backup schedule of the original asset. Use the latest backup to perform the bare metal recovery. For details, see "Windows Bare Metal Protection and Recovery" on page 1207.

2 After recovering the replica's data to the new Windows asset, delete the replica from the appliance and from the hypervisor (if applicable). For instructions, see "To tear down a Windows replica" on page 1114.

Using the live replica as a permanent replacement for the original Windows asset

If the replica will permanently replace the original asset, do these steps after you have booted the replica in live mode:

1 Determine whether to continue protecting the replica with the backup schedules of the original asset or whether to run virtual machine backups for the VM.

2 Delete the replica from the appliance as described in "Tearing down a Windows replica" on page 1114 . Be sure to delete the replica from the appliance only, as you have the option to delete it from the hypervisor as well.

3 If you will be switching to VMware or Hyper-V backups, create the VM backup schedule and remove the asset from the original schedule. See "Backup Administration and Procedures" on page 425 for details.

Notes:

- The replica VM is not automatically added to any existing VM backup schedule.
- It can take several minutes for the replica VM to show up in the list of VMs to protect with VMware or Hyper-V backups. To refresh the list of discovered VMs, click the Gear icon in the upper-right of the UI and select **Inventory Sync**.

Do not cancel an active replica restore job

Do not cancel an active replica restore job. Instead, disable virtual restores.

Each successful backup of the original Windows asset is applied to the replica as soon as the backup completes. The appliance applies the backup by running a replica restore job, which displays on the Active Jobs tab as shown here:

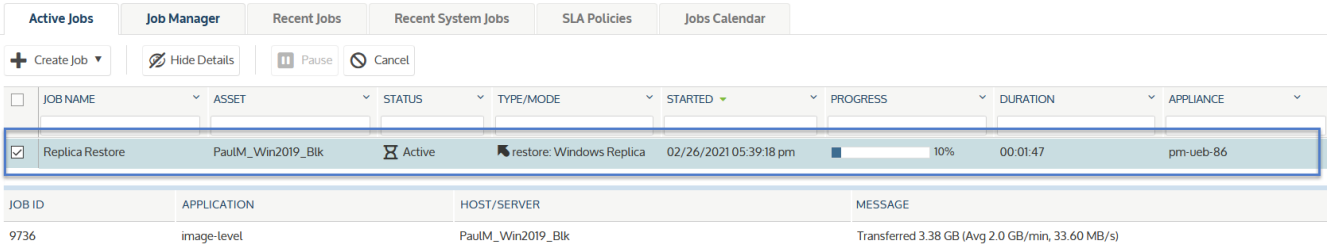

If you cancel a replica restore job by using the Cancel button on the Active Jobs page, the appliance automatically creates a new job to replace the one you canceled. Instead, to temporarily stop applying backups, uncheck the replica's *Enable virtual restores to this Windows Replica VM* box (as described in "Editing a Windows image-level replica" on page 1103). Check this box to start applying backups again. Note that all backups that ran while virtual restores were disabled are applied to the replica upon checking this box. You cannot skip applying a specific backup to a replica.

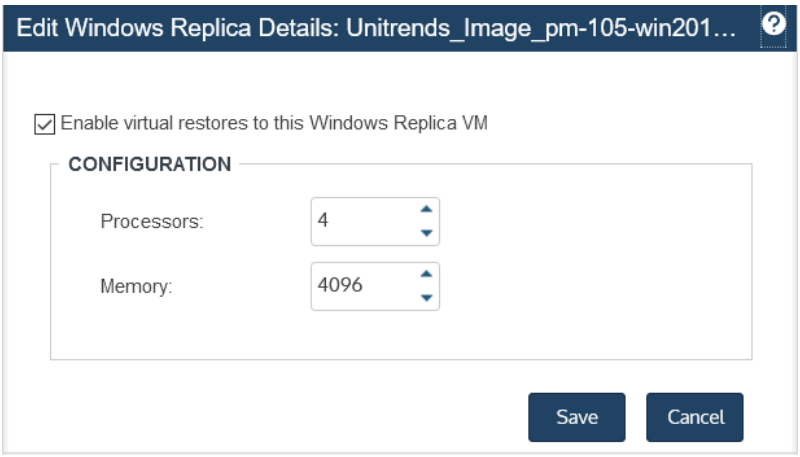

Tearing down a Windows replica

This section provides instructions for deleting a Windows replica.

For a replica running on a Unitrends appliance, you should delete the replica as soon as you have recovered the original asset to new physical hardware, to free up appliance resources.

For a replica running on a hypervisor, you have these options:

- Delete the replica from the appliance only – Select this option to use the replica VM as a permanent replacement for the failed Windows asset.
- Delete the replica from the appliance and delete the replica VM from the hypervisor itself – Select this option if you have recovered the original asset to new physical hardware and will not be using the replica VM as a permanent replacement.

To tear down a Windows replica

IMPORTANT! If you are running a live replica as a temporary replacement for an asset that has failed, do not tear down the replica until you have recovered the original asset to new physical hardware. (For details, see "Using the live replica as a temporary for the original Windows asset" on page 1112.)

1. Select **Recover**, then click the **Replicas** tab.
2. Select the replica, then click **Tear Down**.

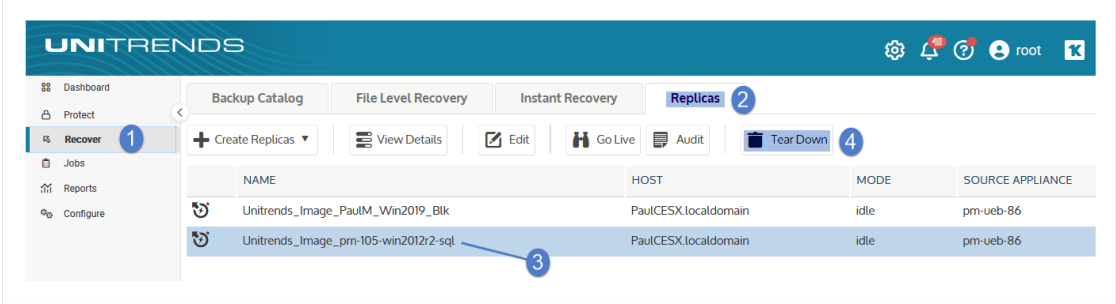

3 Do one of the following (options differ by replica location):

- Replica residing on a Recovery Series or Recovery MAX appliance – Click **Delete** to delete the replica.
- Replica residing on an external hypervisor – A box displays with options to delete the replica from the appliance only or from both the appliance and the hypervisor. Select the desired option and click **Tear Down**.

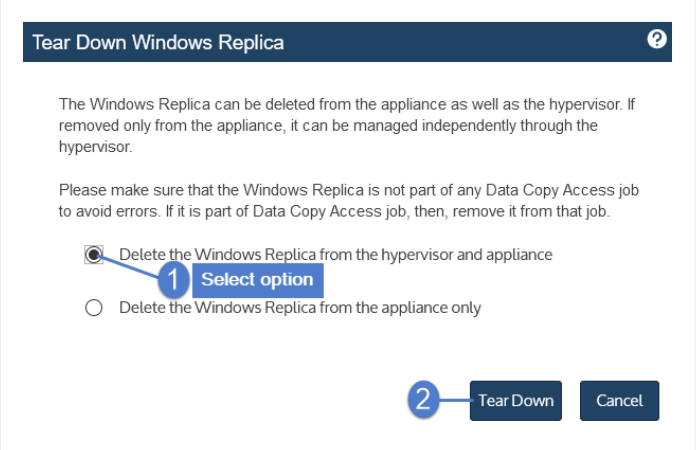

It can take several minutes for the appliance to purge all information about the replica. If you need to create a new replica for the original asset, you must wait for this information to purge. If it has not yet purged, the original asset does not display in the list of assets for which you can create a replica.

Monitoring Windows image-level replicas

Use these procedures to check the status and details of existing Windows replicas:

- "To view all Windows replicas"
- "To view image-level replica details" on page 1116
- "Windows replica modes" on page 1118

To view all Windows replicas

1 Select **Recover**, then click the **Replicas** tab.

2 All Windows image-level, file-level, and VM replicas display in a list on the Replicas tab.

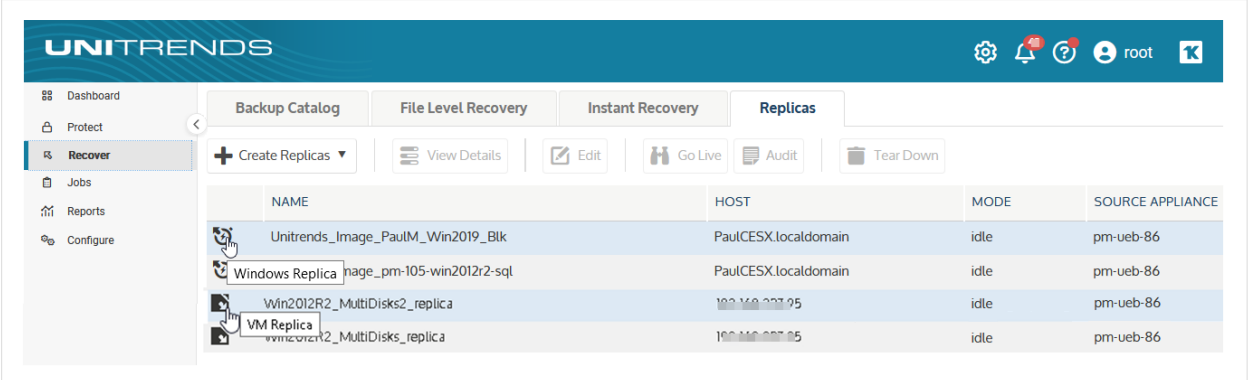

The following information is given for each replica:

Column	Description
Replica type icon	Indicates the replica type: Windows or VM.
Name	Replica name. By default, the replica is named Unitrends _<WindowsAssetName>.
Host	Host where the replica resides: • For ESXi servers, *localhost.localdomain* or server IP address. • For Hyper-V servers, hostname of the Hyper-V server asset. • For Unitrends appliances, *Unitrends appliance*.
Mode	Replica mode. Examples: new, audit, restore, or idle. See "Windows replica modes" on page 1118 for additional details.
Source Appliance	Appliance where the replica was created.
Alert icon	Indicates that an alert has been generated for the replica. Hover over the icon for details.

To view image-level replica details

1. Select **Recover**, then click the **Replicas** tab.
2. Select the replica, then click **View Details**. Tabs containing replica details display. (See the table below for a description of each tab.)

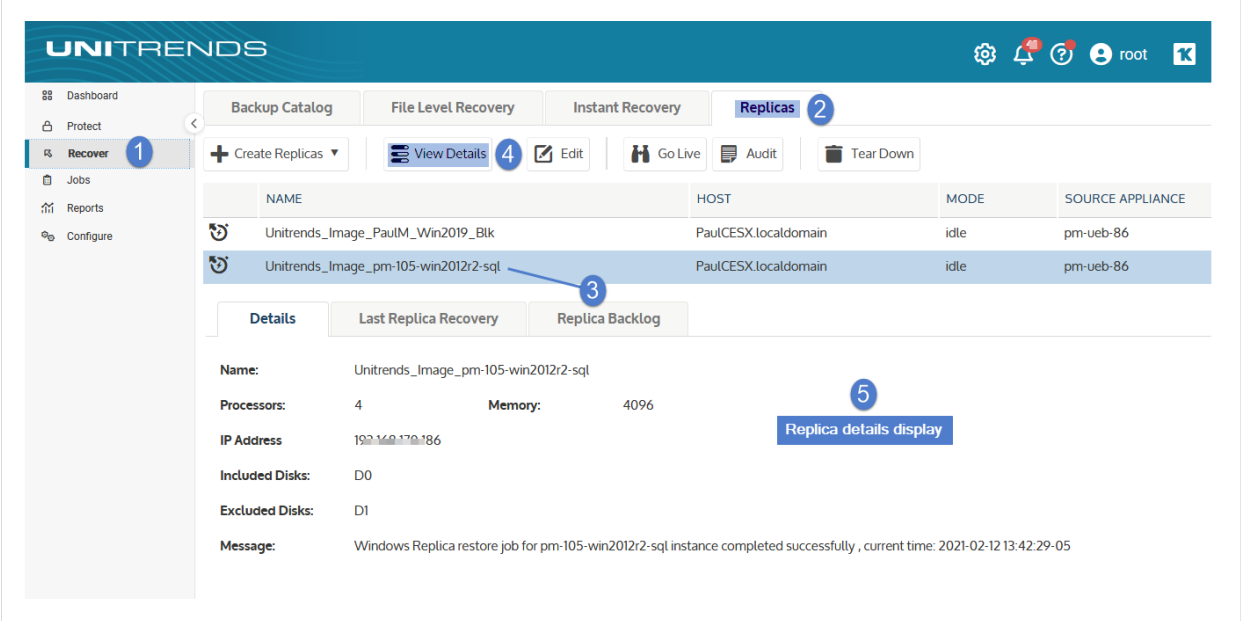

Details tab	Description
Details	Shows the replica name, number of processors, amount of memory, IP address (if applicable), disks/volumes included and excluded, and a message describing the last replica operation. Example: 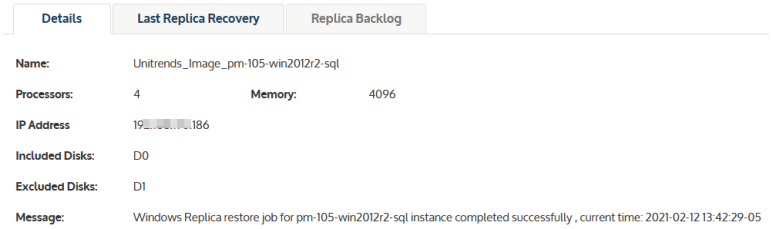
Last Replica Recovery	Shows details of the most recent backup that has been applied to the replica. Example: 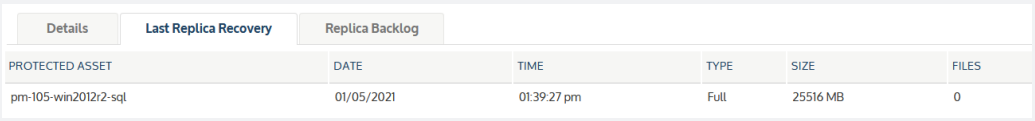
Replica Backlog	List of backups that have not yet been applied to the replica. No backups are listed if the replica is up to date. Example:

Details tab	Description
	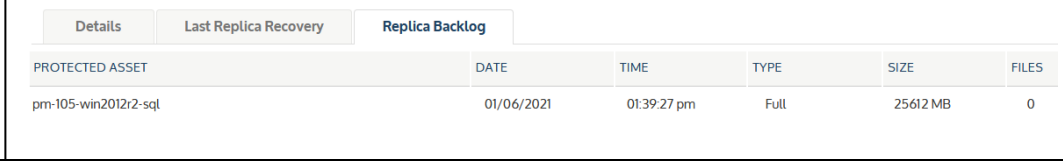

3 (Optional) Click **Hide Details** to stop displaying details for the selected replica.

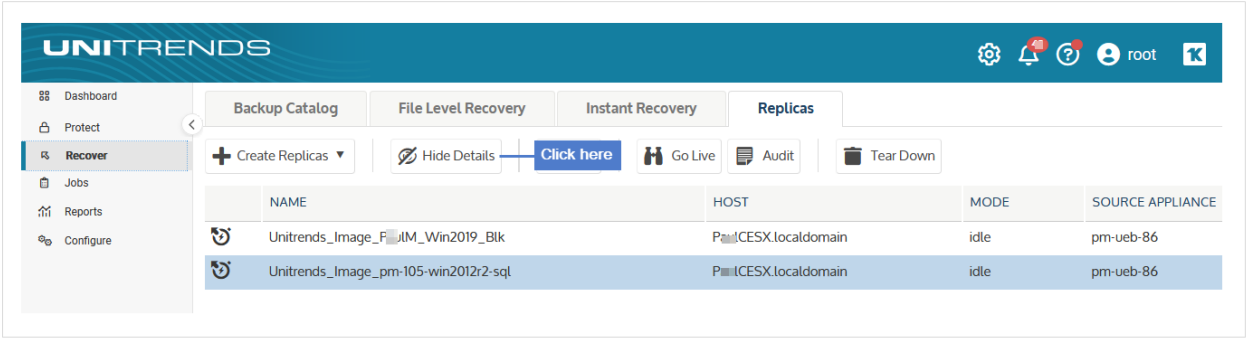

Windows replica modes

You can monitor a Windows replica by checking its mode on the Replicas tab. The mode indicates what is currently happening with the replica (for example, whether it is newly created, whether a backup is being applied, or whether it is in audit mode.)

Windows replica modes are described in the following table:

Mode	Description
New	The replica is new and no backup has been applied. A replica in *new* mode cannot be audited or booted into live mode.
Restore	A backup has completed, and the appliance is applying it to the replica. The replica remains in *restore* mode until the restore completes.
Idle	At least one backup has been applied to the replica, but currently no action is occurring.
Halted	A backup has completed, and the appliance has requested a restore. The replica goes into a *halted* state if the restore cannot be performed. The following can occur when a replica is in this mode: • If the restore could not be performed because the appliance could not reach the replica, it tries again after several minutes, and the mode changes from *halted* to *idle*. After

Mode	Description
	three failed attempts, the replica becomes invalid, and it remains in *halted* mode until a user deletes it. • If the restore could not be performed because a configuration change was made to the original asset, the replica becomes invalid, and it remains in *halted* mode until a user deletes it.
Audit	A user is performing a manual audit and the replica has booted in *audit* mode. For details, see "Auditing a Windows image-level replica" on page 1104.
Verify	A user has enabled verification reports (automated audits). The appliance has booted the replica in *audit* mode to take a screenshot of the replica's login screen. For details about verification reports, see "Automated audits for a Windows image-level replica" on page 1104.
Live	A user has booted the replica in *live* mode to replace the original asset. For details, see "Bringing the replica live in production" on page 1109. Once the replica is live, the only other mode it can enter is *off*.
Off	A user has taken the replica out of *live* mode. Once the replica is *off*, the only other mode it can enter is *live*.

This page is intentionally left blank.

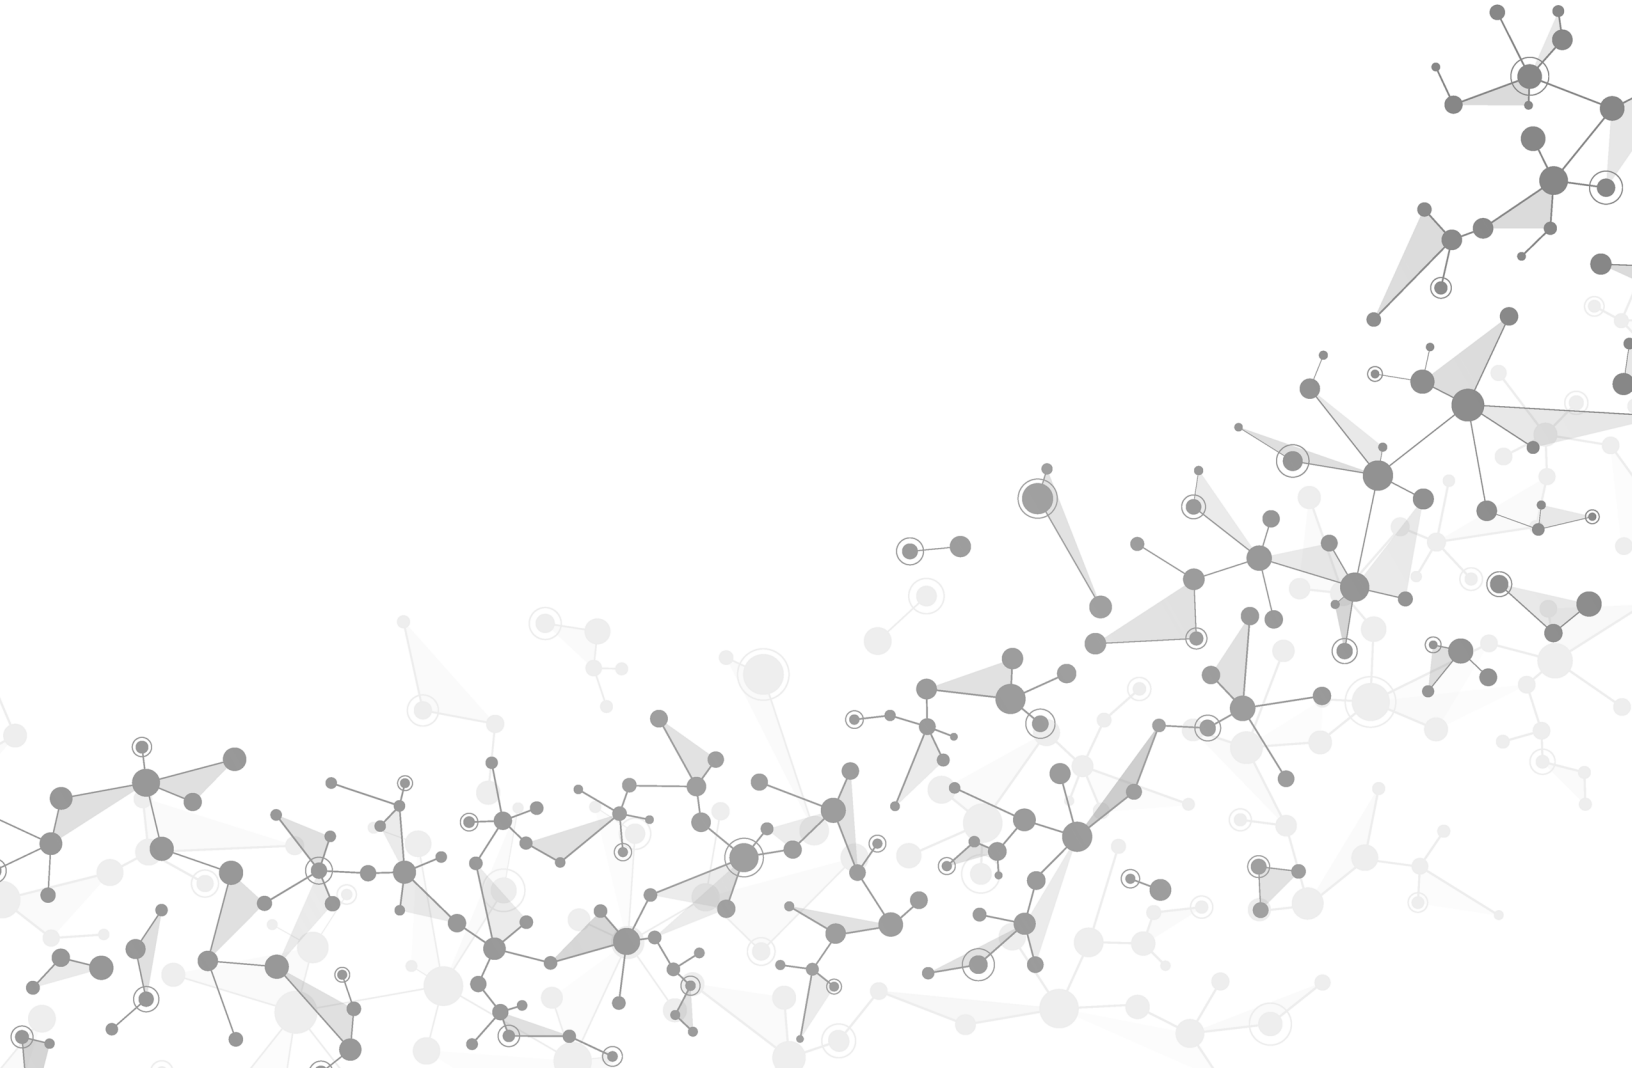

Chapter 17: Recovering NAS Backups

Recovery procedures vary depending on the NAS protocol used to create the backup. See these topics for details:

- "Recovering NAS CIFS or NFS backups" on page 1121
- "Recovering NAS NDMP backups" on page 1140

Recovering NAS CIFS or NFS backups

NAS CIFS and NFS backups have the following recovery options:

- Search Files – Search for specific files in all backups of this NAS that are currently stored on the appliance. Choose files to recover from the search results.
- Recover – Select a specific backup to recover all files in the backup group up to the point in time when the backup ran.
- Recover Files – Search a backup and choose specific directories and/or files to recover.

You can recover from a backup by using the Backup Catalog or the Backup Browser. To recover from an imported backup, from a backup copy, or to search for files across multiple backups, you must use the Backup Catalog. See these topics for details:

- "Specifying the target recovery location" on page 1121
- "To search multiple backups for files to recover" on page 1122
- "To recover an entire NAS CIFS or NFS backup by using the Backup Catalog" on page 1125
- "To recover an entire NAS CIFS or NFS backup by using the Backup Browser" on page 1128
- "To browse one NAS CIFS or NFS backup and recover files by using the Backup Catalog" on page 1132
- "To browse one NAS CIFS or NFS backup and recover files by using the Backup Browser" on page 1136

Specifying the target recovery location

For all recovery procedures, you must specify the target recovery location by selecting an asset and entering the full directory path where you want to recover the files. This location can be the original NAS share, an alternate location on the original share, or another asset that has been added to the Unitrends appliance.

For example, to recover to an alternate location on the original NAS share:

The original backup was of this directory:

/mnt/NAS/folder/subFolder1

and you wish to restore to:

/mnt/NAS/folder/subFolder2

enter the full path in the Directory field:

/mnt/NAS/folder/subFolder2

Note: If you enter */subFolder2* only, the files are recovered to the root mount point on the backup appliance (*/mnt* in our example). This could fill the root mount point and crash the appliance.

To search multiple backups for files to recover

1. Log in to the backup appliance.
2. Click **Recover > Backup Catalog > Search Files**.

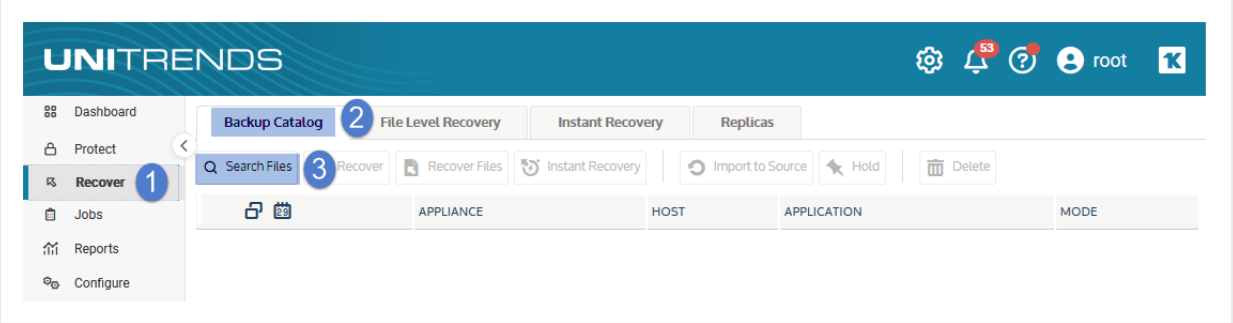

3. For Type, select **Backup**.
4. For Asset, select the NAS whose files you want to recover.
5. Enter additional criteria and click **Search**.

   All backups of this NAS stored on the appliance are searched for matching files.

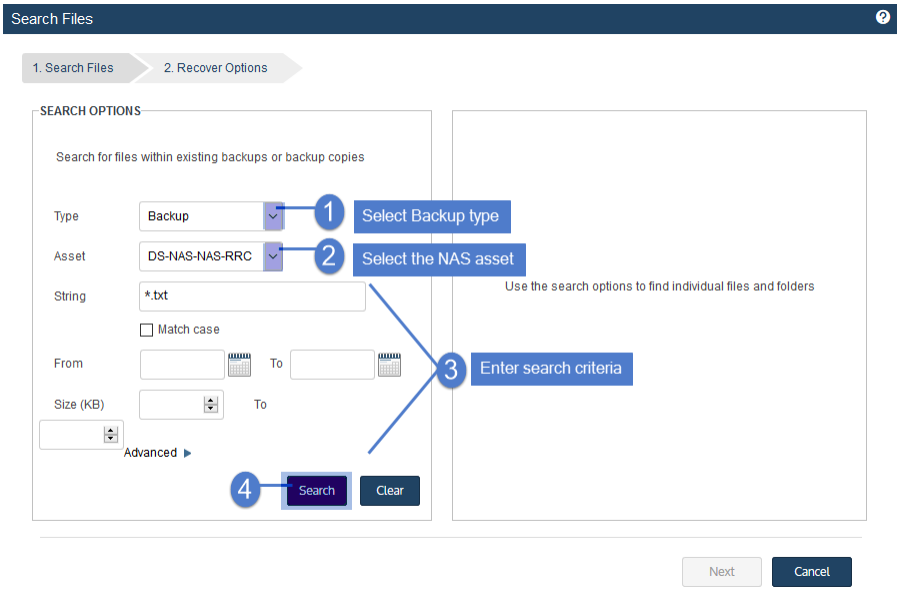

6. In the results list, click to select files to recover.

Note: All files you select must be from a single backup. Check the Backup ID to determine a file's backup. If you select files from multiple backups, the Save button becomes disabled.

7 Click **Next**.

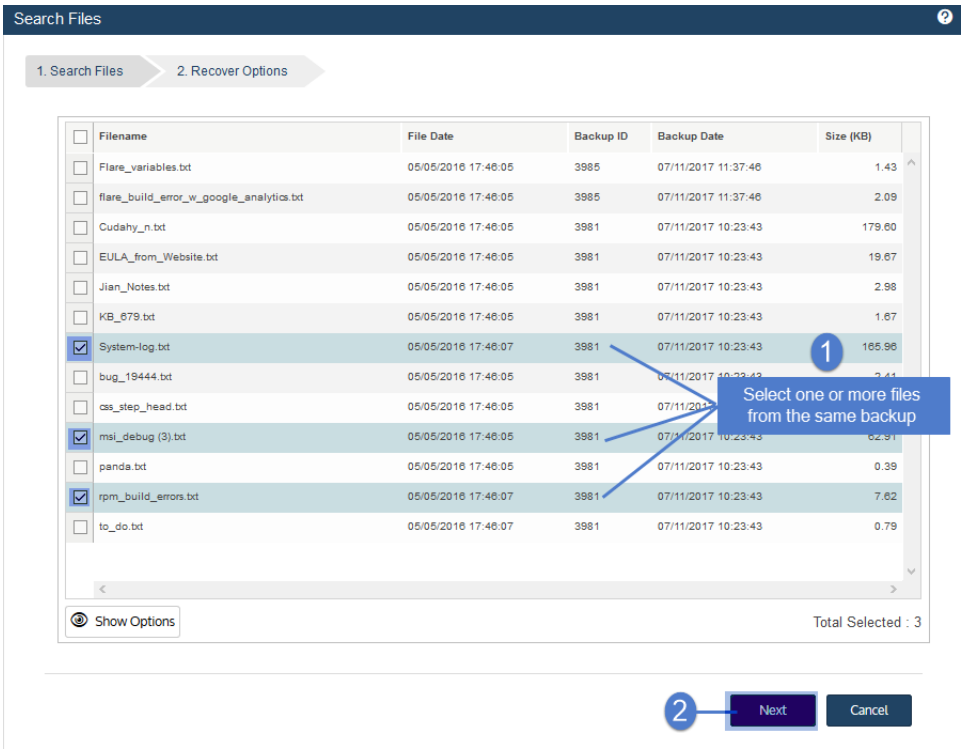

8 Select the asset where the files will be recovered.

9 In the Directory field, enter the full path where the files will be recovered.

For details, see "Specifying the target recovery location" on page 1121.

10 (Optional) Specify Exclusions.

11 (Optional) Specify Advanced Options.

12 Click **Save**.

13 Click **OK** to close the Notice message.

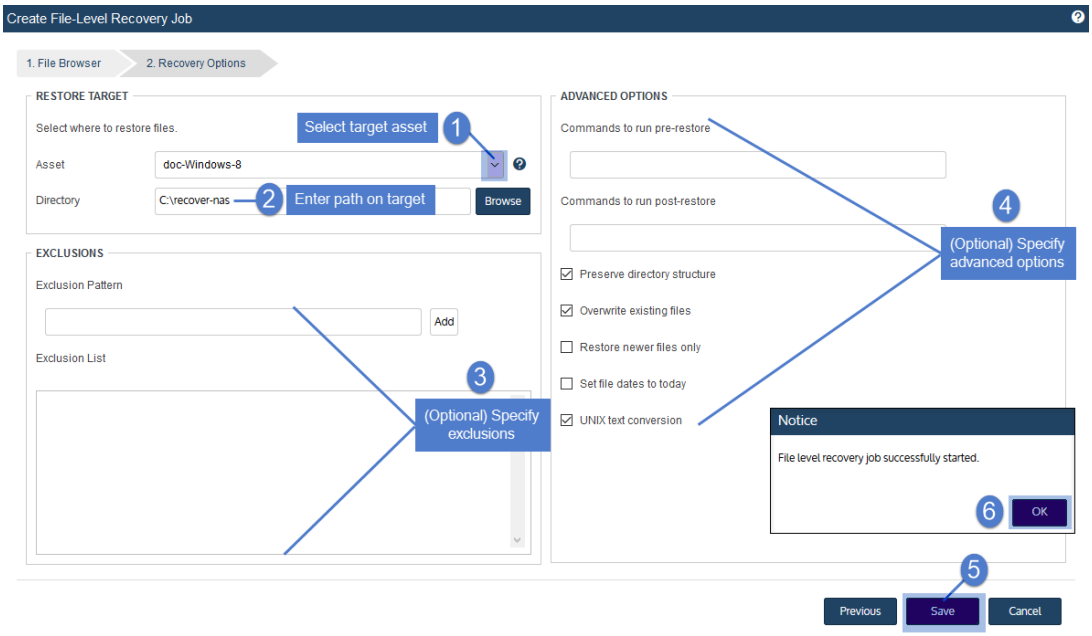

14 To monitor the recovery job:

- Select **Jobs > Active** Jobs.
- Select the job in the list and click **View Details**.

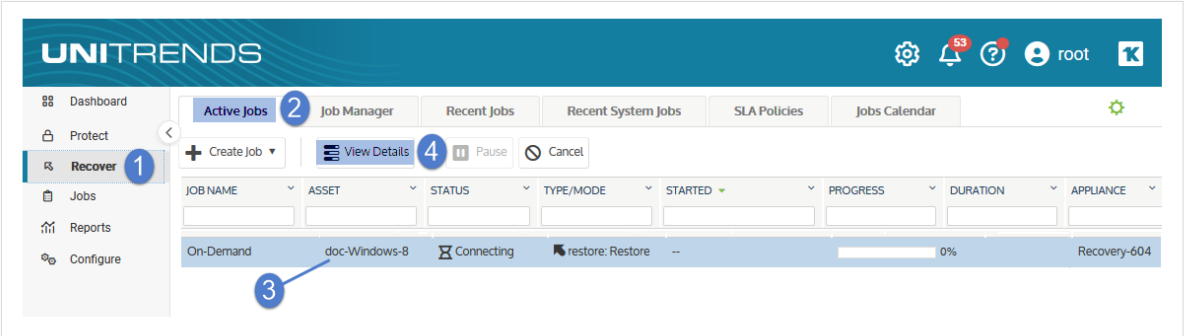

- The recovery is complete when the job's status changes to Success.

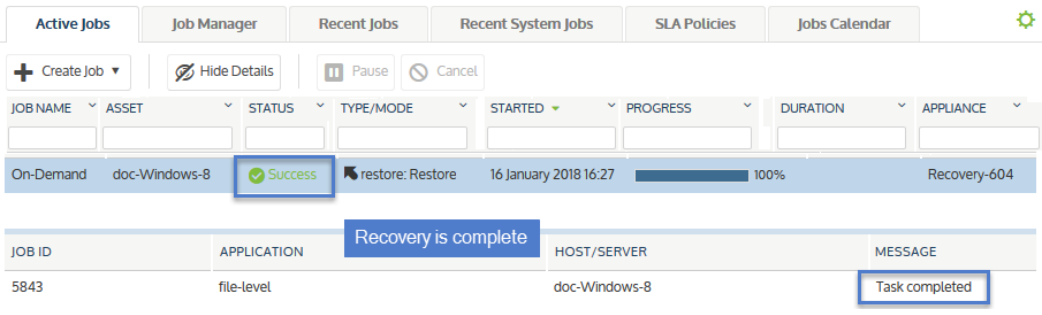

15 Access the recovered files on the recovery target.

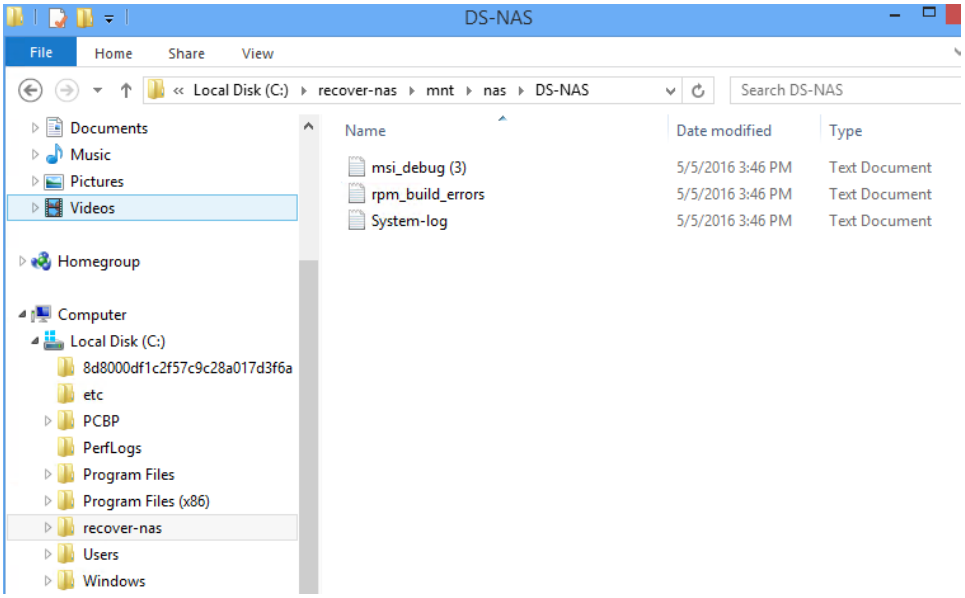

To recover an entire NAS CIFS or NFS backup by using the Backup Catalog

Run this procedure from the backup appliance to recover a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover a hot backup copy.

Note: The file browser contains the backup you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1 Log in to the backup appliance or target appliance (if recovering a hot backup copy).

2 Select **Recover > Backup Catalog**.

(Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

3 Expand the NAS asset and select one of the following:

- A NAS backup.
- An imported NAS backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
- A hot backup copy (supported if performing the recovery on the target appliance where the hot copy resides).

4 Click **Recover**.

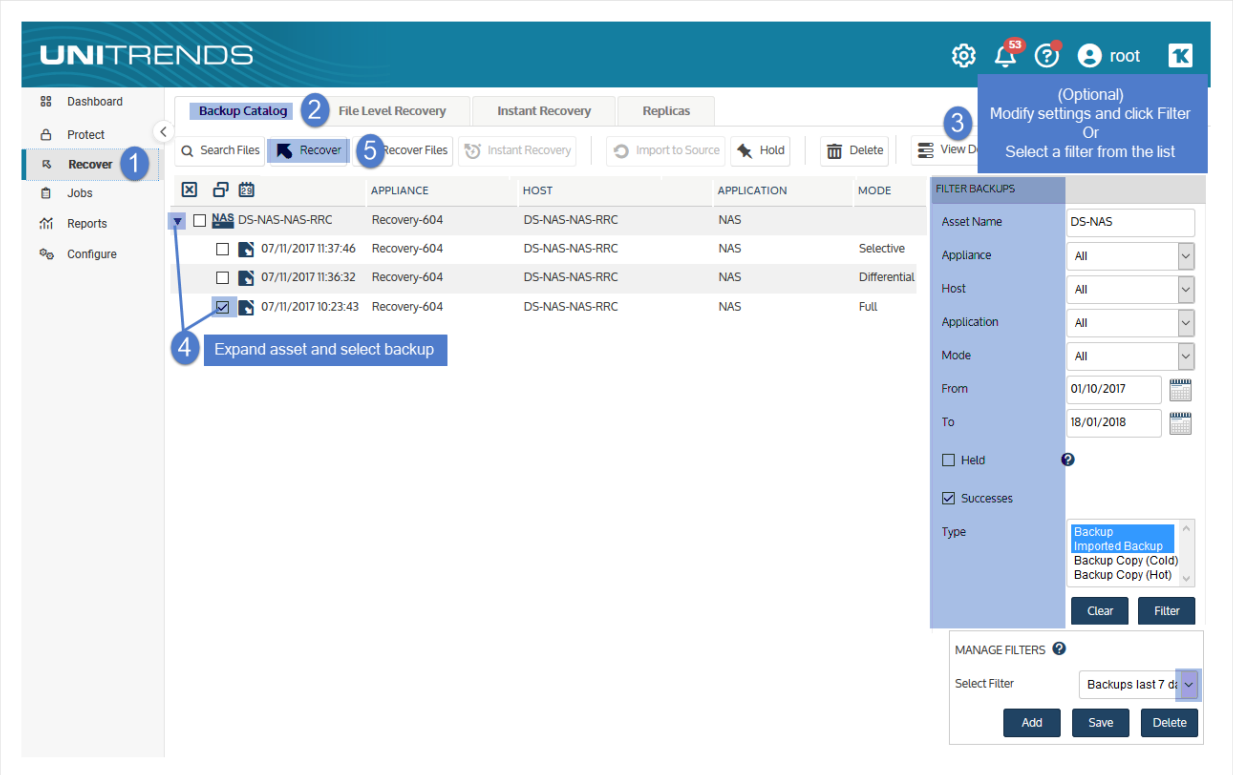

5 Select the Asset where the files will be recovered.

6 In the Directory field, enter the full path where the files will be recovered.

For details, see "Specifying the target recovery location" on page 1121.

7 (Optional) Specify Advanced Options.

8 Click **Next**.

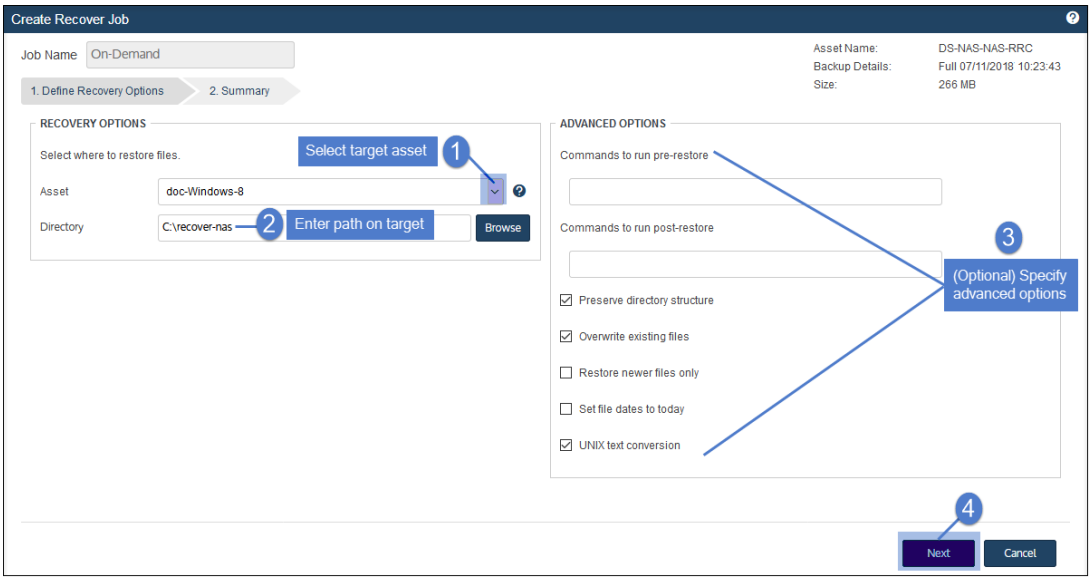

9 Review settings and click **Save**.

10 Click **OK** to close the Notice message.

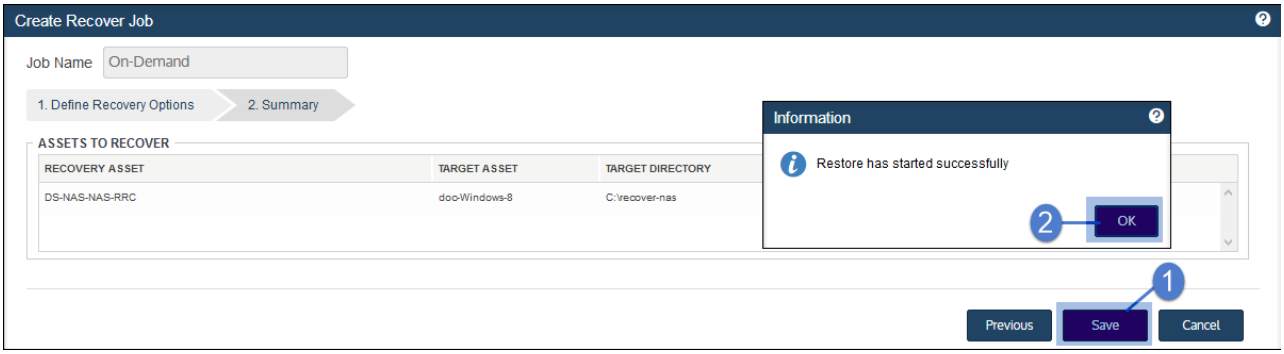

11 To monitor the recovery job:

- Select **Jobs > Active** Jobs.
- Select the job in the list and click **View Details**.

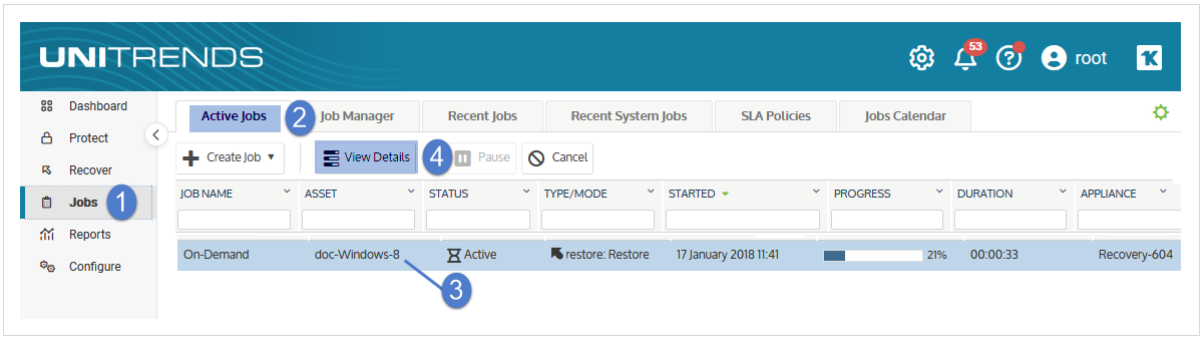

- The recovery is complete when the job's status changes to Success.

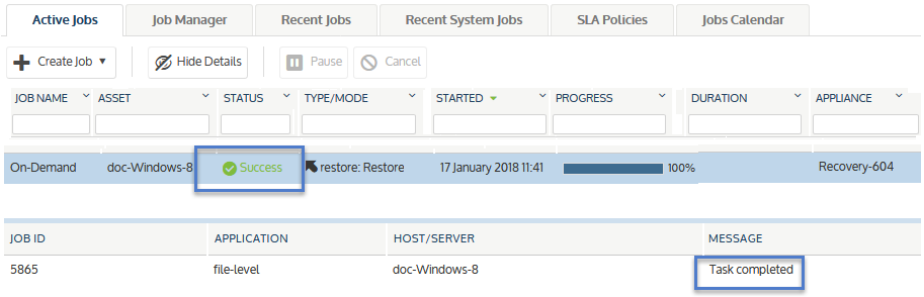

12 Access the recovered files on the recovery target.

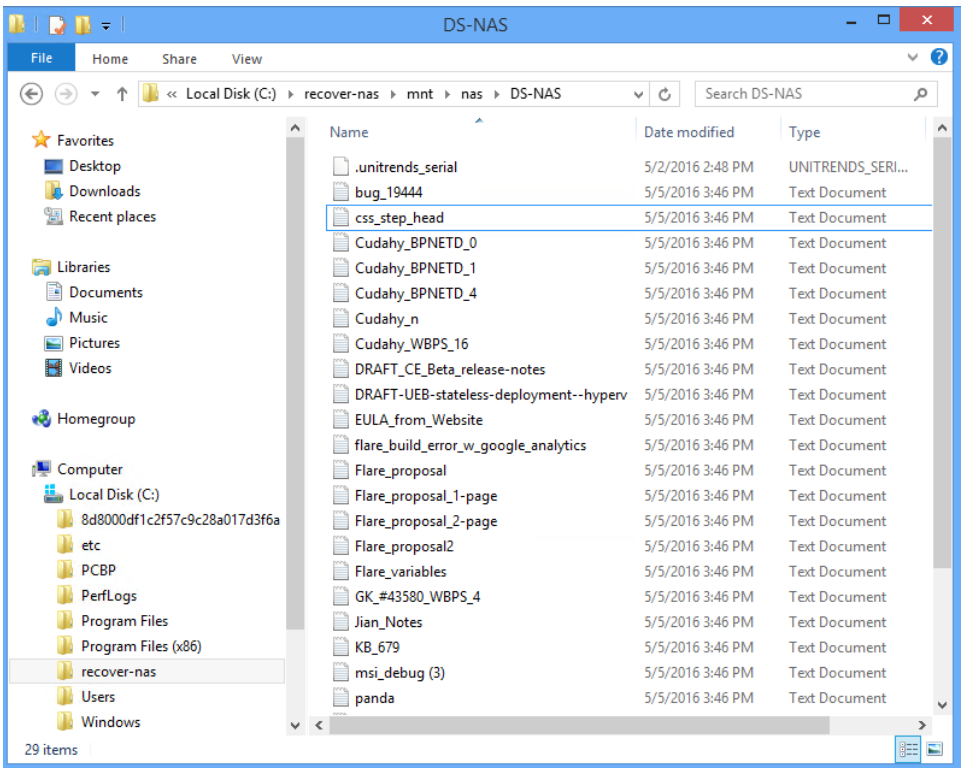

To recover an entire NAS CIFS or NFS backup by using the Backup Browser

Use this procedure to recover the entire backup.

Notes:

- This procedure is not supported for imported backups and backup copies. To recover an imported backup or hot backup copy, use the Backup Catalog procedure above.

- The file browser contains the backup you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

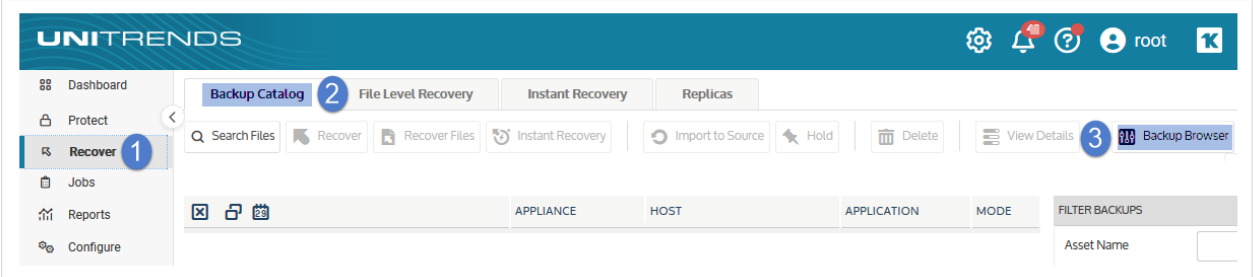

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display:

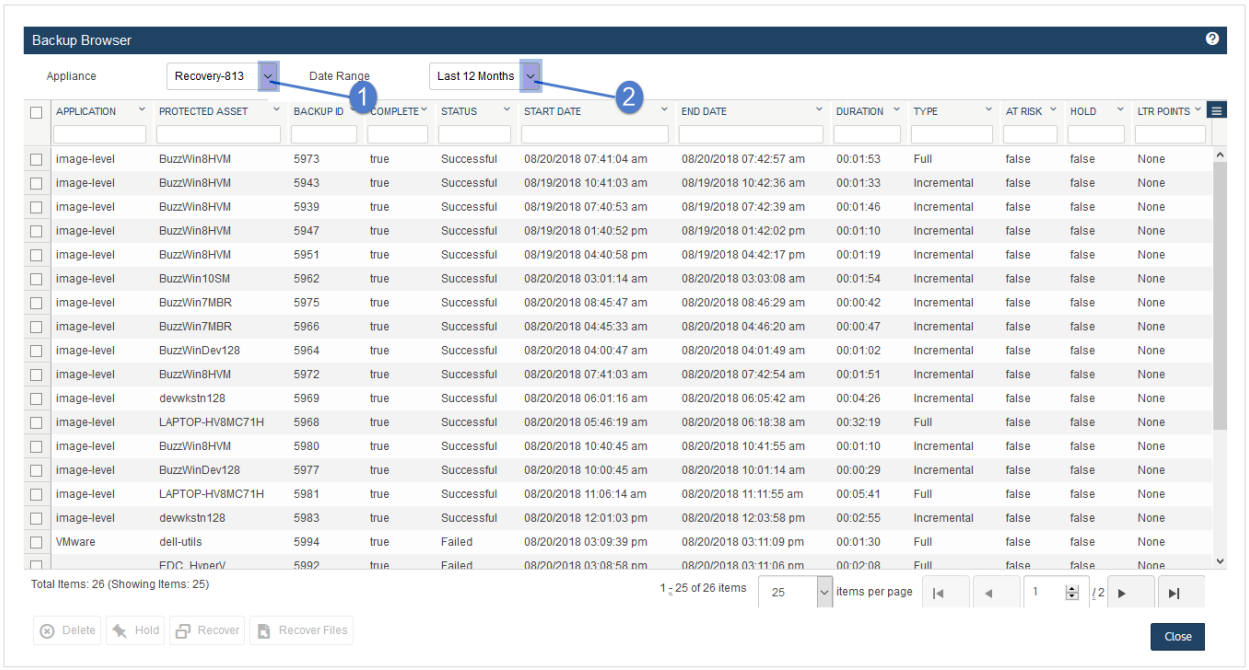

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Select the NAS backup and click **Recover**.

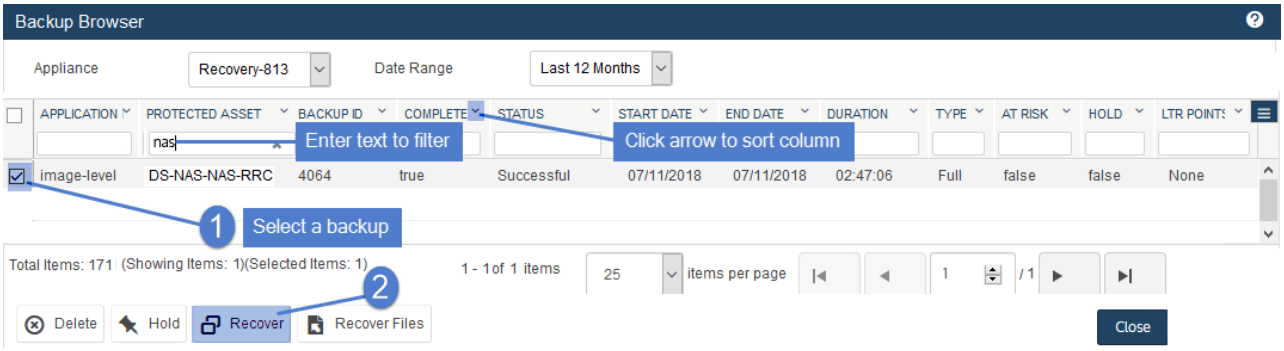

6 Select the Asset where the files will be recovered.

7 In the Directory field, enter the full path where the files will be recovered.

For details, see "Specifying the target recovery location" on page 1121.

8 (Optional) Specify Advanced Options.

9 Click **Next**.

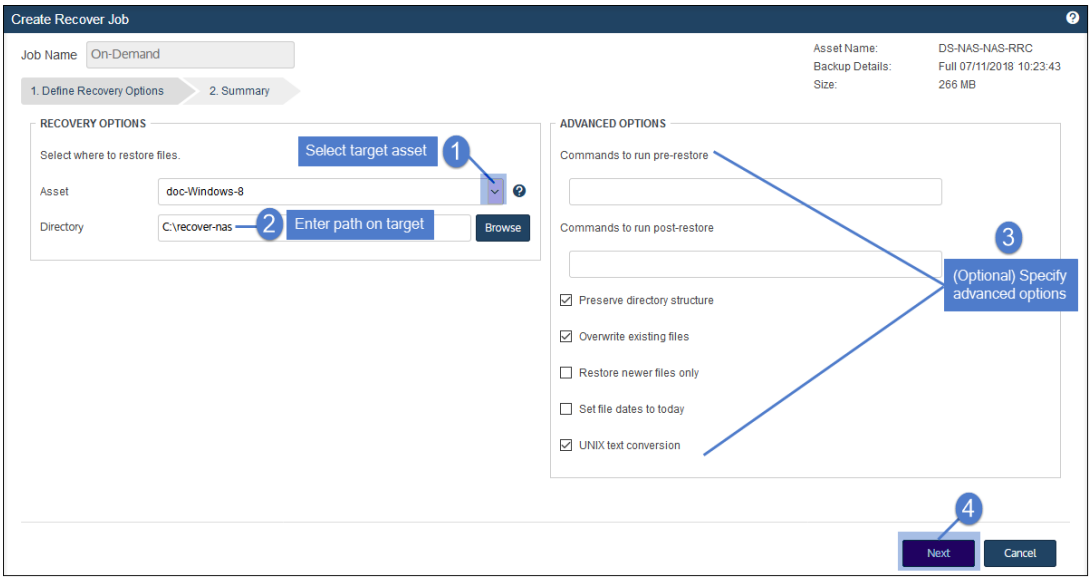

10 Review settings and click **Save**.

11 Click **OK** to close the Notice message.

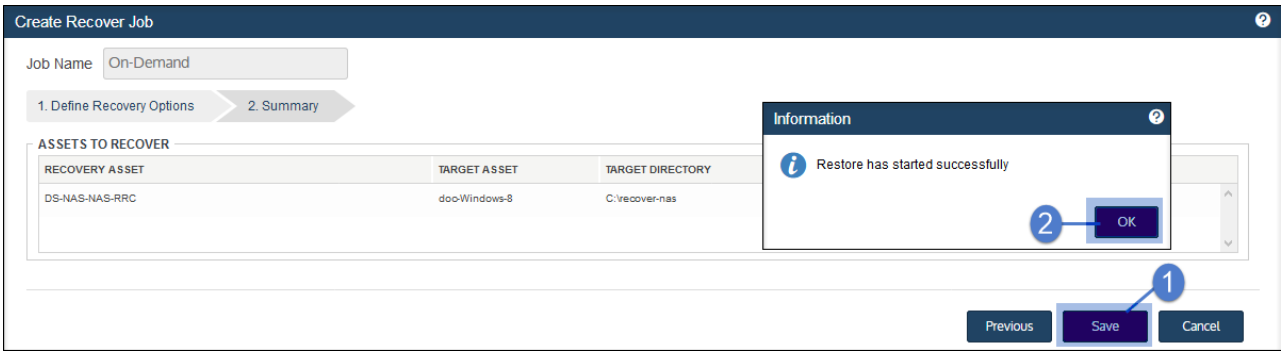

12 To monitor the recovery job:

- Select **Jobs** > **Active** Jobs.
- Select the job in the list and click **View Details**.

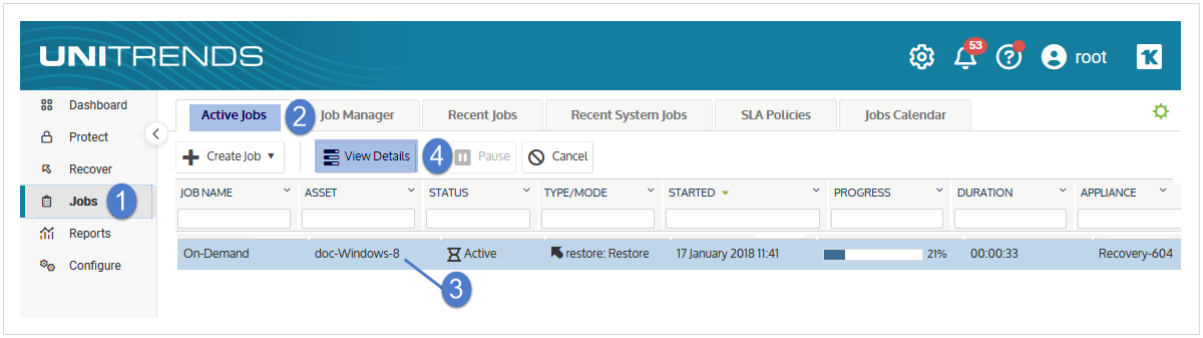

- The recovery is complete when the job's status changes to Success.

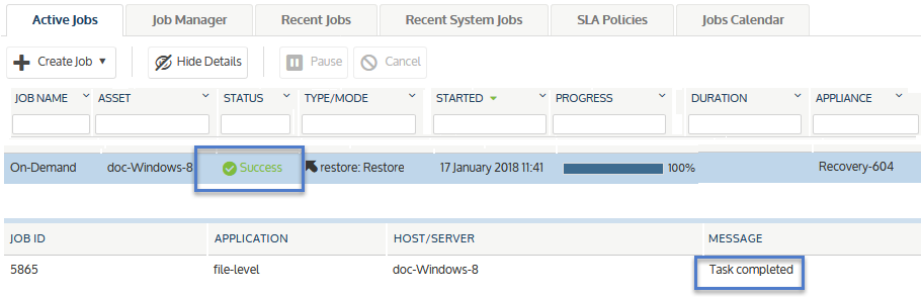

13 Access the recovered files on the recovery target.

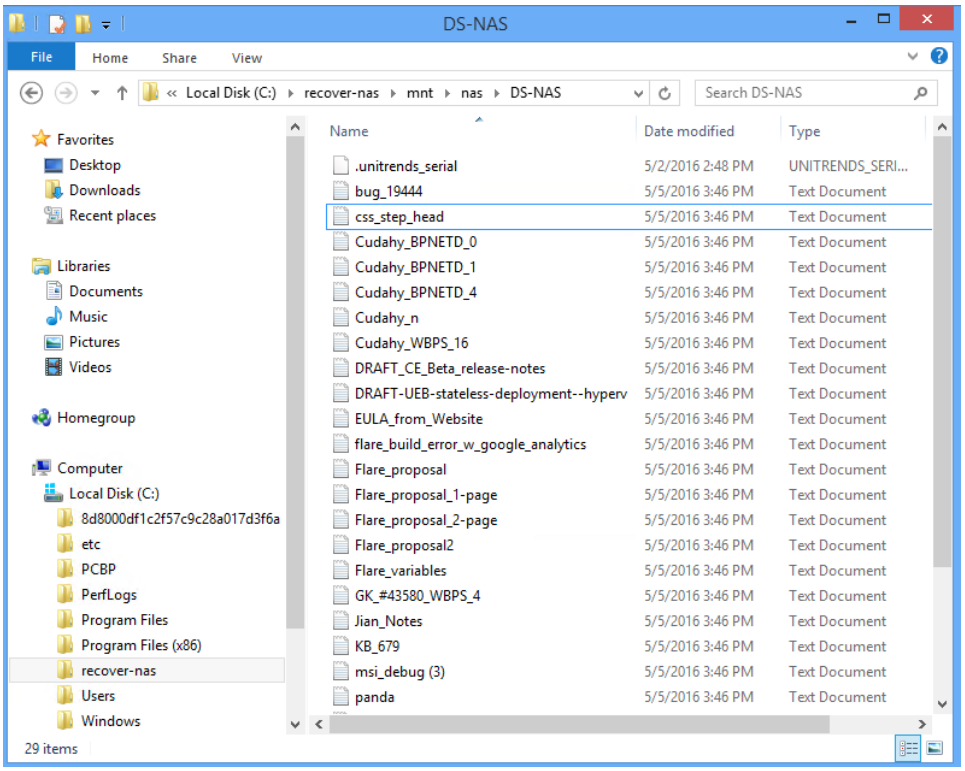

To browse one NAS CIFS or NFS backup and recover files by using the Backup Catalog

Run this procedure from the backup appliance to recover from a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover from a hot backup copy.

Note: The file browser contains the backup you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
2. Select **Recover > Backup Catalog**.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the NAS asset and select one of the following:
   - A NAS backup.
   - An imported NAS backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
   - A hot backup copy (supported if performing the recovery on the target appliance where the hot copy resides).
4. Click **Recover Files**.

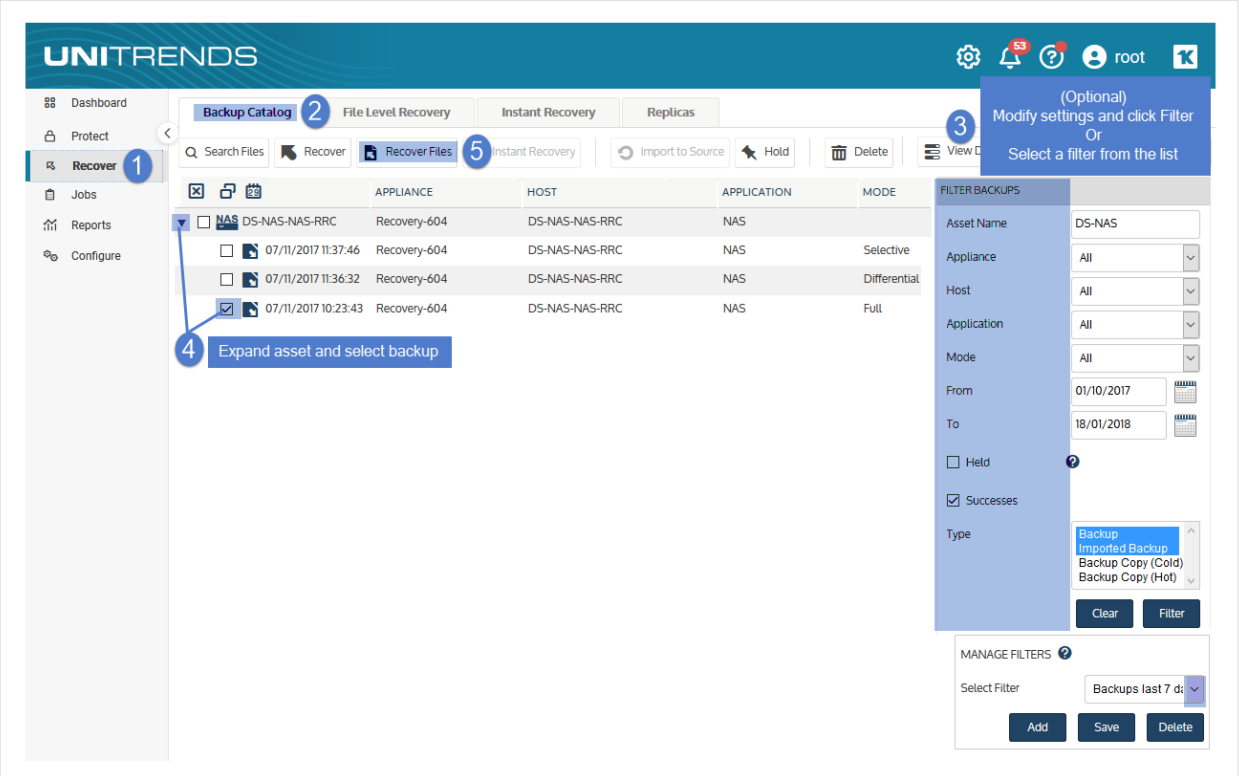

5 In the File Browser, expand folders to view items in the backup.

6 Select or drag files and/or folders to recover.

7 Click **Next**.

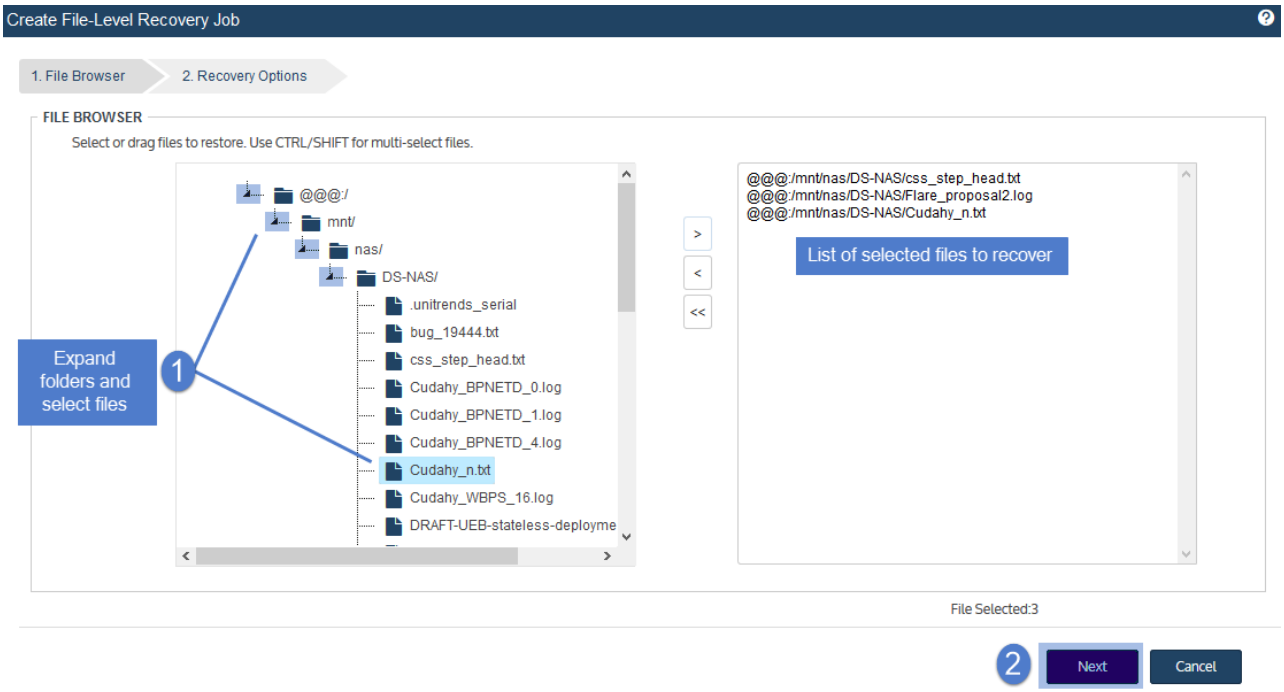

8 Select the asset where the files will be recovered.

9 In the Directory field, enter the full path where the files will be recovered.

For details, see "Specifying the target recovery location" on page 1121.

10 (Optional) Specify Exclusions.

11 (Optional) Specify Advanced Options.

12 Click **Save**.

13 Click **OK** to close the Notice message.

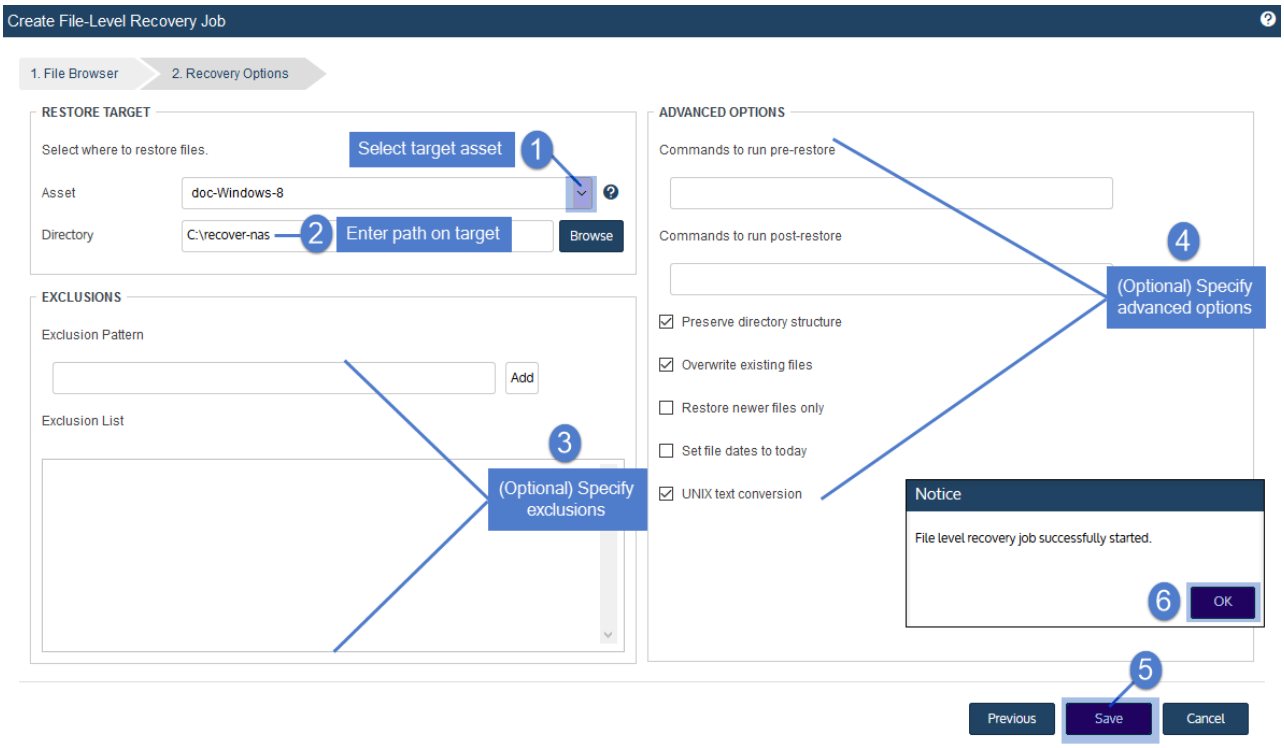

14 To monitor the recovery job:

- Select **Jobs > Active** Jobs.
- Select the job in the list and click **View Details**.

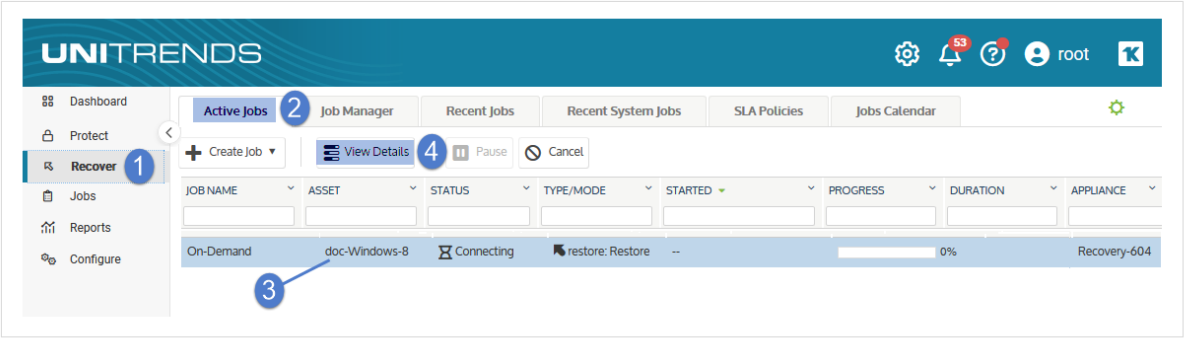

- The recovery is complete when the job's status changes to Success.

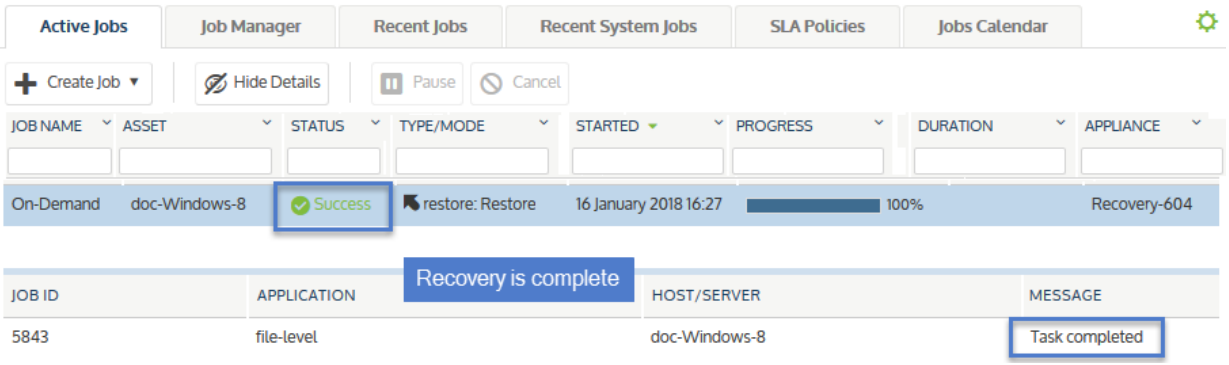

15 Access the recovered files on the recovery target.

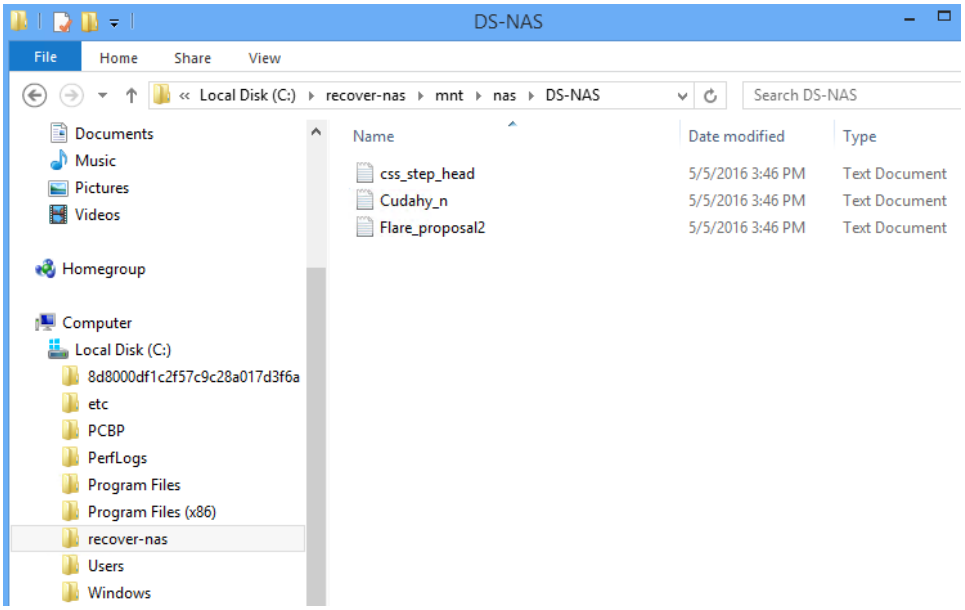

To browse one NAS CIFS or NFS backup and recover files by using the Backup Browser

Use this procedure to browse a NAS backup and recover selected files.

Notes:

- This procedure is not supported for imported backups and backup copies. To recover from an imported backup or hot backup copy, use the Backup Catalog procedure above.
- The file browser contains the backup you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

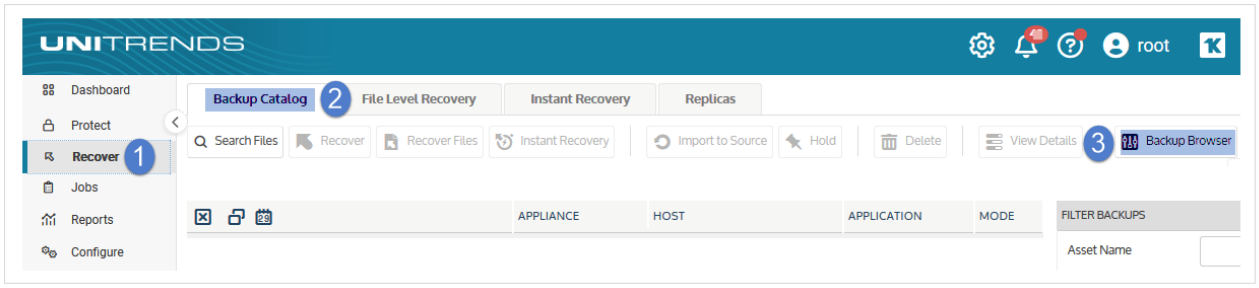

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display:

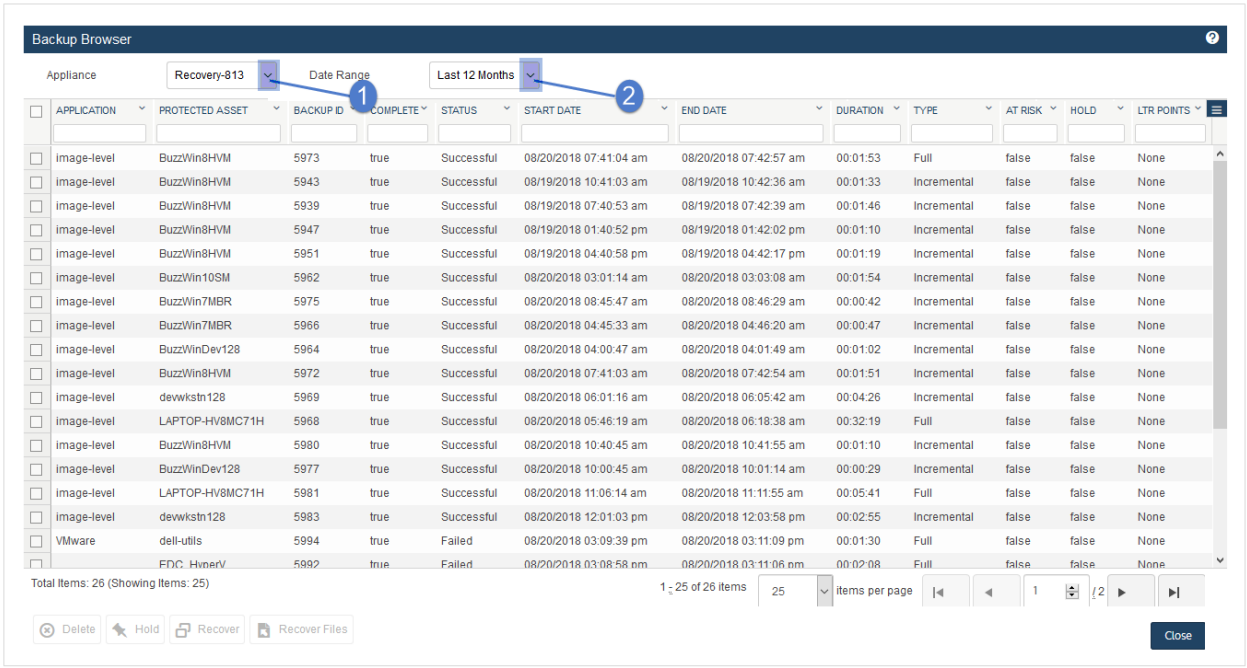

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Select the NAS backup and click **Recover Files**.

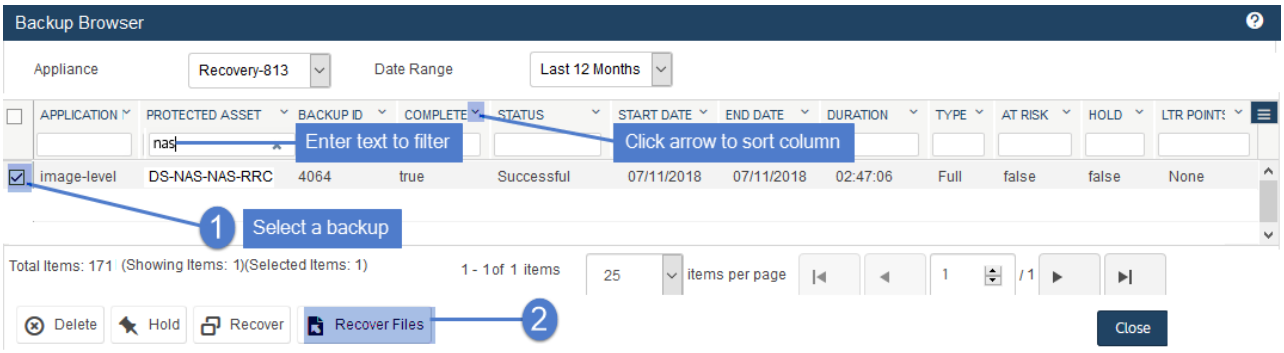

6 In the File Browser, expand folders to view items in the backup.

7 Select or drag files and/or folders to recover.

8 Click **Next**.

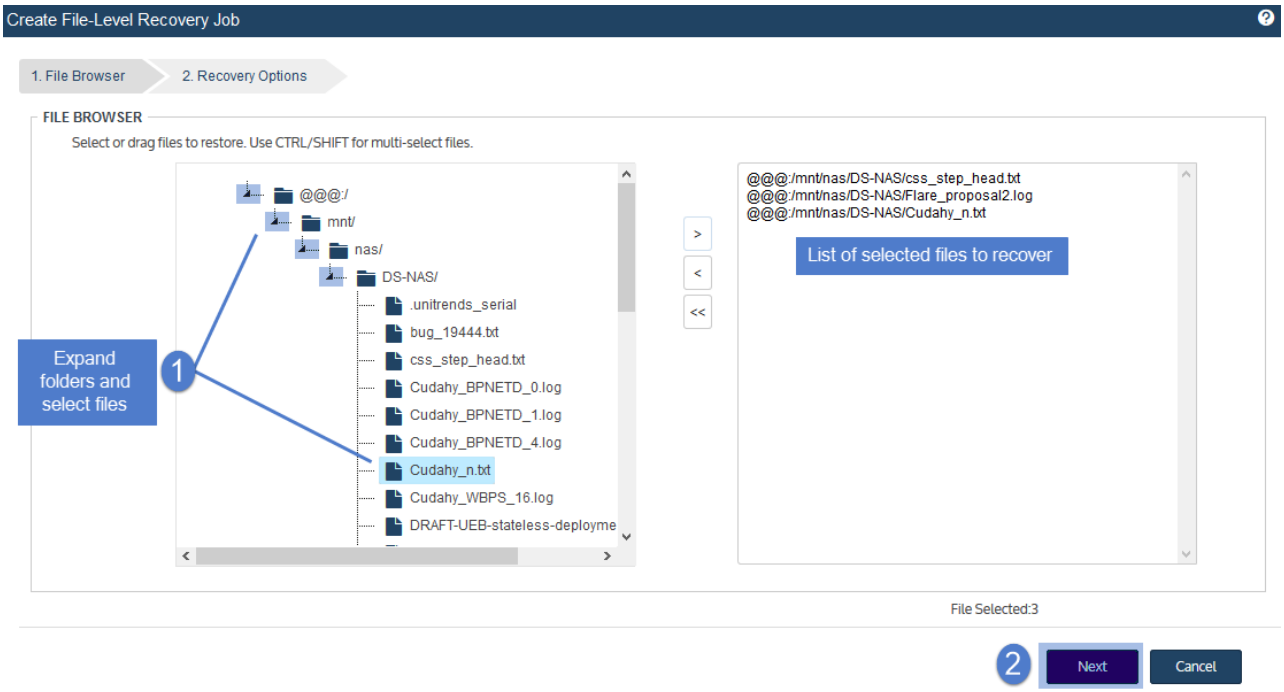

9 Select the asset where the files will be recovered.

10 In the Directory field, enter the full path where the files will be recovered.

For details, see "Specifying the target recovery location" on page 1121.

11 (Optional) Specify Exclusions.

12 (Optional) Specify Advanced Options.

13 Click **Save**.

14 Click **OK** to close the Notice message.

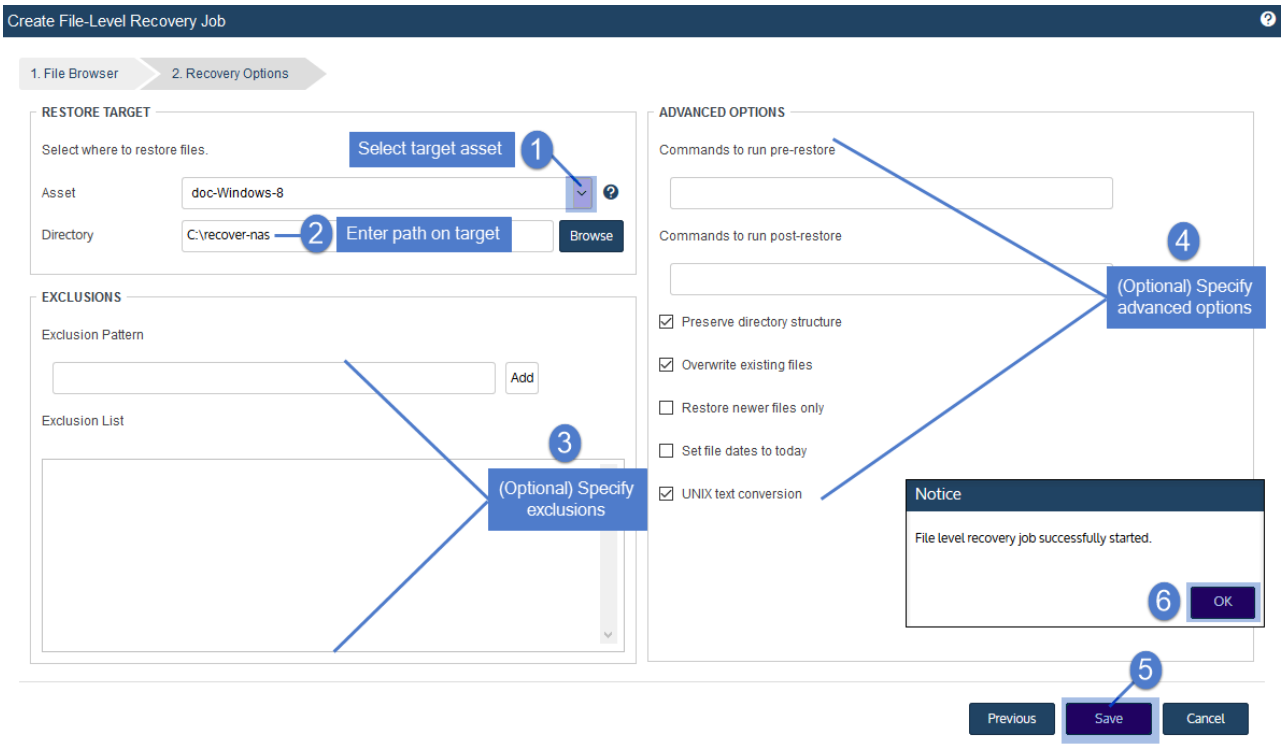

15 To monitor the recovery job:

- Select **Jobs > Active** Jobs.
- Select the job in the list and click **View Details**.

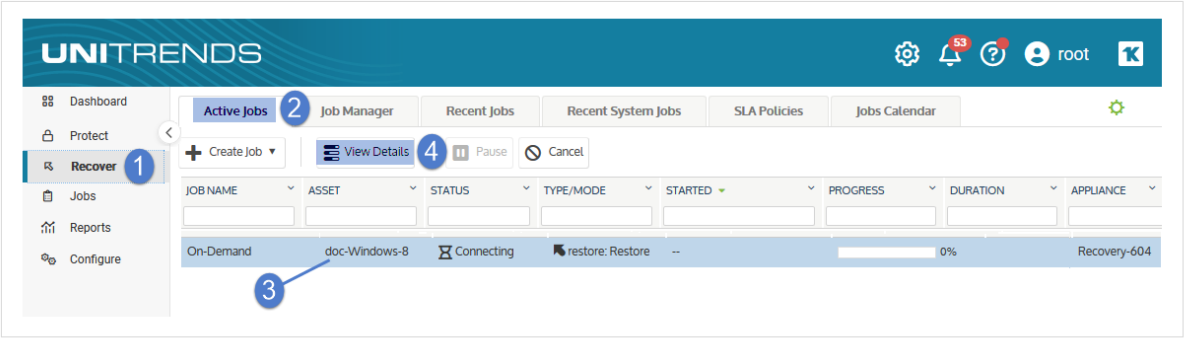

- The recovery is complete when the job's status changes to Success.

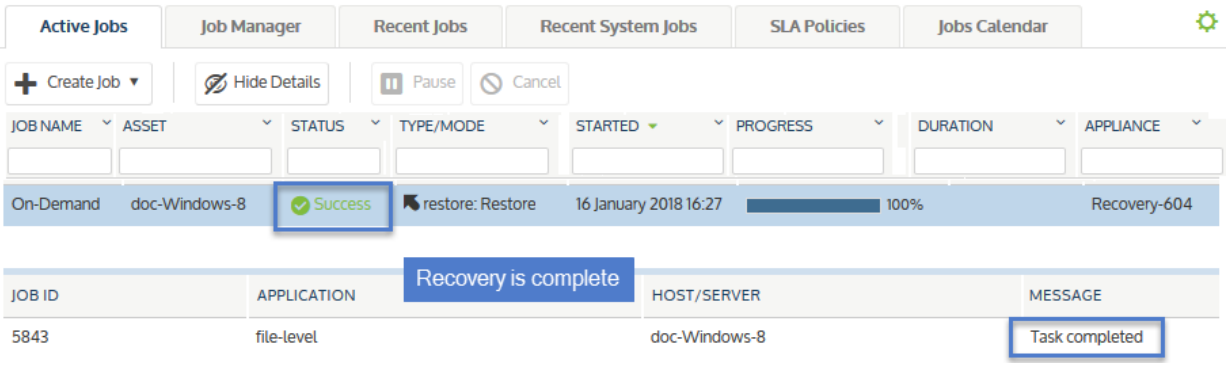

16 Access the recovered files on the recovery target.

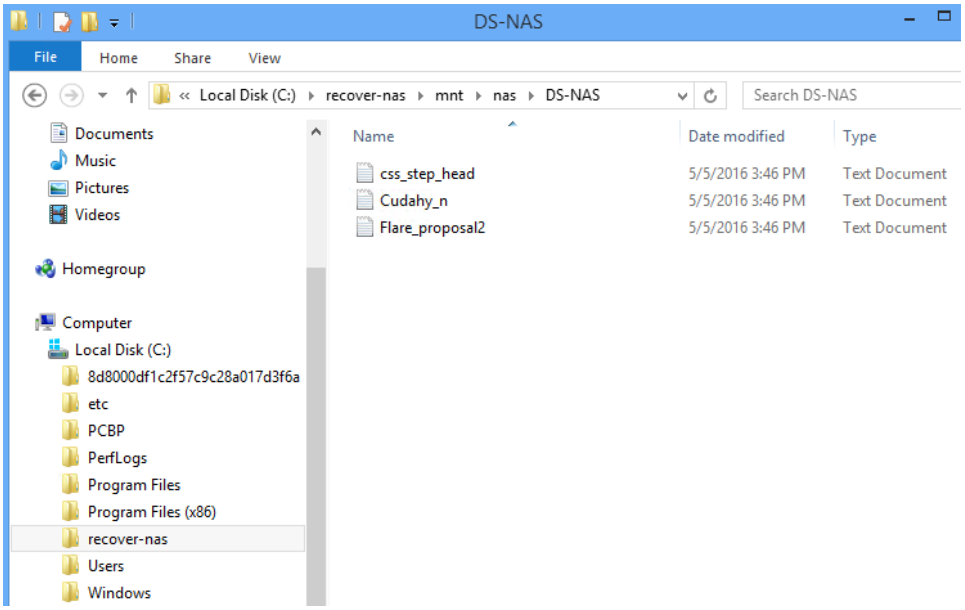

Recovering NAS NDMP backups

The following limitations apply to NAS NDMP recovery:

- NDMP backups can only be recovered to NDMP devices of the same vendor.
- Supported recovery targets vary by vendor.

  Recovery to the original location is supported for all vendors. See the vendor documentation to determine whether you can recover to another location on the original NDMP device, or to another NDMP device that has been added to the appliance as an asset.
- Point-in-time recovery of the entire backup group is supported.
- Recovery of selected files is supported for some NDMP devices from the certified vendors. See the vendor documentation for compatibility limitations.

- When performing point-in-time recovery of an NDMP volume, you cannot specify files to include or exclude. The volume is recovered exactly as it was at the selected recovery point.
- Recovery of selected files that contain non-UTF-8 compatible characters is not supported. Instead you must recover the entire backup.

You can recover from a backup by using the Backup Catalog or the Backup Browser. For imported backups and backup copies, you must use the Backup Catalog. See these topics for details:

- "To recover an entire NDMP backup by using the Backup Catalog" on page 1141
- "To recover an entire NDMP backup by using the Backup Browser" on page 1142
- "To recover files from an NDMP backup by using the Backup Catalog" on page 1143
- "To recover files from an NDMP backup by using the Backup Browser" on page 1144

To recover an entire NDMP backup by using the Backup Catalog

Run this procedure from the backup appliance to recover a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover a hot backup copy.

1. Log in to the backup appliance or target appliance (if recovering a hot backup copy).
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the NDMP asset and select one of the following:
   - An NDMP backup.
   - An imported NDMP backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
   - A hot backup copy (supported if performing the recovery on the target appliance where the hot copy resides).
4. Click **Recover**.
5. Specify a recovery target:

   Note: Some filers only support recovery to the original location on the original NDMP device. See the vendor documentation to determine which options are supported for your filer.

   - Select an NDMP device in the Asset list.

     Only devices of the same vendor display. Additional vendor compatibility limitations apply.
   - Select a target Volume.

     Only directories on the selected device that are online and have enough space for the recovery display.
   - (Optional) Enter the target Directory to which the backup will be recovered.
     - The target directory cannot exceed 255 characters.

- If the directory entered does not exist, it is created within the selected volume.
- If this field is left blank, the backup is recovered to the target volume.

6 Click **Next**.

7 Review the job details, then click **Save**.

The recovery job is queued immediately.

8 Click **OK** to close the Information message.

To view the running job, select **Jobs > Active Job**.

If you encounter a permissions error when attempting to access an NDMP backup that was recovered to an environment with different permissions, unmount the target volume and remount it with the appropriate permissions.

To recover an entire NDMP backup by using the Backup Browser

Use this procedure to recover the entire backup.

Notes:

- This procedure is not supported for imported backups and backup copies. To recover an imported backup or hot backup copy, use the Backup Catalog procedure above.
- The file browser contains the backup you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display:

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Select the NDMP backup and click **Recover**.

6 Specify a recovery target:

Note: Some filers only support recovery to the original location on the original NDMP device. See the vendor documentation to determine which options are supported for your filer.

- Select an NDMP device in the Asset list.

  Only devices of the same vendor display. Additional vendor compatibility limitations apply.

- Select a target Volume.

  Only directories on the selected device that are online and have enough space for the recovery display.

- (Optional) Enter the target Directory to which the backup will be recovered.
  - The target directory cannot exceed 255 characters.
  - If the directory entered does not exist, it is created within the selected volume.
  - If this field is left blank, the backup is recovered to the target volume.

7 Click **Next**.

8 Review the job details, then click **Save**.

The recovery job is queued immediately.

9 Click **OK** to close the Information message.

To view the running job, select **Jobs > Active Job**.

If you encounter a permissions error when attempting to access an NDMP backup that was recovered to an environment with different permissions, unmount the target volume and remount it with the appropriate permissions.

To recover files from an NDMP backup by using the Backup Catalog

Run this procedure from the backup appliance to recover selected files from a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover files from a hot backup copy.

Notes:

- Recovering files is not supported for all NDMP filers. See the vendor documentation to determine whether recovering files is supported.
- Recovery of files that contain non-UTF-8 compatible characters is not supported.
- The file browser contains the backup you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1 Log in to the backup appliance or target appliance (if recovering a hot backup copy).

2 Select **Recover** and click the **Backup Catalog** tab.

(Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

3 Expand the NDMP asset and select one of the following:

- An NDMP backup.
- An imported NDMP backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
- A hot backup copy (supported if performing the recovery on the target appliance where the hot copy resides).

4 Click **Recover Files**.

5 In the File Browser, expand folders to view items in the backup.

6 Select or drag files and/or folders to recover.

7 Click **Next**.

8 Specify a recovery target:

> **Note:** Some filers only support recovery to the original location on the original NDMP device. See the vendor documentation to determine which options are supported for your filer.

- Select an NDMP device in the Asset list.

  Only devices of the same vendor display. Additional vendor compatibility limitations apply.

- Select a target Volume.

  Only directories on the selected device that are online and have enough space for the recovery display.

- (Optional) Enter the target Directory to which the backup will be recovered.
  - The target directory cannot exceed 255 characters.
  - If the directory entered does not exist, it is created within the selected volume.
  - If this field is left blank, the backup is recovered to the target volume.

9 Click **Save**.

The recovery job is queued immediately.

10 Click **OK** to close the Information message.

To view the running job, select **Jobs > Active Job**.

If you encounter a permissions error when attempting to access files that were recovered to an environment with different permissions, unmount the target volume and remount it with the appropriate permissions.

To recover files from an NDMP backup by using the Backup Browser

Use this procedure to recover selected files from an NDMP backup.

Notes:

- This procedure is not supported for imported backups and backup copies. To recover from an imported backup or hot backup copy, use the Backup Catalog procedure above.
- Recovering files is not supported for all NDMP filers. See the vendor documentation to determine whether recovering files is supported.
- Recovery of files that contain non-UTF-8 compatible characters is not supported.

- The file browser contains the backup you select plus all dependent data in the backup group. For example, selecting an incremental enables you to browse the incremental plus its parent full and any other dependent incrementals in the chain. (For details, see "Backup groups" on page 98.)

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog** and click **Backup Browser**.
3. Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.
4. (Optional) Refine the search:
   - Enter text in any column field to filter the display.
   - Click an arrow to sort by column.
   - Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
   - For a description of each column, see "Backup Browser column descriptions".
5. Select the NDMP backup and click **Recover Files**.
6. Select or drag files and/or folders to recover.
7. Click **Next**.
8. Specify a recovery target:

   Note: Some filers only support recovery to the original location on the original NDMP device. See the vendor documentation to determine which options are supported for your filer.

   - Select an NDMP device in the Asset list.

     Only devices of the same vendor display. Additional vendor compatibility limitations apply.
   - Select a target Volume.

     Only directories on the selected device that are online and have enough space for the recovery display.
   - (Optional) Enter the target Directory to which the backup will be recovered.
     - The target directory cannot exceed 255 characters.
     - If the directory entered does not exist, it is created within the selected volume.
     - If this field is left blank, the backup is recovered to the target volume.
9. Click **Save**.

   The recovery job is queued immediately.
10. Click **OK** to close the Information message.

    To view the running job, select **Jobs > Active Job**.

If you encounter a permissions error when attempting to access files that were recovered to an environment with different permissions, unmount the target volume and remount it with the appropriate permissions.

This page is intentionally left blank.

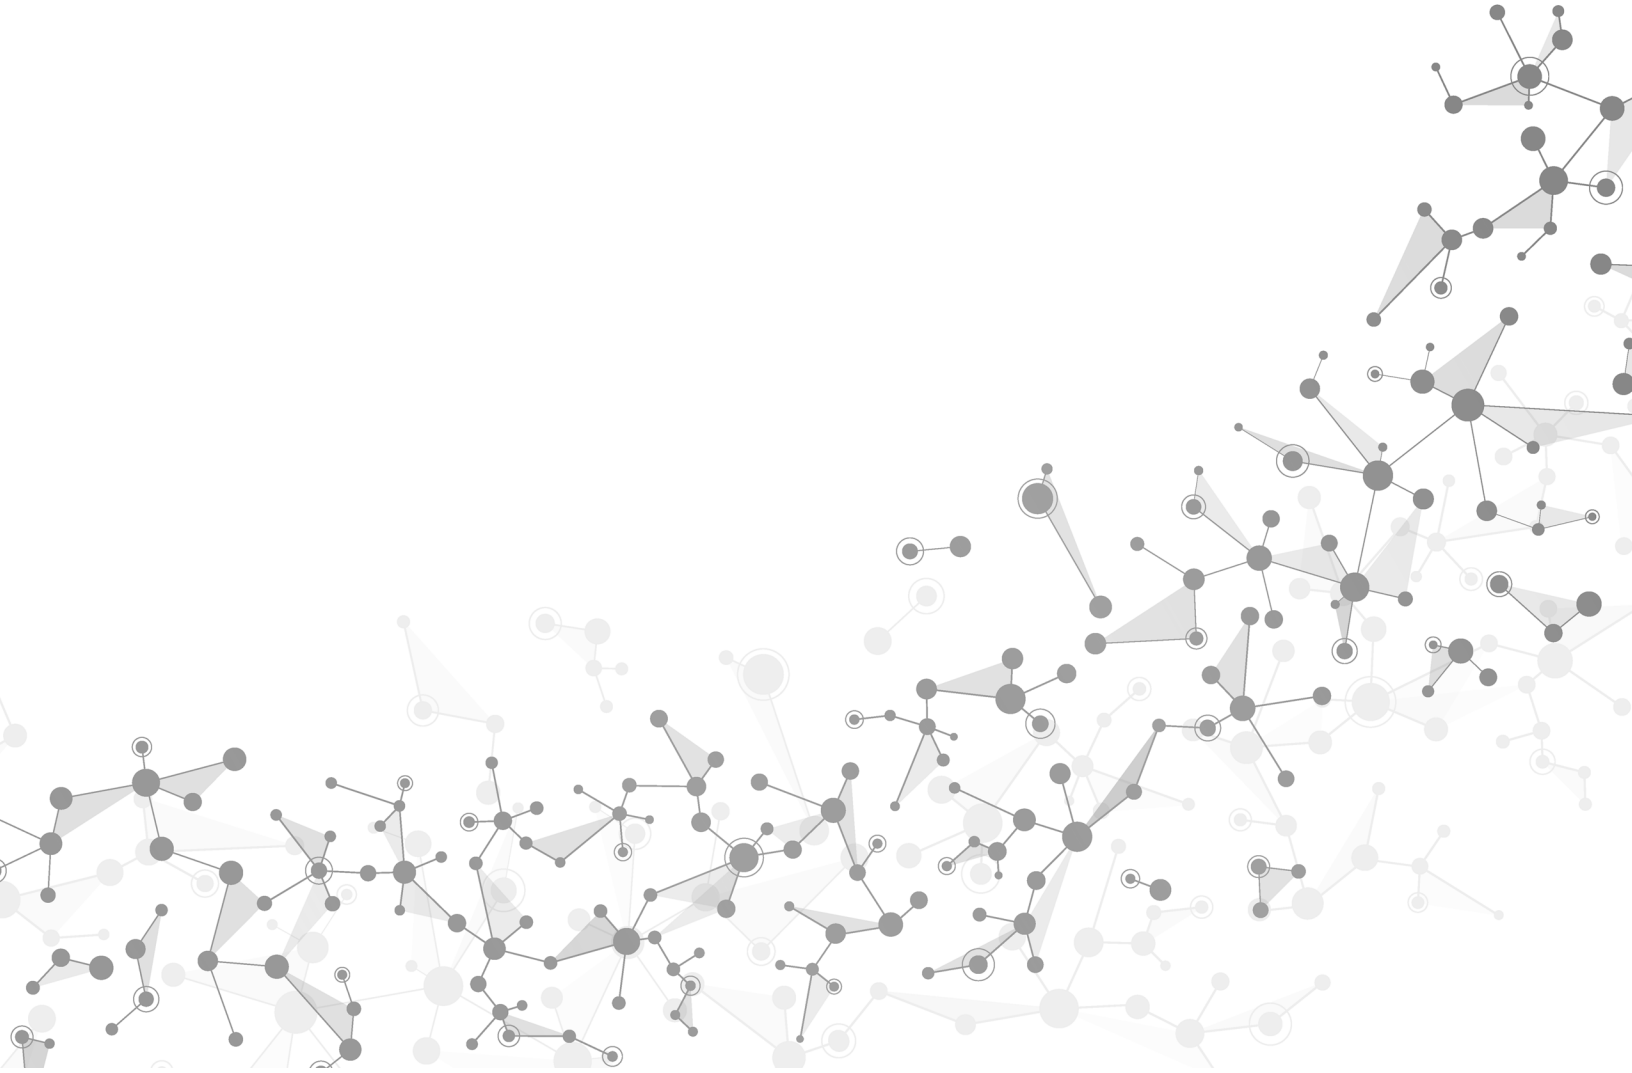

Chapter 18: Recovering Application Backups

This section provides requirements, procedures, and details for recovering applications. The recovery process varies by application type. See the following for details on a specific application:

Recovering Exchange backups

Use the recovery feature to recover an entire database, storage group, or selected items from Exchange backups. See the following topics for details:

Preparing to recover Exchange backups

Before recovering, see these topics for details specific to your Exchange environment:

About recovering Exchange 2016, 2013, and 2010 from a backup

Exchange 2016, 2013, and 2010 use recovery databases. Each server has a recovery database, and there can only be a single mounted recovery database at a time.

Once you have created the recovery database, recover the backup to it. Then use the Microsoft Exchange Management Shell to extract mailbox data from the information store into the local *mail.pst* file. You can also merge the extracted data back into the currently active information store.

Recovery databases differ from the RSG (Recovery Storage Group) mechanism in Exchange 2007 and 2003.

About recovering Exchange 2007 from a backup

Exchange 2007 uses the RSG (Recovery Storage Group) mechanism. RSG enables you to mount a second copy of an Exchange information store on any Exchange server that belongs to the same Exchange Administrative Group as the

original, while the original information store is still active. This feature enables you to recover data from the backup copy of the information store without interfering with the on-going operation of the Exchange server.

Once you have created the RSG, you first recover the backup to it and then use the Microsoft Exchange Management Shell in Exchange 2007 to extract mailbox data from the information store into the local mail .pst file. Optionally, you may also merge the extracted data back into the currently active information store.

RSG differs from the recovery database mechanism in Exchange 2016, 2013, and 2010.

About recovering Exchange 2003 from a backup

Direct recovery to the RSG (Recovery Storage Group) is not permitted for Exchange 2003 backups. Instead only recovering to the original location or to an alternate location is supported.

Recovering an Exchange database or storage group

Use the procedures in this section to recover an Exchange database or storage group to a specific target. Before the recovery, verify the recovery target is set up as required and that any restrictions have been met. Choose from the following recovery targets:

- "Recovering to the original Exchange server" on page 1148
- "Recovering to a recovery area" on page 1153
- "Recovering to an alternate location" on page 1156

Recovering to the original Exchange server

Recovery to the Exchange server is the default recovery type. All database and transaction log files recover directly to the original location. To perform a successful recovery to the original Exchange server, the following conditions must be met:

Condition	Explanation
Database name and file name must remain unchanged from the time the backup was performed.	The database name is the symbolic or displayable name of the database. For example, *Mailbox Database* or *Mailbox1*. The actual database file name, for example *Mailbox1.edb*, must also be unchanged since the backup was run. Note that the location of the database files and transaction log files may be changed after the backup has been performed, if needed. If the log files or database files must be moved to another volume or disk, the actual names of the database files must be preserved.
Databases must be dismounted.	For Exchange 2003 and 2007, this includes all databases contained in the storage group. For Exchange 2016, 2013, and 2010, only the database being recovered must be dismounted.
Database must be in a Clean Shutdown state.	If the database is in a Dirty Shutdown state, you can recover the backup, but need to bring the database into a Clean Shutdown state to mount the database. After recovering, if you cannot mount the database, see this Microsoft article to determine

Condition	Explanation
	whether this is the problem: Exchange Database Is in a Dirty Shutdown State.
Databases must be marked as **overwrite allowed on restore**.	All databases must have the **overwrite allowed on restore** flag set. This task can be performed using the Exchange server administrative console or the appropriate Exchange server command line utility. If this is not the case, the recovery fails.
Remove all existing database and transaction log files.	Unitrends recommends that all database and transaction log files be removed from the recovery location. Recovering a differential, incremental, or a full backup recovers the server to a specific point-in-time. To ensure that the storage group or database can be remounted without integrity errors, any existing database and transaction log files should be removed before performing the recovery.

You can recover from a backup by using the Backup Catalog or the Backup Browser. To recover from an imported backup or hot backup copy, you must use the Backup Catalog. See these topics for details:

- "To recover a database or storage group to the original Exchange server by using the Backup Catalog"
- "To recover a database or storage group to the original Exchange server by using the Backup Browser" on page 1151

To recover a database or storage group to the original Exchange server by using the Backup Catalog

Run this procedure from the backup appliance to recover an entire backup or imported backup copy. Run this procedure from the backup copy target appliance to recover an entire hot backup copy.

1. Verify that all prerequisites in "Recovering to the original Exchange server" on page 1148 have been met.
2. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
3. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
4. Expand the Exchange asset and select one of the following:
   - An Exchange backup.
   - An imported Exchange backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
   - A hot backup copy (supported if performing the recovery on the target appliance where the hot copy resides).
5. Click **Recover**.

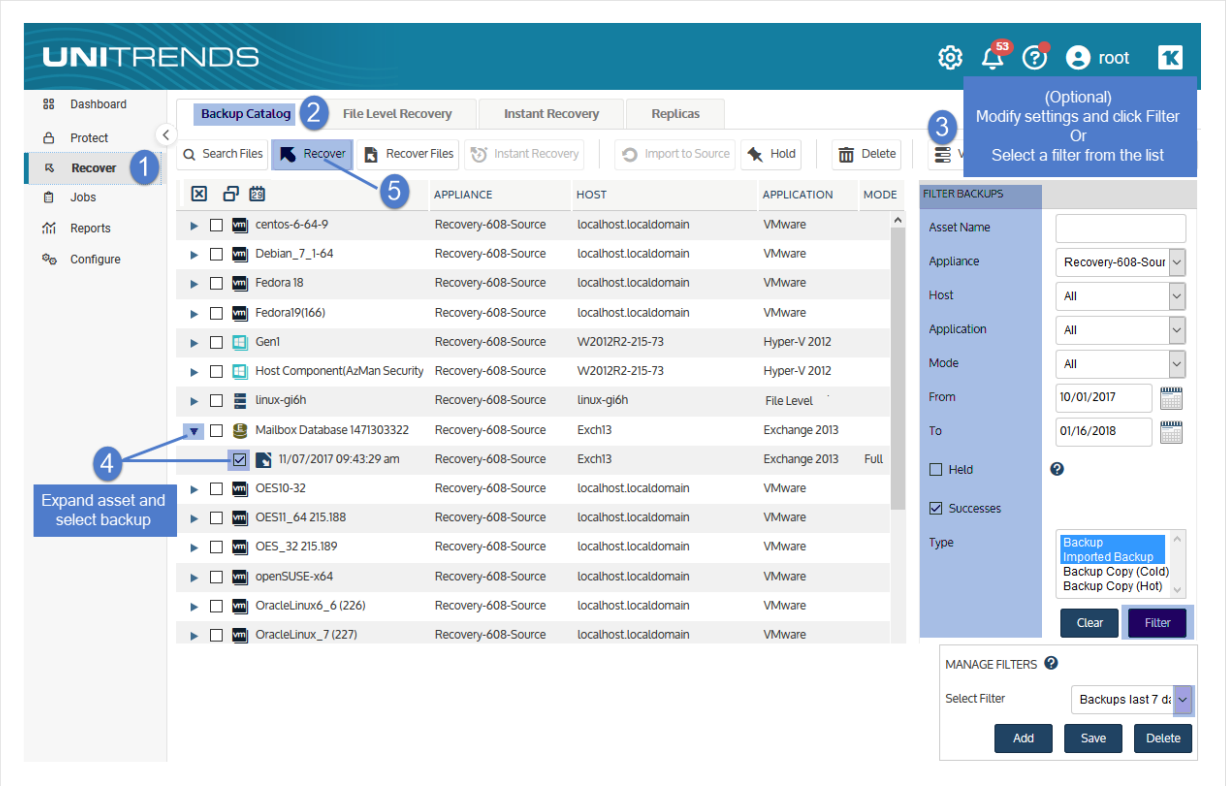

6 Select these Recovery Options:

- **Exchange Server** from the Recovery Target list.
- **Exchange** from the Asset list.
- **Original Location** for the Recovery Location(Optional) Specify asset-side commands to run by entering any system command or user script in **Commands to run pre-restore** and **Commands to run post-restore**.

7 Click **Next**.

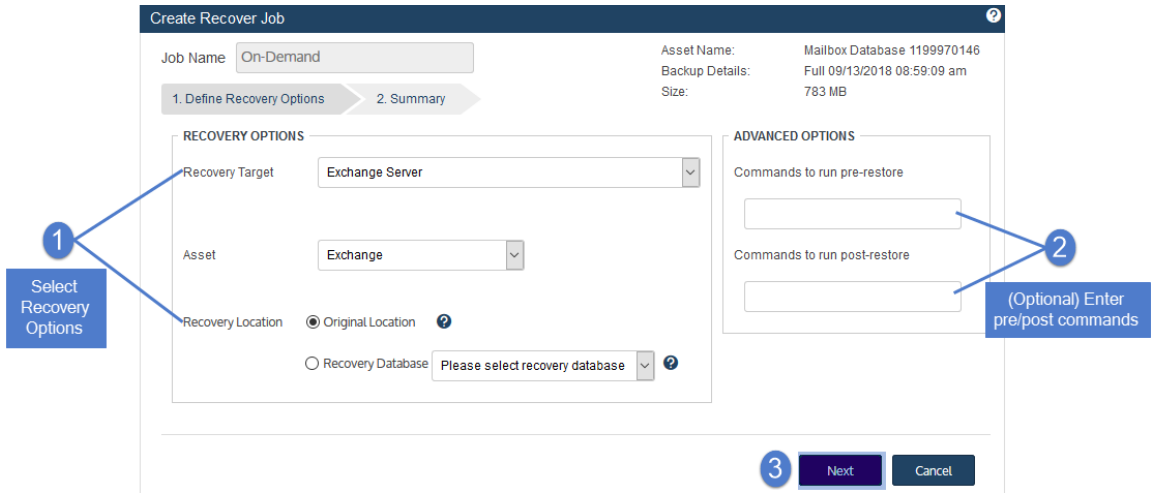

8 Click **Save**.

- The recovery job is queued immediately. To view the running job, select **Jobs > Active Job**.
- All database and transaction log files are recovered directly to the original location.

9 Re-mount any databases you dismounted for the recovery.

To recover a database or storage group to the original Exchange server by using the Backup Browser

Use this procedure to recover a backup to the original location.

Note: This procedure is not supported for imported backups or hot backup copies. To recover an imported backup or backup copy, use the Backup Catalog procedure above.

1 Verify that all prerequisites in "Recovering to the original Exchange server" on page 1148 have been met.

2 Log in to the backup appliance.

3 Select **Recover > Backup Catalog** and click **Backup Browser**.

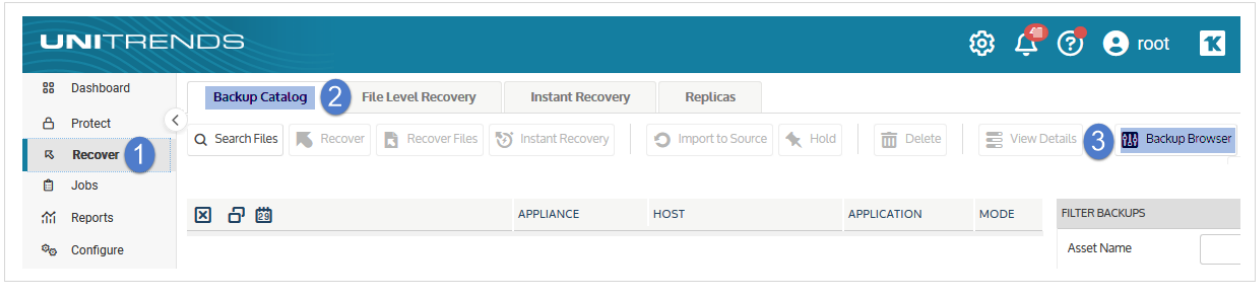

4 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.

5 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

6 Select the Exchange backup and click **Recover**.

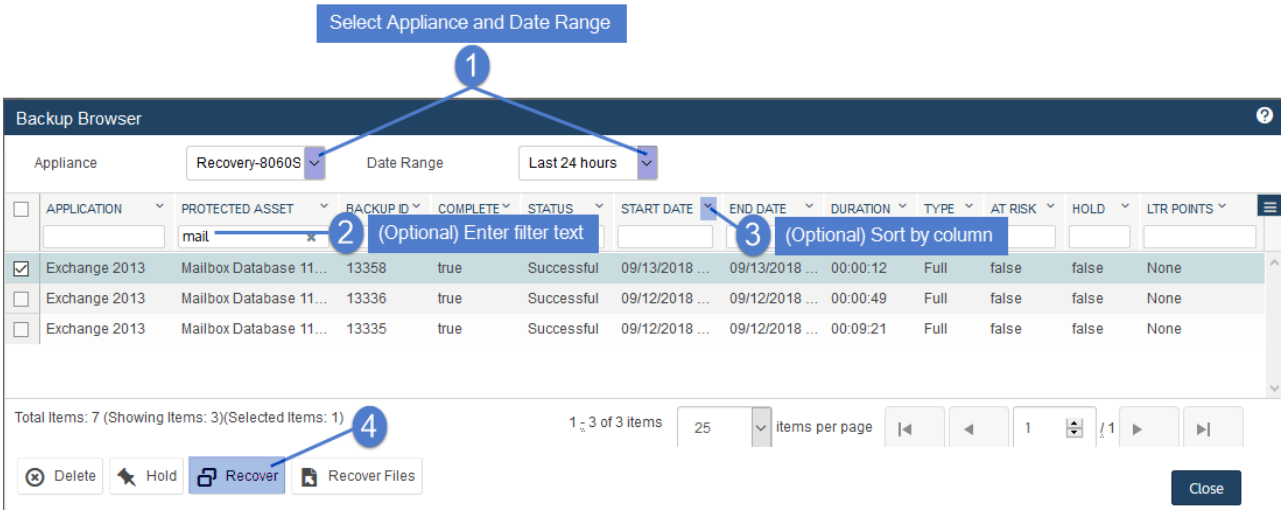

7 Select these Recovery Options:

- **Exchange Server** from the Recovery Target list.
- **Exchange** from the Asset list.
- **Original Location** for the Recovery Location(Optional) Specify asset-side commands to run by entering any system command or user script in **Commands to run pre-restore** and **Commands to run post-restore**.

8 Click **Next**.

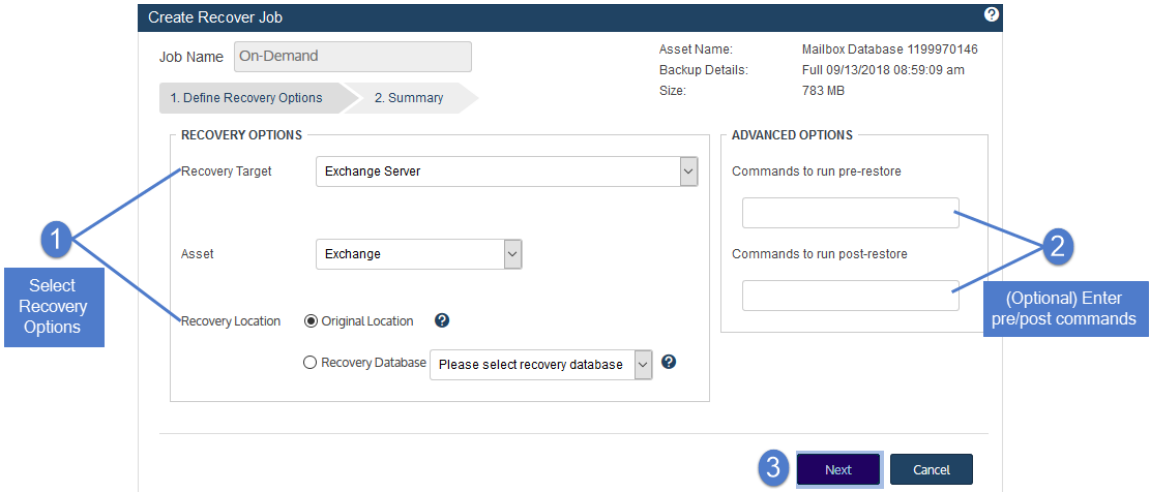

9 Click **Save**.

- The recovery job is queued immediately. To view the running job, select **Jobs > Active Job**.
- All database and transaction log files are recovered directly to the original location.

10 Re-mount any databases you dismounted for the recovery.

Recovering to a recovery area

Recovery to a recovery area is supported only for Exchange 2016/2013/2010 (a recovery database) and Exchange 2007 (an RSG, or recovery storage group). It is not supported for Exchange 2003 or earlier versions. Additionally, it is only available if there is a recovery database or recovery storage group available in the backup.

To perform a successful recovery to a recovery database or recovery storage group, the following conditions must be met:

Condition	Explanation
Exchange 2016/2013/2010 recovery database or Exchange 2007 RSG	Exchange 2003 or earlier versions are not supported.
Databases must be dismounted.	For Exchange 2007, this includes all databases contained in the storage group. For Exchange 2016, 2013 and 2010, only the recovery databases must be dismounted.
Database must be in a Clean Shutdown state.	If the database is in a Dirty Shutdown state, you can recover the backup but need to bring the database into a Clean Shutdown state to mount the database. After recovering, if you cannot mount the database, see this Microsoft article to determine whether this is the problem: Exchange Database Is in a Dirty Shutdown State.
Databases must be marked as **overwrite allowed on restore**.	All databases must have the **overwrite allowed on restore** flag set. This task can be performed using the Exchange server administrative console or the appropriate Exchange server command line utility. If this is not the case, the recovery fails.
Remove all existing database and transaction log files.	Unitrends recommends that all database and transaction log files be removed from the recovery location. Recovering a differential, incremental, or a full backup recovers the server to a specific point-in-time state. To ensure that the storage group or database can be remounted without integrity errors, any existing database and transaction log files should be removed before performing the recovery.
[Exchange 2007 only] The RSG must contain the same number of mailbox databases and public	Each database filename (e.g., mailbox1.edb, publicfolder.edb) created in the recovery storage group must match the corresponding database file name in the original storage group that is being recovered. Creating recovery storage groups using the Exchange 2007 Administrative Console enforces this rule.

Condition	Explanation
folder databases as the original storage group	

You can recover from a backup by using the Backup Catalog or the Backup Browser. To recover from an imported backup or hot backup copy, you must use the Backup Catalog. See these topics for details:

- "To recover a database or storage group to a recovery area by using the Backup Catalog"
- "To recover a database or storage group to a recovery area by using the Backup Browser" on page 1155

To recover a database or storage group to a recovery area by using the Backup Catalog

Run this procedure from the backup appliance to recover an entire backup or imported backup copy. Run this procedure from the backup copy target appliance to recover an entire hot backup copy.

1. Verify that all prerequisites in "Recovering to a recovery area" on page 1153 have been met.
2. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
3. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
4. Expand the Exchange asset and select one of the following:
   - An Exchange backup.
   - An imported Exchange backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
   - A hot backup copy (supported if performing the recovery on the target appliance where the hot copy resides).
5. Click **Recover**.

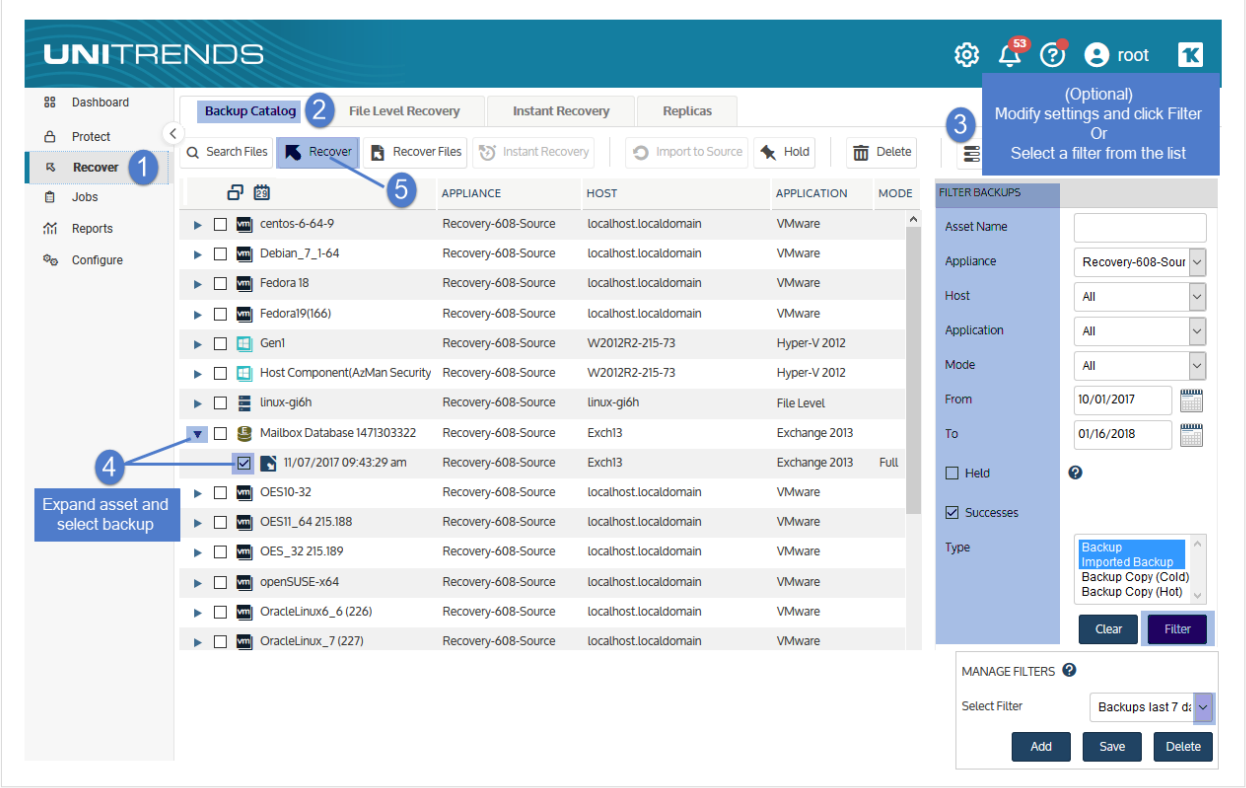

6 Select **Original Exchange Server** as the Recovery Target.

7 Select a Recovery Database from the drop-down.

8 (Optional) Specify asset-side commands to run by entering any system command or user script in **Commands to run pre-restore** and **Commands to run post-restore**.

9 Click **Next**.

10 Click **Save**.

The recovery job is queued immediately. To view the running job, select **Jobs > Active Job**.

To recover a database or storage group to a recovery area by using the Backup Browser

Use this procedure to recover a backup to a recovery area.

Note: This procedure is not supported for imported backups or hot backup copies. To recover an imported backup or backup copy, use the Backup Catalog procedure above.

1 Verify that all prerequisites in "Recovering to a recovery area" on page 1153 have been met.

2 Log in to the backup appliance.

3 Select **Recover > Backup Catalog** and click **Backup Browser**.

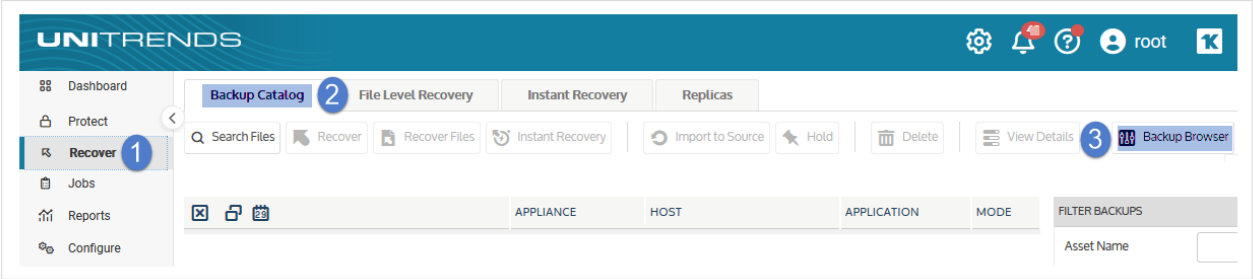

4 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.

5 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

6 Select the Exchange backup and click **Recover**.

7 Select **Original Exchange Server** as the Recovery Target.

8 Select a Recovery Database from the drop-down.

9 (Optional) Specify asset-side commands to run by entering any system command or user script in **Commands to run pre-restore** and **Commands to run post-restore**.

10 Click **Next**.

11 Click **Save**.

The recovery job is queued immediately. To view the running job, select **Jobs > Active Job**.

Recovering to an alternate location

This option recovers the Exchange information store to be recovered to a location other than the original location where it resided when the backup occurred. The alternate location can be either to the same Exchange server host, a different Windows protected asset, the network share of your Unitrends appliance, or any CIFS/NFS network storage. To recover to CIFS/NFS network storage, you must first add the storage to the Unitrends appliance as a NAS asset. For details, see "Managing NAS assets" on page 296.

The following specifies the difference between a full, differential, and incremental recovery to an alternate location:

Backup type	Explanation
Full	All data associated with the Exchange information store is recovered to the specified location.

Backup type	Explanation
Differential	Only the data contained in the differential backup is recovered to the named location; the associated full backup is not recovered. A differential recovery should be used only if certain files within the backup are required or a third-party tool is used for individual mailbox or item recovery, e.g. Ontrack PowerControls. This type of recovery can be performed to any Windows-based protected asset, and the server is not required to have Microsoft Exchange Server installed. This type of recovery may also be done to the backup appliance.
Incremental	Incremental backup chains can become very lengthy, creating a large recovery. To simplify the recovery process, the appliance creates a single recovery job on the recovery point chosen. This is not a synthetic backup, it is simply a concentration of the available backups, sent in backed-up order.

You can recover from a backup by using the Backup Catalog or the Backup Browser. To recover from an imported backup or hot backup copy, you must use the Backup Catalog. See these topics for details:

- "To recover a database or storage group to an alternate location by using the Backup Catalog"
- "To recover a database or storage group to an alternate location by using the Backup Browser" on page 1161

To recover a database or storage group to an alternate location by using the Backup Catalog

Run this procedure from the backup appliance to recover an entire backup or imported backup copy. Run this procedure from the backup copy target appliance to recover an entire hot backup copy.

1. Verify that all prerequisites in "Recovering to an alternate location" on page 1156 have been met.
2. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
3. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
4. Expand the Exchange asset and and select one of the following:
   - An Exchange backup.
   - An imported Exchange backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
   - A hot backup copy (supported if performing the recovery on the target appliance where the hot copy resides).
5. Click **Recover**.

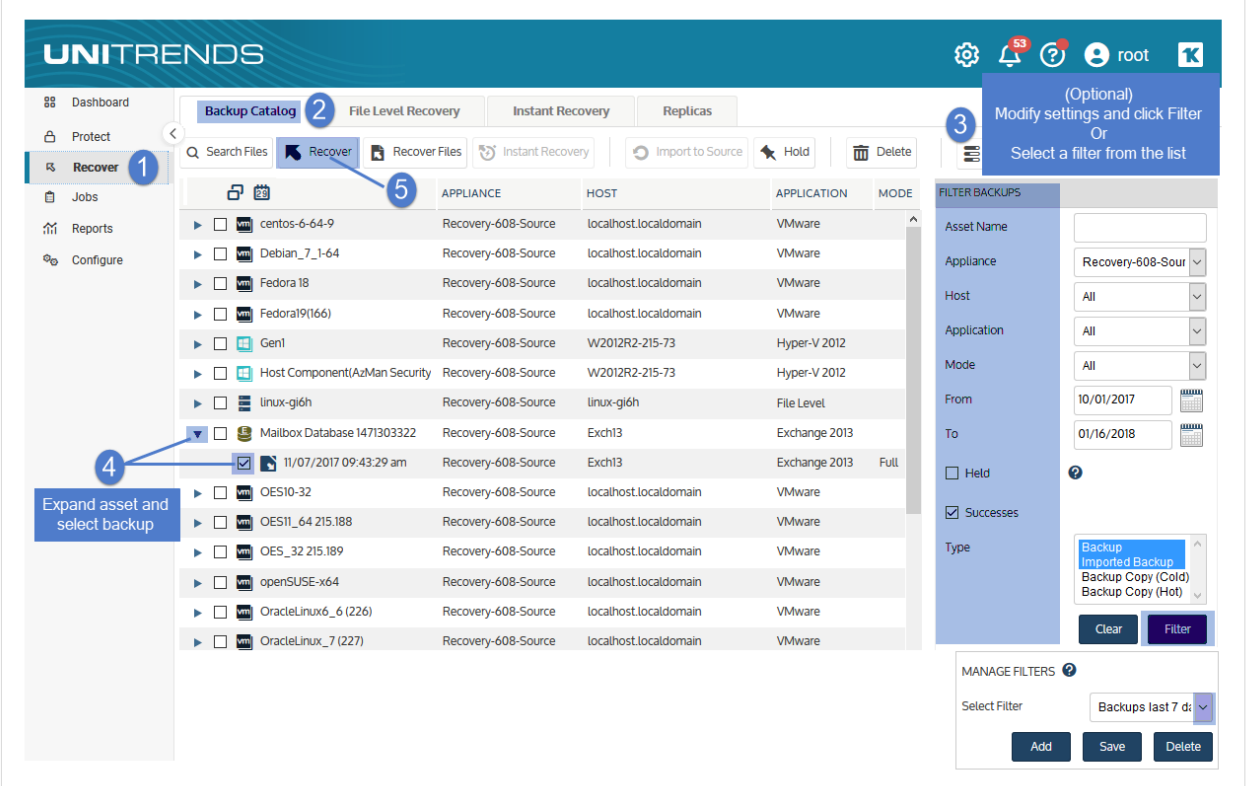

6 Select **Alternate Location** as the Recovery Target.

7 Select the asset where the backup will be recovered.

8 Enter the Directory Path where the backup will be recovered.

9 (Optional) Specify Advanced Options. To run client-side commands, enter any system command or user script in the **Commands to run pre-restore** and **Commands to run post-restore** fields.

10 Click **Next**.

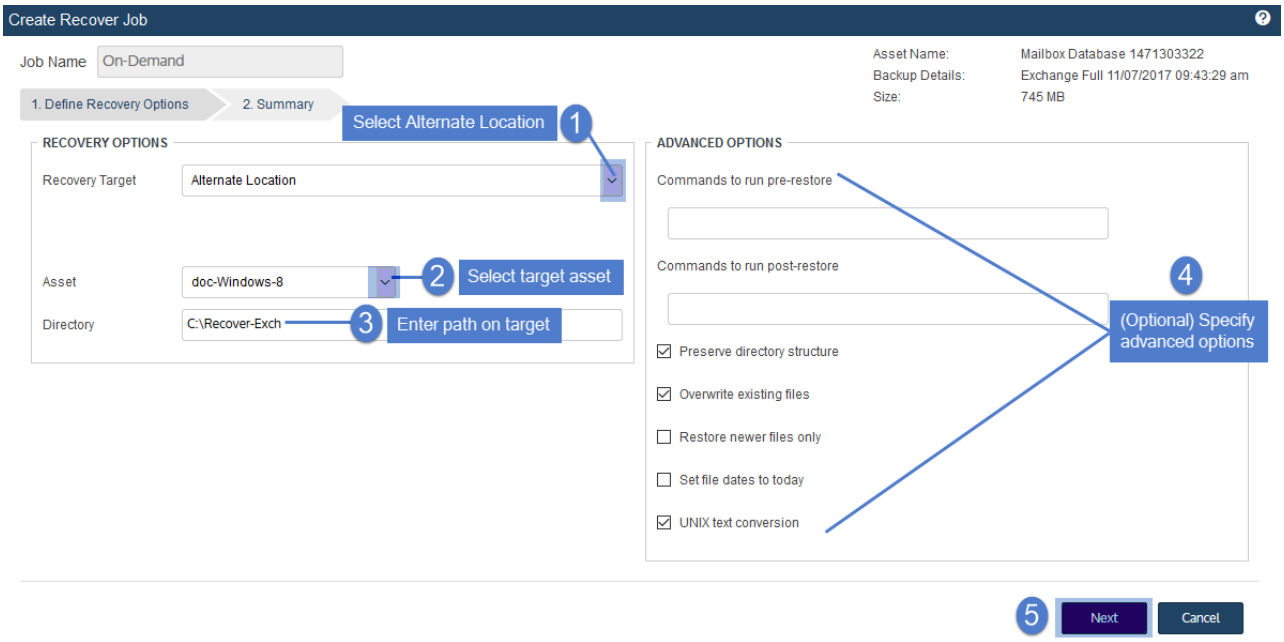

11 Review settings and click **Save**. Click **OK** to close the Information message.

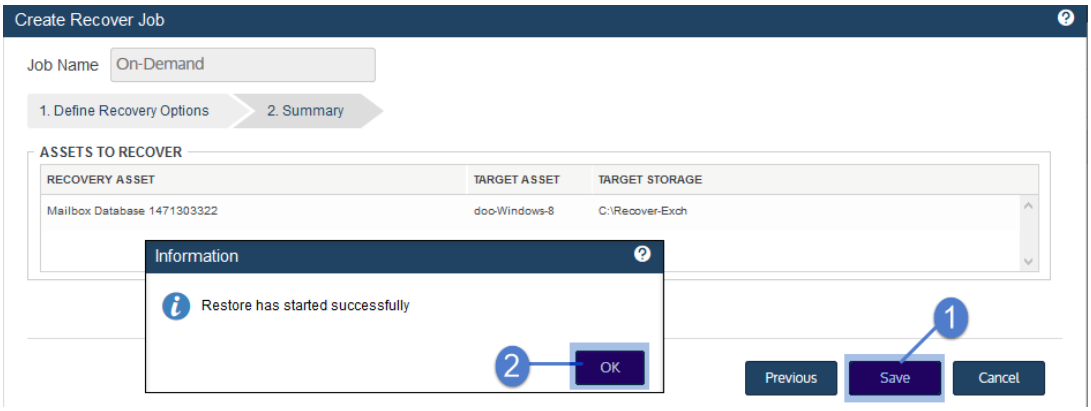

12 To monitor the recovery job:

- Select **Jobs > Active** Jobs.
- Select the recovery job in the list and click **View Details**.

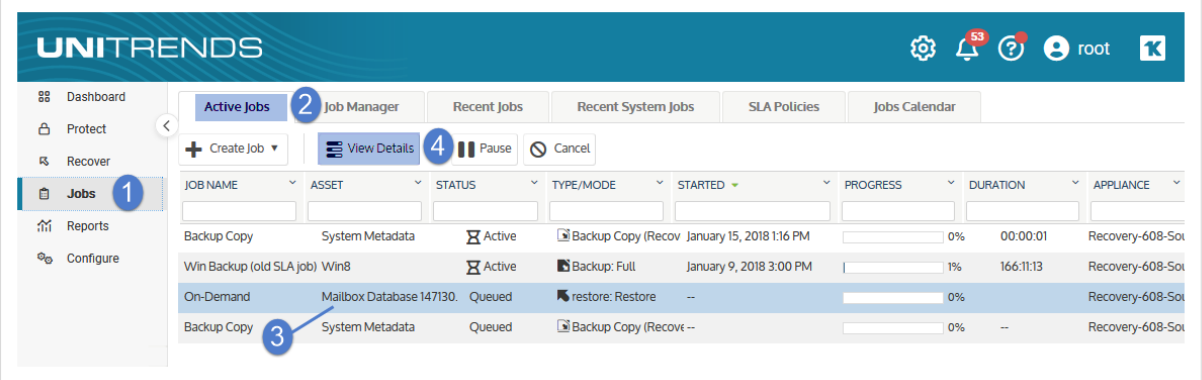

- The recovery is complete when the job's status changes to Success.

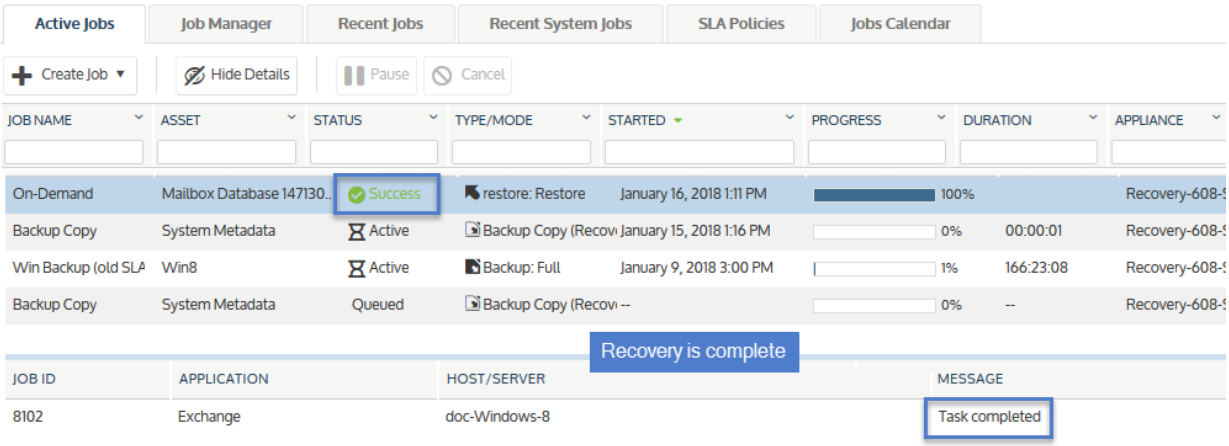

13 Access the recovered data on the recovery target.

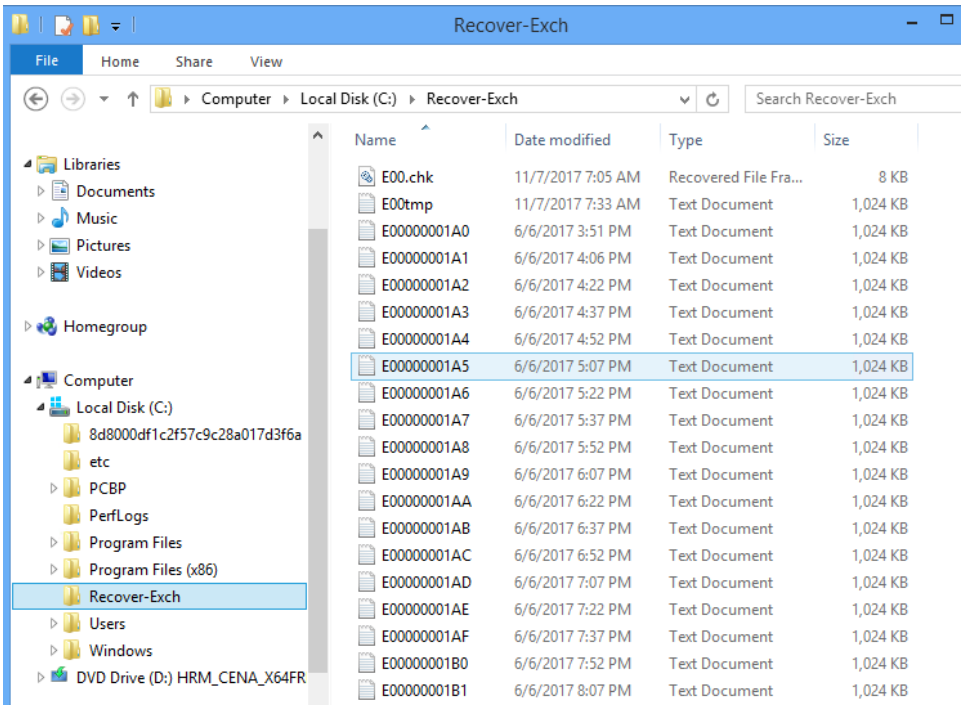

To recover a database or storage group to an alternate location by using the Backup Browser

Use this procedure to recover a backup to an alternate location.

Note: This procedure is not supported for imported backups or hot backup copies. To recover an imported backup or backup copy, use the Backup Catalog procedure above.

1. Verify that all prerequisites in "Recovering to an alternate location" on page 1156 have been met.
2. Log in to the backup appliance.
3. Select **Recover > Backup Catalog** and click **Backup Browser**.

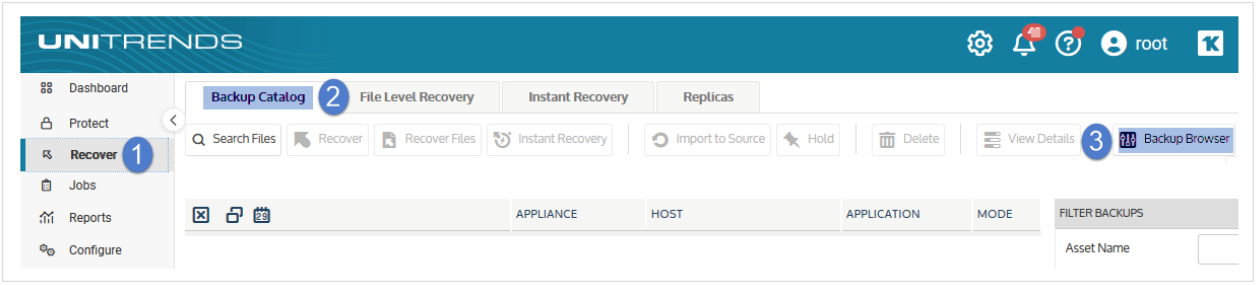

4. Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.
5. (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

6 Select the Exchange backup and click **Recover**.

7 Select **Alternate Location** as the Recovery Target.

8 Select the asset where the backup will be recovered.

9 Enter the Directory Path where the backup will be recovered.

10 (Optional) Specify Advanced Options. To run client-side commands, enter any system command or user script in the **Commands to run pre-restore** and **Commands to run post-restore** fields.

11 Click **Next**.

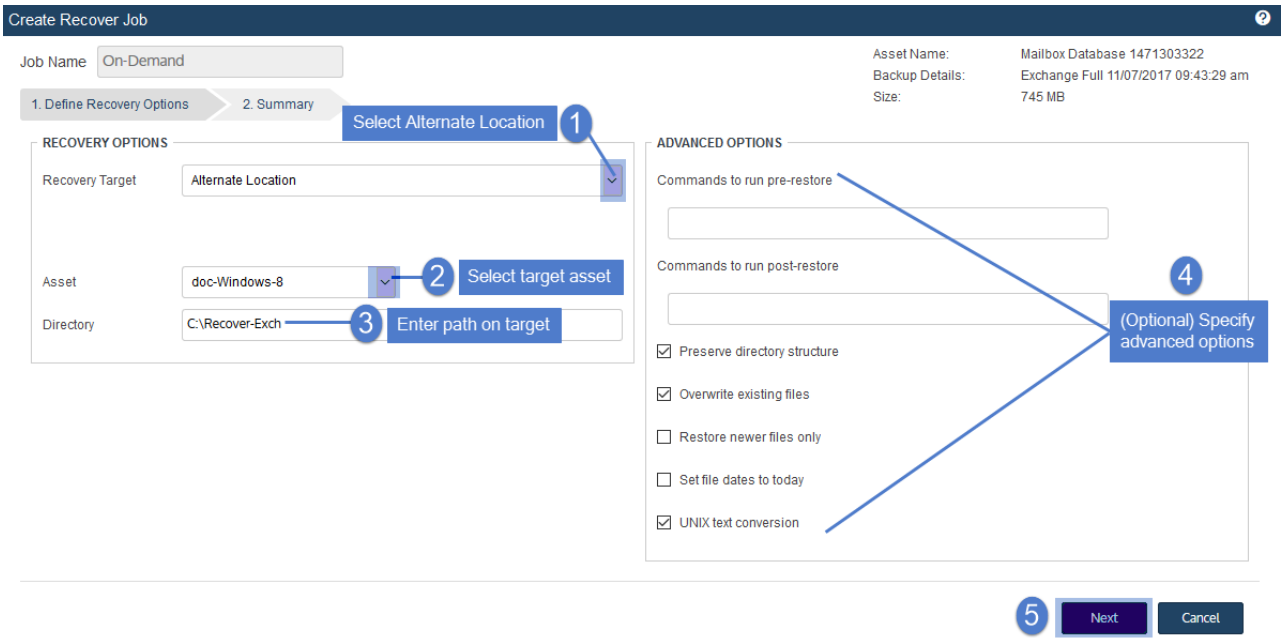

12 Review settings and click **Save**. Click **OK** to close the Information message.

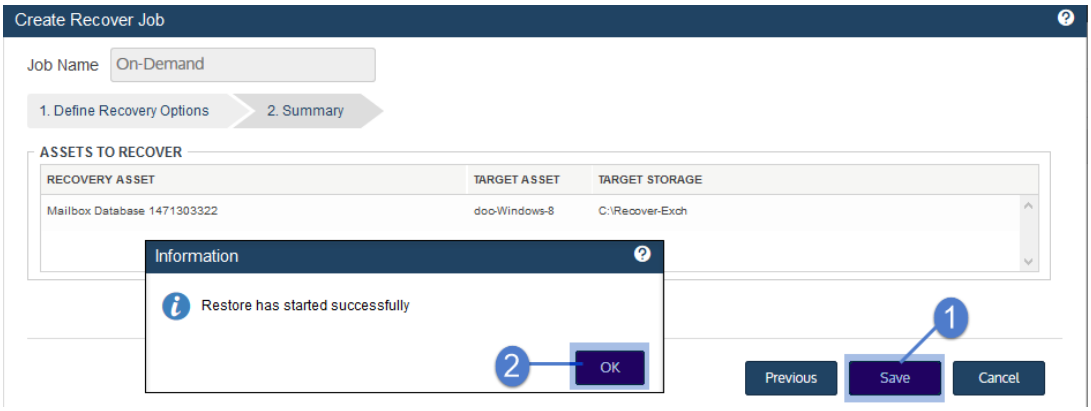

13 To monitor the recovery job:

- Select **Jobs > Active** Jobs.
- Select the recovery job in the list and click **View Details**.

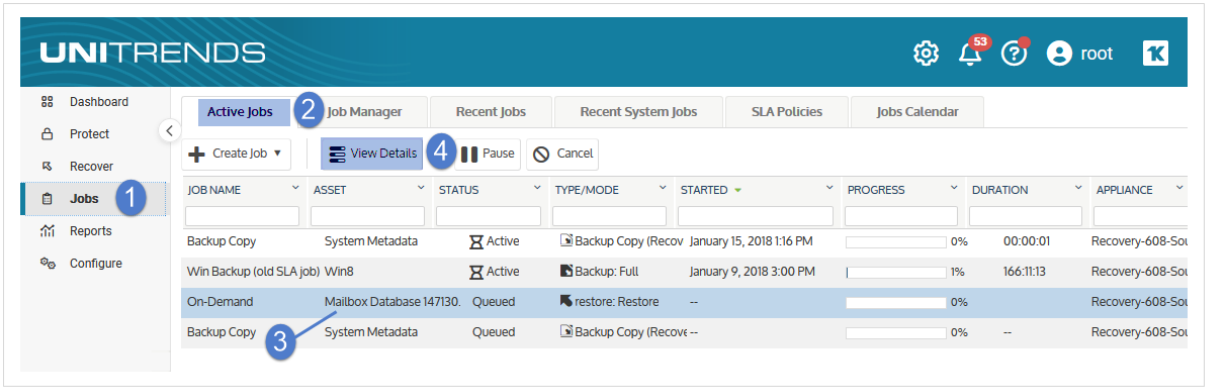

- The recovery is complete when the job's status changes to Success.

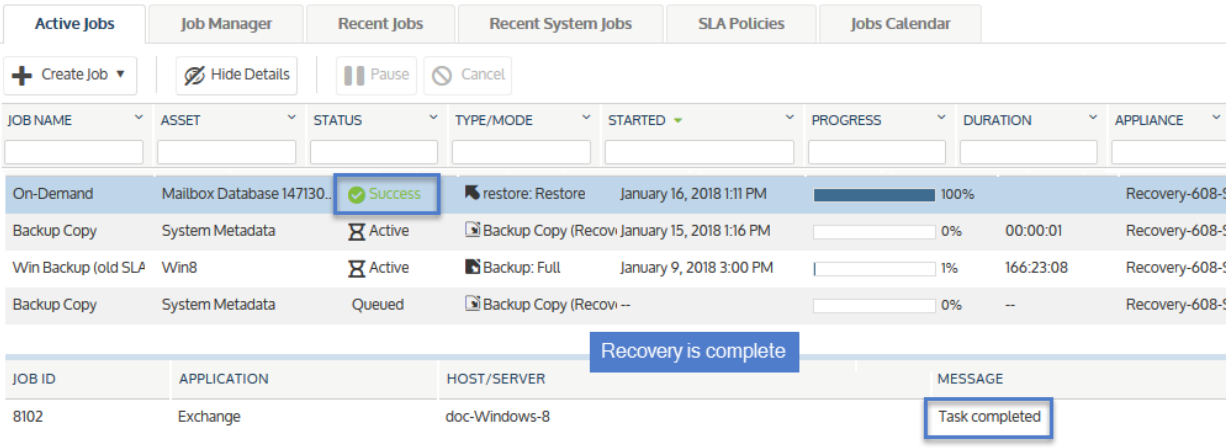

14 Access the recovered data on the recovery target.

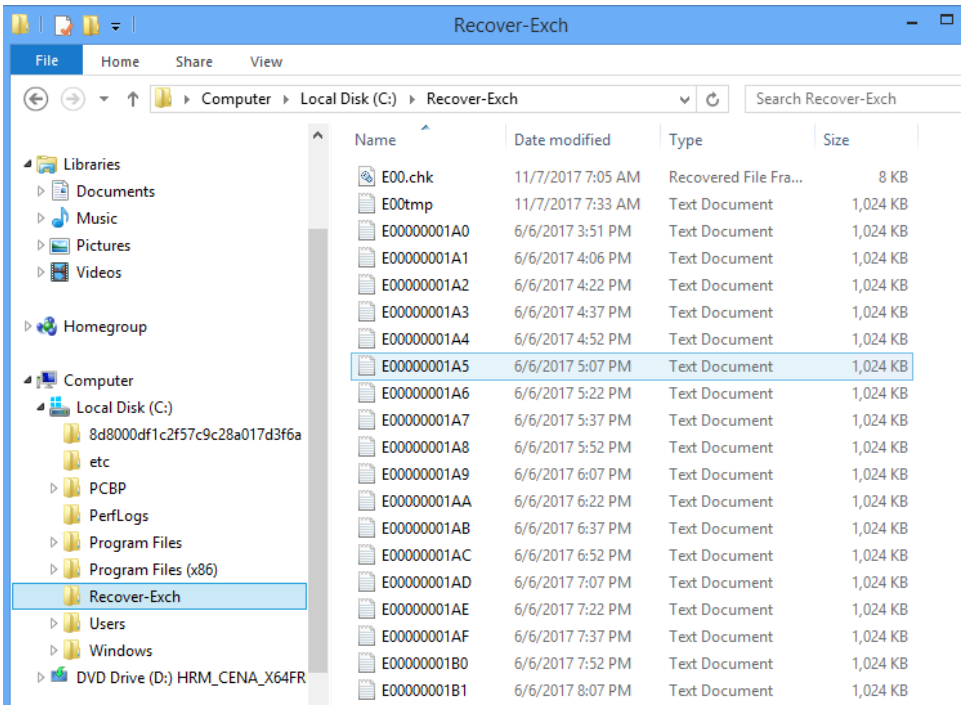

Recovering Exchange items

In addition to giving you the ability to recover an entire Exchange database or selected Exchange storage groups, Unitrends provides EQR (Exchange Quantum Recovery) which allows granular items, down to the individual mail item, to be recovered.

Unitrends offers and supports Ontrack PowerControls to recover individual items from an Exchange backup.

Note: Ontrack PowerControls can be used with 32-bit versions of Outlook only. 64-bit Outlook versions are not supported. For a complete overview of Ontrack PowerControls for Exchange, including procedures and limitations, see the Ontrack PowerControls User Guide for Exchange and the Ontrack PowerControls ExtractWizard Guide.

There are two fundamental ways that this tool may be used:

Recover from	Explanation
Directly from the Exchange backup	Use to perform all of the functions associated with Ontrack PowerControls directly from the Exchange backup without having to first perform the recovery of an Exchange backup.
A previously recovered Exchange backup	After an Exchange backup has been recovered (see "Recovering an Exchange database or storage group" on page 1148), Ontrack PowerControls or a third-party tool (e.g., Lucid8) can be used to search and recover individual Exchange items. Note that unlike Ontrack PowerControls, third-party tools are certified and supported by third parties and not by

Recover from	Explanation
	Unitrends.

Recovering Exchange items directly from a backup

Unitrends offers an optional feature that allows individual Exchange items to be recovered directly from the Exchange backup. This means that you can recover individual Exchange items without first having to perform the recovery of an Exchange backup. This option provides the fastest recovery time possible.

You can recover from a backup by using the Backup Catalog or the Backup Browser. To recover from an imported backup or hot backup copy, you must use the Backup Catalog. See these topics for details:

- "To recover individual Exchange items by using the Backup Catalog"
- "To recover individual Exchange items by using the Backup Browser" on page 1166

To recover individual Exchange items by using the Backup Catalog

Use this procedure to recover items directly from an Exchange backup. Run this procedure from the backup appliance to recover from a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover from a hot backup copy.

1. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the Exchange asset and select one of the following:
   - An Exchange backup.
   - An imported Exchange backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
   - A hot backup copy (supported if performing the recovery on the target appliance where the hot copy resides).
4. Click **Recover Files**.

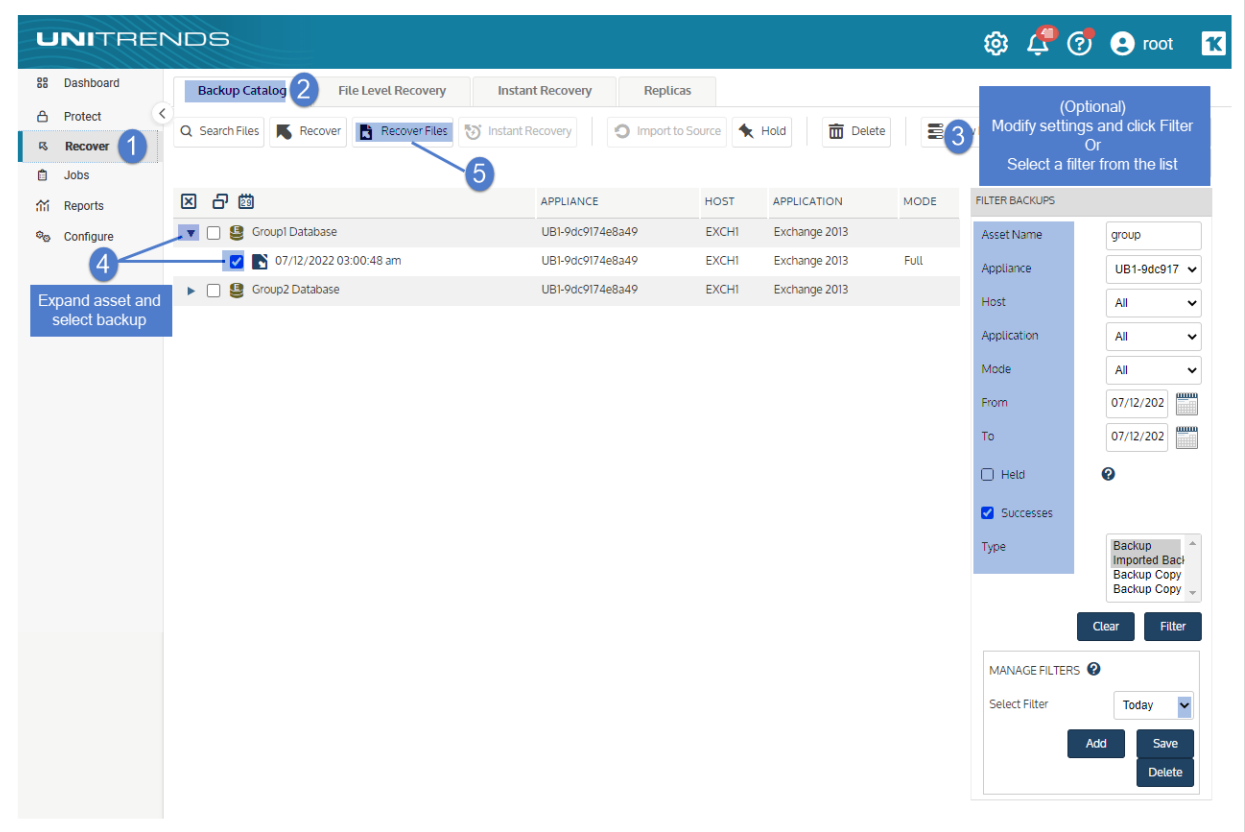

5 Click **Confirm** to create the recovery point.

> **Note:** Creating the recovery point object can take some time. If you go to Ontrack PowerControls and do not see any available items, check back later.

6 Use Ontrack PowerControls to recover Exchange items. See "Recovering items with Ontrack PowerControls for Exchange" on page 1168 for details.

7 Tear down the recovery object. For instructions, see "To view or tear down Exchange recovery objects" on page 1167.

To recover individual Exchange items by using the Backup Browser

Use this procedure to recover items directly from an Exchange backup.

> **Note:** This procedure is not supported for imported backups and backup copies. To recover from an imported backup or hot backup copy, use the Backup Catalog procedure above.

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

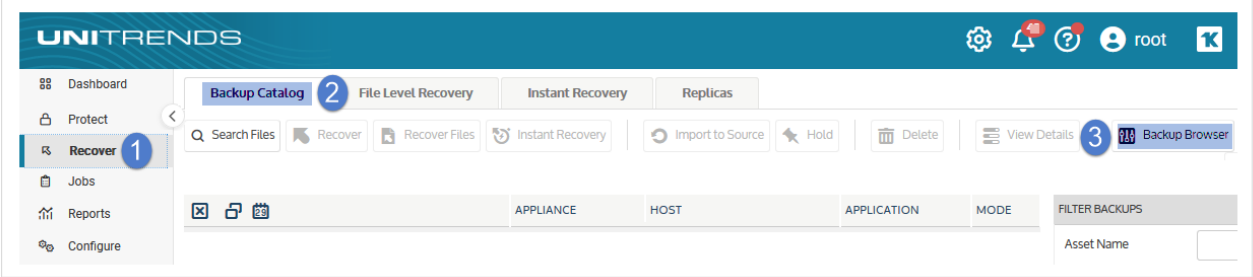

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Select the Exchange backup and click **Recover Files**.

6 Click **Confirm** to create the recovery point.

> **Note:** Creating the recovery point object can take some time. If you go to Ontrack PowerControls and do not see any available items, check back later.

7 Use Ontrack PowerControls to recover Exchange items. See "Recovering items with Ontrack PowerControls for Exchange" on page 1168 for details.

8 Tear down the recovery object. For instructions, see "To view or tear down Exchange recovery objects" on page 1167.

About the Exchange recovery session

After files have been recovered, the session remains until you tear it down. Because appliance resources are used to maintain the session, it is important to tear it down to ensure optimal performance.

To view or tear down Exchange recovery objects

1 If tearing down the object, disconnect from the network drive you mounted on the Ontrack PowerControls system.

2 Select **Recover** and click the **Backup Catalog** tab.

3 Select **File Level Recovery**.

4 Select the Exchange recovery object in the list.

5 Click **Remove** to tear down the object.

Recovering Exchange items from a previously recovered backup

If the option of recovering an individual Exchange item directly from the Exchange backup is not available, then the recovery may be performed from a previously recovered Exchange backup. After an Exchange backup has been recovered, Ontrack PowerControls or a third-party tool (e.g., Lucid8) may be used to search and recover individual Exchange items. Note that unlike Ontrack PowerControls, third-party tools are certified and supported by third parties and not Unitrends.

There are two classes of recovery targets in this situation: the recovery of the Exchange backup may be performed to the Unitrends system or the recovery of the Exchange backup may be performed to a customer's Windows system. The advantage to recovering the Exchange backup to the Unitrends appliance is that typically the recovery is faster because there is no network bandwidth overhead incurred.

Recovering items with Ontrack PowerControls for Exchange

After completing one of the procedures in "Recovering Exchange items" on page 1164, use the following procedure to recover individual items.

Note that to copy to a mailbox other than the one you logged in under, Full Mailbox Access must be set to Allow. For additional information, see *About Restoring Messages to a Microsoft Exchange Server* in the Ontrack® PowerControls™ User Guide.

To recover items using Ontrack PowerControls for Exchange

Note: Creating the recovery object can take some time. If you do not see any available items in Ontrack PowerControls, check back later.

1. Log in to your Windows machine with Ontrack PowerControls installed. Run Ontrack PowerControls for Exchange.
2. On the Welcome screen, click **Next**.
3. Next to the **Source File** field, click **Browse**. Browse to the *exchange_restore* share on your Unitrends appliance and double click the Exchange *.edb* file.
   If recovering from a full backup, the *.edb* file should be located in the *backup0* folder. If recovering from a differential or incremental backup, the *.edb* file should be in the *merged* folder.
4. Back on the Source Path Selection screen, click **Next**.
5. Choose to recover to a PST file or directly back to a live Exchange environment.
   If recovering back to a live Exchange environment, supply administrative credentials to any mailbox to which you want to recover.
6. Click **Next**.
   If creating a PST file, click **Next** and make a selection on the compatibility of the file.
7. To recover items, do one of the following:
   - To recover items to a PST file or live Exchange environment, navigate to the items you want to recover in the Source pane on top, select them, then drag and drop them to the node you want to recover them to in the Target pane on bottom.
   - To recover items to a network location, navigate to the items you want to recover in the Source pane on top, select them, right click and select **Export**, select a message format and recovery location, and click **Export**.

8 After recovering all the items you want to recover, close Ontrack PowerControls for Exchange.

9 Tear down the recovery object as described in "To view or tear down Exchange recovery objects" on page 1167.

Recovering SQL backups

Unitrends supports recovery of full, differential, and transaction log backups of SQL databases. See the following topics for details:

- "Requirements for recovering SQL backups" on page 1169
- "SQL recovery procedures" on page 1172

Requirements for recovering SQL backups

Requirements vary depending on what type of database and backup are being recovered. Review the following table to determine which requirements are applicable to your environment. Ensure these requirements have been met before you perform the "SQL recovery procedures" on page 1172.

Recovery type	Requirements and considerations
All recovery operations	The following apply to all SQL recovery operations. See the remainder of this table for additional requirements by backup mode and database type. • The entire database is recovered in a live state with each recovery operation. Unitrends does not support granular recovery of SQL database records. • A database must be recovered to a SQL instance that is the same version or later than that of the original SQL instance. Databases cannot be recovered to an older SQL version. • During recovery, you select a SQL asset and instance where the database will be recovered. Eligible targets that have been added to the appliance display in the UI. For SQL, the Recovery Target list includes all SQL server and/or SQL cluster assets that are hosting at least one eligible database instance. (If needed, you can add a SQL server or cluster asset before you recover the backup. For details, see "Start protecting non-clustered SQL environments" on page 754 or "Start protecting SQL clusters and availability groups" on page 749.) • Any existing database of the same name that resides on the recovery target is overwritten during the recovery operation (even if you specify a different path). To retain the original database, you must modify the recovery database name in the Create Recover Job dialog.
Full backup of a user database	The following apply to recovering full backups of SQL user databases: • A full backup can be recovered to the original location (the instance where the backup was taken) or to an alternate location.

Recovery type	Requirements and considerations
	• If applicable, see "Recovery to a clustered instance", "Availability groups", "Stretch databases", or "Always Encrypted databases" on page 1171 for additional requirements.
Differential or transaction log backup of a user database	The following apply to recovering SQL differential or transaction log backups: • A local backup can only be recovered to the original location (the instance where the backup was taken). • An imported backup copy can only be recovered as the original name and to the original location (the instance where the backup was taken). **Note:** Clustered instances and availability groups can move between cluster nodes. To ensure that the original location is available, it is best to verify that the cluster and all nodes have been added to the appliance before you run the recovery. (See "Start protecting SQL clusters and availability groups" on page 749 for details.) • During recovery, all previous backups in the group are also recovered. This means that when recovering a transaction backup, all previous transaction backups, the latest differential (if any), and the parent full are also recovered. Each backup is recovered as a separate job, and all jobs in the group are queued automatically. • Transaction log backups – It is highly recommended that you synchronize the date and time of the SQL server with that of the Unitrends appliance before you start the recovery. • If applicable, see "Recovery to a clustered instance", "Availability groups", "Stretch databases", or "Always Encrypted databases" on page 1171 for additional requirements.
System databases	The following apply to recovering system databases: • The master, model, and msdb system databases can only be recovered to their original SQL instances and names. The recovery job overwrites the existing database (even if you specify a different path). • The master, model and msdb databases must be recovered individually. • To recover the master database, you must first stop the SQL instance. See the Microsoft TechNet article How to Stop an Instance of SQL Server for details.
Recovery to a clustered	In a SQL cluster, a SQL Server failover cluster instance is installed into a Windows Server Failover Cluster (WSFC). The cluster's databases reside on shared storage that is

Recovery type	Requirements and considerations
instance	accessible to each server node. When recovering a SQL backup to a clustered instance, you must recover to a clustered storage resource that is accessible to all nodes in the cluster, and is a dependency of the SQL clustered instance you are restoring to. As long as you leave the Specify Path field empty (on the Create Recover Job dialog), the backup is recovered to the original clustered storage resource. If you enter a path, be sure to specify a clustered storage resource that is a dependency of the SQL clustered instance you are restoring to. (If you specify a local volume, the recovery fails with a SQL VSS Writer error.) **Note:** Because it is the instance that is clustered, and not the hosted database, a backup taken on a clustered instance is no different than one taken on a non-clustered instance. As with all user database backups, a full can be recovered to any eligible instance (which can be clustered or non-clustered), and a differential or log backup must be recovered to the original location.
Availability groups	These additional requirements apply to recovering availability group databases. Before you recover, ensure that these SQL prerequisites have been met: • The database must not be present in an availability group on the target server. If the database already exists, you must remove it from the availability group, then delete the mirrored copies from all secondary replicas. • A database of the same name that is in the *restoring* state must not be present on the target instance. If found, you must delete the database before running the recovery. • A differential or transaction log backup must be recovered to the original instance, and this instance must have an availability group that is configured with the same listener IP address as the backup's Availability Group asset. (For more on SQL assets, see "Protecting SQL clusters and availability groups" on page 748.) After recovery, you must manually reconfigure the mirrors and add the recovered database to the availability group. See Microsoft's documentation for details.
Stretch databases	A stretch database can only be recovered to its original SQL instance and name. The recovery job overwrites the existing database (even if you specify a different path). Only local SQL data is included in the backups. After recovering the local SQL data, you must reconnect the local recovered database to the remote Azure database to reconcile the recovered data. See "To recover a Stretch database backup by using the Backup Catalog" on page 1182 for details.
Always Encrypted databases	The SQL Column Master Keys (CMKs) must be available on the recovery target so you can access the recovered data. The keys are stored in a certificate on the SQL server. If the keys are not available on the recovery target, you must install them after you recover the backup. In most environments:

Recovery type	Requirements and considerations
	• You will not need to install CMKs if recovering to the original database or to another database on the original instance. • You will need to install CMKs if recovering to a different SQL server. • You may need to install CMKs if recovering to a different instance on the original server. • See Microsoft's documentation for instructions on installing the CMKs.

SQL recovery procedures

You can recover from a backup by using the Backup Catalog or the Backup Browser. To recover from an imported backup or hot backup copy, you must use the Backup Catalog. After you have reviewed the "Requirements for recovering SQL backups" on page 1169, use the following procedures to recover a SQL backup or imported backup copy:

- "To recover one SQL full, differential, or transaction log backup by using the Backup Catalog"
- "To recover one SQL full, differential, or transaction log backup by using the Backup Browser" on page 1176
- "To recover multiple SQL full, differential, or transaction log backups by using the Backup Catalog" on page 1181
- "To recover a Stretch database backup by using the Backup Catalog" on page 1182
- "To recover a Stretch database backup by using the Backup Browser" on page 1185

To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.

To recover one SQL full, differential, or transaction log backup by using the Backup Catalog

Use this procedure to recover one backup to the original location or to an alternate location. Run this procedure from the backup appliance to recover an entire backup or imported backup copy. Run this procedure from the backup copy target appliance to recover an entire hot backup copy.

1. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
2. Select **Recover**, then click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

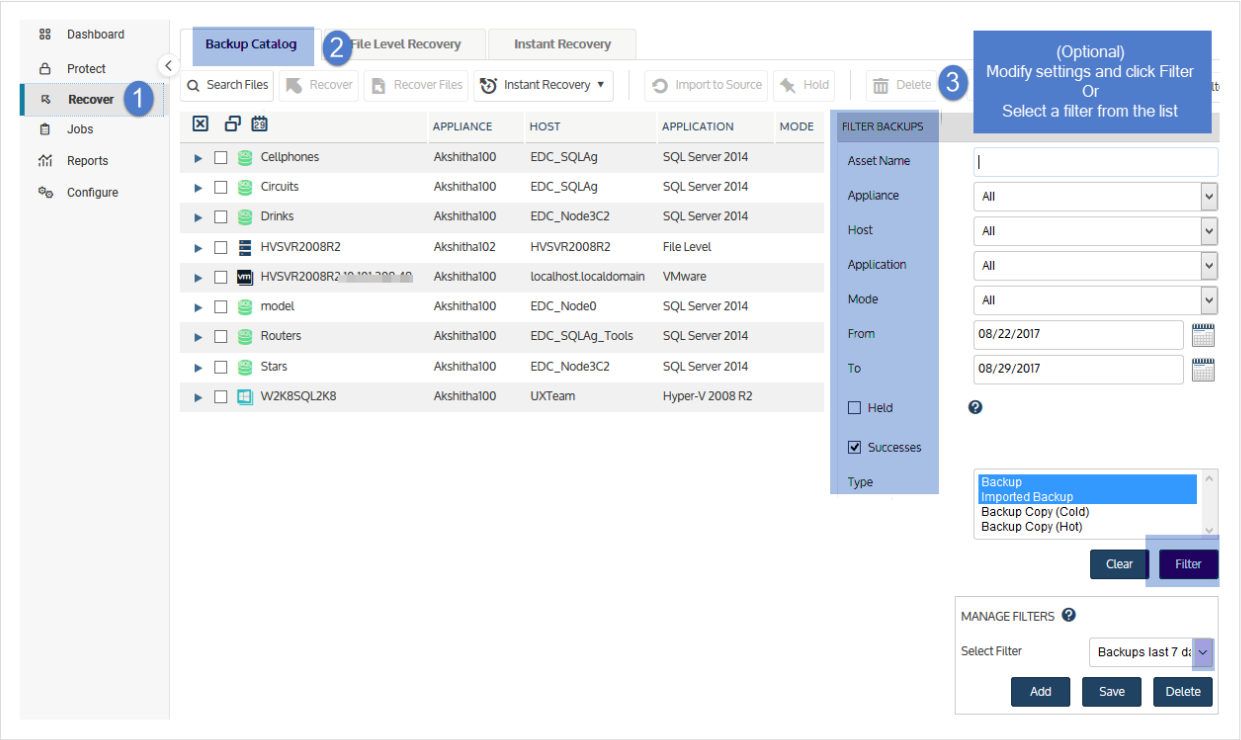

3 Expand the SQL database asset and select one of the following to use for the recovery:

- A SQL backup.
- An imported SQL backup copy. (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
- A hot backup copy (supported if performing the recovery on the target appliance where the hot copy resides).

4 Click **Recover**.

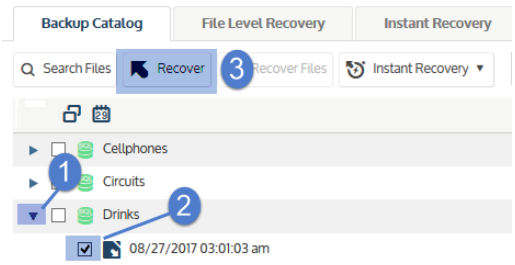

5 Enter Recovery Options and Advanced Options, described in the following table:

Recovery Options	Description
Recovery Target	Select the asset where the database will be recovered. **Notes:** • If recovering from a local backup, only eligible targets that have been added to the appliance display in the Recovery Target list. For SQL, the list includes all SQL server and/or SQL cluster assets that are hosting at least one eligible database instance. • Eligible recovery targets vary by backup and database type. If you don't see the desired target, it may not be supported for the backup you selected. For details, see "Requirements for recovering SQL backups" on page 1169. • If the Recovery Target list is empty, the asset where the backup was taken needs to be added to the backup appliance. Check the upper-right corner of the Create Recover Job dialog to identify the original instance or availability group, then add its host asset to the appliance. • Clustered instances and availability groups can move between cluster nodes. To ensure that the original location is available, it is best to verify that the cluster and all nodes have been added to the appliance before you run the recovery. (See "Start protecting SQL clusters and availability groups" on page 749 for details.) • Differential and transaction log backups must be recovered to the original location where the backup was taken. If recovering from an *imported backup* or from a *hot backup copy*, additional targets may display in the Recovery Target list because the appliance cannot determine the original location. Use care to select the original location. If you select any target other than the original, the recovery fails.
Recovery Instance	Select the SQL instance where the database will be recovered.
Database	Enter a name for the recovered database. • To create a new database during recovery, be sure to enter a unique database name (one that does not yet exist on this recovery target). Any existing database of the same name is overwritten during recovery. • If the original database resides on the recovery target, you must modify the database name to retain the original database (regardless of whether you specify a new path).

Recovery Options	Description
	Note: For system databases, the database name cannot be changed. The existing database is overwritten during recovery.
Specify Path	Use this field to specify a path on the recovery target. You can enter a path manually or by browsing the target. (Check **Specify Path** to enable the **Browse** button.) A path is required in some cases and optional in others: • If a path is required, you are prompted to enter one before you can start the recovery. • If you recover to the original instance and leave the Specify Path field empty, files are recovered to their original locations. • If you recover to an alternate non-clustered instance and leave the Specify Path field empty, files are recovered to *<Vol>:\UnitrendsRestore*, where *<Vol>* is the volume with the most free space. • If you recover to an alternate clustered instance and leave the Specify Path field empty, files are recovered to one of the following: – *<Vol1>:\<ClusterStorage>\<Volume1>:\UnitrendsRestore*, where *<ClusterStorage>* is the dependent cluster resource with the most free space. – *<Vol2>:\UnitrendsRestore*, where <Vol2> is mounted to a cluster storage resource and is the non-dependent cluster storage resource with the most free space. (This option is used only if no dependent cluster storage resources are defined.) • To specify a path to an alternate clustered instance, you must supply a path on a clustered storage resource that is accessible to all nodes in the cluster. (If you specify a local volume, the recovery fails with a SQL VSS Writer error.)
Point-in-time Recovery	(Optional) This option is available for transaction log backups only. For additional recovery points, check the **Point-in-time Recovery** box and select the desired recovery point by moving the Earliest/Latest slider.
Commands to run pre-restore	(Optional) Enter commands to be run before the recovery.
Commands to run post-restore	(Optional) Enter commands to be run after the recovery.

6 Click **Next**.

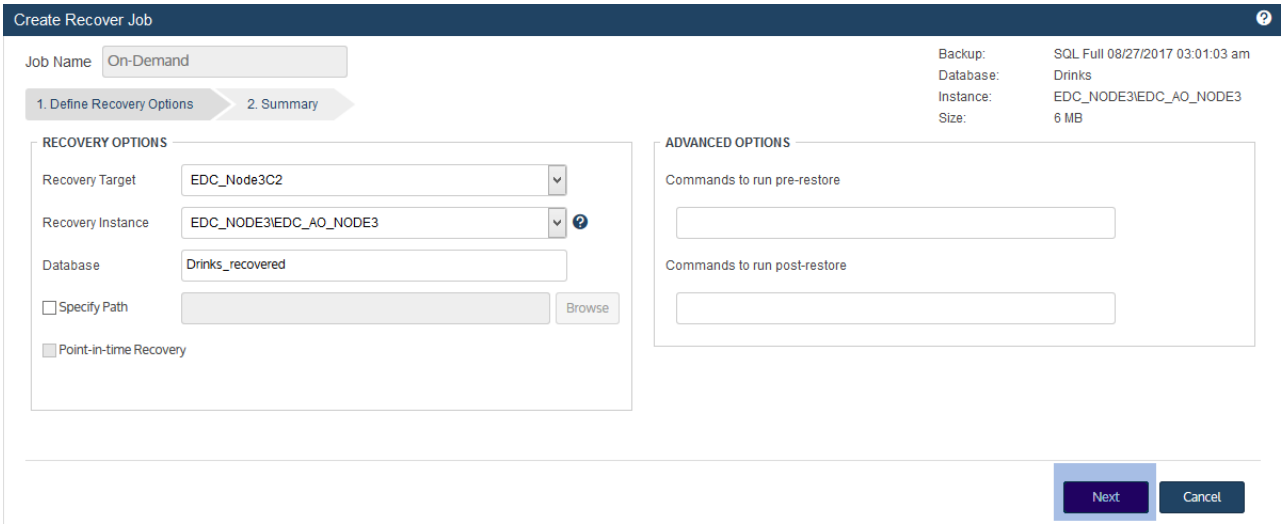

7 Read the *I understand...* message, then check the box to indicate that you understand any database of the same name will be overwritten by the recovery. Click **Save** to start the recovery.

Note: If you recovered an availability group database, you must manually reconfigure the mirrors and add the recovered database to the availability group. See Microsoft's documentation for details.

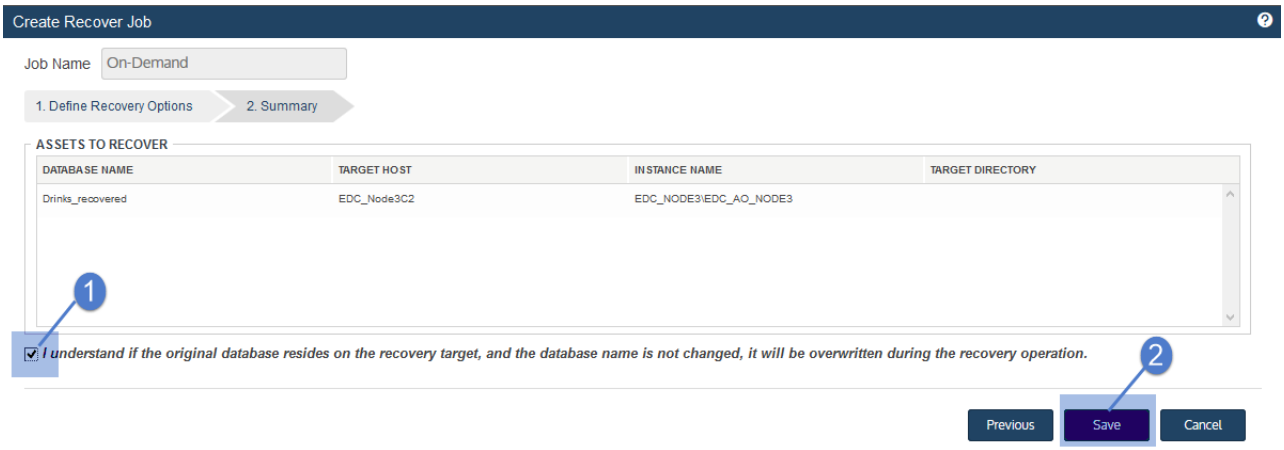

The recovery job is queued immediately. To view the running job, select **Jobs > Active Job**.

To recover one SQL full, differential, or transaction log backup by using the Backup Browser

Use this procedure to recover one backup to the original location or to an alternate location.

Note: This procedure is not supported for imported backups and backup copies. To recover an imported backup or hot backup copy, use the Backup Catalog procedure above.

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog** and click **Backup Browser**.

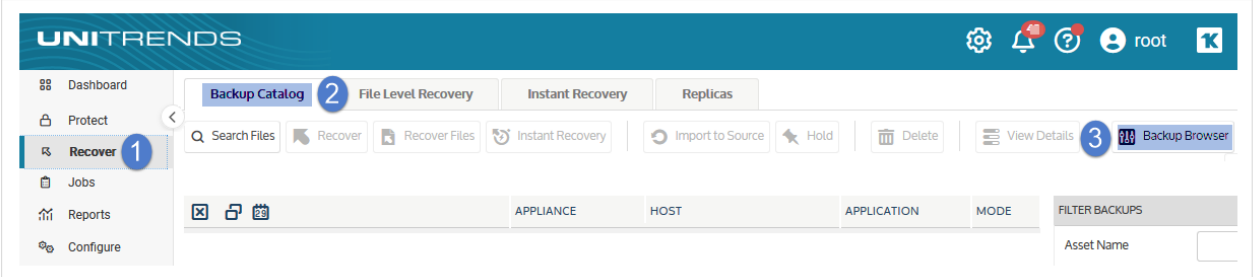

3. Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.
4. (Optional) Refine the search:
   - Enter text in any column field to filter the display.
   - Click an arrow to sort by column.
   - Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
   - For a description of each column, see "Backup Browser column descriptions".
5. Select the SQL backup and click **Recover**.

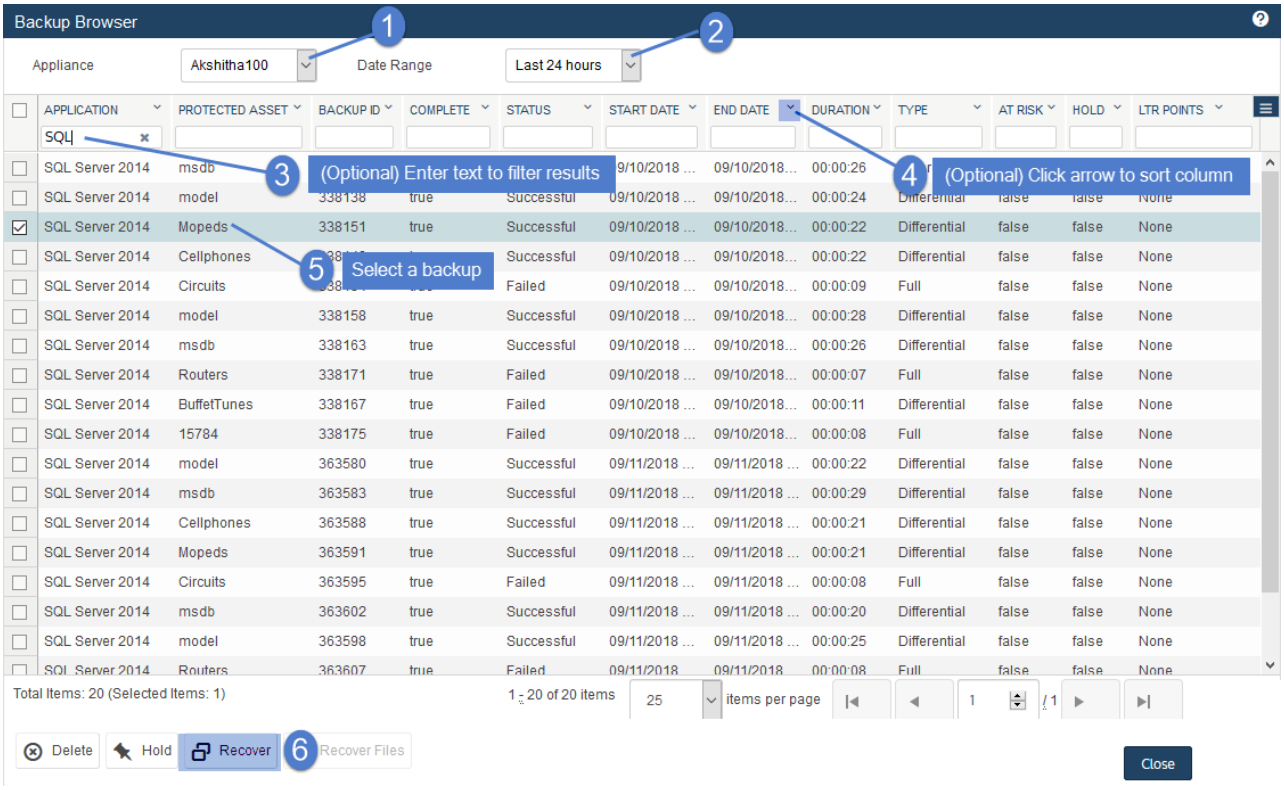

6 Enter Recovery Options and Advanced Options, described in the following table:

Recovery Options	Description
Recovery Target	Select the asset where the database will be recovered. **Notes:** • If recovering from a local backup, only eligible targets that have been added to the appliance display in the Recovery Target list. For SQL, the list includes all SQL server and/or SQL cluster assets that are hosting at least one eligible database instance. • Eligible recovery targets vary by backup and database type. If you don't see the desired target, it may not be supported for the backup you selected. For details, see "Requirements for recovering SQL backups" on page 1169. • If the Recovery Target list is empty, the asset where the backup was taken needs to be added to the backup appliance. Check the upper-right corner of the Create Recover Job dialog to identify the original instance or availability group, then add its host asset to the appliance. • Clustered instances and availability groups can move between cluster nodes. To ensure that the original location is available, it is best to verify that the cluster and all nodes have been added to the appliance before you run the recovery. (See "Start protecting SQL clusters and availability groups" on page 749 for details.) • Differential and transaction log backups must be recovered to the original location where the backup was taken. If recovering from an *imported backup* or from a *hot backup copy*, additional targets may display in the Recovery Target list because the appliance cannot determine the original location. Use care to select the original location. If you select any target other than the original, the recovery fails.
Recovery Instance	Select the SQL instance where the database will be recovered.
Database	Enter a name for the recovered database. • To create a new database during recovery, be sure to enter a unique database name (one that does not yet exist on this recovery target). Any existing database of the same name is overwritten during recovery. • If the original database resides on the recovery target, you must modify the database name to retain the original database (regardless of whether you

Recovery Options	Description
	specify a new path). **Note:** For system databases, the database name cannot be changed. The existing database is overwritten during recovery.
Specify Path	Use this field to specify a path on the recovery target. You can enter a path manually or by browsing the target. (Check **Specify Path** to enable the **Browse** button.) A path is required in some cases and optional in others: • If a path is required, you are prompted to enter one before you can start the recovery. • If you recover to the original instance and leave the Specify Path field empty, files are recovered to their original locations. • If you recover to an alternate non-clustered instance and leave the Specify Path field empty, files are recovered to *<Vol>:\UnitrendsRestore*, where *<Vol>* is the volume with the most free space. • If you recover to an alternate clustered instance and leave the Specify Path field empty, files are recovered to one of the following: – *<Vol1>:\<ClusterStorage>\<Volume1>:\UnitrendsRestore*, where *<ClusterStorage>* is the dependent cluster resource with the most free space. – *<Vol2>:\UnitrendsRestore*, where <Vol2> is mounted to a cluster storage resource and is the non-dependent cluster storage resource with the most free space. (This option is used only if no dependent cluster storage resources are defined.) • To specify a path to an alternate clustered instance, you must supply a path on a clustered storage resource that is accessible to all nodes in the cluster. (If you specify a local volume, the recovery fails with a SQL VSS Writer error.)
Point-in-time Recovery	(Optional) This option is available for transaction log backups only. For additional recovery points, check the **Point-in-time Recovery** box and select the desired recovery point by moving the Earliest/Latest slider.
Commands to run pre-restore	(Optional) Enter commands to be run before the recovery.
Commands to	(Optional) Enter commands to be run after the recovery.

Recovery Options	Description
run post-restore	

7 Click **Next**.

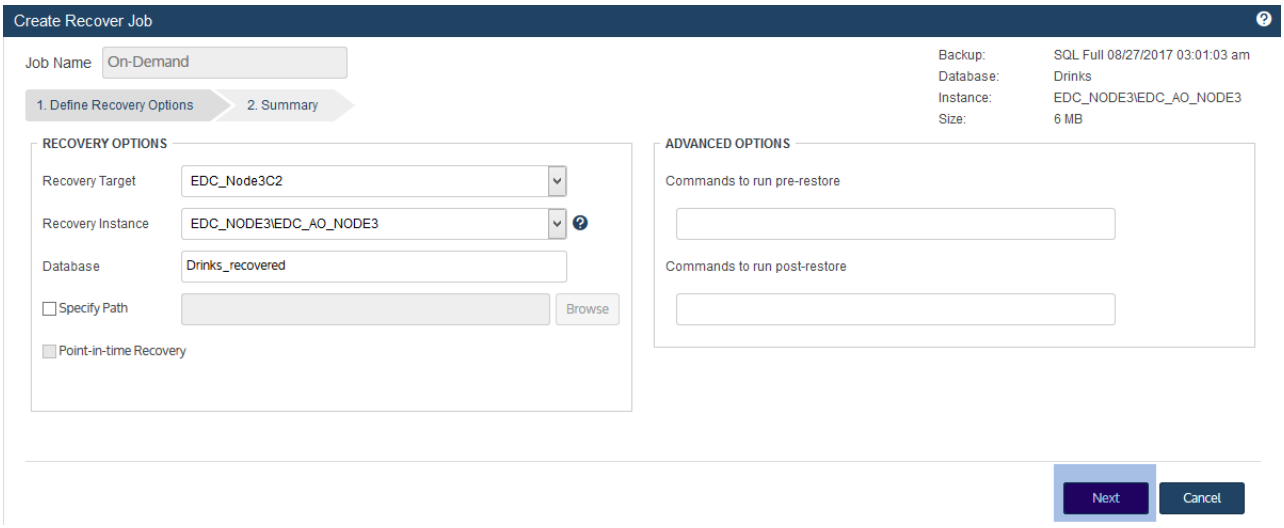

8 Read the *I understand...* message, then check the box to indicate that you understand any database of the same name will be overwritten by the recovery. Click **Save** to start the recovery.

Note: If you recovered an availability group database, you must manually reconfigure the mirrors and add the recovered database to the availability group. See Microsoft's documentation for details.

The recovery job is queued immediately. To view the running job, select **Jobs > Active Job**.

To recover multiple SQL full, differential, or transaction log backups by using the Backup Catalog

Use this procedure to recover multiple user databases to their original locations with their original names. Any existing data will be overwritten. (To recover to a different location, or to recover a system or availability group database, see "To recover one SQL full, differential, or transaction log backup by using the Backup Catalog" or "To recover one SQL full, differential, or transaction log backup by using the Backup Browser" on page 1176.)

Run this procedure from the backup appliance to recover backups and/or imported backup copies. Run this procedure from the backup copy target appliance to recover hot backup copies.

1. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
2. Select **Recover**, then click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

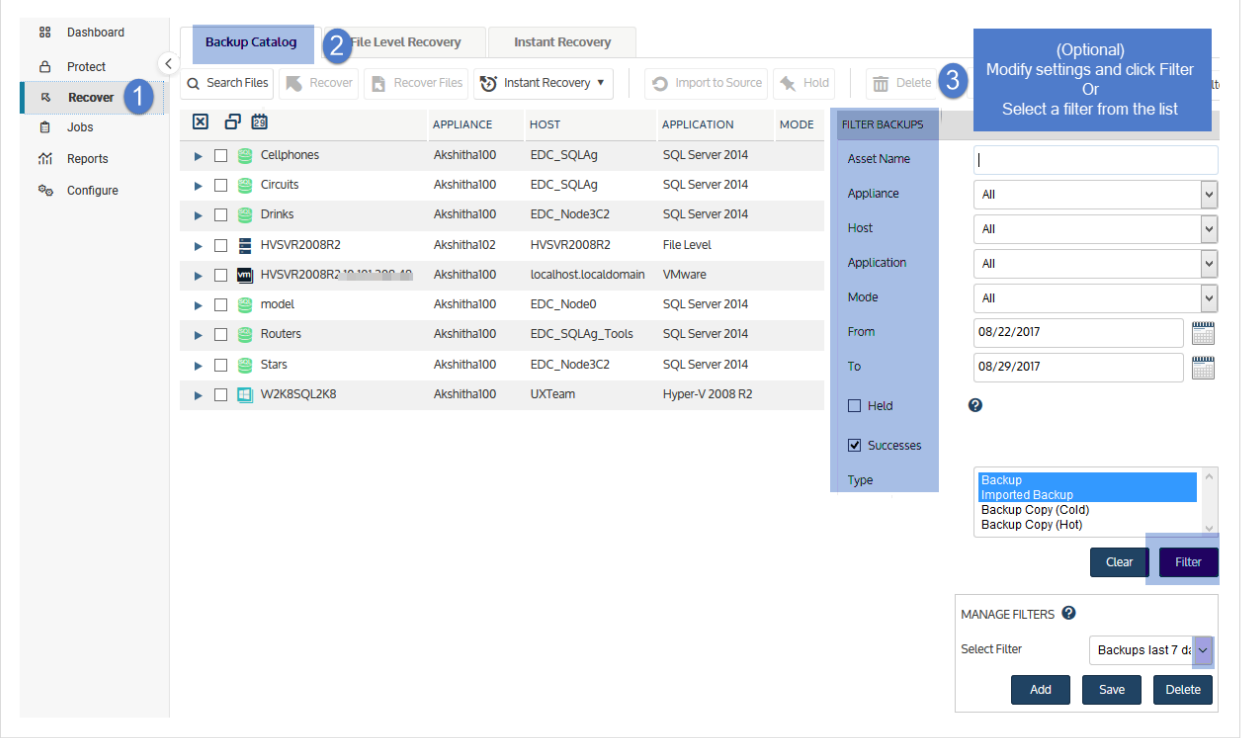

3. Expand the SQL database asset, then select the backup, imported backup, or hot backup copy to recover.
4. Repeat step 3 to select additional backups.
5. Click **Recover**.

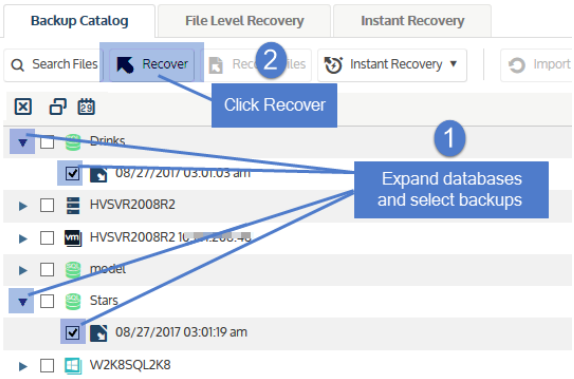

6 Click **Save** to start the recovery.

- Each backup you select is recovered as a separate job.
- Databases are recovered to their original locations with their original names. Any existing data is overwritten.

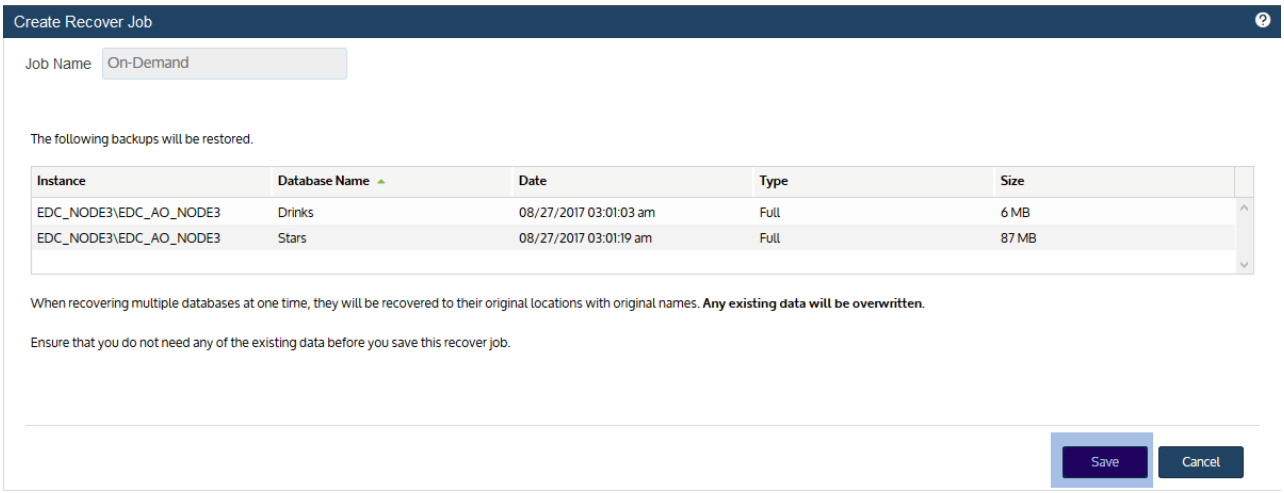

The recovery job is queued immediately. To view the running job, select **Jobs > Active Job**.

To recover a Stretch database backup by using the Backup Catalog

A stretch database can only be recovered to its original SQL instance and name. The recovery job overwrites the existing database.

Because only local SQL data is included in the backup, you must reconnect the local recovered database to the remote Azure database to reconcile the recovered local data with data that has been migrated to Azure.

Run this procedure from the backup appliance to recover a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover a hot backup copy.

1 Log in to the backup appliance or target appliance (if recovering from a hot backup copy).

2 Select **Recover**, then click the **Backup Catalog** tab.

(Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

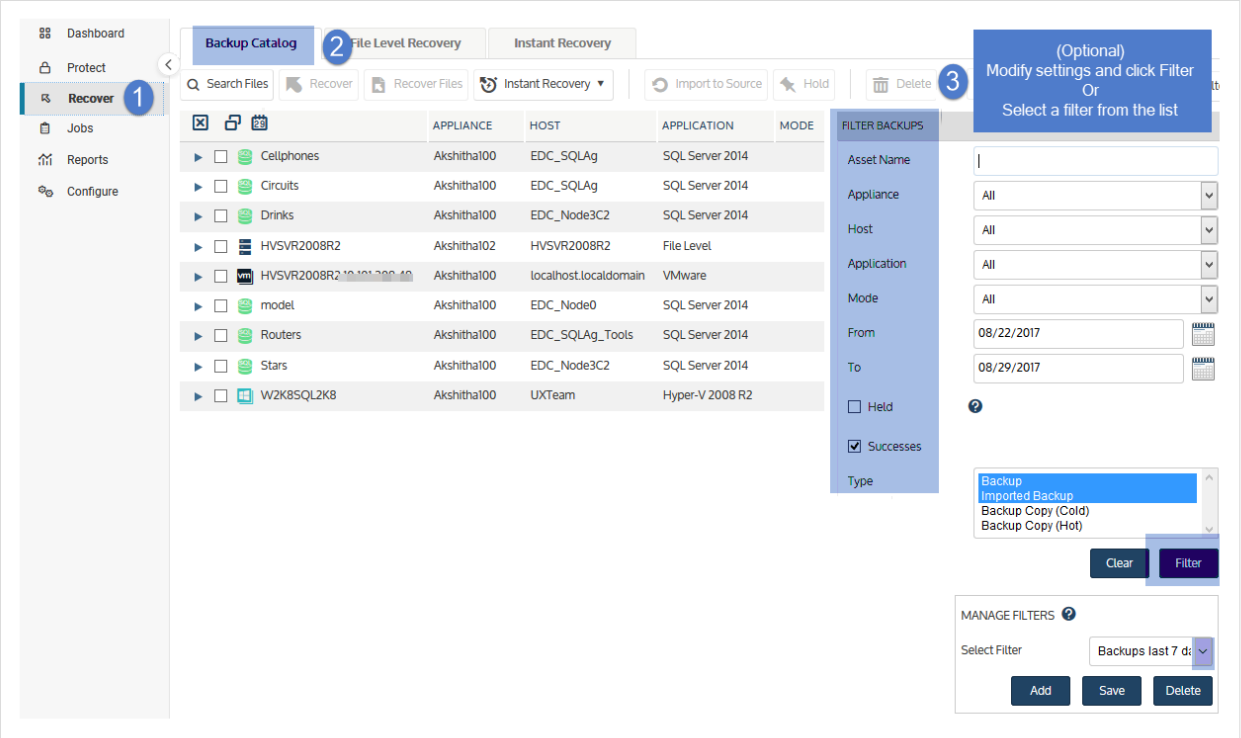

3 Expand the SQL database asset and select the backup, imported backup, or hot backup copy to recover.

4 Click **Recover**.

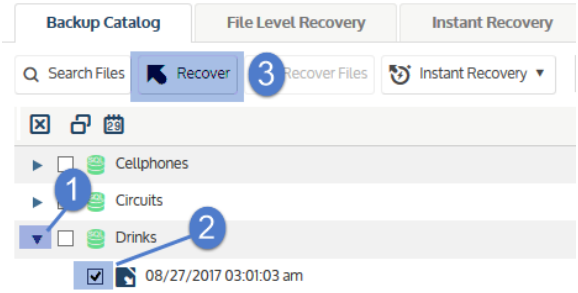

5 Select these Recovery Options and Advanced Options:

Recovery Options	Description
Recovery Target	Select the original SQL server asset (where the backup was taken).

Recovery Options	Description
Recovery Instance	Select the original SQL instance (where the backup was taken).
Database	Enter the name of the original database. Note that existing database files are overwritten during the recovery.
Specify Path	Leave this field empty. Files will be recovered to their original locations.
Commands to run pre-restore	(Optional) Enter commands to be run before the recovery runs.
Commands to run post-restore	(Optional) Enter commands to be run after the recovery runs.

6 Click **Next**.

7 Read the *I understand...* message, then check the box to indicate that you understand any database of the same name will be overwritten by the recovery. Click **Save** to start the recovery.

- The recovery job is queued immediately. To view the running job, select **Jobs > Active Job**.
- The backup is recovered to the original location, creating a new SQL server database.

8 Follow the instructions in Microsoft's Backup and restore Stretch-enabled databases article to connect the newly recovered database to Azure and reconcile the local recovered data with the Azure database. Note the following:

- For credentials, you will need to supply the Azure SQL database FQDN credentials that were created when the original SQL database was stretched. (Do not use the Azure SQL Server administrator credentials.)
- Stretch database credentials are stored on the SQL server and are encrypted with a SQL Master Key. If credentials have been lost, you must recreate them manually using the master key before you can connect the recovered database to the remote Azure instance.

- Connecting the newly recovered database causes a new copy of the remote Azure database to be created.

To recover a Stretch database backup by using the Backup Browser

A stretch database can only be recovered to its original SQL instance and name. The recovery job overwrites the existing database.

Because only local SQL data is included in the backup, you must reconnect the local recovered database to the remote Azure database to reconcile the recovered local data with data that has been migrated to Azure.

Note: This procedure is not supported for imported backups and backup copies. To recover an imported backup or hot backup copy, use the Backup Catalog procedure above.

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog** and click **Backup Browser**.

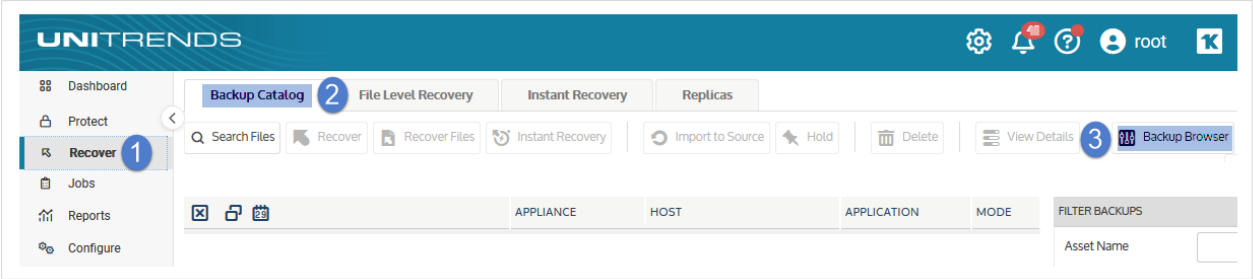

3. Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.
4. (Optional) Refine the search:
   - Enter text in any column field to filter the display.
   - Click an arrow to sort by column.
   - Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
   - For a description of each column, see "Backup Browser column descriptions".
5. Select the SQL backup and click **Recover**.
6. Select these Recovery Options and Advanced Options:

Recovery Options	Description
Recovery Target	Select the original SQL server asset (where the backup was taken).
Recovery Instance	Select the original SQL instance (where the backup was taken).
Database	Enter the name of the original database. Note that existing database files are overwritten during the recovery.

Recovery Options	Description
Specify Path	Leave this field empty. Files will be recovered to their original locations.
Commands to run pre-restore	(Optional) Enter commands to be run before the recovery runs.
Commands to run post-restore	(Optional) Enter commands to be run after the recovery runs.

7 Click **Next**.

8 Read the *I understand...* message, then check the box to indicate that you understand any database of the same name will be overwritten by the recovery. Click **Save** to start the recovery.

- The recovery job is queued immediately. To view the running job, select **Jobs > Active Job**.
- The backup is recovered to the original location, creating a new SQL server database.

9 Follow the instructions in Microsoft's Backup and restore Stretch-enabled databases article to connect the newly recovered database to Azure and reconcile the local recovered data with the Azure database. Note the following:

- For credentials, you will need to supply the Azure SQL database FQDN credentials that were created when the original SQL database was stretched. (Do not use the Azure SQL Server administrator credentials.)
- Stretch database credentials are stored on the SQL server and are encrypted with a SQL Master Key. If credentials have been lost, you must recreate them manually using the master key before you can connect the recovered database to the remote Azure instance.
- Connecting the newly recovered database causes a new copy of the remote Azure database to be created.

Recovering SharePoint backups

The SharePoint agent supports recovery of full and differential backups. With SharePoint backups, the agent leverages STSADM or PowerShell (SharePoint 2013 and higher) to perform recovery operations. Recoveries occur in two phases. In the first phase, backups are unfolded to a local share or backup appliance. In the second phase, the agent invokes STSADM or PowerShell commands to recover to the SharePoint assets.

See the following topics for details:

- "SharePoint recovery considerations" on page 1187
- "About the SharePoint recovery items session" on page 1187
- "SharePoint recovery procedures" on page 1188
- "Recovering items with Ontrack PowerControls" on page 1192

SharePoint recovery considerations

Consider the following when recovering SharePoint environments:

- Free space equivalent to twice the size of the backup is required on the local share for recovery processing. If adequate space is not available, the recovery fails.
- Recoveries are to the original farm only.
- Full catastrophic farm recovery can only be performed for SharePoint 2013 and 2010 deployments where the installation type is *single server*. For *full farm* (all SharePoint releases) or *single server farm* (SharePoint 2016) installations, you must recover items instead. To check your installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 757.
- Granular recovery of farm items is supported on both single-server and multi-server farms through the use of a Windows or another third-party tool. Granular recovery can be performed from full backups only.
- For a recovery to succeed, all nodes in the farm must be online and available.
- Only one recovery or backup operation per farm can run at any given time.
- For a given farm, any backups initiated while a recovery is in progress fails. Once the recovery completes, backups can be run for the farm.
- For a given farm, any recovery initiated while a backup is in progress fails. Once the backup completes, recoveries can be run for the farm.
- For granular, item-level recoveries, the backup is unfolded to a local share on the backup appliance. From here you can recover items using a Windows or third-party tool. When you are finished recovering, you must tear down the share. Subsequent backup or recovery operations for the given farm fail until the recovery share is torn down.

About the SharePoint recovery items session

Prior to starting the recovery, check the shares to see if one exists for the farm. If no share exists for this instance, start the recovery procedure.

If a share exists for this farm, it is either in use by another backup or recovery process, or it has not been torn down after a prior item recovery. You cannot perform the item recovery until the share becomes available.

If you are sure no active job is using the share, tear it down. Be sure to disconnect any network drive mappings to this share before tearing down.

To view or tear down SharePoint recovery objects

1. Select **Recover** and click the **File Level Recovery** tab.
2. Select the recovery object to tear down.
3. Click **Remove**.
4. Click **Yes** to confirm the removal.

SharePoint recovery procedures

You can recover from a backup by using the Backup Catalog or the Backup Browser. To recover from an imported backup or hot backup copy, you must use the Backup Catalog. See these topics for details:

- "To recover a SharePoint farm, site, service, or item by using the Backup Catalog"
- "To recover a SharePoint farm, site, service, or item by using the Backup Browser" on page 1189
- "To recover the entire farm on a standalone SharePoint 2013 or 2010 server by using the Backup Catalog" on page 1190
- "To recover the entire farm on a standalone SharePoint 2013 or 2010 server by using the Backup Browser" on page 1191

Note: To perform a farm, site, or service recovery, see the SharePoint Central Administration web site for recommendations and best practices.

To recover a SharePoint farm, site, service, or item by using the Backup Catalog

Run this procedure from the backup appliance to recover from a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover from a hot backup copy.

1. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the SharePoint Farm and select the backup, imported backup copy, or hot backup copy.

   (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
4. Click **Recover Files**.
5. Click **Confirm** to create the file recovery object.

6 Click **OK**.

7 The system creates a share for this farm instance and starts the recovery. A row for this recovery displays in the grid.

8 Select the share and click **Show Details**.

9 The details form displays the full path of the share, \\<SourceSystemIP>\<ClientName>-Farm.

Note or copy the Network Path.

10 On the workstation used to recover items, map a network drive to the following location: *<SourceSystemIP>**<ClientName>*-*<Instance>*

11 Launch Explorer; right-click **Computer** and select **Map Network Drive**.

12 In the Folder field, enter the share displayed in the Network Path field.

13 Click **Finish**.

14 Recover the items by using Windows or another third-party tool. For details on using Ontrack PowerControls, see "Recovering items with Ontrack PowerControls" on page 1192.

> **Note:** Creating the recovery object can take some time. If you do not see any available items, check back later.

15 Disconnect the network share once items have been recovered by right-clicking the share and selecting **Disconnect**.

16 On the backup appliance, tear down the recovery object as described in "To view or tear down SharePoint recovery objects" on page 1188.

> **IMPORTANT!** Tear down the share as soon as possible. Subsequent backups and recoveries cannot run for this farm until the share has been manually torn down.

To recover a SharePoint farm, site, service, or item by using the Backup Browser

> **Note:** This procedure is not supported for imported backups and backup copies. To recover from an imported backup or hot backup copy, use the Backup Catalog procedure above.

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Select the backup.

6 Click **Recover Files**.

7 Click **Confirm** to create the file recovery object.

8 Click **OK**.

9 The system creates a share for this farm instance and starts the recovery. A row for this recovery displays in the grid.

10 Select the share and click **Show Details**.

11 The details form displays the full path of the share, \\<SourceSystemIP>\<ClientName>-Farm.

Note or copy the Network Path.

12 On the workstation used to recover items, map a network drive to the following location: \\<*SourceSystemIP*>\<*ClientName*>-<*Instance*>

13 Launch Explorer; right-click **Computer** and select **Map Network Drive**.

14 In the Folder field, enter the share displayed in the Network Path field.

15 Click **Finish**.

16 Recover the items by using Windows or another third-party tool. For details on using Ontrack PowerControls, see "Recovering items with Ontrack PowerControls" on page 1192.

> **Note:** Creating the recovery object can take some time. If you do not see any available items, check back later.

17 Disconnect the network share once items have been recovered by right-clicking the share and selecting **Disconnect**.

18 On the backup appliance, tear down the recovery object as described in "To view or tear down SharePoint recovery objects" on page 1188.

> **IMPORTANT!** Tear down the share as soon as possible. Subsequent backups and recoveries cannot run for this farm until the share has been manually torn down.

To recover the entire farm on a standalone SharePoint 2013 or 2010 server by using the Backup Catalog

> **IMPORTANT!** This procedure is for full catastrophic farm recovery only. Recovering the entire farm removes the existing farm. This procedure is supported for SharePoint 2013 or 2010 servers deployed with the *single server* installation type only. (To check your installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 757.) Once a catastrophic farm recovery is complete, you must reconfigure all farm accounts and settings. Before performing catastrophic farm recovery, try to recover items using "To recover a SharePoint farm, site, service, or item by using the Backup Catalog" on page 1188.

Run this procedure from the backup appliance to recover from a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover from a hot backup copy.

1 Log in to the backup appliance or target appliance (if recovering from a hot backup copy).

2 Select **Recover** and click the **Backup Catalog** tab.

(Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.

3 Expand the SharePoint Farm and select a backup, imported backup copy, or hot backup copy.

(To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)

4 Click **Recover Files**.

5 Define the recovery options.

6 Click **Next**.

7 Select **I understand** and click **Confirm**.
The Recovery Status page indicates whether the recovery has been queued successfully. Click **Okay**.

To monitor the status of the recovery, select **Jobs > Active Jobs**. The recovery job displays in the grid. In a successful recovery, status changes from *Queued* to *Active* to *Successful*.

To recover the entire farm on a standalone SharePoint 2013 or 2010 server by using the Backup Browser

IMPORTANT! This procedure is for full catastrophic farm recovery only. Recovering the entire farm removes the existing farm. This procedure is supported for SharePoint 2013 or 2010 servers deployed with the *single server* installation type only. (To check your installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 757.) Once a catastrophic farm recovery is complete, you must reconfigure all farm accounts and settings. Before performing catastrophic farm recovery, try to recover items using "To recover a SharePoint farm, site, service, or item by using the Backup Catalog" on page 1188.

Run this procedure from the backup appliance to recover from a backup.

Note: This procedure is not supported for imported backups and backup copies. To recover from an imported backup or hot backup copy, use the Backup Catalog procedure above.

1 Log in to the backup.

2 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.

3 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

4 Select the backup.

5 Click **Recover Files**.

6 Define the recovery options.

7 Click **Next**.

8 Select **I understand** and click **Confirm**.
The Recovery Status page indicates whether the recovery has been queued successfully. Click **Okay**.

To monitor the status of the recovery, select **Jobs > Active Jobs**. The recovery job displays in the grid. In a successful recovery, status changes from *Queued* to *Active* to *Successful*.

Recovering items with Ontrack PowerControls

Using Ontrack PowerControls for SharePoint, you can recover individual items or a group of items from a Unitrends SharePoint full backup. (It is not possible to perform item-level recovery from SharePoint differential backups. Recover from either a SharePoint full or a SQL backup instead.)

Items can be recovered to the same SharePoint instance, a different SharePoint instance, or a network location. Ontrack PowerControls for SharePoint and the PowerControls ExtractWizard must be installed on a workstation in your network, and a valid Ontrack PowerControls license must be applied. Speak with your Unitrends sales representative for information about obtaining an Ontrack PowerControls license.

The following procedure walks you through a typical item-level recovery. For a complete overview of Ontrack PowerControls for SharePoint, including procedures and limitations, see the Ontrack PowerControls SharePoint user guide available at https://download.ontrack.com/downloads/OPCSP_Manual.pdf.

To recover items using Ontrack PowerControls

Note: Creating the recovery object can take some time. If you do not see any available items in KOP, check back later.

1 Run one of the following procedures: "To recover a SharePoint farm, site, service, or item by using the Backup Catalog" on page 1188 or "To recover a SharePoint farm, site, service, or item by using the Backup Browser" on page 1189.

2 From your Ontrack PowerControls workstation, run the Ontrack PowerControls ExtractWizard. On the Welcome screen, click **Next**.

3 Choose the Direct Method of extraction, and click **Next**.

4 Select **Extract from Disk**, and click **Browse**. Browse to the network path supplied by the Unitrends appliance.

5 Locate the file *spbackup.xml* (the location varies by SharePoint version). Select spbackup.xml, and click **Open**, then click **Next**.

6 Select **Catalog SharePoint backup datasets only**, and click **Next**.

7 Browse to the content databases from which you want to recover data, and click **Next**.

8 From the Destination Folder, browse to or type a path to a convenient temporary working directory, such as a folder on your desktop. Click **Next**.

9 After the databases have been extracted, click **Finish**.

10 Open Ontrack PowerControls for SharePoint.

11 On the Welcome screen, click **Next**.

12 On the Source Path Selection screen, click **Add**. Browse to the folder you created above in step 8. Locate the .mdf and .ldf files of the databases from which you want to recover. Select both, click **Open**, then **Next**.

13 On the Target Server Selection screen, supply the URL and administrative credentials for the SharePoint site to which you want to recover, and click **Next**.

14 The Source pane at the top of the screen represents your Unitrends backup. The Target pane at the bottom of the screen represents your live SharePoint environment. Do one of the following:

- To recover items back to the SharePoint site, browse to items you want to recover in the Source pane. Drag and drop into the desired nodes in the Target pane.
- To recover items to a local folder or network location, right click and select Export. **Deselect Maintain message**, browse to where you want to save the files, and click **Finish**.

15 Disconnect the network share once items have been recovered by right-clicking the share and selecting **Disconnect**.

16 On the backup appliance, tear down the recovery object as described in "To view or tear down SharePoint recovery objects" on page 1188.

Recovering Oracle backups

Unitrends supports recovery of full and incremental backups. As with Oracle backups, the agent leverages RMAN to perform recovery operations. Oracle recoveries occur in two phases. First, the backup is extracted to the server's storage and exposed as a CIFS share (/backups/rae/<asset_name>/ <instance>). In the second phase, the asset accesses the exposed CIFS share, and the RMAN is invoked to recover back to the Oracle database.

Requirements and considerations

Consider the following before recovering Oracle data from the backup system:

- Free space equivalent to twice the size of the backup is required on the Unitrends appliance for recovery processing. If adequate space is not available, the recovery fails.
- Recoveries are performed to the original database only. If you are recovering from a Unitrends appliance backup copy target and the original database is not available, you can recover from the target after performing a bare metal recovery of the Oracle asset. This procedure is only supported for Oracle on Windows platforms running Oracle 19c, 18, 12c, or 11g. See "Windows Bare Metal Protection and Recovery" on page 1207 for details.
- Only one recovery or backup operation per database can run at any given time.

- For a given database, any backups initiated while a recovery is in progress fails. Once the recovery completes, backups can be run for the given database.
- For a given database, any recoveries initiated while a backup is in progress will fail. Once the backup completes, recoveries can be run for the given database.
- For item-level recovery, the backup is unfolded to a Unitrends appliance. From here you can recover items using an Oracle or third-party tool. After recovery, you must tear down the recovery object. See "About the Oracle recovery object" on page 1198 for procedures and details. Subsequent backup or recovery operations for the given instance fail until the share has been manually torn down.
- For Oracle on Windows, recovery requires that the underlying file system has the same structure as when the database was initially backed up. For details, see Oracle Database: Failed to Create File.

Recovering an Oracle backup

The following recovery options are available:

- Recover – Select a backup to recover all data in the backup group up to the point in time when the backup ran.
- Recover Files – Search a selected backup and choose specific files to recover.

You can recover from a backup by using the Backup Catalog or the Backup Browser. To recover from an imported backup copy, you must use the Backup Catalog. See these topics for details:

- "To recover an Oracle backup by using the Backup Catalog"
- "To recover an Oracle backup by using the Backup Browser" on page 1195
- "To recover items from an Oracle backup by using the Backup Catalog" on page 1196
- "To recover items from an Oracle backup by using a Backup Browser" on page 1196

To recover an Oracle backup by using the Backup Catalog

This procedure recovers a database to the original location. Note that the existing database is deleted from the Oracle instance as part of the recovery process.

Run this procedure from the backup appliance to recover a backup or imported backup copy. To recover from a hot backup copy, see "Oracle recovery from a Unitrends appliance backup copy target" on page 1198.

1. Log in to the backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the Oracle asset and select the backup or imported backup copy.

   (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
4. Click **Recover**.

5 Select the desired recovery options.

6 Click **Save**.

To monitor the recovery, select **Jobs > Active Jobs**.

Notes:

- In a successful recovery, status changes from *Queued* to *Active* to *Successful*.
- If the status displays as *Canceled* with a message *Share is unavailable*, the recovery cannot run because a share already exists for this instance. To determine what process is using the share and how to proceed, see "Oracle share is unavailable" on page 1197.

To recover an Oracle backup by using the Backup Browser

This procedure recovers a database to the original location. Note that the existing database is deleted from the Oracle instance as part of the recovery process.

Note: This procedure is not supported for imported backups and backup copies. To recover an imported backup, use the Backup Catalog procedure above. To recover from a hot backup copy, see "Oracle recovery from a Unitrends appliance backup copy target" on page 1198.

1 Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Select the Oracle backup.

6 Click **Recover**.

7 Select the desired recovery options.

8 Click **Save**.

To monitor the recovery, select **Jobs > Active Jobs**.

Notes:

- In a successful recovery, status changes from *Queued* to *Active* to *Successful*.
- If the status displays as *Canceled* with a message *Share is unavailable*, the recovery cannot run because a share already exists for this instance. To determine what process is using the share and how to proceed, see "Oracle share is unavailable" on page 1197.

To recover items from an Oracle backup by using the Backup Catalog

Use this procedure to unfold the backup to a share on the backup appliance. Once unfolded, use an Oracle or third-party tool to recover desired items.

Run this procedure from the backup appliance to recover items from a backup or imported backup copy. To recover from a hot backup copy, see "Oracle recovery from a Unitrends appliance backup copy target" on page 1198.

1. Log in to the backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the Oracle asset and select the backup or imported backup copy.

   (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
4. Click **Recover Files**.
5. Click **Confirm** to create the recovery object and start the recovery.
6. Click **OK** to clear the **Notice** dialog.

   A row for this recovery displays on the **File Level Recovery** tab.
7. Select the Oracle recovery row and click **Show Details** to display the full path of the share: \\<*ApplianceIP*>\<*InstanceID*>.

   Note: Record the Network Path to access files to recover later.
8. On the workstation to be used to recover files, map a network drive to \\<*ApplianceIP*>\<*InstanceID*>.
9. Recover the desired items using an Oracle or third-party tool.
10. Disconnect the network share once files are recovered. To disconnect the share, select the share and click **Remove**.
11. On the backup appliance, tear down the recovery object.

IMPORTANT! Tear down the objects soon as possible. Subsequent backups and recoveries cannot run for this instance until the object has been manually torn down.

To recover items from an Oracle backup by using a Backup Browser

Use this procedure to unfold the backup to a share on the backup appliance. Once unfolded, use an Oracle or third-party tool to recover desired items.

Note: This procedure is not supported for imported backups and backup copies. To recover items from an imported backup, use the Backup Catalog procedure above. To recover from a hot backup copy, see "Oracle recovery from a Unitrends appliance backup copy target" on page 1198.

1. Log in to the backup appliance.

2 Select **Recover > Backup Catalog** and click **Backup Browser**.

3 Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.

4 (Optional) Refine the search:

- Enter text in any column field to filter the display.
- Click an arrow to sort by column.
- Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
- For a description of each column, see "Backup Browser column descriptions".

5 Select the Oracle backup.

6 Click **Recover Files**.

7 Click **Confirm** to create the recovery object and start the recovery.

8 Click **OK** to clear the **Notice** dialog.
A row for this recovery displays on the **File Level Recovery** tab.

9 Select the Oracle recovery row and click **Show Details** to display the full path of the share: \\<*ApplianceIP*>\<*InstanceID*>.

> **Note:** Record the Network Path to access files to recover later.

10 On the workstation to be used to recover files, map a network drive to \\<*ApplianceIP*>\<*InstanceID*>.

11 Recover the desired items using an Oracle or third-party tool.

12 Disconnect the network share once files are recovered. To disconnect the share, select the share and click **Remove**.

13 On the backup appliance, tear down the recovery object.

> **IMPORTANT!** Tear down the objects soon as possible. Subsequent backups and recoveries cannot run for this instance until the object has been manually torn down.

Oracle share is unavailable

Only one recovery or backup job per instance can run at any given time. If an Oracle backup or recovery operation fails with a share is unavailable message, a share already exists for this database. Review the following for additional information:

- If a share exists for this database, it is in use by another backup or recovery process, or it has not been torn down after a previous recovery.
- If no active job is using the share, tear down the share. Disconnect any network drive mappings to the share before tearing it down. To tear down the share, select **Recover**, click the **File Level Recovery** tab, select the share, and click **Remove**.
- To view details on active jobs, select **Jobs** and click the **Active Jobs** tab.

About the Oracle recovery object

After files are recovered, the object remains until you tear it down. Because appliance resources are used to maintain the object, it is important to tear it down to ensure optimal performance.

To view Oracle recovery image

Select **Recover** and click the **File Level Recovery** tab.

To tear down the Oracle recovery object

1. Disconnect any network drive mappings to the share before tearing it down.
2. Select **Recover** and click the **File Level Recovery** tab.
3. Select the recovery image to tear down.
4. Click **Remove**.

Oracle recovery from a Unitrends appliance backup copy target

Use this procedure if you are unable to recover from the backup appliance. This procedure is only supported on Oracle on Windows platforms running Oracle 19c, 18, 12c, or 11g.

Because Oracle recovery to an alternate server is not supported, this procedure requires that you perform a disaster recovery (DR) of the Windows asset to a new asset that is directly attached to the backup copy target. See "Windows Bare Metal Protection and Recovery" on page 1207 for details.

Once the asset has been recovered, you perform an Oracle granular recovery.

Considerations and procedures for Oracle recovery from a Unitrends appliance backup copy target

Consider the following before recovering from a Unitrends appliance backup copy target:

- Free space equivalent to twice the size of the backup is required on the local share for recovery processing. If adequate space is not available, the recovery fails.
- A backup copy of the source asset taken after the database was deployed is required. To perform DR using Windows integrated recovery, the backup copy must have been run using version 7.3 or higher for BIOS-based assets, or version 7.4 or higher for UEFI-based assets. To perform DR using Windows image-based bare metal, the backup must be a bare metal backup.
- A backup copy of the database to recover is required.
- Only one recovery or backup operation per database can run at any given time. Any backup or recovery initiated while another is in progress fails.
- The Oracle backup is unfolded to a remote share. When you are finished recovering, you must tear down the share. Subsequent backup or recovery operations for the given instance will fail until the share has been torn down.

To recover an Oracle database from a Unitrends appliance backup copy target

1 Perform a disaster recovery of the Oracle asset from the Unitrends appliance backup copy target. See "Windows Bare Metal Protection and Recovery" on page 1207 for details.

2 Select **Recover** and click the **File Level Recovery** tab to determine if a share exists for this database. Review the following before proceeding:

- If no share exists for this database, proceed to the next step in this procedure.
- If a share exists for this database, it is in use by another backup or recovery process, or it has not been torn down after a previous recovery. You cannot perform the item recovery until the share is available.
- If no active job is using the share, tear down the share. Disconnect any network drive mappings to the share before tearing it down. To tear down the share, select it and click **Remove**.
- To view details on active jobs, select **Jobs** and click the **Active Jobs** tab.

3 Select **Recover**, and click the **Backup Catalog** tab.

4 In the **Filter Backups** window on the right, select **Backup Copy**.

5 (Optional) Select other filter options. For details, see "Working with custom filters" on page 67.

6 Click **Filter**.

7 Expand the list of backups under the Oracle instance and select a backup to recover.

8 Click **Recover Files**.

9 Click the **File Level Recovery** tab.

10 Select the share and click **Show Details**.

Record or copy the CIFS path as it displays in the File Level Recovery Details window.

11 Log in to the target you created in step 1 on page 1199.

12 Open Windows Explorer.

13 In Windows Explorer, navigate to the CIFs path recorded in step 10.

14 Open the file `unitrends-<database>.env`. Note the following values:

- Oracle SID
- Oracle Home
- Backup #

15 Open a command-line prompt.

16 Execute the following commands:

```
# set ORACLE_SID=<SIDfromLastStep>
# set ORACLE_HOME=<HomefromLastStep>
# %ORACLE_HOME%\bin\rman target /
```

```
# shutdown immediate;
# startup nomount;
# restore controlfile from '<pathFromStep8>\unitrends-<database>.ctf';
# alter database mount;
# crosscheck backup;
# catalog start with '<pathFromStep8>\unitrends';
# list backup tag 'unitrends-<backup#FromLastStep>';
```

17 Run the following commands:

```
# restore database;
# recover database;
# alter database open resetlogs;
# quit;
```

The Oracle recovery is complete.

18 On the Unitrends appliance backup copy target, disconnect any network mapping to the share and tear down the restore share. To tear down the share, select it and click **Remove**.

IMPORTANT! Tear down the share as soon as possible. Subsequent backups and restores cannot run for this instance until the share has been manually torn down.

Recovering Cisco UCS service profile backups

Unitrends leverages the native Cisco XML API to recover from service profile backups. Use the Recover option to recover the entire backup group up to the point in time when the backup ran or to recover selected items.

See the following topics for details:

- "Preparing to recover service profile backups"
- "Recovery procedures" on page 1201

Preparing to recover service profile backups

Before recovering, review the following requirements and considerations

- Service profiles, templates, pools, and policies must be recovered using the original name to prevent namespace collisions.
- Recovering a profile, template, pool, or policy overwrites the original if it exists on the UCS.
- Recovering an active profile takes down that service profile. You must restart the service profile after the recovery completes.
- You can recover the entire backup or selected items to the original asset or to an alternate UCS manager asset that has been added to the appliance.

- Before recovering an entire backup, it is recommended to view its contents to be sure you want to recover all items. During the recovery procedure, you can expand the file browser to view these items. For UCS file naming conventions, see "Identifying files in UCS service profile backups" on page 1202.
- Only one recovery or backup operation per UCS manager can run at any given time. Any subsequent jobs are queued and started once the last run completes.

Recovery procedures

You can recover from a backup by using the Backup Catalog or the Backup Browser. To recover from an imported backup or hot backup copy, you must use the Backup Catalog. See the following topics for details:

- "To recover from a service profile backup by using the Backup Catalog"
- "To recover from a service profile backup by using the Backup Browser" on page 1202

To recover from a service profile backup by using the Backup Catalog

Use this procedure to recover the entire backup or selected service profiles, templates, pools, and policies.

Run this procedure from the backup appliance to recover from a backup or imported backup copy. Run this procedure from the backup copy target appliance to recover from hot backup copy.

1. Log in to the backup appliance or target appliance (if recovering from a hot backup copy).
2. Select **Recover** and click the **Backup Catalog** tab.

   (Optional) Use Filter Backups to the right to customize the backups that display. For details, see "Working with custom filters" on page 67.
3. Expand the UCS asset and select a backup, imported backup, or hot backup copy to use for the recovery.

   (To import a backup copy, see "To import a cold backup copy" on page 786 or "To import a hot backup copy" on page 780.)
4. Click **Recover**.
5. In the File Browser, expand folders to view items in the backup.

   To expand all folders, press * on your keyboard.
6. Select or drag files to recover. To recover the entire backup, select the top-level folder.
7. Click **Next**.
8. Specify a Restore Target by selecting an asset.

   This must be the original UCS manager or another UCS manager that has been added to the appliance.
9. Click **Save**.
10. Click **OK** to close the Notice message.

    Selected items are recovered to the target location. To view the running job, select **Jobs > Active Job**.
11. Once the job completes, use the UCS manager to restart any active service profiles that have been recovered.

To recover from a service profile backup by using the Backup Browser

Use this procedure to recover the entire backup or selected service profiles, templates, pools, and policies.

Note: This procedure is not supported for imported backups and backup copies. To recover an imported backup or hot backup copy, use the Backup Catalog procedure above.

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog** and click **Backup Browser**.
3. Select the **Appliance** and **Date Range** of backups to search. Backups that ran during the date range display.
4. (Optional) Refine the search:
   - Enter text in any column field to filter the display.
   - Click an arrow to sort by column.
   - Click the accordion icon to add or remove columns from the display. (Not all columns display by default.)
   - For a description of each column, see "Backup Browser column descriptions".
5. Select the backup.
6. Click **Recover**.
7. In the File Browser, expand folders to view items in the backup.

   To expand all folders, press * on your keyboard.
8. Select or drag files to recover. To recover the entire backup, select the top-level folder.
9. Click **Next**.
10. Specify a Restore Target by selecting an asset.

    This must be the original UCS manager or another UCS manager that has been added to the appliance.
11. Click **Save**.
12. Click **OK** to close the Notice message.

    Selected items are recovered to the target location. To view the running job, select **Jobs > Active Job**.
13. Once the job completes, use the UCS manager to restart any active service profiles that have been recovered.

Identifying files in UCS service profile backups

Each UCS service profile backup contains all supported service profiles, templates, pools, and policies present on the UCS at the time the backup ran. When recovering, it may be necessary to select specific items. Use the following naming conventions table to identify items in a UCS profile backup.

Note: The following objects are not included in Unitrends UCS profile backups: BIOS defaults, IPMI access policies, management firmware policies (deprecated, replaced by host firmware packages), and iSCSI authentication profiles.

Supported Cisco UCS objects	File prefix naming convention
Service profiles and templates	*ls-**
Adapter policies	*eth-profile** or *fc-profile**
BIOS policies	*bios-prof-**
Boot policies	*boot-policy-**
Host firmware packages	*fw-host-pack-**
Local disk configuration policies	*local-disk-config-**
Maintenance policies	*maint-**
Power control policies	*power-policy-**
Scrub policies	*scrub-**
Serial over LAN policies	*sol-**
Server pool policies	*compute-pool-**
Server pool policy qualifications	*blade-qualifier-**
Threshold policies	*thr-policy-**
vNIC/vHBA placement policies	*vcon-profile-**
Sub-organizations	*org-**

This page is intentionally left blank.

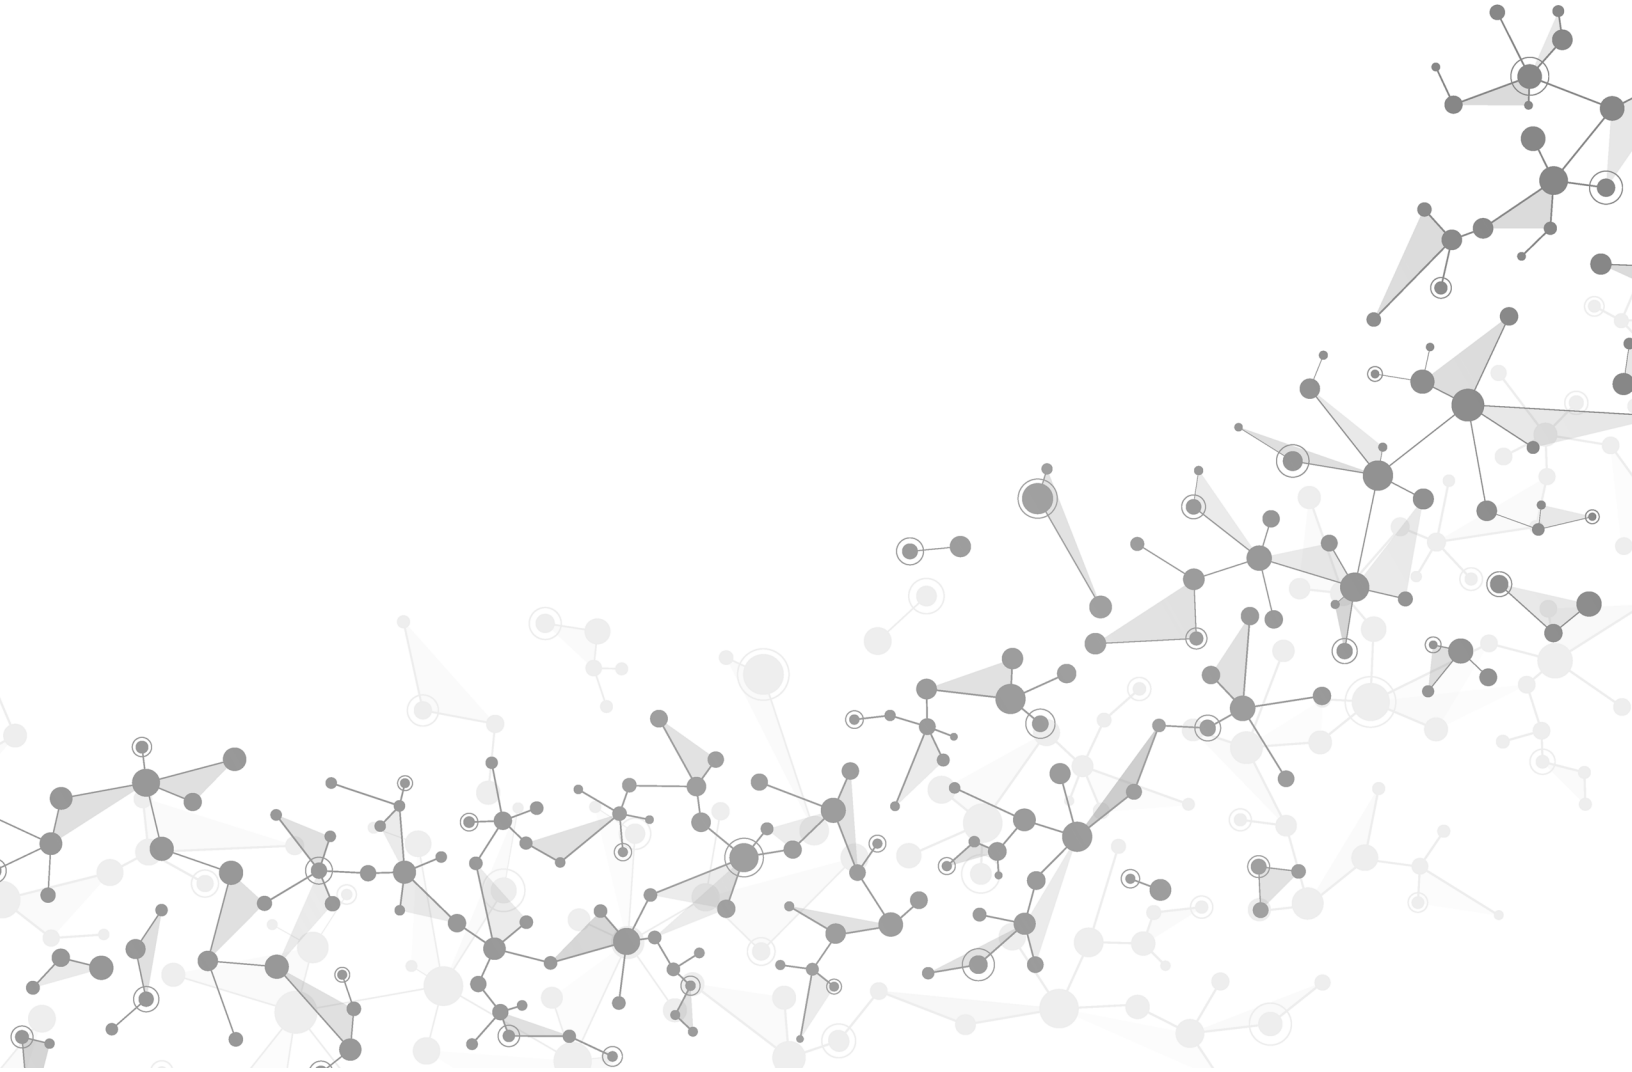

Chapter 19: Recovering iSeries Backups

Use the recovery feature to recover the entire backup or selected items to the original iSeries or to an alternate iSeries server. See the following topics for details:

Preparing to recover from iSeries backups

Before performing the recovery, ensure that these requirements have been met:

- No other jobs are running on the iSeries server. To check for active jobs, log in to the OS400 operating system and issue the **WRKACTJOB** command.
- The user performing the recovery must, at a minimum, have *SECADM privileges. The recovery is run by the user associated with the profile of the backup you are recovering.
- Files to recover must have read-write attributes. To assign read-write attributes, log in to the OS400 operating system and grant object authority to the user performing the recovery. For example, enter the following to modify security privileges in the QGPL and QUSRSYS libraries for user QSECOFR:

```
# GRTOBJAUT OBJ(QGPL/*ALL) OBJTYPE(*ALL) USER(QSECOFR) AUT(*ALL)
# GRTOBJAUT OBJ(QUSRSYS/*ALL) OBJTYPE(*ALL) USER(QSECOFR) AUT(*ALL)
```

To recover from an iSeries backup

After the requirements in "Preparing to recover from iSeries backups" on page 1205 have been met, use this procedure to recover from the iSeries backup:

1. Using a terminal emulator, such as PuTTY, connect to the appliance using the following:
   - Appliance IP address
   - Port 22
   - SSH connection type
2. Log in as user *root*.
3. Enter the following command to access the console menu:

```
# dpuconfig
```

4. Select option **4** for Advanced Options.
5. Select option **2** for IBM iSeries Backup and Recovery.
6. Select option **3** for Restore iSeries.
7. Select the iSeries server whose backup you want to recover.

8 Select the desired backup.

9 Select the recovery type: full or selective restore.

Notes: Full restore –

- A full restore recovers everything in the backup.
- The /Security Data object contains the save file from the **SAVSECDTA** command. If the /Security Data object has not been excluded from the backup, it is the first object recovered (via the **RSTUSRPRF** command), then a **RSTAUT** command is executed after everything else has been recovered.
- The /System Configuration object contains the save file from the **SAVCFG** command. If the /Security Configuration object has not been excluded from the backup, it is recovered before any other objects, except /Security Data, in the backup file. It is recovered using the **RSTCFG** command.

10 Follow the prompts to select desired options and start the recovery job.

For selective restore:

- Enter each file, library, or IFS object to recover, one entry per line.
- To recover to the original location, do not enter a destination path.
- To recover to an alternate location, enter the full path of the destination.
- When your list is complete, type **q** to end your list and launch the recovery job.

iSeries disaster recovery

To recover the iSeries from a catastrophic state, you must first recover the OS and Licensed Internal Code from your supplemental system backup. Next, recover the full backup to obtain data, as described in "To recover from an iSeries backup" on page 1205.

Chapter 20: Windows Bare Metal Protection and Recovery

Bare metal technology is used for disaster recovery of protected assets. Use the procedures in this chapter to set up bare metal protection for your Windows assets and to recover failed Windows assets. (To recover the Unitrends appliance itself, see "Appliance Disaster Recovery" on page 1295. To set up bare metal protection for assets running other operating systems, see the Bare Metal Protection and Recovery Guide.)

There are two options for hot bare metal recovery (BMR) of Windows agent-based assets: Windows unified BMR (formerly known as *integrated BMR*) and Windows image-based BMR.

With Windows unified BMR, Unitrends provides Unified Bare Metal™ protection that enables you to perform disaster recovery (DR) right from a file-level or image-level backup. This reduces recovery time, provides additional recovery points, increases on-appliance retention by eliminating the need for bare metal backups, and simplifies the Windows DR process. You perform unified BMR by using the Unified Bare Metal Recovery wizard and standard 32-bit and 64-bit ISO images, eliminating the need to create bare metal ISOs for each protected asset and keep them on-hand in case disaster strikes.

With image-based BMR, you must run bare metal backups and create a separate bare metal ISO for each Windows asset you want to protect. You perform image-based BMR by booting from the asset's bare metal ISO. Image-based BMR can protect older versions of Windows that are not supported by unified BMR.

Note: If you have Windows virtual machines, you can protect them by running host-level backups that use hypervisor snapshots or by installing the Windows agent and running file-level backups. If you are running agent-based file-level backups for a VM, use the hot bare metal procedures in this chapter for disaster recovery. If you are running host-level backups for a VM, see "Recovering a virtual machine" on page 796.

Which bare metal method should I use?

It is recommended to use unified BMR where possible. The following table provides a high-level comparison of unified and image-based hot bare metal recovery. To set up bare metal protection, see the following topics:

- "Windows unified bare metal recovery" on page 1209
- "Windows image-based bare metal recovery" on page 1246

Note: Unified BMR does not support recovering a Windows 2003 Server to dissimilar hardware. In this case, you must use image-based BMR instead.

Item	Unified BMR	Image-based BMR
Recovery Time Objective (RTO)	Faster recovery time than with image-based BMR.	Slower recovery time than with unified BMR.
Recovery Point Objective (RPO)	More recovery points available since you	Fewer recovery points since you restore from a bare metal backup only.

Item	Unified BMR	Image-based BMR
	restore from any eligible file-level or image-level backup.	
Recovery types	Supports physical-to-virtual (P2V), virtual-to-physical (V2P), physical-to-physical (P2P), and virtual-to-virtual (V2V) DR.	Supports physical-to-virtual (P2V), virtual-to-physical (V2P), physical-to-physical (P2P), and virtual-to-virtual (V2V) DR.
Recovery of Windows Server 2022	Yes, recovering Windows Server 2022 to identical or dissimilar hardware is supported.	No, recovering Windows Server 2022 assets is not supported.
Recovery of Windows Server 2019	Yes, recovering Windows Server 2019 to identical or dissimilar hardware is supported.	No, recovering Windows Server 2019 assets is not supported.
Recovery of Windows Server 2016	Yes, recovering Windows Server 2016 to identical or dissimilar hardware is supported.	No, recovering Windows Server 2016 assets is not supported.
Dissimilar recovery of Windows Server 2012	Yes, recovering Windows Server 2012 to identical or dissimilar hardware is supported.	Yes, recovering Windows Server 2012 to identical or dissimilar hardware is supported.
Dissimilar recovery of Windows Vista/Server 2008	Yes, recovering Windows Vista/Server 2008 to identical or dissimilar hardware is supported.	Yes, recovering Windows Vista/Server 2008 to identical or dissimilar hardware is supported.
Dissimilar recovery of Windows Server 2003	No, recovering Windows Server 2003 to dissimilar hardware is not supported. Recovery to identical hardware is supported.	Yes, recovering Windows Server 2003 to dissimilar hardware is supported for some distributions. See the Compatibility and Interoperability Matrix for details. Recovery to identical hardware is supported.
Dissimilar recovery of Windows XP	No, recovering Windows XP to dissimilar hardware is not supported. Recovery to identical hardware is supported.	No, recovering Windows XP to dissimilar hardware is not supported. Recovery to identical hardware is supported.

Item	Unified BMR	Image-based BMR
On-appliance retention	More on-appliance retention due to eliminating bare metal backups.	Less on-appliance retention due to bare metal backup storage.
ISO image/boot disk	Standard 32-bit and 64-bit ISO images used for most Windows assets; available on the Unitrends appliance.	Separate ISO required for each Windows asset; ISOs must be created manually with the Unitrends bare metal agent.
Bare Metal Interface	Simplified wizard interface enables DR to the desired point-in-time using a single process, decreasing overall recovery time. Leverages WinPE 10.0 for all Windows assets.	Two dialog-based interfaces (one WinPE 1.5 for older assets, one WinPE 2.0 for newer assets). Cannot perform DR in a single process.
Target disk size	For file-level backups, supports recovery of original Windows asset to a smaller disk size. (For details, see "Prerequisites for file-level backups" on page 1210.) For image-level backups, must recover to a disk of an equal or greater size than that of the original asset. (For details, see "Prerequisites for image-level backups" on page 1215.)	Must recover to a disk of an equal or greater size than that of the original asset.
UEFI-based assets	Supports recovery of UEFI-based assets.	Cannot recover UEFI-based assets.
GPT-partitioned assets	Supports recovery of GPT-partitioned assets.	Cannot recover GPT-partitioned assets.

Windows unified bare metal recovery

With Windows unified bare metal recovery (BMR), you can protect a Windows asset's operating system without running a bare metal backup or creating a custom ISO image. File-level and image-level backups capture the disk metadata necessary for the recovery, and you perform DR using a standard 32-bit or 64-bit ISO image provided on the Unitrends backup appliance. The target for the recovery can be a physical or virtual machine.

When you boot the target machine from the standard ISO, it boots into WinPE (a minimal version of Windows used for installations) and the Windows Unified Bare Metal Recovery (UBMR) wizard launches to guide you through the recovery. Depending on your operating system and hardware, it might be necessary to add drivers during the recovery. You can use the wizard to add any drivers.

See the following topics for details about protecting your Windows assets with unified bare metal recovery:

- "Implementing Windows unified bare metal protection"
- "Performing unified bare metal recovery" on page 1218

Implementing Windows unified bare metal protection

For best results, it is recommended that you plan your strategy for disaster recovery before an asset fails. Following is a high-level overview of the steps you must complete to implement unified bare metal protection for your Windows assets. It identifies steps to complete before and after an asset fails.

Perform the following before an asset fails

Step 1: Check the Windows asset to determine the following:

- Its operating system version and whether it is 32-bit or 64-bit. For instructions, see the Microsoft document Which Windows operating system am I running? (Recovery of image-level backups is supported for 64-bit assets only.)
- Its firmware interface type (BIOS or UEFI). Check this by viewing the system information or volumes on the Windows machine.

Step 2: Review the "Prerequisites for Windows unified bare metal recovery" on page 1210 to verify that the operating system is supported and the asset meets all requirements.

Step 3: Run file-level or image-level backups that include all critical system information. (For details see "An eligible file-level backup" or "An eligible image-level backup".)

Step 4: (Optional/recommended) Perform a test recovery using the procedures in "Performing unified bare metal recovery" on page 1218.

To recover a failed asset

Step 5: Perform unified bare metal recovery (BMR) using the procedures in "Performing unified bare metal recovery" on page 1218.

Prerequisites for Windows unified bare metal recovery

Consider the prerequisites for unified BMR as you plan your disaster recovery strategy. For Windows operating systems not supported by unified BMR, see "Windows image-based bare metal recovery" on page 1246.

Requirements differ by backup type. See these topics for details:

- "Prerequisite for the Unitrends backup appliance"
- "Prerequisites for file-level backups"
- "Prerequisites for image-level backups" on page 1215

Prerequisite for the Unitrends backup appliance

For increased security, the UniView Portal provides an option to block local access to the Unitrends appliance.

To perform unified BMR, local access must be unblocked on the Unitrends appliance whose backup you will use for the recovery. If needed, temporarily unblock local access until you have recovered your Windows asset. For details, see Blocking or unblocking local access to an appliance in the UniView Portal Guide.

Prerequisites for file-level backups

The following requirements must be met to recover from a file-level backup.

Requirement	Description
Operating systems	Recovery to identical hardware and virtual machines is supported for the client operating systems listed below. (Additional version limitations apply. See the Compatibility and Interoperability Matrix for details.) Recovery to dissimilar hardware is NOT supported for XP and is supported for the others in this list. • Windows XP, 32-bit and 64-bit (SP2 and later) • Windows Vista, 32-bit and 64-bit (SP2) • Windows 7, 32-bit and 64-bit • Windows 8, 32-bit and 64-bit • Windows 8.1, 32-bit and 64-bit • Windows 10, 64-bit • Windows 11, 64-bit Recovery to identical hardware and virtual machines is supported for the server operating systems listed below. (Additional version limitations apply. See the Compatibility and Interoperability Matrix for details.) Recovery to dissimilar hardware is NOT supported for 2003/2003 R2 and is supported for the others in this list. • Windows 2003, 32-bit and 64-bit (SP2) • Windows 2003 R2, 32-bit and 64-bit • Windows Small Business Server 2003 and later, 32-bit and 64-bit • Windows 2008, 32-bit and 64-bit • Windows 2008 R2, 64-bit • Windows 2012, 64-bit • Windows 2012 R2, 64-bit • Windows 2016, 64-bit • Windows 2019, 64-bit • Windows 2022, 64-bit
An eligible file-level backup	The backup used for recovery must meet these requirements: • It is successful. • It is a full, differential, or incremental file-level backup that contains disk metadata (known as the asset's *system state*). Disk metadata is captured in file-level backups unless you have opted to exclude critical volumes or to exclude the

Requirement	Description
	system state. To check whether the system state is included in a backup, run the "Backup History report" on page 1320 and select the backup in the list. In the Backup Status: Report Entry dialog, check the Output area for System State *Excluded* or *Included*. If system state was excluded, you need to modify the backup job to include all critical volumes to create a backup that can be used for unified BMR. For details, see "To create a file-level backup job" on page 437.
Unified BMR ISO images	For the recovery, you must use the 32-bit or 64-bit unified BMR ISO image provided on the your Unitrends appliance. The ISOs contain WinPE (a minimal version of Windows used for installations) and the Unitrends Unified Bare Metal Recovery wizard that guides you through the recovery. To prepare for DR, it is recommended that you do the following: • Create bootable CDs of these ISOs and store them in a safe place (so that you can quickly recover to a physical machine target). • Save the ISOs to your hypervisor (so you can quickly recover to a virtual machine target). For details, see "Access the unified bare metal recovery ISO image" on page 1218.
Firmware interface type	Supported for BIOS- and UEFI-based assets. The firmware interface type (BIOS or UEFI) of the recovery target machine must match that of the failed asset.
Disk configuration	• GPT disks are supported. • Dynamic disks are not supported. • iSCSI disks are not supported. Recover the critical (non-iSCSI) volumes as described in "Performing unified bare metal recovery" on page 1218. Once the critical volumes have been restored, recover data on the iSCSI volumes as described in "Recover from backups or imported backup copies" on page 926.
Software RAID volumes	Software RAID configurations are not supported.
Network adapter	Wireless network adapters cannot be used for the recovery.
Drivers	Drivers needed for unified BMR are determined by the operating system of the asset you are recovering and the operating system and hardware of the target recovery machine. You might need to add drivers during different stages of the recovery.

Requirement	Description
	• Loading WinPE drivers for unified bare metal recovery — The Unified BMR wizard uses WinPE 10.0 for the recovery. WinPE 10.0 is based on Windows 11 or Server 2022. If WinPE 10.0 cannot detect a network adapter or storage disks, you must load Windows 11 or Server 2022 drivers into WinPE during recovery. • Injecting drivers into the recovered operating system — After the critical volumes have been recovered, you must inject drivers into the recovered operating system if you are recovering to dissimilar hardware or to a virtual machine. – If recovering to a physical machine, required drivers vary by hardware and operating system. It is recommended that you verify whether the operating system of the asset you are recovering requires additional drivers to run on the hardware of the recovery target machine before you begin the recovery. – If recovering to a virtual machine, you must inject ESX, Hyper-V, or XenServer guest storage drivers, depending upon your virtual environment. These drivers are included in the unified BMR ISO image.
Processor features on the recovery target	WinPE requires that these processor features are enabled on the recovery target machine: NX, PAE, and SSE2. You might need to enable these features on the recovery target machine before booting from the ISO image. For instructions, see Error 0x0000005D (Unsupported Processor) when booting Release 7.4 Integrated Bare Metal restore media. Machines that do not have these processor features cannot be used for the recovery.
Disk space on recovery target	Make sure the target machine has enough disk space for the recovery. To recover from a file-level backup, adhere to these requirements: • The recovery target can have smaller disks than the failed asset, but the recovery fails if the disks do not have enough space for the data on the critical volumes. • After an asset has failed, there is no way to determine the size of its critical volumes. You can determine the size of an asset's full backup by viewing the backup details in the "Backup History report". However, the size of the critical volumes will be smaller than the full backup if it also contains non-critical volumes. If you are unsure about the size of the critical volumes, it is recommended that you recover to destination disks that are the same size as the original disks or larger. • If you are recovering to new disks, any existing data on the destination disks is overwritten or deleted during the recovery, even if the disks have more than enough space. Before performing a recovery, make sure you have additional copies of any data on the destination disks. • If you are recovering to the original disk, only the recovered volumes are

Requirement	Description
	overwritten. Other volumes on the original disk are not impacted by the recovery.
Supported recovery scenarios	Use the procedures in "Performing unified bare metal recovery" on page 1218 for the following recovery scenarios: • Recover to the same physical hardware as the failed asset. • Recover a failed physical asset to dissimilar hardware. Supported for Windows Vista/Server 2008 and higher. • Recover a failed physical asset to dissimilar hardware with fewer disks. Supported for Windows Vista/Server 2008 and higher. • Recover a failed physical asset to hardware with smaller or larger disks. • Recover a failed asset BIOS/MBR configuration to a dissimilar BIOS/MBR configuration. Supported for Windows Vista/Server 2008 and higher. • Recover a failed asset UEFI/GPT configuration to a dissimilar UEFI/GPT configuration. Supported for Windows Vista/Server 2008 and higher. • Recover multi-boot configured BIOS servers. • Recover a failed physical asset to a virtual machine (VM). See "Recovery to a VM is supported for these virtual hosts:" below for supported virtual hosts. • Recover a failed VM from a file-level backup to a VM or to a physical asset: – If you have opted to protect a VM by installing the Windows agent and running file-level backups, you can recover the failed VM by using unified BMR. – Recovering a VM to a physical asset is supported for Windows 7/ Server 2008 R2 and higher. – Recovering a VM to a VM is supported for the hosts listed below in "Recovery to a VM is supported for these virtual hosts:". If you are also protecting the VM with host-level backups, it is easier to recover from the host-level backup as described in "Recovering a virtual machine" on page 796. • Recovery to a VM is supported for these virtual hosts: – VMware ESX/ESXi versions 5.0 and higher that are listed in the Compatibility and Interoperability Matrix. – All versions of Hyper-V listed in the Compatibility and Interoperability Matrix. – All versions of Citrix XenServer listed in the Compatibility and Interoperability

Requirement	Description
	Matrix. The virtual host must support the operating system (OS) of the Windows asset you are recovering. (See the VMware, Microsoft, or Citrix documentation for details. For Hyper-V, see this Microsoft article: Should I create a generation 1 or 2 virtual machine in Hyper-V?) For example, you cannot recover Windows 2016 to ESXi 5.1 or Hyper-V 2008 R2.

Prerequisites for image-level backups

The following requirements must be met to recover from an image-level backup.

Requirement	Description
Operating systems	Recovery to identical physical hardware and virtual machines is supported for the operating systems listed below. Recovery to dissimilar hardware is NOT supported. Supported client operating systems: • Windows 7 with SP1, 64-bit only • Windows 8, 64-bit only • Windows 8.1, 64-bit only • Windows 10, 64-bit only • Windows 11, 64-bit only Supported server operating systems: • Windows 2008 R2 with SP1, 64-bit only • Windows 2012, 64-bit only • Windows 2012 R2, 64-bit only • Windows 2016, 64-bit only • Windows 2019, 64-bit only • Windows 2022, 64-bit
An eligible image-level backup	The backup used for recovery must meet these requirements: • It is successful. • It is a full or incremental image-level backup that contains all critical system volumes. **Notes:**

Requirement	Description
	• By default, image-level backups include all system information needed for unified bare metal recovery. If you opt to exclude volumes from backup, use care not to exclude the boot and critical system (OS) volumes. • The recovered asset is created based on the backup you select. Volumes that were excluded from backup are not recovered. • When you recover the entire asset, any existing data on the target is overwritten or deleted. Volumes on the target disk that were excluded from backup may also be overwritten. • For SQL, the master, model, and msdb system databases must also be present in the image-level backup of the Windows asset. (These are included by default. If you want the recovered asset to include a hosted SQL application, use care not to exclude these system databases from the image-level backup.)
Unified BMR ISO image	For the recovery, you must use the 64-bit unified BMR ISO image provided on the your Unitrends appliance. The ISO contains WinPE (a minimal version of Windows used for installations) and the Unitrends Unified Bare Metal Recovery wizard that guides you through the recovery. To prepare for DR, it is recommended that you do the following: • Create a bootable CD of the ISO and store it in a safe place (so that you can quickly recover to a physical machine target). • Save the ISO to your hypervisor (so you can quickly recover to a virtual machine target). For details, see "Access the unified bare metal recovery ISO image" on page 1218.
Firmware interface type	Supported for BIOS- and UEFI-based assets. The firmware interface type (BIOS or UEFI) of the recovery target machine must match that of the failed asset.
Disk configuration	• GPT disks are supported. • Dynamic disks are not supported. • iSCSI disks are not supported. Recover the critical (non-iSCSI) volumes as described in "Performing unified bare metal recovery" on page 1218. Once the critical volumes have been restored, recover data on the iSCSI volumes as described in "Recovering files from Windows image-level backups" on page 1033.
Software RAID volumes	Software RAID configurations are not supported.

Requirement	Description
Network adapter	Wireless network adapters cannot be used for the recovery.
Drivers	Drivers needed for unified BMR are determined by the operating system and hardware of the asset you are recovering and whether you are recovering to a physical or virtual machine. You might need to add drivers during different stages of the recovery. • Loading WinPE drivers for unified bare metal recovery — The Unified BMR wizard uses WinPE 10.0 for the recovery. WinPE 10.0 is based on Windows 11 or Server 2022. If WinPE 10.0 cannot detect a network adapter or storage disks, you must load Windows 11 or Server 2022 drivers into WinPE during recovery. • Recovering to a virtual machine — After the critical drives have been recovered, you must inject ESX, Hyper-V, or XenServer guest storage drivers. These drivers are included in the unified BMR ISO image.
Processor features on the recovery target	WinPE requires that these processor features are enabled on the recovery target machine: NX, PAE, and SSE2. Ensure that these features are enabled on the recovery target machine before booting from the ISO image. For instructions, see Error 0x0000005D (Unsupported Processor) when booting Release 7.4 Integrated Bare Metal restore media. Machines that do not have these processor features cannot be used for the recovery.
Disk space on recovery target	Make sure the target machine has enough disk space for the recovery. To recover from an image-level backup, adhere to these requirements: • You must recover to destination disks that are the same size as the original disks or larger. Recovery fails if the disk does not have enough space. • Any existing data on the destination disks is overwritten or deleted during the recovery, even if the disk has more than enough space. Before performing a recovery, make sure you have additional copies of any data on the destination disk. • Deduplicated volumes - Data is recovered in its non-deduplicated form. Ensure that the target disk has enough capacity to house this non-deduplicated data.
Supported recovery scenarios	Use the procedures in "Implementing Windows unified bare metal protection" on page 1210 for the following recovery scenarios: • Recover to the same physical hardware as the failed asset. • Recover a failed physical asset to a virtual machine (VM). Recovery to a VM is supported for these virtual hosts: – VMware ESX/ESXi versions 5.0 and higher that are listed in the Compatibility and Interoperability Matrix.

Requirement	Description
	– All versions of Hyper-V listed in the Compatibility and Interoperability Matrix. – All versions of Citrix XenServer listed in the Compatibility and Interoperability Matrix. The virtual host must support the operating system (OS) of the Windows asset you are recovering. (See the VMware, Microsoft, or Citrix documentation for details.) For example, you cannot recover Windows 2016 to ESXi 5.1 or Hyper-V 2008 R2.

Performing unified bare metal recovery

Use the following procedures to perform unified bare metal recovery. For a successful recovery, be sure to run the procedures in order. Before you start, it is recommended that you read "Implementing Windows unified bare metal protection" on page 1210 for prerequisites and supported recovery scenarios.

Step 1: "Access the unified bare metal recovery ISO image"

Step 2: "Prepare the recovery target machine" on page 1219

Step 3: "Run the Unified Bare Metal Recovery wizard" on page 1227

Step 4: "Complete the unified bare metal recovery" on page 1244

Step 1: Access the unified bare metal recovery ISO image

For the recovery, you must use the 32-bit or 64-bit unified BMR ISO image provided on the Unitrends backup appliance.

To access the unified bare metal recovery ISO images

1 Mount a working asset to the *virtual_failover* share on your Unitrends appliance by entering **\\<ApplianceIP>\virtual_failover** in the Windows File Explorer.

2 Download *winbm.iso* (64-bit assets) or *winbm32.iso* (32-bit assets).

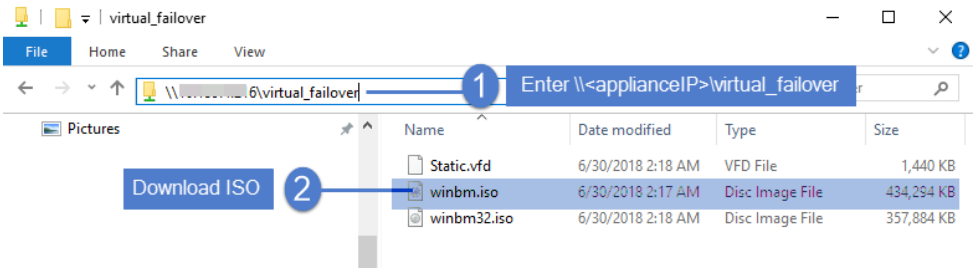

3 Do one of the following:

- If recovering to a physical machine, burn the ISO to CD. Refer to the documentation for the burner you are using to walk you through the process of creating a bootable CD from an ISO image, which is not the same as burning an ISO image to a CD.

- If recovering to a virtual machine, save the ISO in a location that you can access from your hypervisor.

4 Proceed to "Step 2: Prepare the recovery target machine".

Step 2: Prepare the recovery target machine

You can recover a failed asset to a physical or virtual machine (VM). Use one of the following procedures to prepare the recovery target machine:

- "To prepare a physical machine for unified bare metal recovery"
- "To prepare a VM for unified bare metal recovery"

To prepare a physical machine for unified bare metal recovery

1 Ensure that the target machine meets the following requirements:

Requirement	Description
Processor features	WinPE requires these processor features to be enabled on the recovery target machine: NX, PAE, and SSE2. Be sure these features are enabled before you boot from the ISO image. For instructions, see Error 0x0000005D (Unsupported Processor) when booting Release 7.4 Integrated Bare Metal restore media. Machines that do not have these processor features cannot be used for the recovery.
Firmware interface type	Ensure that the machine's firmware interface type (BIOS or UEFI) matches the firmware interface type of the failed asset. The recovered asset will not boot if you recover to a dissimilar interface type.
Disk space	Make sure the machine has enough disk space for the recovery. See these rows below for requirements by backup type: "File-level backup" and "Image-level backup".
File-level backup	To recover from a file-level backup, adhere to these requirements: • The recovery target can have smaller disks than the failed asset, but the recovery fails if the disks do not have enough space for the data on the critical volumes. • After an asset has failed, there is no way to determine the size of its critical volumes. You can determine the size of an asset's full backup by viewing the backup details in the "Backup History report". However, the size of the critical volumes will be smaller than the full backup if it also contains non-critical volumes. If you are unsure about the size of the critical volumes, it is recommended that you recover to destination disks that are the same size as the original disks or larger. • If you are recovering to new disks, any existing data on the destination disks is

Requirement	Description
	overwritten or deleted during the recovery, even if the disks have more than enough space. Before performing a recovery, make sure you have additional copies of any data on the destination disks. • If you are recovering to the original disk, only the recovered volumes are overwritten. Other volumes on the original disk are not impacted by the recovery.
Image-level backup	To recover from an image-level backup, adhere to these requirements: • You must recover to destination disks that are the same size as the original disks or larger. Recovery fails if the disk does not have enough space. • Any existing data on the destination disks is overwritten or deleted during the recovery, even if the disk has more than enough space. Before performing a recovery, make sure you have additional copies of any data on the destination disk. • Deduplicated volumes - Data is recovered in its non-deduplicated form. Ensure that the target disk has enough capacity to house this non-deduplicated data.

2 Load the disk containing the bootable ISO image into the machine's CD/DVD drive. (For details on creating the ISO disk, see "Step 1: Access the unified bare metal recovery ISO image".)

3 Proceed to "Step 3: Run the Unified Bare Metal Recovery wizard" on page 1227.

To prepare a VM for unified bare metal recovery

A VMware, Hyper-V, or XenServer virtual machine can be used as the recovery target. You can create a new VM or edit the settings of an existing VM.

1 Ensure that the target VM meets the following requirements:

Requirement	Description
Firmware interface type	Ensure the VM's firmware interface type (or UEFI) matches the firmware interface type of the failed asset. The recovered asset will not boot if you recover to a dissimilar interface type.
Memory	Add enough memory to satisfy Microsoft's support guidelines for the operating system being recovered. The recovery requires at least 1 GB of memory.
Disk space	The VM's virtual hard disks must have adequate space for the recovery. See these rows below for requirements by backup type: "File-level backup" and "Image-level

Requirement	Description
	backup".
File-level backup	To recover from a file-level backup, adhere to these requirements: • The recovery target can have smaller disks than the failed asset, but the recovery fails if the disks do not have enough space for the data on the critical volumes. • After an asset has failed, there is no way to determine the size of its critical volumes. You can determine the size of an asset's full backup by viewing the backup details in the "Backup History report". However, the size of the critical volumes will be smaller than the full backup if it also contains non-critical volumes. If you are unsure about the size of the critical volumes, it is recommended that you recover to destination disks that are the same size as the original disks or larger. • If you are recovering to new disks, any existing data on the destination disks is overwritten or deleted during the recovery, even if the disks have more than enough space. Before performing a recovery, make sure you have additional copies of any data on the destination disks. • If you are recovering to the original disk, only the recovered volumes are overwritten. Other volumes on the original disk are not impacted by the recovery.
Image-level backup	To recover from an image-level backup, adhere to these requirements: • You must recover to destination disks that are the same size as the original disks or larger. Recovery fails if the disk does not have enough space. • Any existing data on the destination disks is overwritten or deleted during the recovery, even if the disk has more than enough space. Volumes on the target disk that were excluded from backup may also be overwritten. Before performing a recovery, make sure you have additional copies of any data on the destination disk.
Secure Boot (Windows 10 and 11, Server 2016, 2019, and 2022)	The Windows Secure Boot option is not supported. Ensure that Secure Boot is set to OFF before you perform the bare metal recovery.
VMware NIC	For recovery to a VMware virtual machine, make sure you are using the E1000 NIC. This NIC is only required during the recovery. After rebooting the recovered VM, you can use a different NIC.

Requirement	Description
	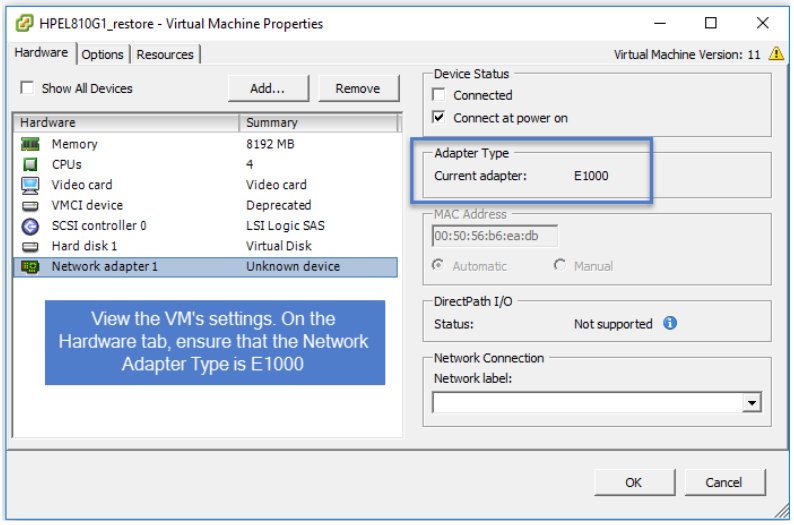

2 Edit the VM settings to boot from the bare metal ISO. (For details on accessing the ISO, see "Step 1: Access the unified bare metal recovery ISO image".)

Procedures vary by hypervisor. A VMware example is given here:

- Power off the VM.
- Right-click the VM and select **Edit Settings**:

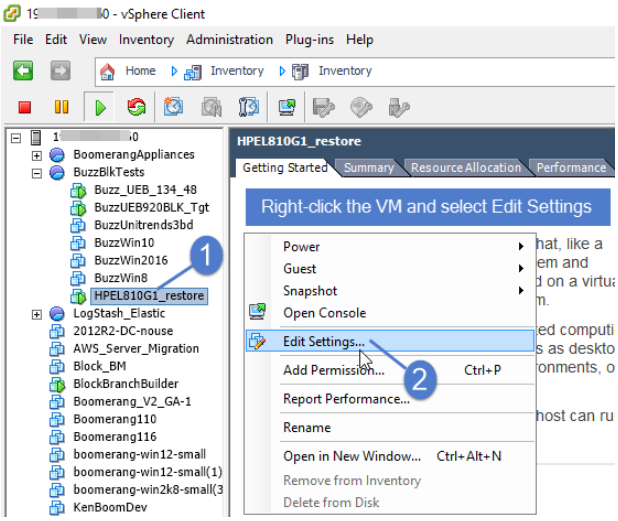

- On the Hardware tab, click **Add**, select **CD/DVD Drive** and click **Next**:

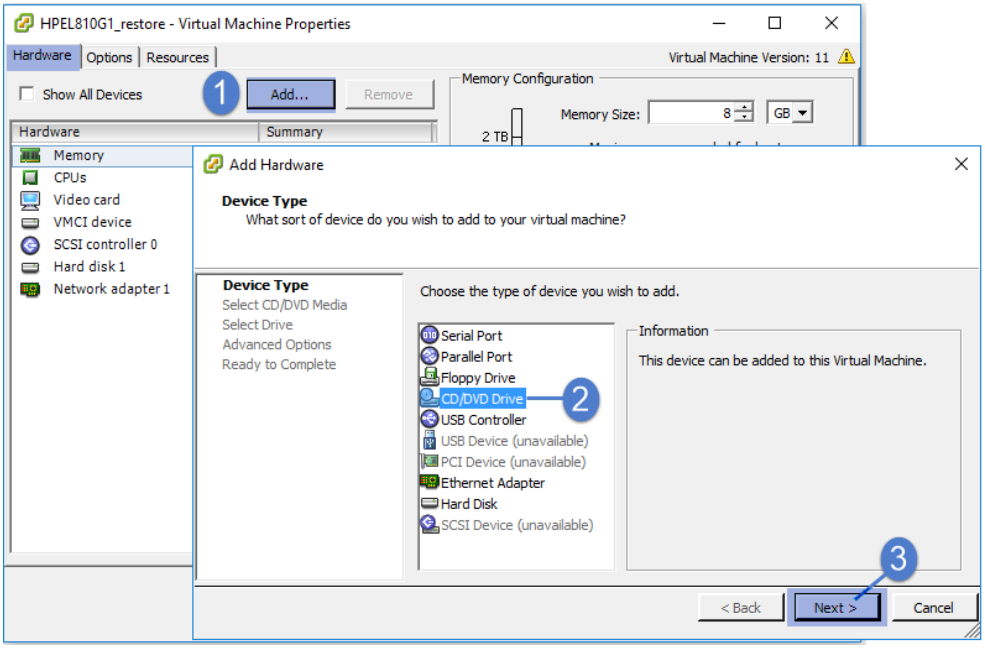

- Select **Use ISO image** and click **Next**:

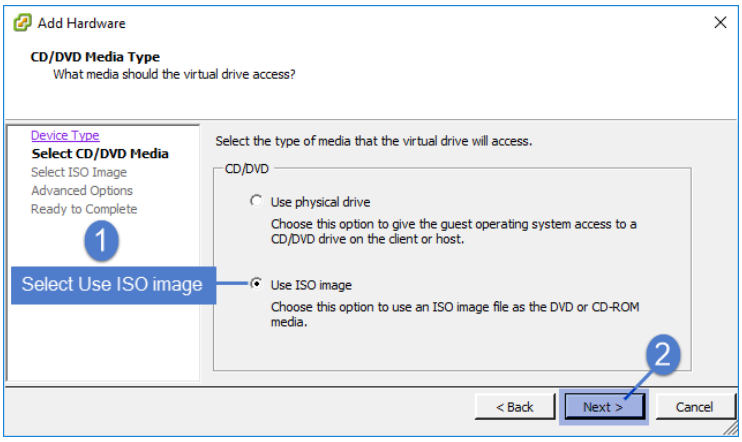

- Browse or enter a path to select the unified bare metal recovery ISO, check the **Connect at power on** box, and click **Next**:

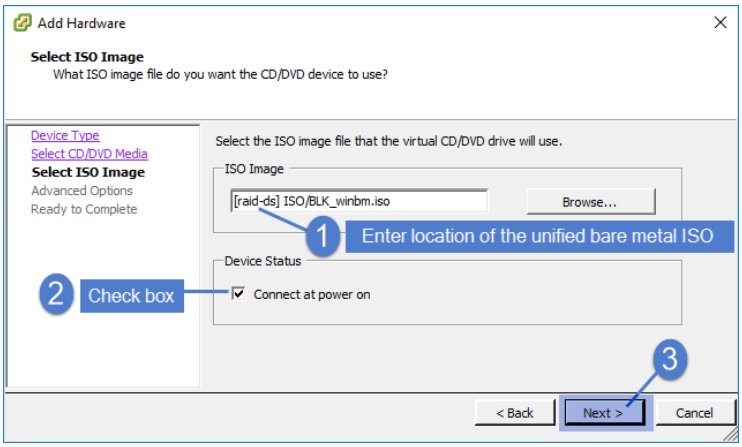

- Click **Next** to continue:

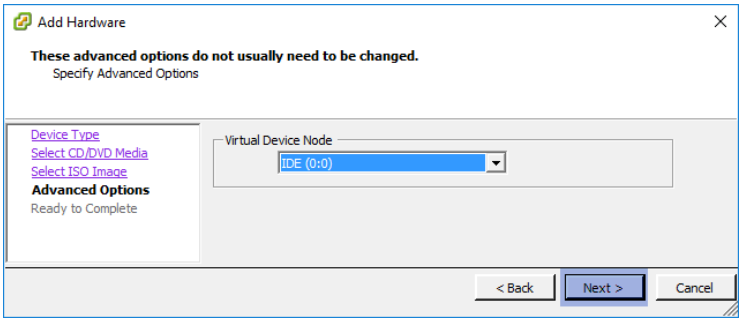

- Click **Finish** to add the device:

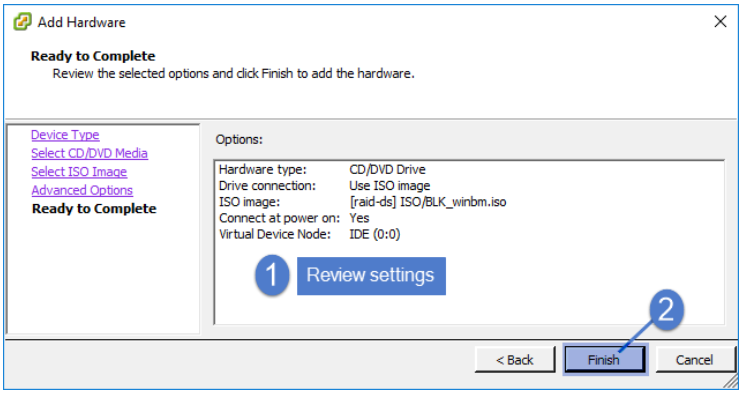

- Click **OK** to save:

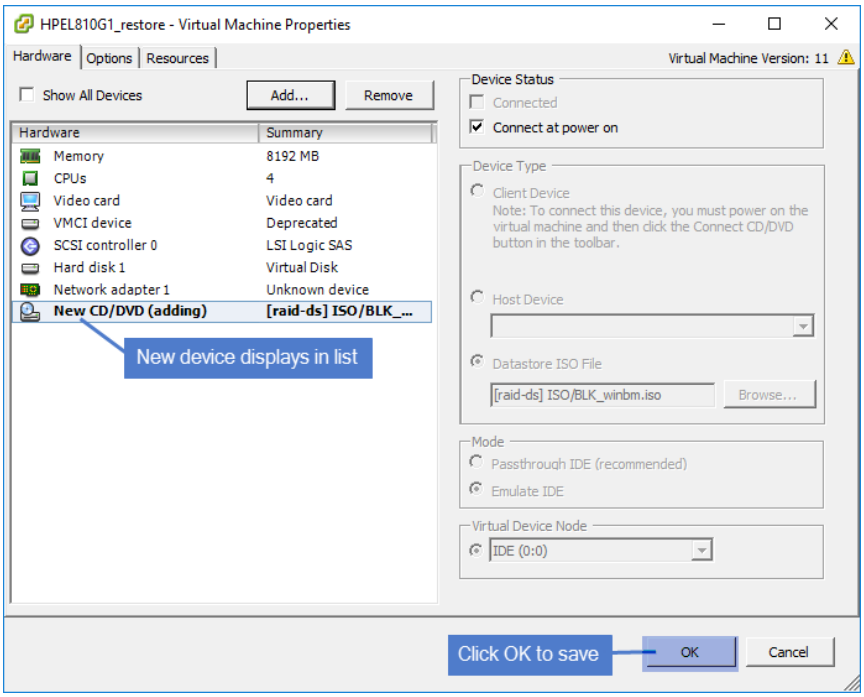

- Do one of the following:
  - If the target VM does NOT have an operating system installed, proceed to "Step 3: Run the Unified Bare Metal Recovery wizard".

  OR

  - If the target VM has an operating system installed, you must boot into BIOS and change the boot order to force the VM to boot from the bare metal ISO. See the remaining steps in this procedure for a VMware example.

    Note: The remaining steps provide an example of how to boot into BIOS and change the VM's boot order. This example is run on VMware. Details will vary by hypervisor.

- Ensure that the VM is powered off.
- Right-click the VM and select **Edit Settings**.
- On the **Options** tab, select **Boot Options**, check the **Force BIOS Setup** box, then click **OK**:

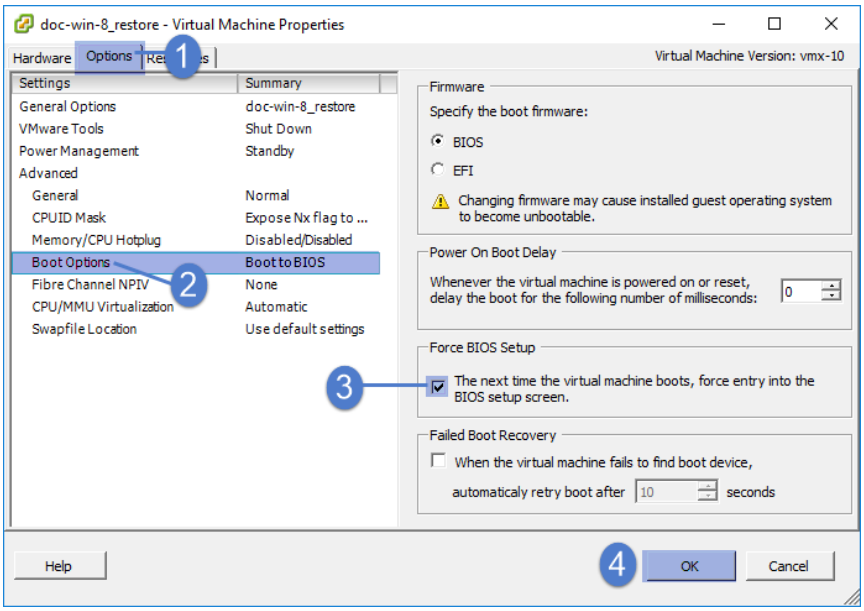

- Power on the VM. The VM boots into BIOS.
- Launch the VM console.
- On the Boot menu in the BIOS Setup Utility, select **CD-ROM Drive** and press **+** to move it to the top of the list. Press **F10** to save.
- Press **Enter** to save and exit.

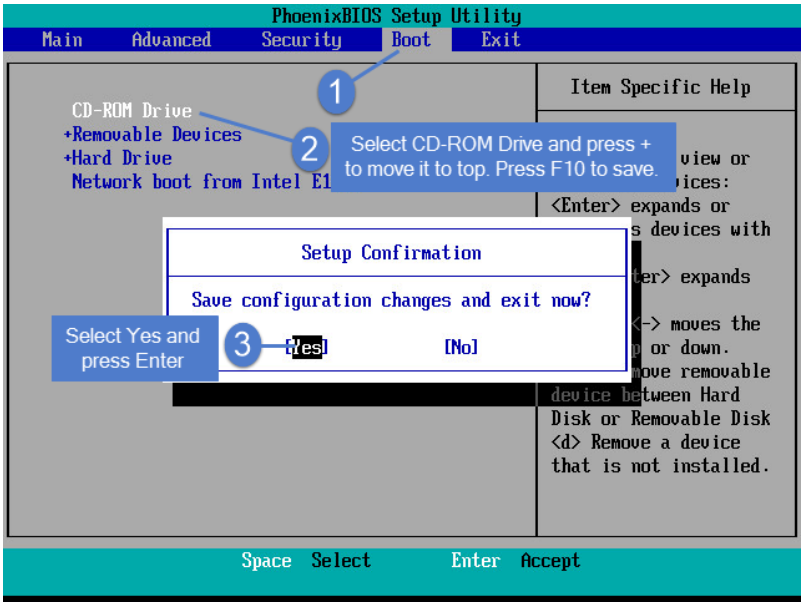

- Power off the VM.

3 Proceed to "Step 3: Run the Unified Bare Metal Recovery wizard".

Step 3: Run the Unified Bare Metal Recovery wizard

Perform the recovery by using the Unified Bare Metal Recovery (UBMR) wizard. See the following topics for details:

- "Considerations for performing a test unified bare metal recovery"
- "Perform the recovery by running the Unified Bare Metal Recovery wizard"

Considerations for performing a test unified bare metal recovery

Before an asset fails, you can perform a test unified bare metal recovery without impacting the original asset. As long as you assign the recovered asset a unique IP address and rename it, the test recovery does not cause any network conflicts with the original asset. IP conflicts will occur if you perform a test recovery and do not assign the recovered asset a unique name and IP address.

To perform the test recovery, use the "Perform the recovery by running the Unified Bare Metal Recovery wizard" procedures.

Perform the recovery by running the Unified Bare Metal Recovery wizard

Use the procedures in this section to perform the recovery by running the UBMR wizard. These procedures apply to any supported target machine (a physical machine, a virtual machine, dissimilar hardware, etc.).

For a successful recovery, you must run the procedures in the following order:

- "Step 1: Boot into WinPE and set up the local environment"
- "Step 2: Select an appliance and the backup to recover"
- "Step 3: Map drives and volumes"
- "Step 4: Start the recovery"

Notes:

- You must prepare the recovery target machine before running the UBMR wizard. (See "Step 2: Prepare the recovery target machine" for details.)
- If you exit the UBMR wizard before you are finished with the recovery, you are taken to a command window. To return to the wizard from this window, run the following command: **`z:\pcpb\Restore.exe`**.

Step 1: Boot into WinPE and set up the local environment

These steps ensure that the recovery target machine can communicate with the appliance that is storing the backup or hot backup copy you will use for the recovery.

1 Boot the recovery target machine from the bare metal ISO image. The machine boots into WinPE and launches the Unified Bare Metal Recovery wizard.

Notes:

- If you see a message stating that you must set up networking to continue or that no disks are detected on the local system, you might need to load drivers into WinPE. Click **OK** to allow the boot to continue. You can load drivers later in this procedure.

- If you are recovering to a VM and the Windows login screen displays, you must boot into BIOS and change the boot order to force the VM to boot from the bare metal ISO. For a VMware example, see "To prepare a VM for unified bare metal recovery".

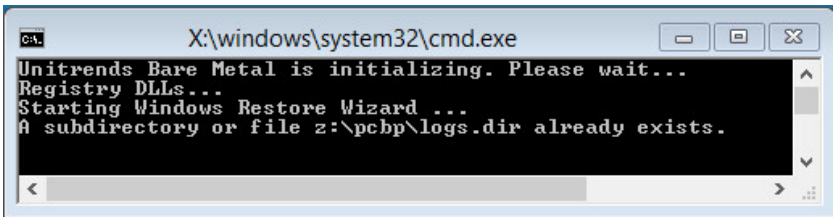

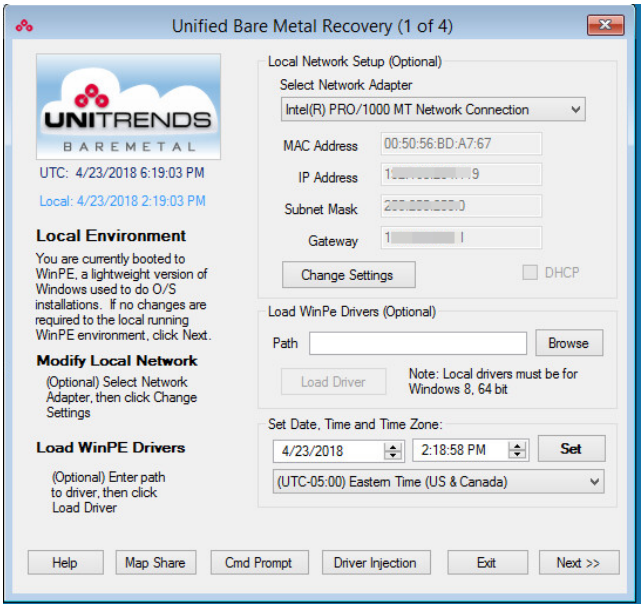

2 Select a Network Adapter. If the machine has more than one adapter, the default adapter displays first.

If a network adapter does not display, WinPE cannot detect one. To resolve this issue, perform the following:

- Ensure that the network cable is plugged in to an active port.
- If the adapter is connected to the network and WinPE is unable to detect it, you must load a network driver into WinPE. In the Load WinPE Drivers section, specify a driver by browsing or entering a path, then click **Load Driver**.

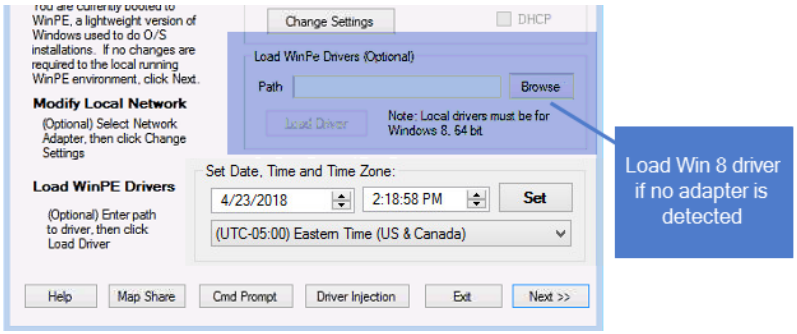

3 Review network settings and modify as needed.

- If DHCP is configured for your network, network settings are assigned automatically.
- If DHCP is not configured, or if you want to configure network settings for the target machine manually, click **Change Settings**. Then enter a unique IP address, the Subnet Mask, and the Gateway. Network settings do not need to match those of the original asset. The only requirement is that the machine can communicate with the appliance that is storing the backup you will use for recovery.

> **Note:** The network settings that you configure during this step are used only for the recovery. They are not applied to the network adapter when you boot into the recovered operating system. Before connecting the recovered asset to your network, you will reconfigure the asset's network settings.

4 Select the time zone of the Unitrends appliance storing the backup or hot backup copy that you will use for recovery.

5 Click **Next**.

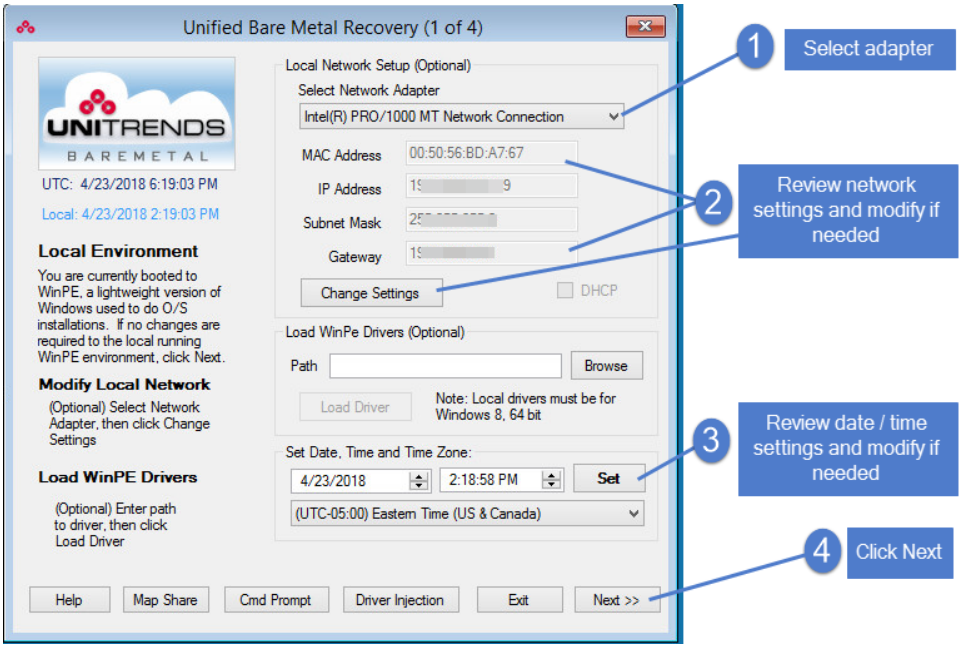

6 Proceed to "Step 2: Select an appliance and the backup to recover".

Step 2: Select an appliance and the backup to recover

1 For Backup Source, specify the Unitrends appliance that is storing the backup or hot backup copy you will use for the recovery. You can either:

- Select the appliance from the Backup Source list. This list contains all appliances on the same subnet as the recovery target machine.
- Enter the IP address of an appliance on a different subnet.

2 Enter the Username and Password credentials of a UI user account that has superuser or administrator privileges. Click **Login**.

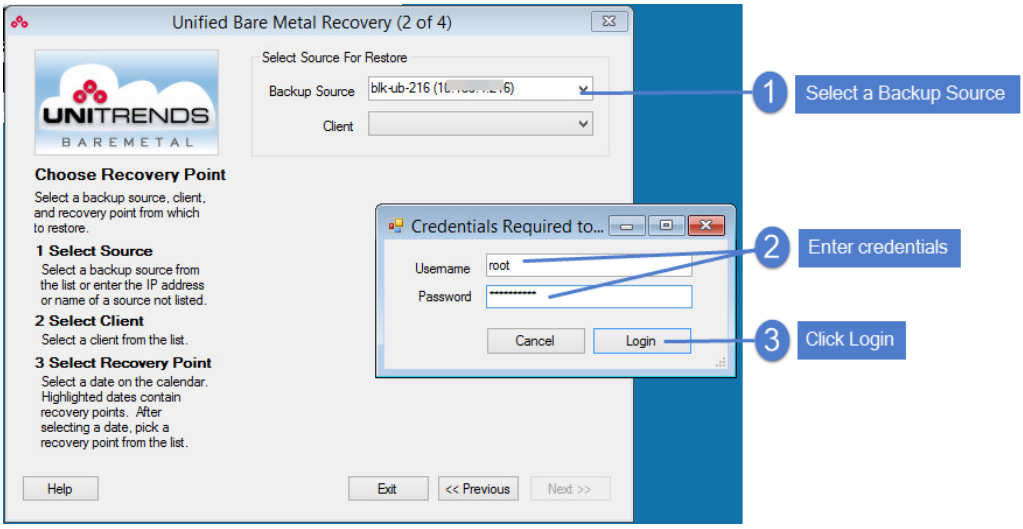

3 Select the Windows asset from the Client list.

> **Note:** Assets that have eligible file-level and/or image-level backups display in the list. (See "An eligible file-level backup" on page 1211 or "An eligible image-level backup" on page 1215 for details.)

4 Select a recovery point in the calendar.

- Days with a backup display in bold on the calendar.
- If multiple backups exist on a given day, the different times for these backups display in the Recovery Point list.
- If the recovery times do not match the backup times on the Unitrends appliance, verify that you selected the time zone of the appliance that you are using for the recovery (in "Step 1: Boot into WinPE and set up the local environment").
- The selected Recovery Point shows the backup date/time, type (*Img* for image-level or *File* for file-level), and mode (*Full* or *Incr* for incremental).

5 Click **Next**.

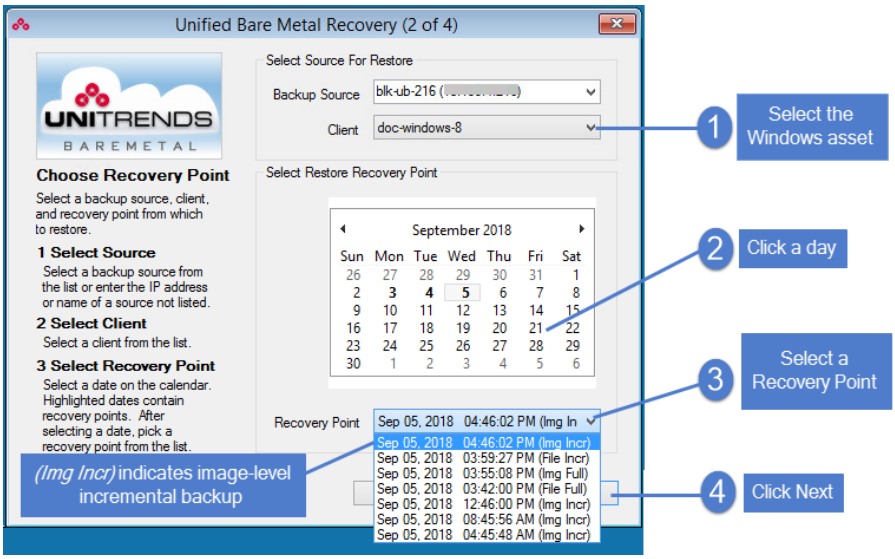

6 Proceed to "Step 3: Map drives and volumes".

Step 3: Map drives and volumes

After selecting an appliance, asset, and recovery point, you must map the failed asset's disks and/or volumes to disks on the recovery target machine. This procedure differs by backup type. Examples are given for file-level and image-level recovery.

1 Map drives and/or volumes as described below in "File-level backup example" or "Image-level backup example" on page 1233.

2 Click **Next**.

3 Click **Yes** to confirm.

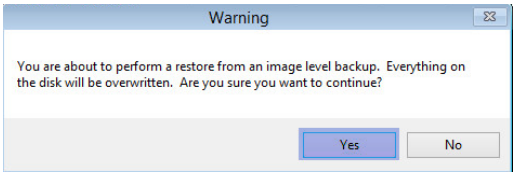

4 Proceed to "Step 4: Start the recovery".

File-level backup example

On the Drive/Volume Mapping screen, select the critical volumes to recover by highlighting the disk(s) that contain them.

- If the recovery target is similar to the original machine, the wizard automatically maps the backed up volumes to the destination disk.
- Disks and critical volumes that can be recovered from the backup display in the Source Disks/Volumes

File-level backup example

area at the top of the screen. For most newer versions of Windows, two volumes display: a boot volume and a system volume. For most older versions, the boot and system files are on a single volume. Some assets require additional critical volumes for booting, and these also display.

- Non-critical volumes do not display because they cannot be recovered through the UBMR wizard. (You can recover these later in "Step 4: Complete the unified bare metal recovery".)
- The wizard does not require that you select all critical volumes to proceed, but for the recovery to succeed, you must select the volumes needed to boot the operating system. If you are unsure which are required, select all critical volumes.
- To perform the mapping manually, uncheck the **Restore to original system (automatic mapping)** box, then manually add volumes to the Destination Disks area at the bottom of the screen by highlighting the destination disk and clicking **Add**. (To remove a volume, select the volume and click **Remove**.)

Consider the following when selecting a destination disk:

- If you are recovering to a new disk, any existing volumes on the disk are deleted during the recovery and new volumes are created.
- If you are recovering to the original disk(s), only the recovered volumes are overwritten. Other volumes on the original disk are not impacted by the recovery.
- It is recommended that the destination disk be the same size as the original disk or larger to ensure that there is enough space for the recovery. You can use a smaller destination disk as long as it is large enough for the critical volumes; otherwise the recovery runs until the disk is full, and then it fails.
- Volumes are assigned letters during recovery that do not necessarily match the numbers from the original disk.
- For dissimilar recovery of multi-boot configured BIOS servers, the boot and system volumes must be recovered to the same disk numbers used on the original server.

File-level backup example

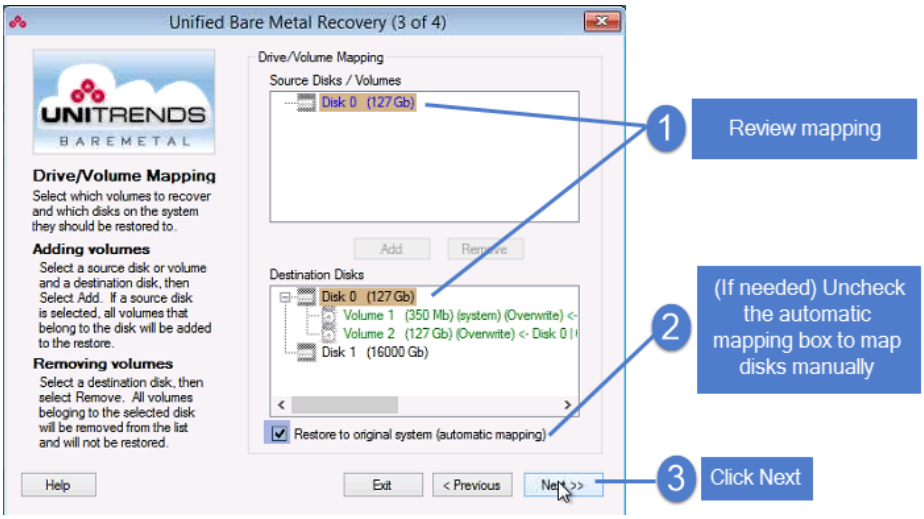

Image-level backup example

On the Disk Mapping screen, select the disks to recover.

- If the recovery target is similar to the original machine, the wizard automatically maps the backed up disks to the destination disks.
- You must select system critical disks. Other disks are optional.

  Note: It is recommended to select all disks. The only way to recover all data on a disk is by selecting the disk during this mapping step. After completing the bare metal recovery, you can opt to recover individual files from the backup (as described in "Recovering files from Windows image-level backups" on page 1033), but you cannot recover entire disks or volumes.

- The wizard does not require that you select all critical disks to proceed, but for the recovery to succeed, you must select the disks needed to boot the operating system. If you are unsure which are required, select all disks.
- To perform the mapping manually, uncheck the **Restore to original system (automatic mapping)** box, then manually add disks to the Destination Disks area at the bottom of the screen by highlighting the destination disk and clicking **Add**. (To remove a disk, select the disk and click **Remove**.)

Consider the following when selecting a destination disk:

- Any existing data on the disk is overwritten or deleted during the recovery, even if the disk has more than enough space. Before performing a recovery, make sure you have additional copies of any data on the destination disk.

Image-level backup example

- You must recover to destination disks that are the same size as the original disks or larger. Recovery fails if the disk does not have enough space.
- Disks are assigned numbers during recovery that may not match the numbers from the original disks.
- Volumes are assigned letters during recovery that may not match the letters from the original volumes.

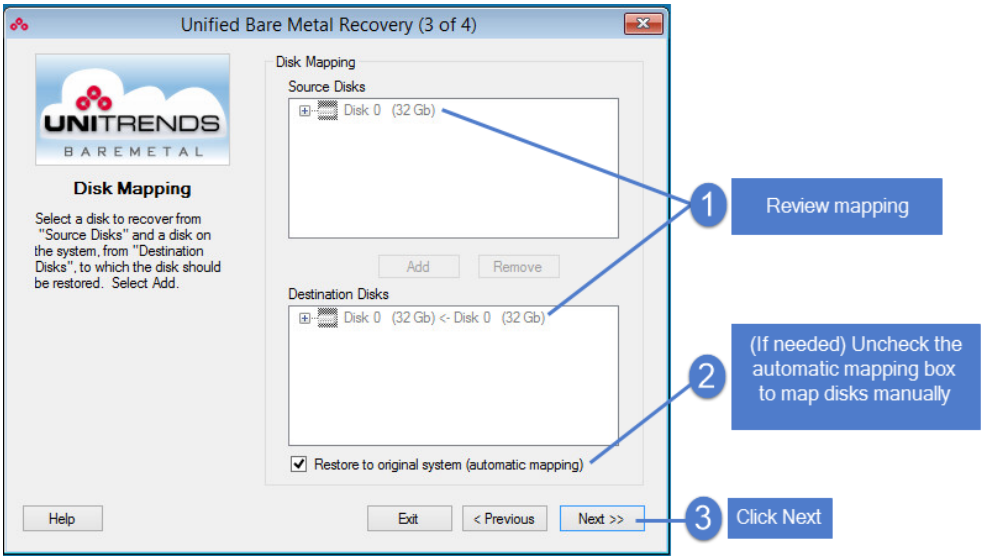

Step 4: Start the recovery

Run one of the following procedures to start the recovery:

- "To recover to a physical machine target"
- "To recover to a virtual machine target" on page 1239

To recover to a physical machine target

1 Click **Restore**. Proceed to the next step when the job completes.

> **Note:** The recovery must complete within 24 hours. If many jobs are running on the appliance, move the UBMR job ahead in the queue by canceling running jobs or pausing queued jobs. For details, see "Managing active jobs" on page 607.

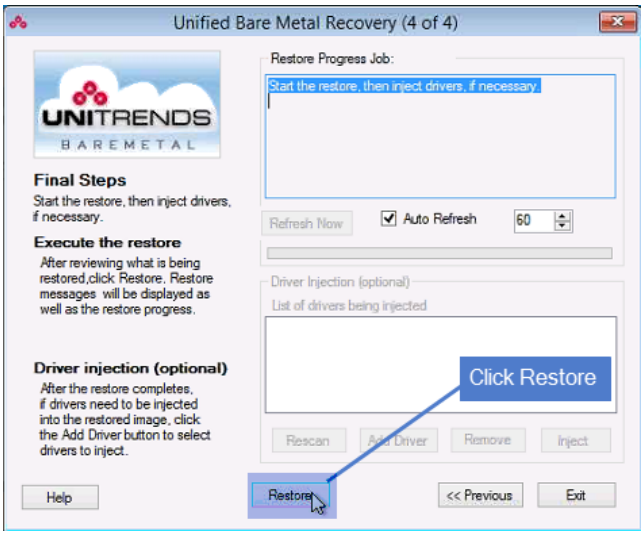

The recovery can take some time. Monitor job progress in the Restore Progress Job area at the top of the screen. Sample messages are given here:

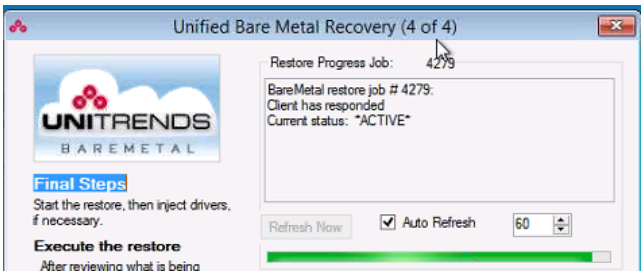

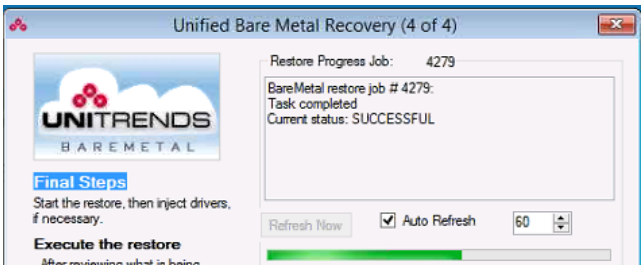

2 Click **OK** to close the Status message:

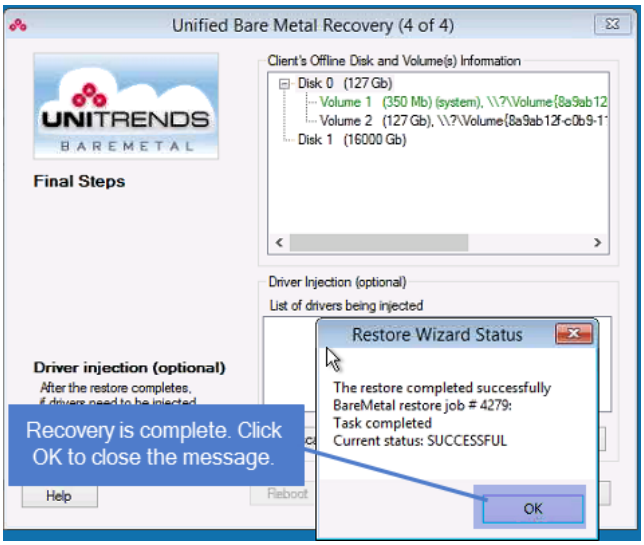

3 Do one of the following:

- If you recovered to identical hardware, remove the bootable ISO CD, then click **Reboot**. After the target machine reboots, proceed to "Step 4: Complete the unified bare metal recovery" on page 1244. (If the target fails to reboot, you need to inject drivers. To inject drivers, boot the machine from the ISO disk, then proceed to step 4.)

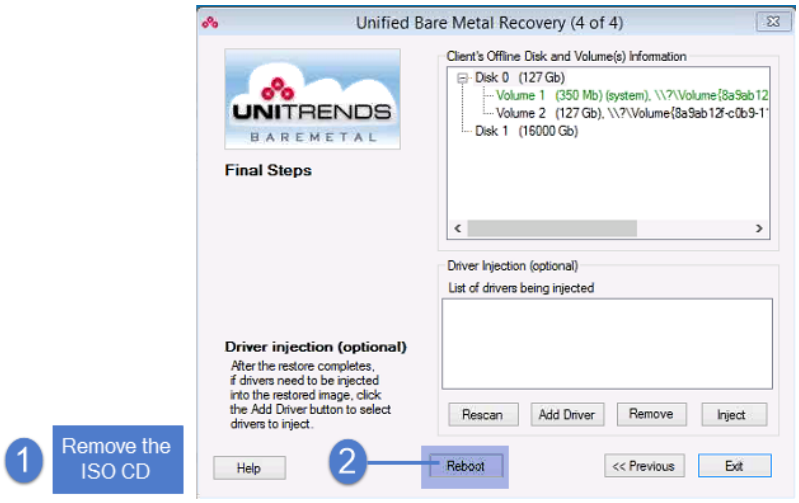

- If you recovered to dissimilar hardware, you must inject storage drivers before rebooting. Proceed to step 4.

4 Inject storage drivers by doing these steps:

- In the UBMR wizard, click one of the following:
  - If you have booted from the bare metal ISO after your recovered operating system failed to boot, you are returned to the first screen of the wizard. Click the **Driver Injection** button on this screen.

    Note: Do not use the Load Driver button under Load WinPE Drivers. This loads drivers into WinPE. Because your operating system has already been recovered, you must inject drivers into the recovered operating system instead.

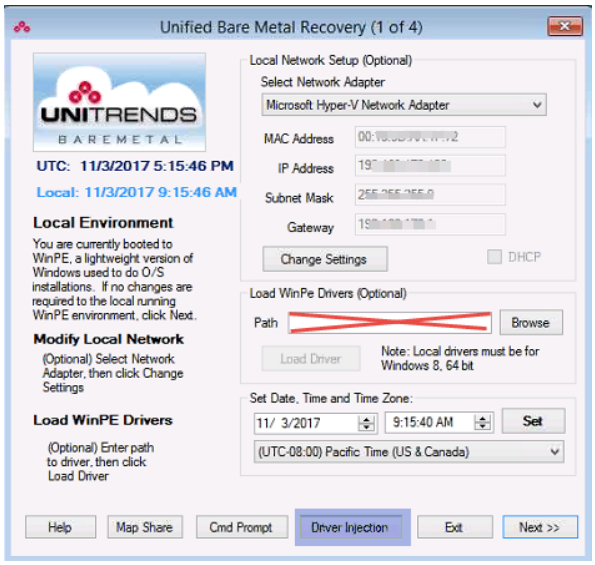

  - If you did not reboot, click **Add Driver** on the screen you used to recover the critical volumes.

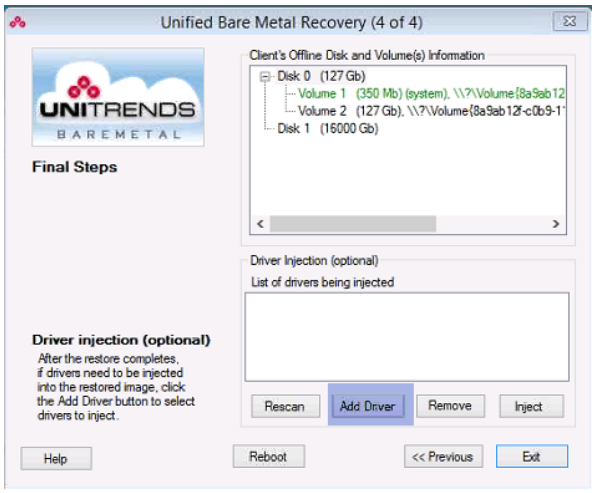

- Browse to select the driver and click **Done** to add it to the list of drivers to inject. Repeat as needed to add more drivers to the list.
- Highlight the volume containing the operating system files in the asset's Offline Disk and Volume(s) Information area. Then highlight a driver and click **Inject**. Repeat this step as needed to inject all the necessary drivers.

5 When you receive a message stating that the driver injection was successful, you are ready to reboot the recovered asset. Do one of the following:

- If you are NOT recovering an Active Directory server, remove the bootable ISO CD, then click **Reboot**.

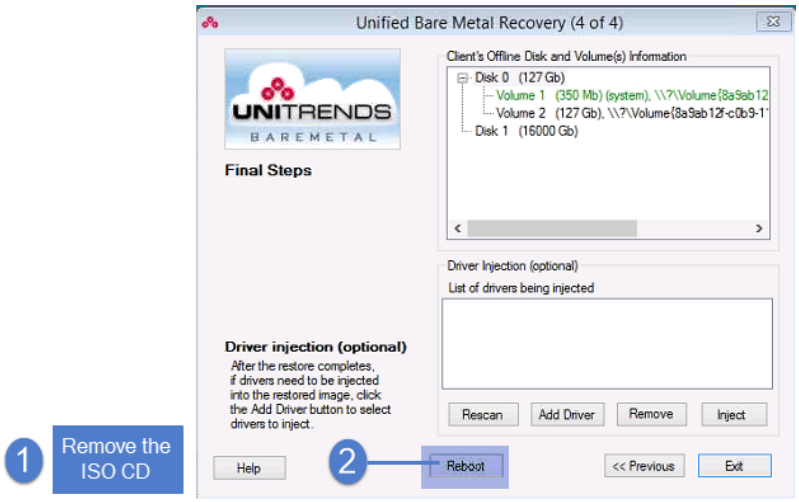

- If you ARE recovering an Active Directory server, you must start the target machine in Directory Services Restore Mode by doing these steps:
  - Remove the bootable ISO CD.
  - Disconnect the server from the network (to ensure the server does not start in normal mode).

- Start the server in Directory Services Restore Mode.
- Do one of the following:
  - If a file-level backup is available, connect the server to the network, then recover the file-level backup as described in "To recover an entire file-level backup by using the Backup Catalog" on page 945.
  - If you do not have a file-level backup, set the **database restored from backup** registry value to **1**.
- Restart the domain controller in normal mode.

6 After the recovered asset boots, proceed to "Step 4: Complete the unified bare metal recovery" on page 1244.

> **Note:** If the asset fails to boot, you might need to add additional drivers.

To recover to a virtual machine target

1 Click **Restore**. Proceed to the next step when the job completes.

> **Note:** The recovery must complete within 24 hours. If many jobs are running on the appliance, move the UBMR job ahead in the queue by canceling running jobs or pausing queued jobs. For details, see "Managing active jobs" on page 607.

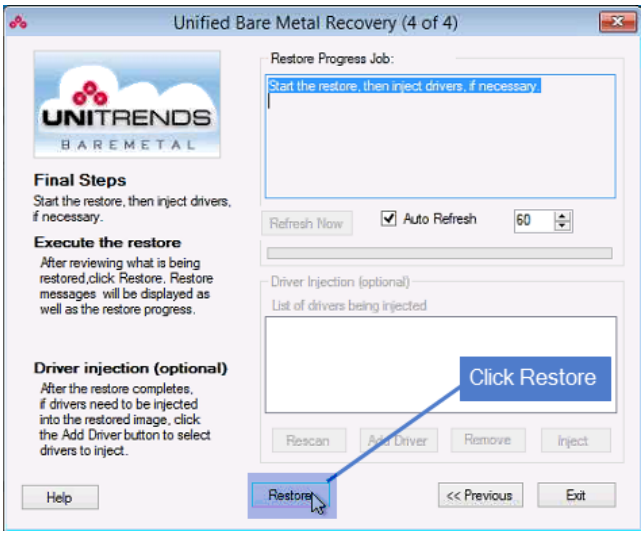

The recovery can take some time. Monitor job progress in the Restore Progress Job area at the top of the screen. Sample messages are given here:

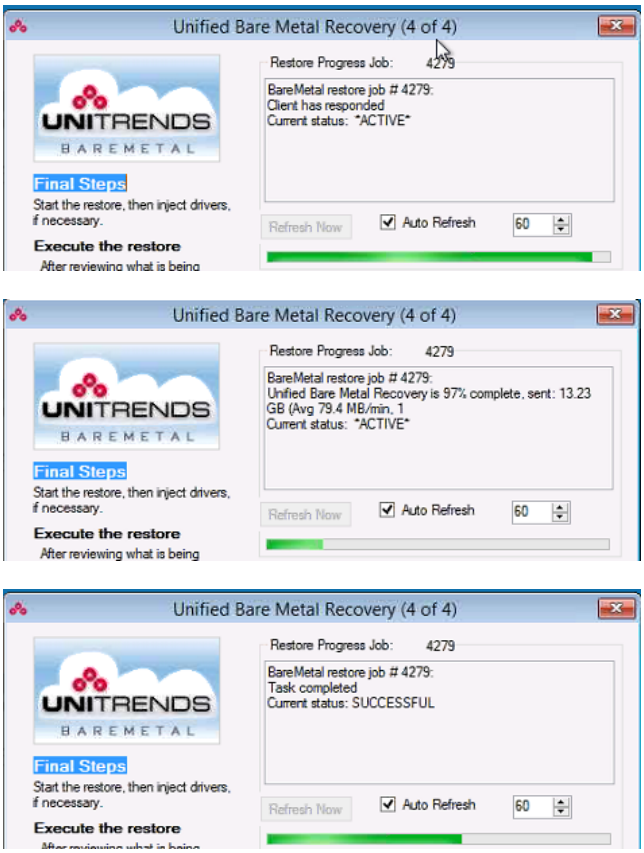

2 Click **OK** to close the Status message:

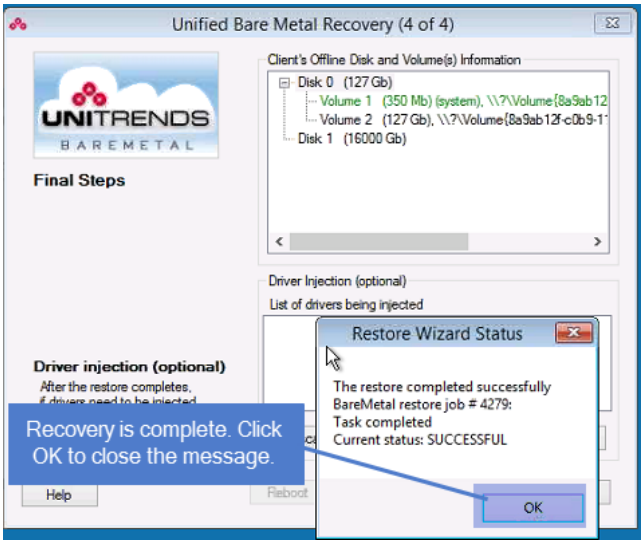

3 In the Restore dialog, click **No**. (You must inject a guest storage driver before you can reboot the recovered VM.)

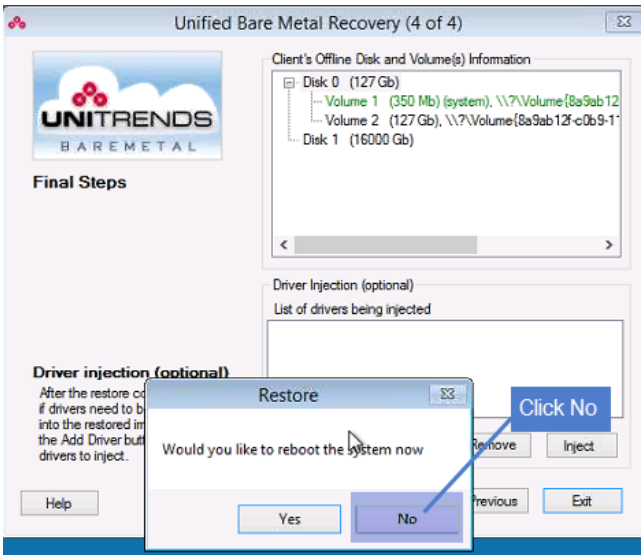

4 Inject an ESX, Hyper-V, or XenServer guest storage driver by doing these steps:

Notes:

If you attempt to reboot without adding the driver, the VM boots to a blue screen. Return to the wizard by booting from the bare metal ISO image. The first wizard screen displays. Click **Driver Injection**. Then inject the necessary driver as described below.

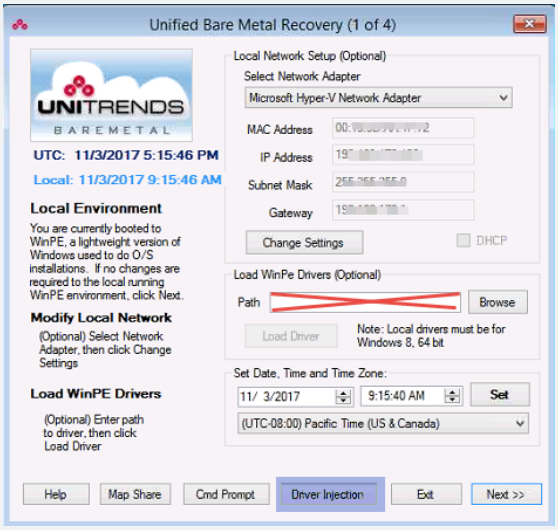

- Click **Add Driver**.

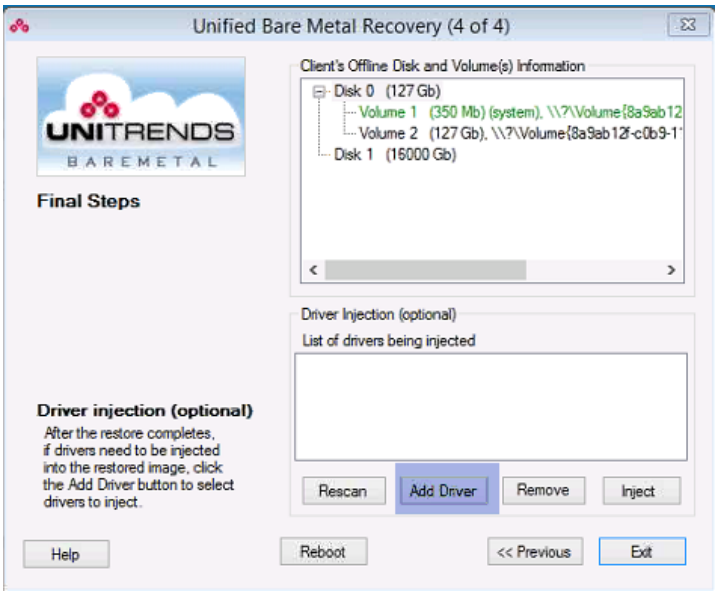

- Select the guest storage driver, and then click **Done** to add the driver to the list of drivers being injected.

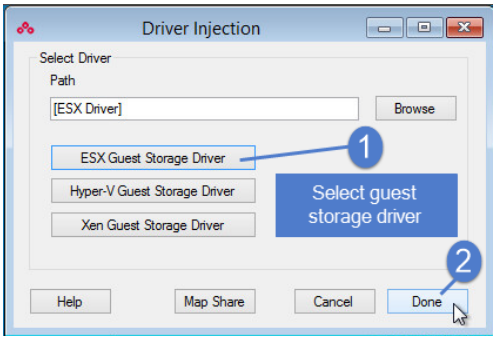

- Highlight the volume containing the operating system files in the asset's Offline Disk and Volume(s) Information area. Then highlight the guest driver and click **Inject**.

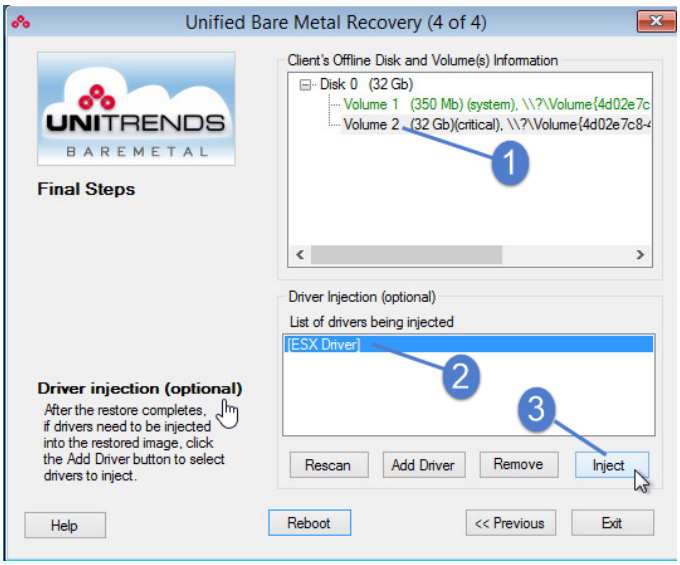

- Click **Yes** to continue.

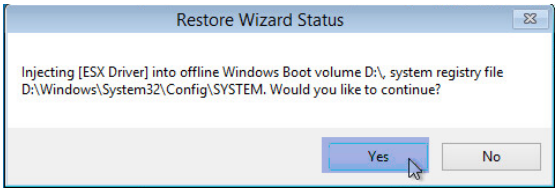

- Click **OK** to close the Restore message.

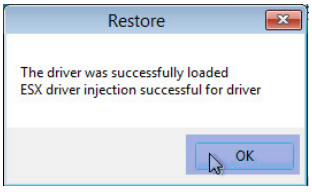

5 When you receive a message stating that the driver injection was successful, you are ready to reboot the VM. Do one of the following:

- If the VM is NOT a Hyper-V domain controller, edit the VM settings so it no longer boots from the bare metal ISO. Click **Reboot**, then click **Yes** to continue.

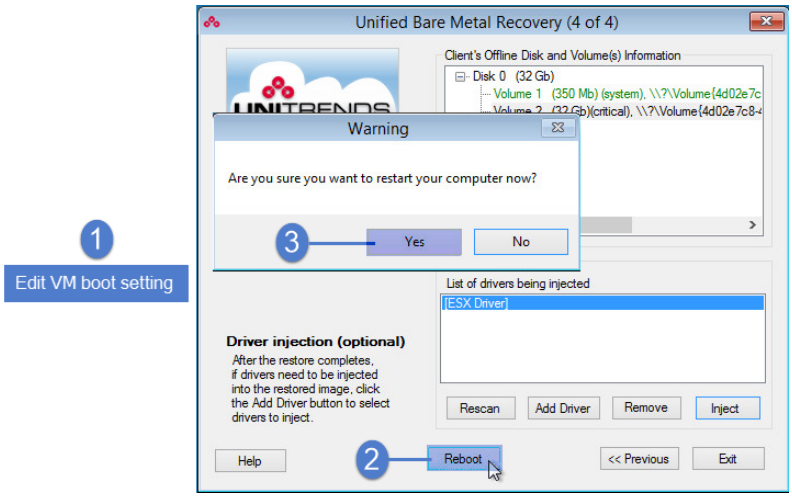

- If the VM IS a Hyper-V domain controller, you must start the VM in Directory Services Restore Mode by doing these steps:
  - Edit the VM settings so it no longer boots from the bare metal ISO.
  - Disconnect the server from the network (to ensure the VM does not start in normal mode).
  - Start the VM in Directory Services Restore Mode.
  - Do one of the following:
    - If a file-level backup is available, connect the VM to the network, then recover the file-level backup as described in "To recover an entire file-level backup by using the Backup Catalog" on page 945.
    - If you do not have a file-level backup, set the **database restored from backup** registry value to **1**. For details on editing this registry value, see this Microsoft article: Backup and Restore Considerations for Virtualized Domain Controllers.
  - Restart the domain controller in normal mode.

6 After the VM boots, proceed to "Step 4: Complete the unified bare metal recovery".

Step 4: Complete the unified bare metal recovery

Use these steps to complete the recovery:

1 Configure network settings for the recovered asset. The network settings that were used for the recovery are not retained after booting into the recovered operating system. Consider the following when configuring network settings:

- If the original asset is still connected to the network, you must assign the recovered asset a unique IP address and rename it before connecting to the network to avoid conflicts.
- If the original asset is no longer connected to the network, you can assign the recovered asset the same IP address as the failed asset.

- If you are using DHCP and you added the original asset to the backup appliance by using only the asset's name, the appliance discovers the recovered asset after you connect it to the network. If the original asset is still connected to the network, rename the recovered asset to prevent conflicts.
- If you configure the recovered asset to use the same name and IP as the original asset, the appliance treats the recovered asset as if it is the original asset. No changes are needed on the Unitrends appliance.
- If you configure the recovered asset to use a different name and IP than the original asset, the appliance treats the recovered asset as new asset. To protect the recovered asset, add it to the appliance and add or modify job schedules.

2 (If needed) Create and format additional drives and volumes:

- If you recovered a file-level backup, only the system critical volumes you selected in "Step 3: Map drives and volumes" on page 1231 have been restored. If backups of the original Windows machine include other volumes, you must create and format those additional volumes.

  IMPORTANT! Recovery of file-level backups fails if these additional volumes do not exist.

- If you recovered an image-level backup, the disks you selected in "Step 3: Map drives and volumes" on page 1231 have been recovered. You can opt to create and format additional disks, but this is not required.

3 (Optional) If you recovered a file-level backup, any data on non-critical volumes has not been restored. To restore this data, recover the failed asset's last file-level backup. For details, see "To recover an entire file-level backup by using the Backup Catalog" on page 945.

4 (If needed) Do these additional steps as needed:

- For Exchange servers – If you are unable to mount Exchange databases after recovery, the databases may be in a Dirty Shutdown state. See this Microsoft article for details: Exchange Database is in a Dirty Shutdown State.
- For Hyper-V servers – After recovery, you must run the following command on the Hyper-V server: **`bcdedit /set hypervisorlaunchtype Auto`**. Then reboot the server.
- For Hyper-V servers recovered to a VMware or XenServer virtual machine – If the Hyper-V server had Integration Services installed, you must remove the Integration Services provider from the recovered VM before running Unitrends backups, as described here:
  - Log in to the recovered VM as administrator and open a command prompt.
  - Enter this command: **vssadmin list providers**
  - In the output, find *Hyper-V IC Software Shadow Copy Provider* and note the Provider ID. Example ID: {74600e39-7dc5-4567-a03b-f091d6c7b092}
  - Using a registry edit tool like **regedit**, remove the provider key and all subkeys: [HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VSS\Providers\{74600e39-7dc5-4567-a03b-f091d6c7b092}]

Windows image-based bare metal recovery

With Windows image-based bare metal recovery (BMR), you protect an asset's operating system by creating a custom ISO-image and running bare metal backups that capture the disk metadata. You perform DR by using the asset's custom ISO and a bare metal backup. The target for the recovery can be a physical or virtual machine.

Note: For most assets, you can recover from regular file-level or image-level backups by using "Windows unified bare metal recovery" on page 1209. It is recommended to use unified bare metal recovery where possible. To determine which to use for your asset, see "Which bare metal method should I use?" on page 1207.

How image-based BMR works

When you boot the recovery target machine from the asset's custom ISO, it boots into WinPE, a minimal version of Windows used for installations, and the Windows Bare Metal Interface displays. You use this interface to perform the recovery. Depending on your operating system and hardware, it might be necessary to add drivers during or after the recovery. The recovery procedures guide you through this process.

Image-based BMR recovers only critical volumes, so to complete the recovery, you need to recover the last file-level backup to restore files that reside on non-critical volumes. After recovering the critical volumes, injecting any necessary drivers, and configuring network settings on the new machine, you use the backup appliance to recover the last file-level backup to the new machine. (If all of your data resides on the critical volumes, it is recovered during BMR and you don't need to recover the last backup by using the backup appliance.)

See the following topics for details about protecting your Windows assets with image-based bare metal recovery:

- "Implementing image-based bare metal protection"
- "Performing image-based bare metal recovery" on page 1252

Implementing image-based bare metal protection

For best results, it is recommended that you plan your strategy for disaster recovery before an asset fails. Following is a high-level overview of the steps you must complete to implement image-based bare metal protection for your Windows assets. It identifies steps to complete before and after an asset fails.

Note: For most assets, you can recover from regular file-level or image-level backups by using "Windows unified bare metal recovery" on page 1209. It is recommended to use unified bare metal recovery where possible. To determine which to use for your asset, see "Which bare metal method should I use?" on page 1207.

Perform the following before a asset fails

Step 1: Review "Prerequisites for Windows image-based bare metal recovery" to verify that the asset's operating system is supported and other requirements have been met.

Step 2: Create the ISO and boot media as described in "Creating the ISO and boot media" on page 1250. (You must create this media for each asset you are protecting with image-based BMR.)

Step 3: Test the boot media as described in "Testing bare metal media for image-based recovery" on page 1251.

Step 4: Run periodic hot bare metal backups. A successful bare metal backup is required to recover the asset. For details on creating a backup job, see "To create a file-level backup job" on page 437.

> **Note:** You must install a separate Windows bare metal agent to run bare metal backups. Install the agent as described in "To install the Windows bare metal agent" on page 372.

Step 5: (Optional/recommended) Perform a test recovery as described in "Testing bare metal media for image-based recovery" on page 1251.

To recover a failed asset

Step 6: Perform image-based BMR using the procedures in "Performing image-based bare metal recovery" on page 1252.

Prerequisites for Windows image-based bare metal recovery

Consider the prerequisites for imaged-base BMR as you plan your disaster recovery strategy.

Requirement	Description
Supported operating systems	Recovery to physical and virtual machines is supported for these operating systems: • Microsoft Windows XP (32-bit and 64-bit, WinPE 1.5). Recovery to identical hardware is supported. Recovery to dissimilar hardware is not supported. • Microsoft Windows Server 2003 (32-bit and 64-bit, WinPE 1.5): – Recovery to identical hardware is supported. – Recovery to dissimilar hardware is supported for some distributions. See the Compatibility and Interoperability Matrix for supported distributions. Recovery to dissimilar hardware is NOT supported for servers with dual-boot or multi-boot configurations. • Microsoft Windows Server 2003 R2 (WinPE 1.5) • Vista (32-bit and 64-bit, WinPE 2.0) • Windows 7 (WinPE 2.0) • Windows 8 (WinPE 2.0) • Windows 8.1 (WinPE 2.0) • Microsoft Windows Server 2008 (32-bit and 64-bit, WinPE 2.0) • Microsoft Windows Server 2008 R2 (WinPE 2.0) • Microsoft Windows Server 2012 (64-bit, WinPE 2.0) • Microsoft Windows Server 2012 R2 (WinPE 2.0)

Requirement	Description
An eligible bare metal backup	These requirements must be met: • The backup used for recovery must be a successful bare metal backup. For details on running a bare metal backup, see"To create a file-level backup job" on page 437. • Local access to the appliance where the backup resides must be unblocked to perform the recovery. If local access has been blocked through the Block Local Access feature in the UniView Portal, temporarily unblock local access until you have recovered your Windows asset. For details, see Blocking or unblocking local access to an appliance in the UniView Portal Guide.
Bare metal ISO and boot media	To recover the asset, you must boot the recovery target machine from the asset's bare metal boot media. Custom media is required for each asset. You cannot create the boot media after an asset has failed. To prepare for image-based BMR, create the media as described in "Creating the ISO and boot media" on page 1250 and keep it in a safe place.
Recovery target machine	You can recover to a physical or virtual target machine. The recovery target must meet the requirements described in the rows below.
Firmware interface type	Supported for BIOS-based assets. (For UEFI-based assets, use "Windows unified bare metal recovery" on page 1209 instead.)
Disk configuration	• Basic disks are supported. Disks must be MBR partitioned. • Dynamic disks and GPT partitions are not supported. **Note:** Assets with UEFI BIOS are automatically partitioned with GPT. Use "Windows unified bare metal recovery" on page 1209 to protect these assets.
Memory	• The target must have enough memory to satisfy Microsoft's support guidelines for the operating system being recovered. • The recovery requires that at least 256MB of RAM is available. • If recovering to a virtual machine, the VM must have at least 2 GB of RAM.
Disk size	The target disk must be at least as large as the source disk you are recovering.

Requirement	Description
Graphics card	The target must have a graphics card supporting a minimum 800X600 resolution.
Network adapter	• Wireless network adapters cannot be used for the recovery. • If recovering to a VMware VM, you must use the E1000 NIC. • If recovering an older asset (Windows 2003/R2 or earlier) to a Hyper-V VM, you must use the Legacy NIC to boot into WinPE 1.5. (For instructions, see "To configure the VM to use the legacy network adapter" below).
Drivers	Drivers needed for image-based BMR are determined by the operating system of the asset you are recovering and the operating system and hardware of the recovery target machine. You might need to add drivers during different stages of the recovery: • If recovering to identical hardware, you do not need lo load drivers. • If recovering to dissimilar hardware, you will need to load drivers to access the network and storage hardware on the recovery target machine. • If recovering to a virtual machine, you will need to load the guest storage driver, During recovery, you boot into WinPE 2.0 (Vista and later operating systems) or WinPE 1.5 (earlier operating systems). If you are recovering to dissimilar hardware you will need to load the drivers described below. • Drivers needed for dissimilar recovery of supported Vista and later operating systems (WinPE 2.0): – The WinPE 2.0 bare metal boot environment requires that you load 32-bit Windows Vista, 2008, or 2012 drivers to access the underlying network and storage hardware. Load the applicable drivers onto a USB drive or CD that you can access the recovery. – Upon booting the bare metal image, the Windows Bare Metal Interface displays. You use this interface to load the network and storage drivers into the WinPE image and then start the recovery. – After recovering the bare metal backup, you will load any additional required drivers. • Drivers needed for dissimilar recovery of supported Windows 2003 operating systems (WinPE 1.5): – The WinPE 1.5 bare metal boot environment requires that you load 32-bit Windows 2003 drivers to access the underlying network and storage hardware. Load the applicable drivers onto a USB drive or CD that you can access the recovery.

Requirement	Description
	– Upon booting the bare metal image, the Windows Bare Metal Interface displays. You use this interface to start the recovery and then load the network and storage drivers. • Recovering to a virtual machine target: You will recover the bare metal backup and then load the networking driver and the guest storage driver (ESX, Hyper-V, or XenServer) right from the Windows Bare Metal Interface.
Dissimilar recovery for Windows 2003	Dissimilar recovery is supported for some Windows 2003 distributions. (See the Compatibility and Interoperability Matrix for details.) These additional requirements apply: • Dual-boot and Multi-boot configurations are not supported. • You must not remove the boot media from the machine once the system has booted. The CD contains important Windows bare metal system files. This is a restriction of Microsoft Windows PE.

Creating the ISO and boot media

You must create a custom ISO image and boot media for each Windows asset. It is recommended to create a new ISO and boot media in these cases:

- After upgrading the Windows agent
- After modifying the operating system and/or system critical volumes.
- After installing or removing programs.

To create the Windows ISO and boot media

1. Log in to the Windows server and launch the Bare Metal Media program:

   From the Start Menu, select **All Programs > Unitrends Agent**, then right-click **Bare Metal Media** and select **Run as Administrator**.

2. Enter the following in the Unitrends System Settings fields:
   - System Name - The hostname of the Unitrends appliance that is protecting this Windows asset.
   - System IP - IP of the Unitrends appliance that is protecting this Windows asset.
   - Select a device in the **Select a backup device** list. If you are storing backups on the default device, select **D2DBackups**.
3. Review the Asset Settings. These are populated by default, and will be the network settings that your server will have after the recovery process.

4 If necessary, check the DHCP checkbox. This will cause the asset to reach out to a DHCP server and grab an available IP address upon booting from the bare metal CD. Leave the firewall and resolve asset IPs boxes unchecked. These features are deprecated and should not be used.

5 Review the path in the Save Windows Bare Metal ISO to area. If desired, you can change this path to save to a different location. The default location is *C:\PCBP_BM\WinBm.dir\cdrom_images*.

6 Review the Save Windows Bare Metal ISO As area to see the name of the ISO that will be created. If desired, you can modify the ISO file name.

7 Click **Create ISO**.

8 On the Continue page, check for the message *All tests are successful*, then click **Yes** to continue.

If you do not see a success message, modify settings as required, then **Create ISO**.

9 Do one of the following (to determine whether WinPE 2.0 or 1.5 is used for your asset, see "Supported operating systems" above):

For...	Procedure
WinPE 2.0 (Vista and later)	The system creates the ISO. Proceed to step 10. **Note:** You will have the ability to inject any drivers while performing recovery.
WinPE 1.5 (2003/R2 and earlier)	You are asked if you would like to insert additional drivers: • Click **Yes** to inject drivers. Browse to your driver(s), select them in the box on the right, and click **Add**. Be sure to check the **Mass storage device** checkbox if drivers are being added for a mass storage device. Click **Continue**. • Click **No** if you do not want to inject drivers.

10 The system creates the ISO and the Success page displays.

- If this is a physical asset (or if you will recover to a physical asset), burn the ISO to CD. Refer to the documentation for the burner you are using to walk you through the process of creating a bootable CD from an ISO image, which is not the same as burning an ISO image to a CD. Store the boot CD in a safe place.
- If this is a virtual machine (VM) asset (or if you will recover to a VM), store the ISO image in a safe place.

11 Test the media as described in "Testing bare metal media for image-based recovery".

Testing bare metal media for image-based recovery

To ensure the boot media you created functions properly, use this procedure to verify that the disk can be used to connect to the Unitrends appliance and recover a bare metal backup.

To test the bare metal media

1 Shut down your Windows server and boot it from its ISO CD.

The server boots and launches the Integrated Bare Metal Recovery Wizard. This can take a few minutes.

2. Click **Bare Metal Hardware Confirmation** in the top right. The Bare Metal Hardware Confirmation dialog displays.
3. Check these boxes to select the test options: **Ping server**, **Quick connect server**, and **Disk read (MBR)**.
4. Click **Start** to begin the test.
5. The results of your test display. If you see *Success*, click **OK** and reboot your server into its operating system. If you see *Windows Bare Metal Quick Test Failed*, do one or all of the following:
   - Ensure that your server and the Unitrends appliance are able to communicate on your network.
   - Verify that the Windows server's hostname and IP address in WinPE matches the settings on the Unitrends appliance:
     - To view settings in WinPE, select **Bare Metal Setup** from the main menu.
     - To view settings on the Unitrends appliance, log in to the appliance, go to the **Configure > Appliances** page, select the appliance and click the **Network** tab below. On the Network tab, select the adapter (typically *eth0*) and click **Edit Hosts File**.
   - Verify that the Windows server uses MBR partitions. To do this, boot into the Windows operating system, open the **Start** menu, right-click **Computer**, click **Manage**, expand **Storage**, and click **Disk Management**. For each of the disks (Disk 0, Disk 1, etc.), right-click and select **Properties > Volumes**, and verify that *Partition Style* is *Master Boot Record (MBR)*.

Performing image-based bare metal recovery

Use the following procedures to perform image-based bare metal recovery. For a successful recovery, be sure to run the procedures in order. Before you start, it is recommended that you read "Windows image-based bare metal recovery" on page 1246 for an overview of the recovery process and review the "Prerequisites for Windows image-based bare metal recovery" on page 1247.

Step 1: "Access the bare metal recovery ISO image"

Step 2: "Prepare the recovery target machine"

Step 3: "Perform image-based bare metal recovery" on page 1254

Step 1: Access the bare metal recovery ISO image

For the recovery, you must use the ISO image that was created for the asset you will recover:

- If you are recovering to a physical machine target, you will need the boot CD containing the failed asset's ISO image.
- If you are recovering to a virtual machine target, save the failed asset's ISO image in a location that you can access from your hypervisor.

Proceed to "Step 2: Prepare the recovery target machine".

Step 2: Prepare the recovery target machine

You can recover a failed asset to a physical or virtual machine (VM). Use one of the following procedures to prepare the recovery target machine:

- "To prepare a physical machine for image-based bare metal recovery"
- "To prepare a VM for image-based bare metal recovery"

To prepare a physical machine for image-based bare metal recovery

1. Check the "Supported operating systems" on page 1247 to verify that the recovery target machine is running a supported operating system version.
2. If you will be recovering to dissimilar hardware, do the following:
   - Check the "Supported operating systems" on page 1247 to verify that recovery to dissimilar hardware is supported for the failed asset's operating system. If dissimilar recovery is not supported, you must recover to a target machine that has identical hardware as that of the failed asset.
   - During recovery, you will need to load drivers into WinPE to access the network and storage hardware on the recovery target machine. Load the applicable drivers onto a USB drive or CD that you can access during the recovery. See "Drivers" on page 1249 for a description of the drivers needed for your operating system.
3. Verify that the machine meets the "Prerequisites for Windows image-based bare metal recovery" on page 1247 and the following requirements:
   - Firmware interface is BIOS-based.
   - Disks are configured as basic disks and the boot disk contains an MBR partition.
   - At least 256MB of RAM is available.
   - The target disk is at least as large as the source disk you are recovering.
   - Has a graphics card supporting a minimum 800X600 resolution.
   - Has a wired network adapter (wireless adapters cannot be used for the recovery).
4. Create copies of any existing data on the target machine's disk if you need to preserve this data. Existing data on the recovery target disks is overwritten or deleted during the recovery.
5. Proceed to "Step 3: Perform image-based bare metal recovery" on page 1254.

To prepare a VM for image-based bare metal recovery

1. Create a new Hyper-V, VMware, or XenServer VM that meets the "Prerequisites for Windows image-based bare metal recovery" on page 1247 and the following requirements:
   - Do not install an OS on the VM.
   - Add enough memory to satisfy Microsoft's support guidelines for the operating system being recovered. The recovery requires at least 2 GB of memory.
   - Assign the VM a virtual hard disk with at least as much space as was available on the failed asset you are recovering. If you give it less space, the recovery fails.

2 Add a NIC to use for the recovery:

- If recovering to a VMware VM, you must use the E1000 NIC.
- If recovering an older asset (Windows 2003/R2 or earlier) to a Hyper-V VM, you must use the legacy NIC to boot into WinPE 1.5.

  To configure the VM to use the legacy network adapter

  - Launch **Hyper-V Manager**, right-click the VM, and select **Shut Down**. The VM shuts down and its State changes to Off.
  - Right-click the VM and select **Settings**.
  - In the Hardware list, select the existing **Network Adapter** and note the current value displayed in the Network drop-down box.
  - Change the value in the Network drop-down box to **Not connected** and click **Apply**.
  - In the Hardware list, select **Add Hardware**.
  - Choose **Legacy Network Adapter** in the list and click **Add**.
  - Change the value in the Network drop-down box to the original value you noted above, and click **OK**.
  - Right-click the VM and select **Start** to power it on.

3 Edit the VM settings to boot from the bare metal ISO that was created for the failed asset.

4 Proceed to "Step 3: Perform image-based bare metal recovery".

Step 3: Perform image-based bare metal recovery

Perform the recovery by using the Integrated Bare Metal Recovery Wizard. See the following topics for details:

- "Considerations for performing a test image-based bare metal recovery"
- "Performing image-based recovery for Vista and later operating systems (WinPE 2.0)"
- "Performing image-based recovery for 2003/R2 and earlier operating systems (WinPE 1.5)" on page 1259

Considerations for performing a test image-based bare metal recovery

Before an asset fails, you can perform a test image-based bare metal recovery without impacting the original asset. To ensure the original asset is not is not impacted, review the considerations in the table below and implement one of the recommended strategies before you perform the test. IP/DNS conflicts will occur if the hostname and IP of the recovered asset match those of the original asset (if the original asset is connected to the network).

To perform the test recovery, use the procedures in one of the following topics:

- "Performing image-based recovery for Vista and later operating systems (WinPE 2.0)"
- "Performing image-based recovery for 2003/R2 and earlier operating systems (WinPE 1.5)" on page 1259.

Consideration	Description
Network considerations	Consider the following before performing the test recovery: • When you boot from the ISO image, the asset connects to the network using the IP and hostname of the asset the ISO was made for. It is critical that the test recovery target machine NOT be connected to the same physical network if the original server is still online as this can cause an IP/DNS conflict. To prevent a conflict, do one of the following: – Remove the production server from the network during the test recovery. After recovering, take the recovery target off of the network and bring the production server back on to the network. – Create a new test ISO image for the production server but assign a different, unused IP address. While creating the bare metal boot disk, just manually change the asset's IP and hostname. (For details, see "Creating the ISO and boot media" on page 1250.) You then need to change the asset's IP/hostname settings on the Unitrends appliance to match these new settings (and change them back after you complete the test recovery). • During recovery, you select a bare metal backup to recover. The Windows Bare Metal interface searches the appliance for these backups by using the asset's IP and/or hostname. If the asset's IP and hostname on the ISO do not match those on the appliance, no bare metal backups display in this list. Modify these settings on the Unitrends appliance to match the settings on the ISO (and change them back after you complete the test recovery). For details, see "To view or edit the hosts file" on page 110.
Active Directory considerations	Assets that are dependent on Active Directory for day to day functions may not function properly when recovered to a test network if a domain controller is not recovered into the test network first. If you plan to recover a domain controller with bare metals, it is extremely critical to perform this recovery into a test network.

Performing image-based recovery for Vista and later operating systems (WinPE 2.0)

Use the procedures in this section to recover assets running Windows Vista and later operating systems listed in "Supported operating systems" on page 1247. These procedures apply to any supported target machine (a physical machine, a virtual machine, dissimilar hardware, etc.).

For a successful recovery, you must run the procedures in the following order:

- "Step 1: Boot into WinPE 2.0 and set up the local environment"
- "Step 2: Start the recovery"
- "Step 3: Boot the recovery target machine into its operating system"
- "Step 4: Complete the image-based bare metal recovery"

Notes:

- You must prepare the recovery target machine before running this procedure. See "Step 2: Prepare the recovery target machine" for details.
- If you are recovering to dissimilar hardware, you will need to load drivers during recovery. WinPE requires that you use 32-bit network and storage drivers. Be sure to load these drivers onto a CD or USB device that is accessible from the recovery target machine before you start.

Step 1: Boot into WinPE 2.0 and set up the local environment

1. Boot the recovery target machine from the bare metal ISO image. The machine boots into WinPE 2.0 and launches the Windows Bare Metal interface.
2. The system attempts to discover the local network hardware. Do one of the following:
   - If the network device is found, proceed to step 3.
   - If a network device cannot be found, select **Yes** to load a driver for the network device. Once the driver is loaded successfully, the boot process continues and the WinPE GUI displays.
3. Verify that local disks display in the disk information area, then proceed to "Step 2: Start the recovery".

   If no disks display, a storage driver is needed to access the local storage devices. Use these steps to load the storage driver:
   - Click **Bare Metal Setup**.
   - Select **Load Driver**.
   - Select the target file, then click **Load Driver**. The selected driver is loaded into the active WinPE image. (If you see a failure message, the driver could not be loaded.)
   - When you see the message indicating the driver was loaded successfully, click **OK** then **Exit** to return to the main menu.
   - Select **Bare Metal Restore**, then **Rescan Disk**. The local disks display in the list. Proceed to "Step 2: Start the recovery".

Step 2: Start the recovery

1. Select the bare metal backup and target disk information, then click **Start Restore**.

   Note: If no bare metal backups display, verify that the asset's hostname and IP address in WinPE are the same as the ones in the hosts file on the backup appliance. If not, modify these settings on the Unitrends appliance to match the ones in WinPE. To check and modify these settings on the Unitrends appliance, see "To view or edit the hosts file" on page 110.
2. The option to view real-time statistics displays in a dialog box. Monitor the progress of the recovery on the Bare Metal Statistics screen.
3. Once the recovery is complete, proceed to "Step 3: Boot the recovery target machine into its operating system".

Step 3: Boot the recovery target machine into its operating system

Run one of the following procedures to boot the recovered machine:

- "To boot a physical recovery target" on page 1257
- "To boot a VM recovery target" on page 1257

To boot a physical recovery target

1. Remove the bare metal CD, then do one of the following:
   - If you are NOT recovering an Active Directory server, reboot your server into its operating system.
   - If you ARE recovering an Active Directory server, you must start the recovered server in Directory Services Restore Mode by doing these steps:
     - Disconnect the server from the network (to ensure the server does not start in normal mode).
     - Start the server in Directory Services Restore Mode.
     - Do one of the following:
       - If a file-level backup is available, connect the server to the network, then recover the file-level backup as described in "To recover an entire file-level backup by using the Backup Catalog" on page 945.
       - If you do not have a file-level backup, set the **database restored from backup** registry value to **1**.
     - Restart the domain controller in normal mode.
2. After booting, do one of the following:
   - If the server boots successfully, proceed to "Step 4: Complete the image-based bare metal recovery" on page 1258.
   - If the server does not boot, you need to inject storage drivers. Continue with the next step in this procedure.
3. Insert the bare metal CD and boot from the ISO.
4. In the Windows Bare Metal interface, select **Bare Metal Restore**, then **Rescan Disk**.
5. Click **Inject Offline Driver**.
6. In the Inject Offline Driver dialog, navigate to the folder containing storage drivers, select the drivers, then click **Inject**.
7. Once the injection completes, a success or failure message displays. (If there is a failure, view the resulting log file to determine the cause.)
8. When you receive a message stating that the driver injection was successful, you are ready to reboot the recovered asset into its operating system. Remove the bare metal CD, then reboot the server.
9. Proceed to "Step 4: Complete the image-based bare metal recovery" on page 1258.

To boot a VM recovery target

Use this procedure to inject the VM guest storage driver and reboot the VM.

1. On the Bare Metal Statistics screen, click **OK** and then **Exit** to return to the Windows Bare Metal Restore screen.
2. Click **Rescan Disk** to scan for newly created partitions and volumes.
3. Click **Inject Offline Driver**.
4. Select your OS volume under Disk information and click either **ESX Guest Storage Driver** or **HyperV or Xen Guest Storage Driver**.
5. Do one of the following to reboot the VM:
   - If the VM is NOT a Hyper-V domain controller, click **Cancel** to close windows until you're back to the Windows Bare Metal menu. Click **Diagnostic Tools** and then **Reboot**.
   - If the VM IS a Hyper-V domain controller, you must start the VM in Directory Services Restore Mode by doing these steps:
     - Disconnect the server from the network (to ensure the VM does not start in normal mode).
     - Start the VM in Directory Services Restore Mode.
     - Do one of the following:
       - If a file-level backup is available, connect the VM to the network, then recover the file-level backup as described in "To recover an entire file-level backup by using the Backup Catalog" on page 945.
       - If you do not have a file-level backup, set the **database restored from backup** registry value to **1**. For details on editing this registry value, see this Microsoft article: Backup and Restore Considerations for Virtualized Domain Controllers.
     - Restart the domain controller in normal mode.
6. After the recovery target machine boots, proceed to "Step 4: Complete the image-based bare metal recovery".

Step 4: Complete the image-based bare metal recovery

1. After booting the recovery target machine into its operating system, verify that the server has network connectivity.

   If the machine is not connected to the network, add the required network drivers. For details, see Microsoft's documentation for the applicable operating system distribution.
2. At this point the Windows boot volume (usually *C:*) has been recovered. Create and format additional volumes as necessary.

   IMPORTANT! If file-level backups of the original Windows machine contain files from volumes outside of the Windows boot volume, you must create and format those additional volumes. Recovery of file-level backups will fail if these additional volumes do not exist.
3. Check the recovered asset's network and hostname settings and make changes as needed:
   - If you are using DHCP and you added the original asset to the backup appliance by using only the asset's name, the appliance discovers the recovered asset after you connect it to the network.
   - If the recovered asset has the same name and IP as the original asset, the appliance treats the recovered asset as if it is the original asset. No changes are needed on the Unitrends appliance.

- If the recovered asset has a different name and IP than the original asset, the appliance treats the recovered asset as new asset. To protect the recovered asset, add it to the appliance and add or modify job schedules.

4 (Optional) Recover the failed asset's last backup to restore data that resides on other volumes. For details, see "To recover an entire file-level backup by using the Backup Catalog" on page 945. If all data resides on the boot volume, all data has already been recovered.

Notes:

- For Exchange servers — If you are unable to mount Exchange databases after recovery, the databases may be in a Dirty Shutdown state. See this Microsoft article for details: Exchange Database is in a Dirty Shutdown State.
- For Hyper-V servers — After recovery, you must run the following command on the Hyper-V server: **`bcdedit /set hypervisorlaunchtype Auto`**. Then reboot the server.

Performing image-based recovery for 2003/R2 and earlier operating systems (WinPE 1.5)

Use the procedures in this section to recover assets running Windows 2003/R2 and earlier operating systems listed in "Supported operating systems" on page 1247. These procedures apply to any supported target machine (a physical machine, a virtual machine, dissimilar hardware, etc.).

For a successful recovery, you must run the procedures in the following order:

- "Step 1: Boot into WinPE 1.5 and run the recovery"
- "Step 2: Boot the recovery target machine into its operating system"
- "Step 3: Complete the image-based bare metal recovery"

Notes:

- You must prepare the recovery target machine before running this procedure. See "Step 2: Prepare the recovery target machine" for details.
- Recovery to dissimilar hardware is supported for some Windows 2003 distributions listed in the Compatibility and Interoperability Matrix. If you are recovering to dissimilar hardware, you will need to load drivers during this procedure. WinPE requires that you use 32-bit Windows 2003 drivers to access the underlying network and storage hardware. Be sure these drivers have been loaded onto a CD or USB device that is accessible from the recovery target machine before you start.

Step 1: Boot into WinPE 1.5 and run the recovery

1 Boot the recovery target machine from the bare metal ISO image. The machine boots into WinPE 1.5 and launches the Windows Bare Metal interface.

2 Verify that network settings display.

Note: If you do not see network settings, you need to load the network driver. To load the driver, boot from the ISO and press **F6** to enter BIOS. Then add the driver.

3 Select **Bare Metal Restore**. This loads the restore GUI which lists the asset's bare metal backups that are currently stored on the backup appliance.

> **Note:** If no bare metal backups display, verify that the asset's hostname and IP address in WinPE are the same as the ones in the hosts file on the backup appliance. If not, modify these settings on the Unitrends appliance to match the ones in WinPE. To check and modify these settings on the Unitrends appliance, see "To view or edit the hosts file" on page 110.

4 Select a bare metal backup to restore (normally the most recent one).

5 Check the **Enable Seek in Restore** box (to speed up the restore).

6 Select the applicable recovery options and click **Start Restore**.

7 Select the backup and a target disk where the backup will be recovered, then click **Add**.

8 Do one of the following:

- If recovering to identical hardware, start the recovery, then skip to step 11.
- If recovering to dissimilar hardware or to a VM, you need to load storage drivers. Proceed to step 9.

9 Select the storage drivers to load by doing these steps:

- Check the **Enable Dissimilar Restore** box.
- Browse to select the drivers to load.
- Select the known platform that you are recovering to from the list.
- Click **OK**.

10 The system verifies the driver files and displays a success or failure message. Do one of the following:

- If the driver file verification is successful, select **OK** to inject the driver files and continue with the recovery.
- If verification fails, select **OK** to select and verify different drivers.

11 The bare metal backup is recovered to the target disk. When the recovery completes, the following message displays:

Quit the Bare Metal Restore GUI and reboot the server with the Windows Bare Metal CD removed from the server.

12 Proceed to "Step 2: Boot the recovery target machine into its operating system".

Step 2: Boot the recovery target machine into its operating system

1 Do one of the following to reboot the recovery target machine:

- If the machine is NOT a Hyper-V domain controller, remove the bare metal CD (physical target) or edit VM settings to no longer boot from the ISO. Then reboot the machine into its operating system by selecting **Diagnostic Tools > Reboot > Yes > Yes**.
- If the machine IS a Hyper-V domain controller, you must start the VM in Directory Services Restore Mode by doing these steps:

- Disconnect the server from the network (to ensure the VM does not start in normal mode).
- Start the VM in Directory Services Restore Mode.
- Do one of the following:
  - If a file-level backup is available, connect the VM to the network, then recover the file-level backup as described in "To recover an entire file-level backup by using the Backup Catalog" on page 945.
  - If you do not have a file-level backup, set the **database restored from backup** registry value to **1**. For details on editing this registry value, see this Microsoft article: Backup and Restore Considerations for Virtualized Domain Controllers.
- Restart the domain controller in normal mode.

2 After the recovery target machine boots, proceed to "Step 3: Complete the image-based bare metal recovery".

> **Note:** If your server does not boot, different storage drivers are needed. Repeat this procedure from the beginning (from "Step 1: Boot into WinPE 1.5 and run the recovery" on page 1259) and add the correct storage drivers.

Step 3: Complete the image-based bare metal recovery

1 After booting the recovery target machine into its operating system, verify that the server has network connectivity.

If the machine is not connected to the network, add the required network drivers. For details, see Microsoft's documentation for the applicable operating system distribution.

2 At this point the Windows boot volume (usually *C:*) has been recovered. Create and format additional volumes as necessary.

> **IMPORTANT!** If file-level backups of the original Windows machine contain files from volumes outside of the Windows boot volume, you must create and format those additional volumes. Recovery of file-level backups will fail if these additional volumes do not exist.

3 Check the recovered asset's network and hostname settings and make changes as needed:

- If you are using DHCP and you added the original asset to the backup appliance by using only the asset's name, the appliance discovers the recovered asset after you connect it to the network.
- If the recovered asset has the same name and IP as the original asset, the appliance treats the recovered asset as if it is the original asset. No changes are needed on the Unitrends appliance.
- If the recovered asset has a different name and IP than the original asset, the appliance treats the recovered asset as new asset. To protect the recovered asset, add it to the appliance and add or modify job schedules.

4 (Optional) Recover the failed asset's last backup to restore data that resides on other volumes. For details, see "To recover an entire file-level backup by using the Backup Catalog" on page 945. If all data resides on the boot volume, all data has already been recovered.

> **Notes:**

- For Exchange servers — If you are unable to mount Exchange databases after recovery, the databases may be in a Dirty Shutdown state. See this Microsoft article for details: Exchange Database is in a Dirty Shutdown State.
- For Hyper-V servers — After recovery, you must run the following command on the Hyper-V server: **`bcdedit /set hypervisorlaunchtype Auto`**. Then reboot the server.

Chapter 21: Recovery Assurance

Recovery assurance addresses the following scenarios for your virtual and physical workloads:

1. Certifying that your backups fulfill the requirements specified in your business continuity plan by testing for RPO/RTO compliance, malware infection, and application behavior in a failover scenario.
2. Orchestrating failover from certified recovery points in the event of a disaster.
3. Accessing production data as fully functional VMs contained within isolated testing labs. These instant labs enable you to scan for malware, perform update testing, run compute intensive reports, and perform business analytics operations without impacting your production environment.

Recovery assurance functionality is provided through the data copy access job type. Data copy access (DCA) jobs run in three modes that directly correspond to the three scenarios described above.

1. Run Test - This is the default mode for data copy access jobs and can be configured to run on a schedule. VMs stood up by this mode are automatically torn down after testing is complete.
2. Failover - This mode is used to orchestrate failover to your recovery network. This mode cannot be configured to run on a schedule.
3. Instant Lab - This mode is used to stand up VMs that remain on your test network until they are manually torn down. This mode cannot be configured to run on a schedule.

Recovery assurance VMs can be run on your virtual host server (ESXi or Hyper-V) or directly on your physical Recovery Series or Recovery MAX appliance:

- On-virtual host DCA – Running recovery assurance VMs on your ESXi or Hyper-V server provides full network capability. You can utilize all job modes (Run Test, Failover, and Instant Lab) and run custom tests for your applications.
- On-box DCA – Running recovery assurance VMs on your physical Unitrends appliance provides no network capability. You can use the Run Test job mode only to stand up and power on VMs. No other job modes or custom tests are supported.

  Note: Network capability for on-box VMs will be added in an upcoming release

See the following table for a comparison of on-virtual host and on-box capabilities:

	On-virtual host recovery assurance	On-box recovery assurance
Supported backup methods	• Agentless host-level backups of Hyper-V and VMware virtual machines • Windows image-level backups (run with the Windows agent) • Windows file-level backups (run with the Windows agent)	• Windows image-level backups (run with the Windows agent)

	On-virtual host recovery assurance	On-box recovery assurance
Supported operating systems	• Any OS running as a VM in a supported version of VMWare with VMWare Tools installed. • Any OS running as a VM in a Windows 2012 or above Hyper-V instance. • Any Windows OS that can be protected by Windows image-level or Windows file-level backups. **Notes:** • See the Unitrends Compatibility and Interoperability Matrix for supported versions of Windows. • For complete requirements, see "Requirements for Windows image-level protection" on page 718 or "Requirements and considerations for file-level backups" on page 703.	Any Windows OS that can be protected by Windows image-level backups. **Notes:** • See the Unitrends Compatibility and Interoperability Matrix for supported versions of Windows. • For complete requirements, see "Requirements for Windows image-level protection" on page 718.
Verification method	The appliance logs into the VM via the hypervisor tools to verify connectivity to the OS and optionally verify application layer services and items within the booted VM. **Note:** There is no screenshot verification.	Screenshot verification: • The desired machine is booted and a screenshot is taken of the login screen after booting. • The screenshot is included in each test report. **Note:** There is no application testing.
Reporting	Recovery Assurance and Compliance reports are available from the desktop and the "Reports" menu. Additionally, certification status is reported in UniView.	Screenshot verification is sent as part of the "DCA Job Notifications report" and the "Compliance report". Recovery Assurance and Compliance reports are

	On-virtual host recovery assurance	On-box recovery assurance
		available from the desktop and the "Reports" menu. Additionally, certification status is reported in UniView.
Testing types	Boot and application level	Boot
Orchestration	Groups of VMs can be tested simultaneously with boot order and boot groups that can be applied.	One VM at a time is booted and tested. The VM is then powered off before booting the next VM. Boot order can also be specified.
Resources used	Resources are shared between the appliance and the hypervisor: • The virtual machine is created in the hypervisor. RAM and CPU is provided by the target hypervisor. • For test/audit functionality, the virtual disk(s) stay on the appliance with only changed data being written to the Internal storage of the appliance's Instant Recovery space. • For Failover functionality, the disk is initially created on the appliance using the appliance's Instant Recovery space. Then, the disk is live-migrated from the appliance to the hypervisor's storage using tools from the hypervisor. See "Virtual machine instant recovery" on page 904 for specifics pertaining to the hypervisor for live storage migration.	All testing resources are provided by the appliance: • Deltas (changes) are written to the Internal storage Instant Recovery space allocated on the appliance. • Appliance CPU and RAM are both used to boot the VM.

For further information on recovery assurance, see the following topics:

- "Recovery assurance requirements and considerations"
- "Recovery assurance procedures" on page 1270
- "Custom tests" on page 1283

Recovery assurance requirements and considerations

Review the following requirements and considerations before proceeding to "Recovery assurance procedures" on page 1270 to create and run data copy access jobs:

Requirement	Description
Appliance license	Data copy access jobs are only enabled on Enterprise Plus licensed appliances. If you are interested in this feature, contact your sales representative or email sales@unitrends.com. **Note:** Upon trial license expiration, data copy access UI objects no longer display.
Supported asset types	Supported asset types vary by where the DCA VMs reside: • On-virtual host DCA jobs can include Windows image-level backups, host-level backups of VMware or Hyper-V VMs, as well as any replicas located on a Hyper-V or ESXi host (VM replicas, Windows image-level replicas, or Windows file-level replicas). **Note:** Replicas located on the backup appliance are not supported. • On-box DCA jobs can include Windows image-level backups. For more on these asset types, see "Windows Image-level Backups Overview" on page 709, "Host-level Backups Overview" on page 653, "Instant recovery of Windows image-level backups" on page 1055, "VM replicas" on page 876, "Windows image-level replicas" on page 1086, and "Windows file-level replicas" on page 993.
Isolated test network	You must configure an isolated test network on your hypervisor prior to running tests and standing up instant labs. Failure to contain tests and instant labs within this VM network may allow them to interfere with your production VMs. • When creating a test network on a VMware host, do not assign a physical adapter to the standard switch. • When creating a test network on a Hyper-V server, use the internal virtual switch type.
VMware requirements	Review the following requirements and considerations before including VMware VMs in your data copy access jobs: • All requirements for VM instant recovery as defined in "Prerequisites for VMware instant recovery" on page 906 must be met. • vSphere versions earlier than 5.5 are not supported. • The following must be installed or configured before taking the first backup intended for inclusion in data copy access jobs:

Requirement	Description
	– VMware Tools must be installed on the VM. **Note:** VMware tools requires Microsoft Visual C++ Redistributable. When upgrading VMware Tools, Visual C++ Redistributable is also updated and a reboot is required. To enable your VM to boot in a DCA job, the Visual C++ Redistributable must be up to date. If not, the DCA job will attempt to install the latest version and the VM will not boot. To prevent this issue, always keep Microsoft Visual C++ 2015-2019 Redistributable updated with the latest version available from Microsoft. For details, see this VMware article: Microsoft Visual C++ Redistributable Requirement for VMware Tools. – Powershell version 3.0 or later must be installed on VMs that are running Windows as a guest OS. – User Account Control must be disabled on VMs that are running Windows as a guest OS. – .NET framework version 2.0 or later must be installed on VMs that are running Windows as a guest OS. – Guest OS credentials must be configured for administrator-level privileges on the VM and added to the backup appliance. For further information on managing credentials, see "Managing asset credentials" on page 322. – For Active Directory domain controllers, guest OS credentials must be configured to use the local system administrator account. A user that is part of the domain administrator group is not sufficient. To run NTDS tests, the local system administrator account is required. Additional considerations: • Failover jobs can only be committed on hosts that are managed by a vCenter. • Re-IP functionality for Linux VMs requires **ifconfig**. **ifconfig** is not installed by default on minimal operating system versions. Be sure to install **ifconfig** to enable re-ip of these operating system versions by running this command: `# yum install net-tools` • Additional requirements must be met to enable re-IP functionality for Windows 7 or Windows Server 2008 R2 VMs. See DCA jobs for Win2K8 R2 may fail for details.
Hyper-V requirements	Review the following requirements and considerations before including Hyper-V VMs in your data copy access jobs:

Requirement	Description
	• All requirements for VM instant recovery as defined in "Prerequisites for Hyper-V instant recovery" on page 907 must be met. • The following must be installed or configured before taking the first backup intended for inclusion in data copy access jobs: – The Unitrends agent must be installed on the Hyper-V server. The latest agent version is recommended, minimum supported version is 10.0.0-1 (see "Installing the Windows agent" on page 362). – The data copy access Hyper-V agent must be installed on the Hyper-V server (see "Installing the data copy access Hyper-V agent" on page 1270). – Hyper-V Integration Services must be installed and updated. – Powershell version 3.0 or later must be installed on Windows VMs. – User Account Control must be set to minimum on Windows VMs. – .NET framework version 2.0 or later must be installed on Windows VMs. – Guest OS credentials must be configured for administrator-level privileges on the VM and added to the backup appliance. For further information on managing credentials, see "Managing asset credentials" on page 322. – For Active Directory domain controllers, guest OS credentials must be configured to use the local system administrator account. A user that is part of the domain administrator group is not sufficient. To run NTDS tests, the local system administrator account is required. Additional considerations: • If re-IP and custom test functionality are required, the data copy access Hyper-V agent must be installed on the Hyper-V server and its guest VMs. For instructions, see "Installing the data copy access Hyper-V agent" on page 1270. • Re-IP and custom tests are not supported for Linux guests and Hyper-V/Windows Server versions earlier than 2012 R2. • Re-IP functionality for Linux VMs requires **ifconfig**. **ifconfig** is not installed by default on minimal operating system versions. Be sure to install **ifconfig** to enable re-ip of these operating system versions by running this command: `# yum install net-tools`
VM replica	A VM replica created on a managed host cannot be moved to a host that is not

Requirement	Description
requirement	managed by a vCenter. For further information, see "VM replicas" on page 876.
Windows file-level replica requirements	Windows file-level replicas on VMware: • Windows Server versions earlier than 2008 are not supported. • The host must be managed by a vCenter. • Data copy access jobs that include Windows replicas and span two hosts require both hosts to be managed by the same vCenter. Windows replicas cannot be moved across hosts that are not managed by the same vCenter. • If the Power On Timeout option is selected, it must be configured for longer than 20 minutes. Windows file-level replicas on Hyper-V: • Powershell version 3.0 must be installed on Windows VMs. • If the Power On Timeout option is selected, it must be configured for longer than 20 minutes. • Hyper-V/Windows Server versions earlier than 2012 are not supported. • If re-IP and custom test functionality are required, the data copy access Hyper-V agent must be installed on the Windows asset from which the replica is created. For instructions, see "Installing the data copy access Hyper-V agent" on page 1270 **Note:** Re-IP and custom tests are not supported for Windows replicas of Windows Server 2012 VMs hosted on Hyper-V/Windows Server 2012 R2 and later.
Image-level backup requirements	Review the following requirements and considerations before including image-level backups in your data copy access jobs. • Image-level backups on VMware – Windows Server versions earlier than 2008 are not supported. • Image-level backups on Hyper-V: – Windows Server versions earlier than 2012 are not supported. – Re-IP and custom tests are not supported for Hyper-V/Windows Server versions earlier than 2012 R2. – If re-IP and custom test functionality are required, the data copy access Hyper-V agent must be installed on the asset before taking the first backup intended for inclusion in data copy access jobs (see "Installing the data copy access Hyper-V agent" on page 1270).

Installing the data copy access Hyper-V agent

To install the data copy access Hyper-V agent

1. Download the agent MSI file from https://helpdesk.kaseya.com/hc/en-gb/articles/4407526882193-Unitrends-Downloads.
2. Save the agent MSI file.
3. Right click the installer file and select **Run as administrator** to start the setup wizard.
4. From the first wizard screen, click **Next**.
5. Click **Finish** when the installation completes.
6. When successful setup is confirmed, click **Close**.

Recovery assurance procedures

- "Creating data copy access jobs "
- "Running data copy access jobs" on page 1279
- "Viewing the results of data copy access jobs " on page 1282

Creating data copy access jobs

Data copy access (DCA) jobs can be created on your source backup appliance, managed appliances, and distributed appliances. The steps required vary by where the DCA VMs will reside:

- On-virtual host DCA job – When creating your first on-virtual host job, you must create a lab profile that enables you to apply consistent settings across multiple DCA jobs. This step can be omitted when subsequent data copy access jobs are created. For details, see "To create an on-virtual host DCA job".
- On-box DCA job – When creating an on-box job, a lab profile is created automatically. There is no need to enter location details since the DCA VMs will reside on the physical appliance itself. For details, see "To create an on-box data copy access job" on page 1276.

To create an on-virtual host DCA job

Use these steps to create a job that will stand up VMs on your ESXi or Hyper-V server.

1. Select **Protect > Data Copy Access**.
2. Enter a job name.
3. Select one of the following in the *What do you want to certify* list: **VMware + Windows Replica** or **Hyper-V + Windows Replica**.

   Notes:

   - A job can contain either a combination of VMware assets, image-level backups, and Windows replicas, or a combination of Hyper-V assets, image-level backups, and Windows replicas. VMware and Hyper-V assets

cannot be included in the same data copy access job.

- Windows replicas must reside on a Hyper-V or ESXi host. Replicas located on the backup appliance are not supported.

4 Check boxes to select assets to include in the job. Click **Next**.

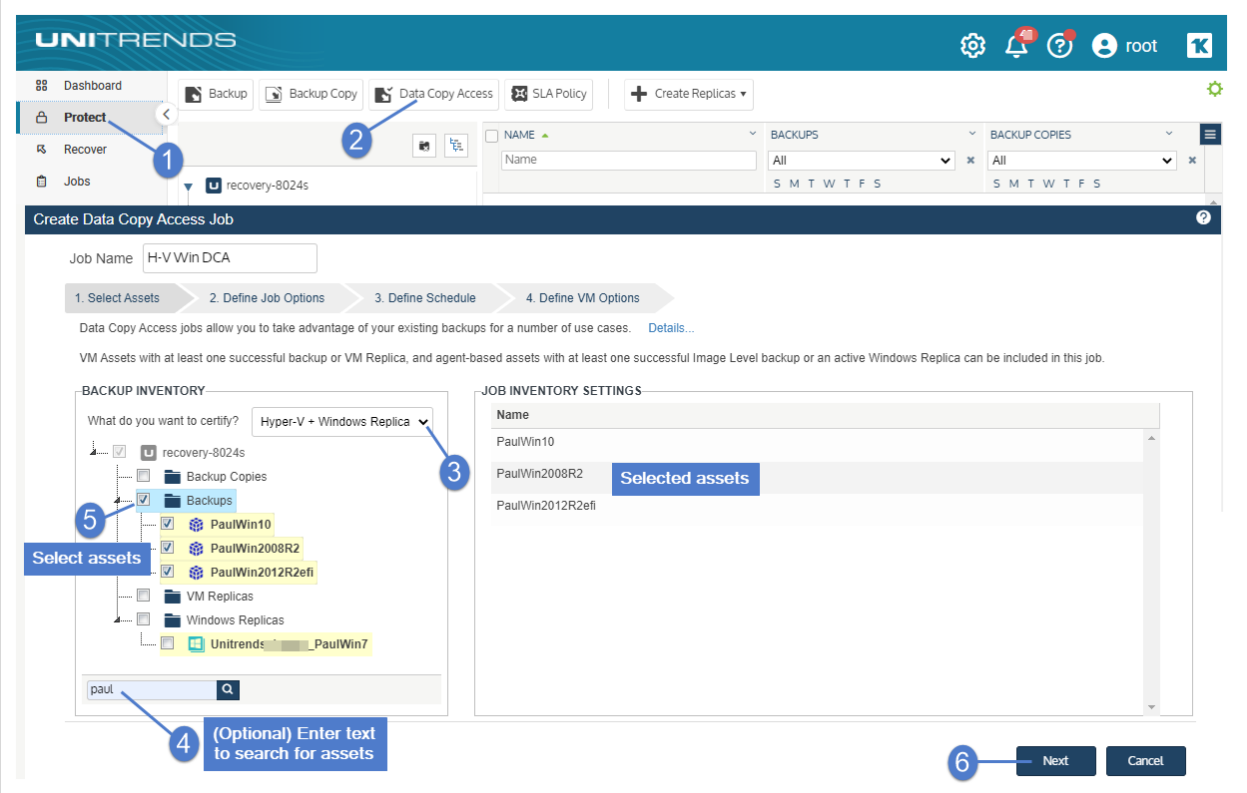

5 Enter location and recovery assurance settings by selecting a lab profile from the list.

If needed, use these steps to create a new lab profile:

- Select **Manage Lab Profiles**. Click **Add**.
- Enter a Profile Name.
- Enter Target Location settings:

Setting	Description
Type	The virtual host type you selected.
Host	Select a virtual host from this list.

Setting	Description
Storage (Hyper-V only)	Select a storage volume.
Resource (VMware only)	Resource pools configured on your selected ESXi host are listed here. You can opt to select a resource pool.
Datastore (VMware only)	Datastores configured on your selected ESXi host are listed here. Select a datastore from the list.
Recovery Network	Networks you have configured via your hypervisor display in this list. The recovery network is where your VMs will be stood up when the data copy access job runs in failover mode.
Test Network	Networks you have configured via your hypervisor display in this list. The test network is where your VMs will be stood up when the data copy access job runs in either Run Test or Instant Lab mode. **WARNING!** The test network must be isolated from your production VMs. Failure to contain tests and instant labs within an isolated test network may allow them to interfere with your production VMs.
Appliance Network	Select the VLAN you are using to connect to your virtual hosts.

- Enter Recovery Assurance settings:

Setting	Description
Enable RPO/RTO tracking	Check this box. RPO/RTO tracking must be enabled in order to certify backups as RPO/RTO compliant recovery points.
Recovery Point Objective (RPO)	The maximum amount of data loss, measured in time, that your business continuity plan can tolerate.
Recovery Time Objective (RTO)	The maximum amount of downtime that your business continuity plan can tolerate.

- Click **Save**.

6 Enter Job Options. Click **Next**.

Setting	Description
Post Custom Script	Enter the filename of a script you wish to run upon conclusion of the data copy access job. **Note:** All custom scripts must output a result.
Post Custom Script Arguments	Enter any arguments for the post custom script.
Power On Timeout	If a VM fails to power on in this amount of time, it is omitted from the data copy access job. **Note:** The inclusion of custom tests and replicas may necessitate longer timeout periods.
Default Suffix	This suffix will be appended to the VM's name when it is stood up on the hypervisor.

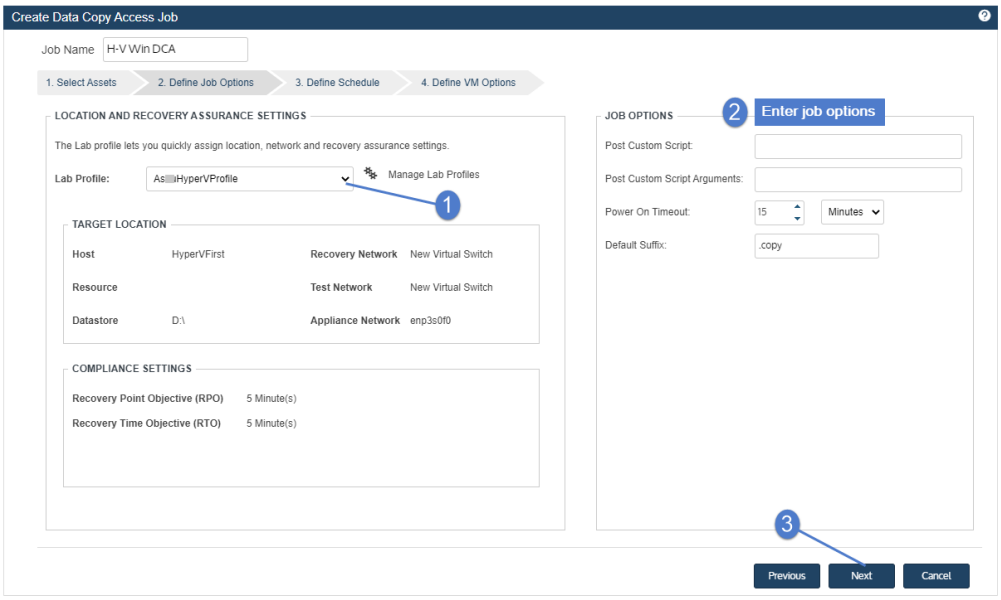

7 Define the test schedule. Click **Next**.

Scheduled data copy access jobs run in Run Test mode.

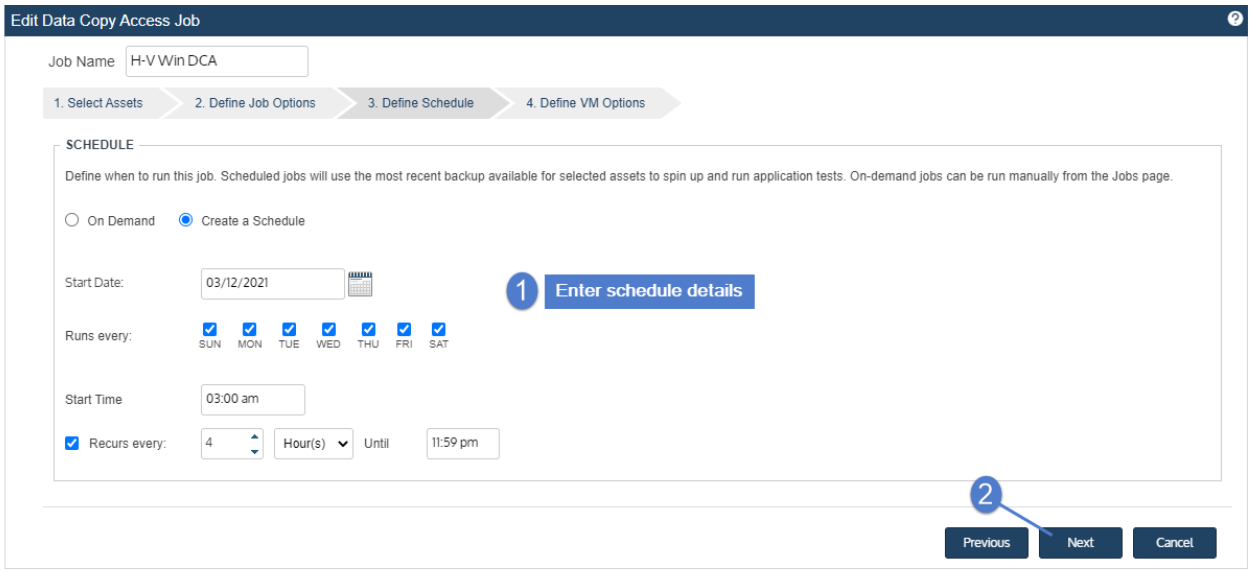

8 Define VM settings. After editing settings for all VMs, click **Save** to create the job.

Click the pencil icon adjacent to each VM in your job and review or edit the settings on each tab. Settings vary based on your individual use case, however guest credentials (on the Options tab) must be configured for each VM. See the tables below for descriptions of these settings.

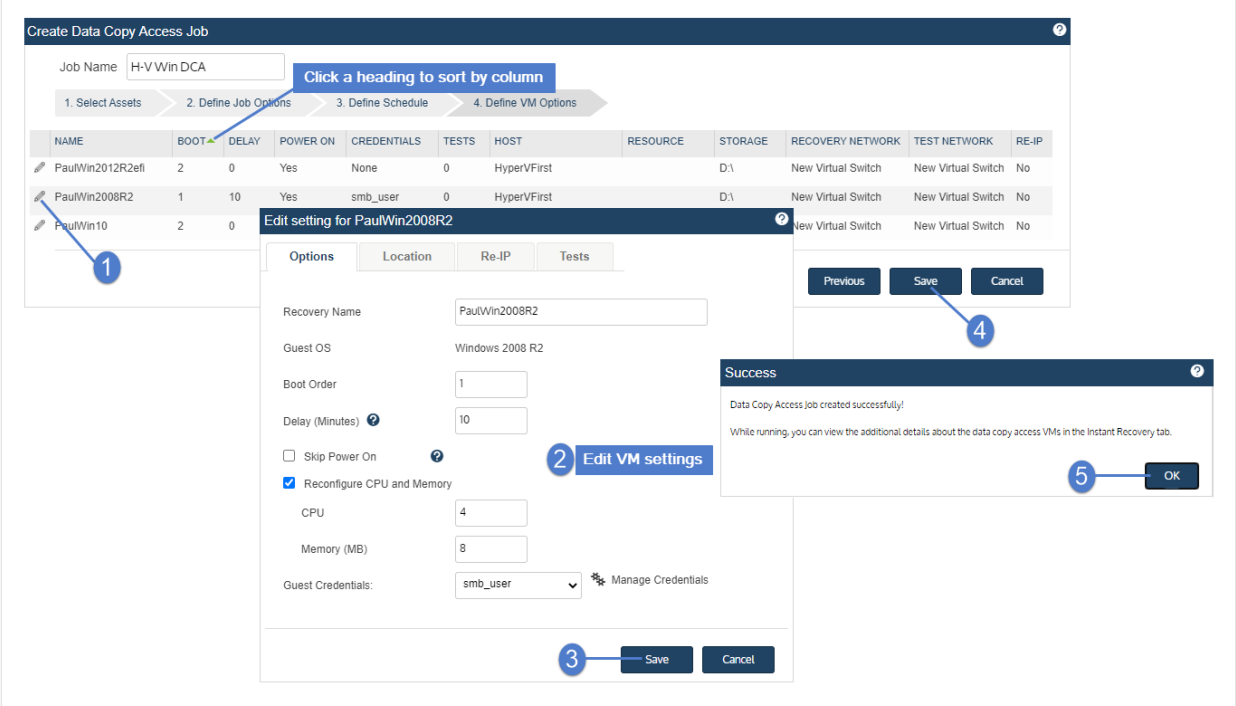

Options Tab

Review and edit settings, then click **Save**.

Setting	Description
Recovery Name	If desired, you can enter a new name for your VM in this box.
Guest OS	Displays the VM's guest OS.
Boot Order	This value determines this VMs position in the boot sequence relative to other VMs in the job. If two or more VMs are assigned the same boot order value, they are sequenced at random relative to each other.
Delay	This value determines the number of minutes to delay custom tests and re-IP after the VM is powered on.
Skip Power On	If this option is selected, the VM is stood up on the test network or recovery network (depending on the job mode) but not turned on.
Reconfigure CPU and Memory	When the data copy access job stands up your VM on either your recovery network or test network, these CPU and memory settings replace the settings the VM was configured with when the backup was taken.
Guest Credentials	Select the guest OS credentials used to log in to this VM. For further information, see "Managing asset credentials" on page 322.

Location Settings Tab

If location settings are entered for this individual VM, they override the location settings specified in your lab profile.

Review and edit settings, then click **Save**.

Re-IP Tab

Click **Add** to assign new IP settings to your VM. When the data copy access job stands up your VM on either your recovery network or test network, these IP settings replace the IP address the VM was configured with when the backup was taken.

Note: Adapter Name refers to the network adapter found within the guest OS of your VM.

Click **Save** after making changes to these settings.

Tests Tab

Add or edit custom tests, then click **Save**. For further instructions on adding custom tests to data copy access jobs, see "Custom tests" on page 1283.

To create an on-box data copy access job

Use these steps to create a job that will stand up VMs on your physical Unitrends appliance.

1 Select **Jobs > Create Job > Data Copy Access**.

2 Enter a job name.

3 Select **Unitrends Appliance** from the *What do you want to certify* list.

4 Check boxes to select Windows image-level assets to include in the job. Click **Next**.

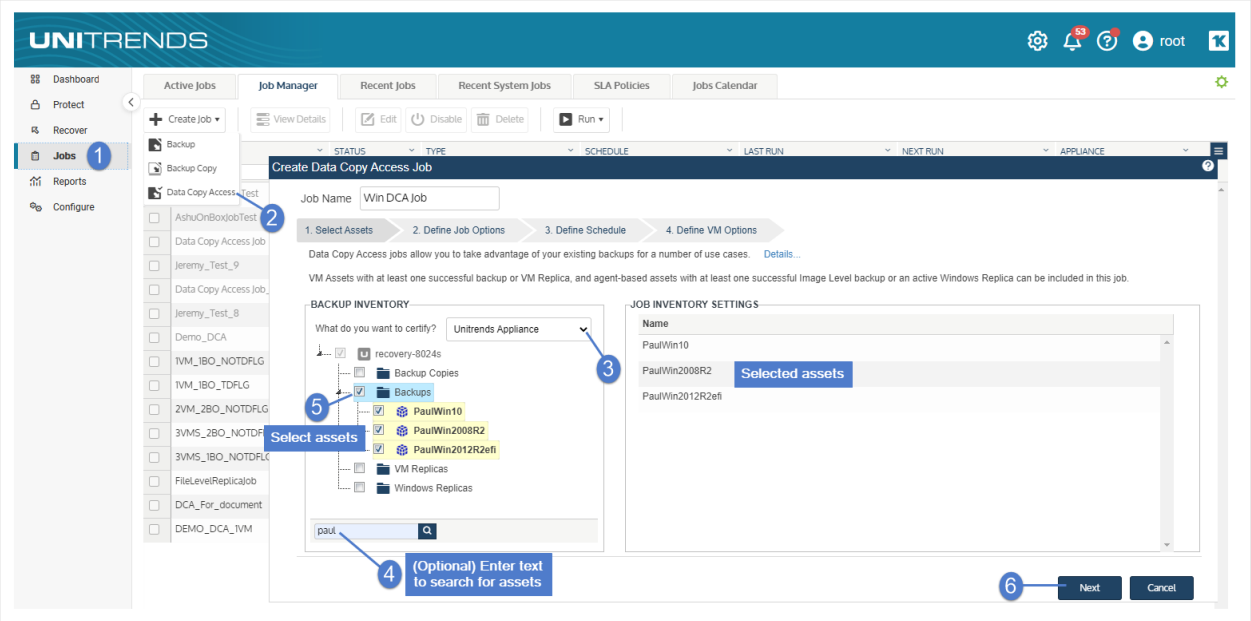

5 Enter job options. Click **Next**.

- Enable RPO/RTO tracking – Check this box. RPO/RTO tracking must be enabled to certify backups as RPO/RTO compliant recovery points.
- Recovery Point Objective (RPO) – The maximum amount of data loss, measured in time, that your business continuity plan can tolerate.
- Recover Time Objective (RTO) – The maximum amount of downtime that your business continuity plan can tolerate.
- Power On Timeout – If a VM fails to power on in this amount of time, it is omitted from the data copy access job.

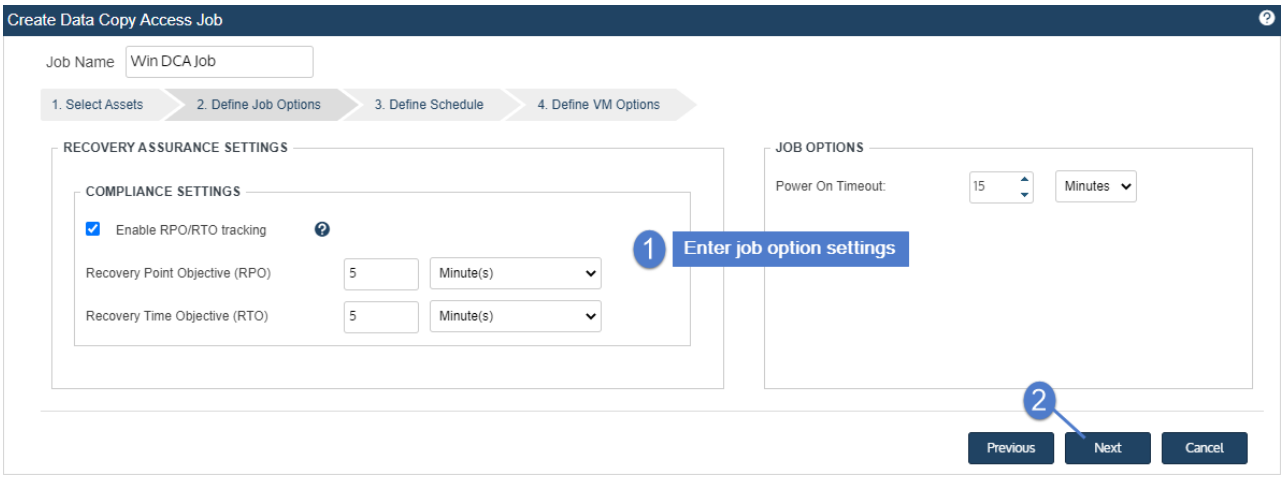

6 Define the test schedule. Click **Next**.

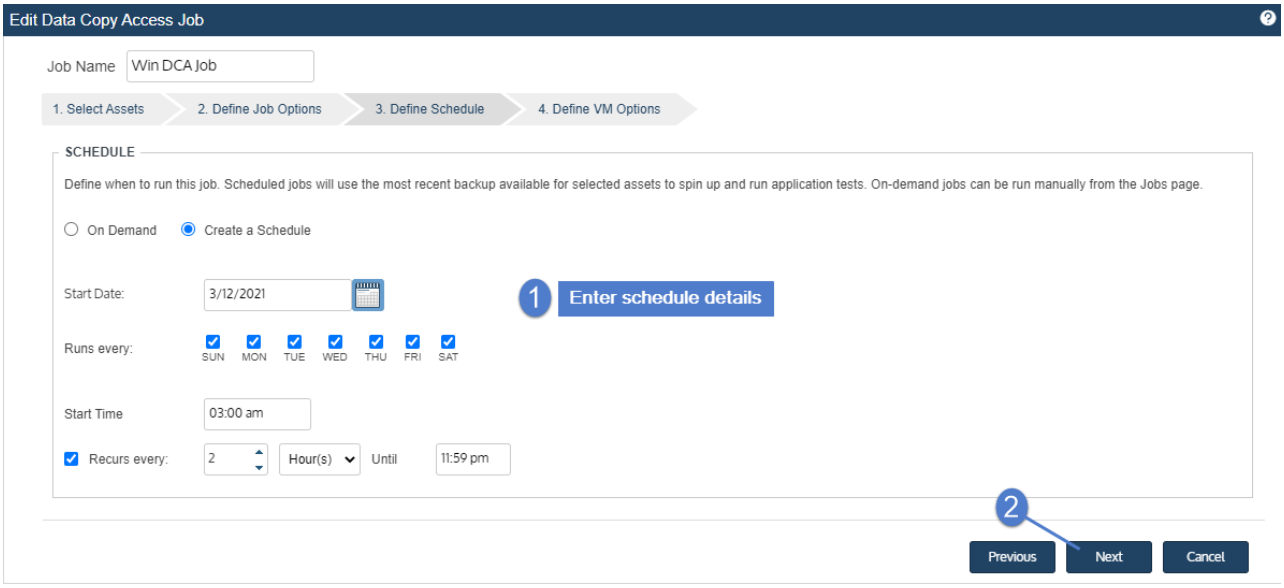

7 Review VM settings. If needed, click the pencil icon to edit a VM's settings. After editing and saving settings for all VMs, click **Save** to create the job.

VM settings

Setting	Description
Recovery Name	If desired, you can enter a new name for your VM in this box.
Guest OS	Displays the VM's guest OS.
Boot Order	This value determines this VMs position in the boot sequence relative to other VMs

Setting	Description
	in the job. If two or more VMs are assigned the same boot order value, they are sequenced at random relative to each other. For on-box DCA jobs, it is recommended to boot VMs serially (one at a time) to reduce the impact to appliance performance.
Delay	Number of minutes the VM will remain powered on before being powered off.
Keep On Until Job Completed	Check this box to keep the VM powered on until the job is finished.
Capture Screenshot	Check this box to capture a screenshot of the VM desktop. This screenshot is included in the emailed "DCA Job Notifications report".
Reconfigure CPU and Memory	When the data copy access job stands up your VM on your test network, these CPU and memory settings replace the settings the VM was configured with when the backup was taken.
Guest Credentials	Select the guest OS credentials used to log in to this VM. For details, see "Managing asset credentials" on page 322.

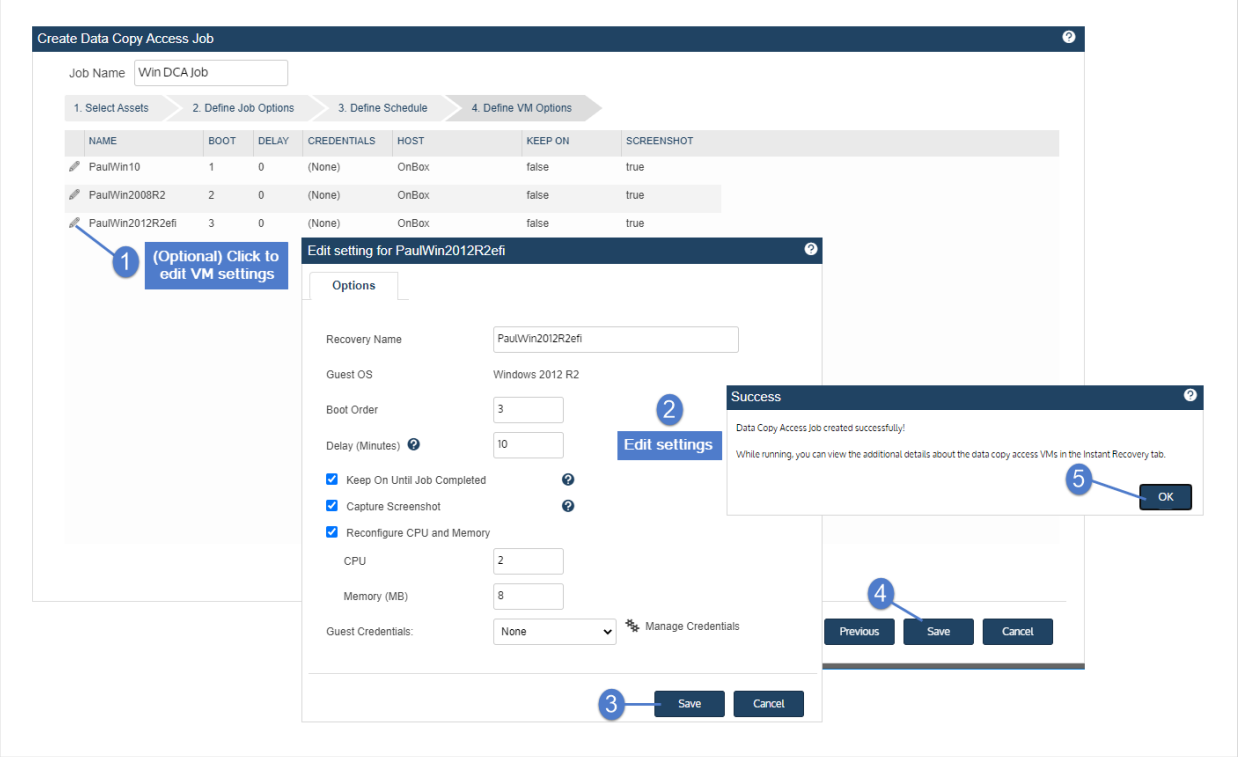

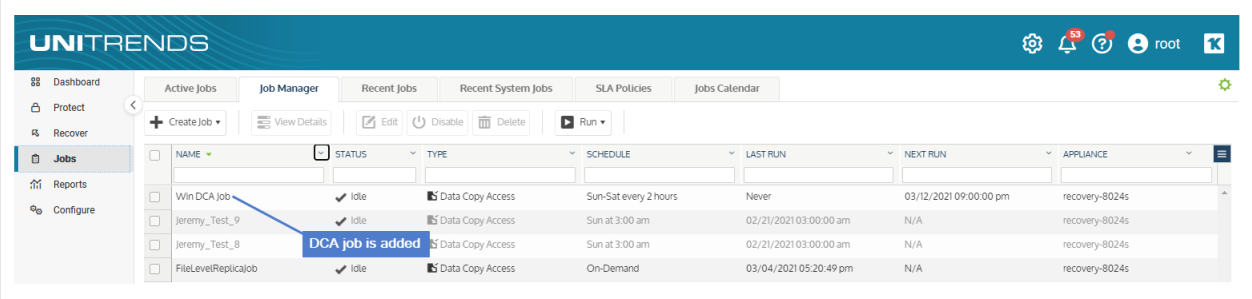

Running data copy access jobs

If you created a schedule for your data copy access job, it will run in Run Test mode at the times specified. This job displays on the **Job Manager** tab, where you can select the job to view details, edit settings, or run it on-demand.

For on-box DCA jobs, use the "To run a DCA job in Run Test mode" procedure to run the job on-demand.

For on-virtual host DCA jobs, you can run the job in any of the three data copy access job modes by using these procedures: "To run a DCA job in Run Test mode", "To run an on-virtual host DCA job in Instant Lab mode" on page 1279, or "To run an on-virtual host DCA job in Failover Mode" on page 1281.

WARNING! Canceling a data copy access job that is in progress can result in orphaned virtual machine instant recovery sessions.

To run a DCA job in Run Test mode

WARNING! If running an on-virtual host DCA job, be sure the job's lab profile is configured to use an isolated test network on your hypervisor. Failure to contain tests and instant labs within an isolated VM test network may allow them to interfere with your production VMs.

1. Select **Jobs > Job Manager**.
2. Select the job.
3. Click **Run**.
4. Select **Run Test**.

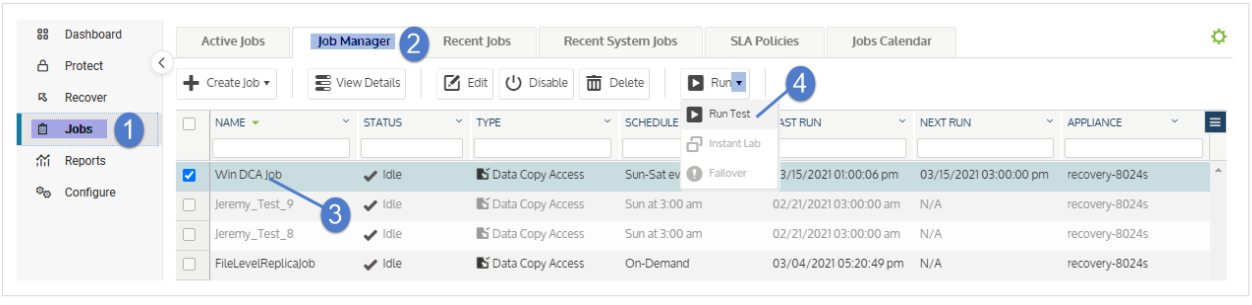

To run an on-virtual host DCA job in Instant Lab mode

This procedure is supported for virtual host targets only.

WARNING! Before running this procedure, be sure the job's lab profile is configured to use an isolated test network on your hypervisor. Failure to contain tests and instant labs within an isolated VM test network may allow them to interfere with your production VMs.

1. Select **Jobs > Job Manager**.
2. Select the job.
3. Click **Run**.
4. Select **Instant Lab**.

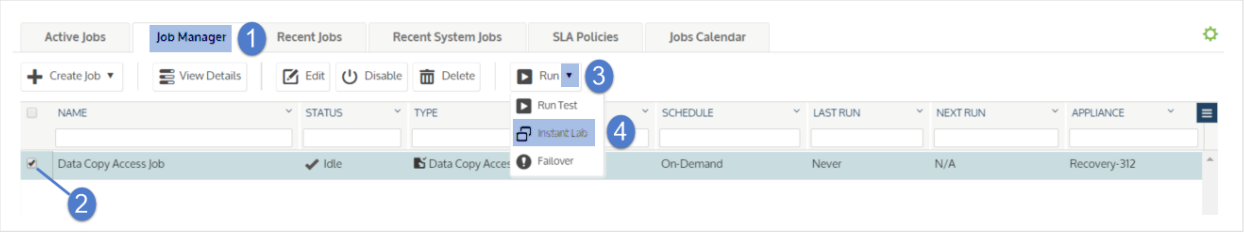

5. Select your desired VMs and recovery points.
   - Green text – This backup has passed all tests specified in the initial Run Test job.
   - Red text – This backup has failed one or more tests specified in the initial Run Test job.
   - Black text – This backup has not been tested.
   - *Failure to latest recovery point* – Select the latest recovery point if recovering from a VM replica or Windows replica.
6. Click **Save**. The dialog closes.

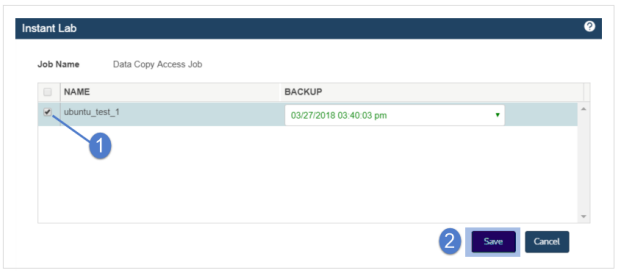

7. Reselect the job.
8. Click **Options**.
9. Select **Commit**. This triggers VM stand-up on your virtual host. Any replicas included in the job will be stood up in audit mode.

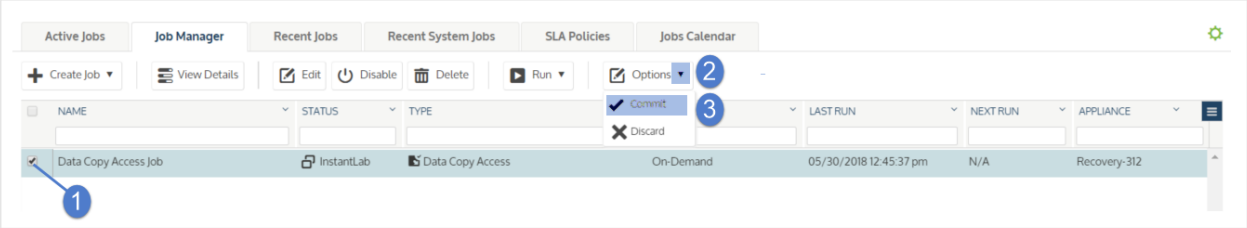

To run an on-virtual host DCA job in Failover Mode

This procedure is supported for virtual host targets only.

WARNING! Before running this procedure, be sure the job's lab profile is configured to use an isolated test network on your hypervisor. Failure to contain tests and instant labs within an isolated VM test network may allow them to interfere with your production VMs.

1. Select **Jobs > Job Manager**.
2. Select the job.
3. Click **Run**.
4. Select **Failover**.

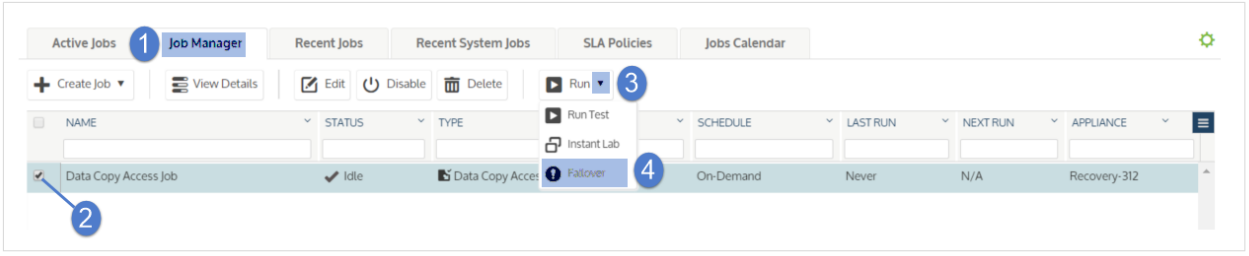

5. Select your desired VMs and recovery points.
   - Green text - This backup has passed all tests specified in the initial Run Test job.
   - Red text - This backup has failed one or more tests specified in the initial Run Test job.
   - Black text - This backup has not been tested.
   - *Failure to latest recovery point* – If recovering from a VM replica or Windows replica, select the latest recovery point.
6. Click **Save**.

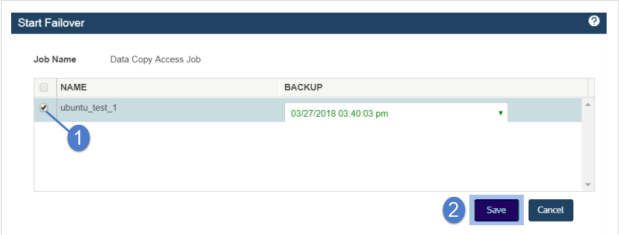

7 Reselect the job.

8 Click **Options**.

9 Select **Commit**. This completes the failover process.

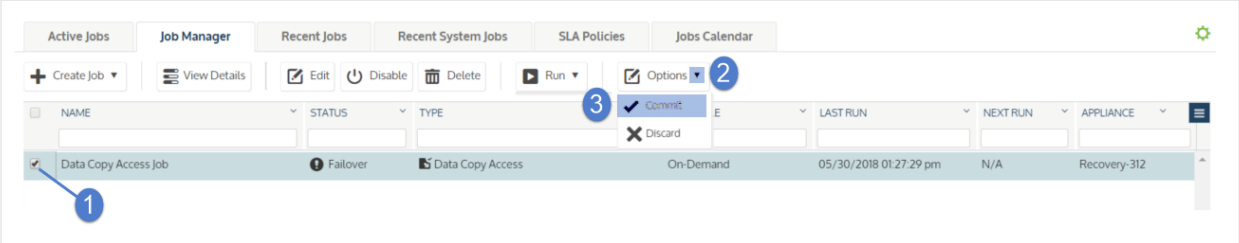

Note: Any replicas included in the job will stand up in audit mode. In a disaster recovery scenario you may wish to put your replicas in live mode. For further information see "Bringing the replica live in production" on page 1020 for Windows replicas and "Bringing the VM replica live in production" on page 895 for VM replicas.

Viewing the results of data copy access jobs

The results of a data copy access job can be viewed from within the Jobs menu or the Recovery Assurance and Compliance reports. For further information, see "Working with reports" on page 1307.

To view the results of a DCA job

1 Select **Jobs > Recent Jobs**.

2 Select job.

3 Click **View Log**.

The tasks included in the data copy access job display in a vertical list.

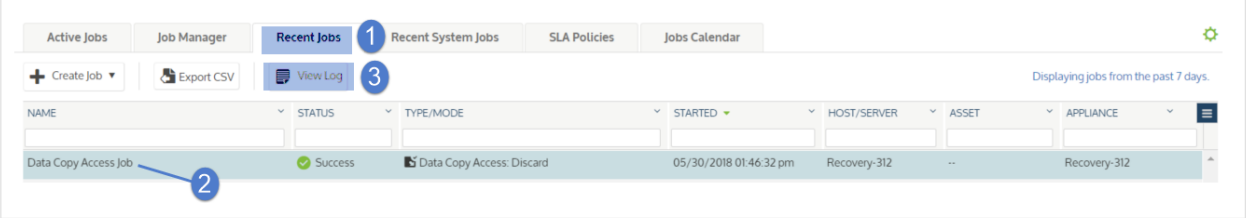

Custom tests

Recovery assurance supports a wide range of tests that can be applied at the VM level. Once configured, these tests are included in the data copy access job.

Notes:

- Custom tests are supported for VMs running on ESXi or Hyper-V only. Custom tests for on-box VMs will be supported in an upcoming release.
- Do not select the Skip Power On option for VMs you wish to test.

Application tests

Application tests enable you to verify that specific applications and processes will function as intended in a disaster recovery scenario.

To include an application test in an on-virtual host DCA job

1. From the Define VM Options tab of the Create Data Copy Access Job dialog, select the pencil icon adjacent to the VM you wish to test.
2. Select the **Tests** tab and click **Add**.
3. Enter test details:

Setting	Description
Type	All tests that are available for this VM display in the dropdown.
Name	Used to identify this test when viewing the results of the data copy access job.
Priority	This value determines the order in which this test will run. If multiple tests are assigned the same priority value, their positions in the sequence respective to each other are chosen at random.
Timeout	If the test fails to complete in this amount of time, it will be flagged as a failure.

Note: Some application tests may have additional settings.

4 Click **Add**.

5 Click **Save** to finish creating the data copy access job.

Custom scripts

Data copy access jobs support custom scripting that enables you to develop your own tests and further customize your disaster recovery plan.

Note: All custom scripts must output a result.

To include a custom script in an on-virtual host DCA job

1 From the Define VM Options tab of the Create Data Copy Access Job dialog, select the pencil icon adjacent to the VM you wish to test.

2 Select the **Tests** tab and click **Add**.

3 Enter test details:

Setting	Description
Type	Select **Run Script** from the dropdown.
Name	Used to identify this test when viewing the results of the data copy access job.
Script Path	This is the file pathway location of your custom script.
Arguments	Enter any script arguments in this box.
Priority	This value determines the order in which this test will run. If multiple tests are assigned the same priority value, their positions in the sequence respective to each other are chosen at random.
Timeout	If the test fails to complete in this amount of time, it will be flagged as a failure.

4 Click **Add**.

5 Click **Save** to finish creating the data copy access job.

Malware scans

Data copy access jobs enable you to scan backups for malware in your test environment rather than impacting your production environment.

Note: Malware scans are not supported on Linux VMs or Windows guest OS versions earlier than Windows 2008 R2.

To include a malware scan in an on-virtual host DCA job

1. From the Define VM Options tab of the Create Data Copy Access Job dialog, select the pencil icon adjacent to the VM you wish to scan.
2. Select the **Tests** tab and click **Add**.
3. Enter test details:

Setting	Description
Type	Select **Scan drive or folder** from the dropdown.
Name	Used to identify this test when viewing the results of the data copy access job.
Drive or Folder	The drive or folder that will be scanned in this test. To scan multiple volumes, create a separate test for each.
Priority	This value determines the order in which this test will run. If multiple tests are assigned the same priority value, their positions in the sequence respective to each other are chosen at random.
Timeout	If the test fails to complete in this amount of time, it will be flagged as a failure. Unitrends recommends a minimum timeout period of 60 minutes for whole drive scans.

4. Click **Add**.
5. Click **Save** to finish creating the data copy access job.

Note: If a potential malware infection is detected, the test is flagged as a failure.

This page is intentionally left blank.

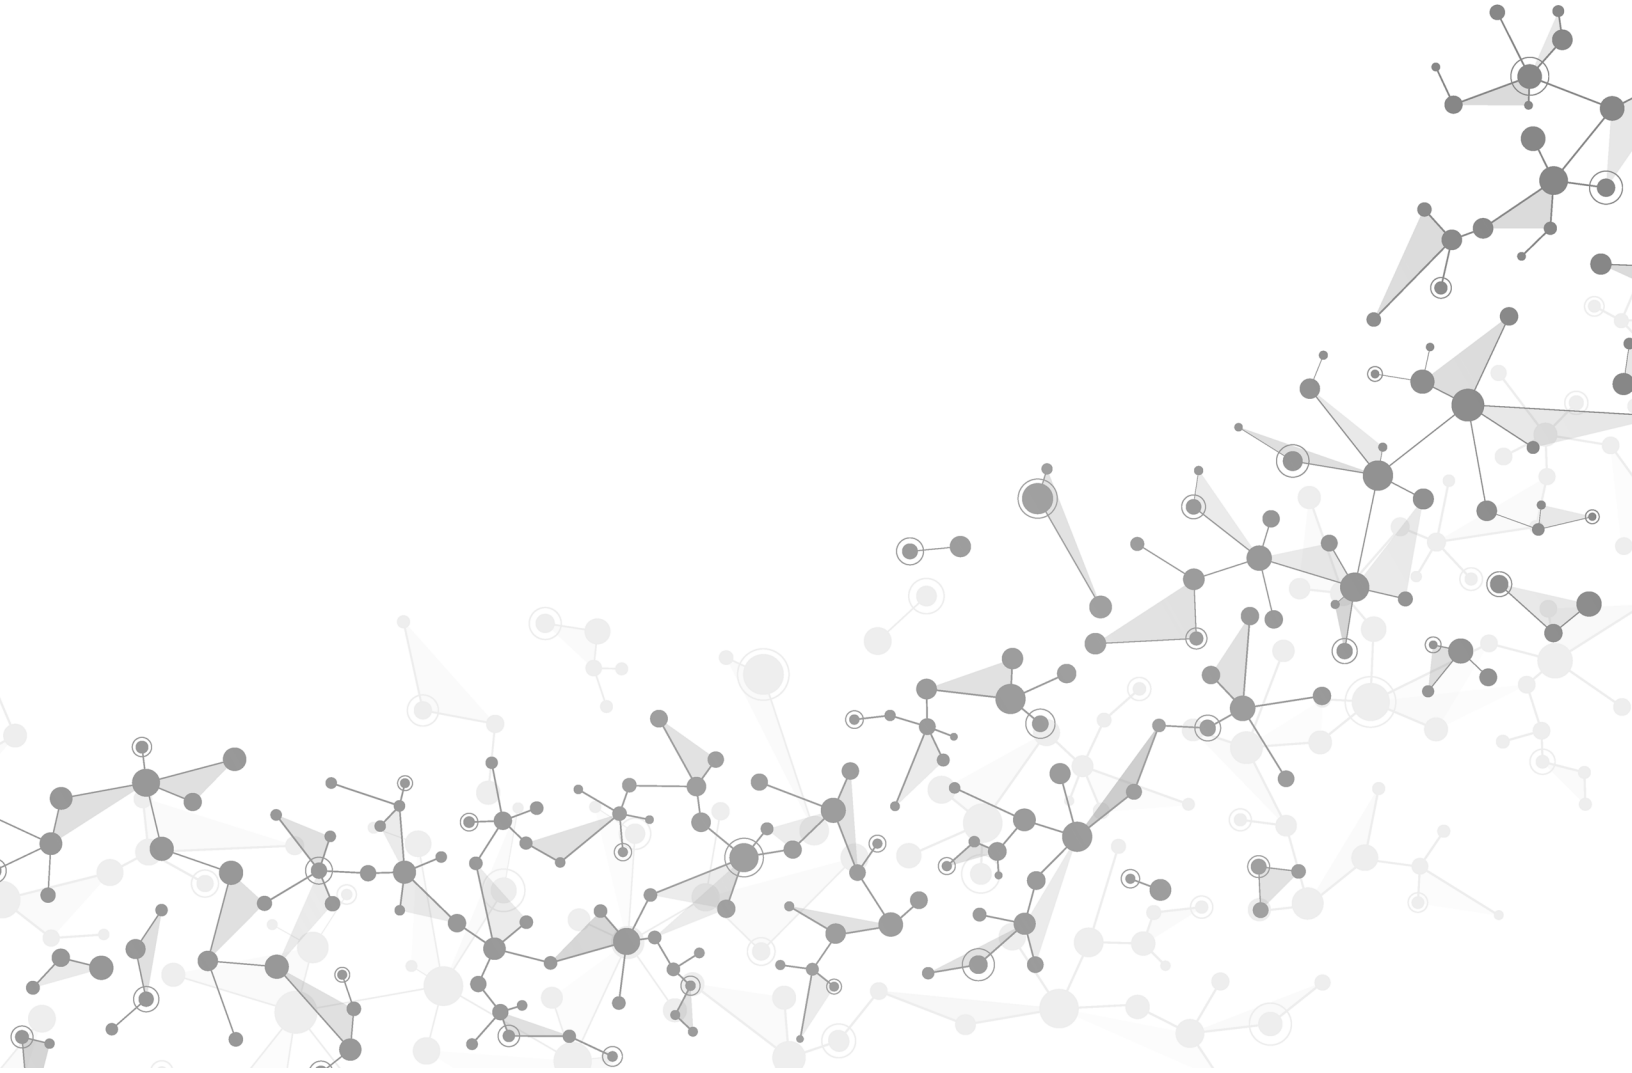

Chapter 22: Helix Self Healing

Helix is an intelligent SaaS remediation platform laser focused on eliminating manual tasks that IT administrators hate performing. Helix uses a SaaS delivery model to keep your Unitrends backup appliances and protected assets healthy, no matter where they are located.

This Helix release targets Unitrends appliance updates, Windows Volume Shadow Copy Services (VSS), and HDD/SSD disk health monitoring leveraging Self-Monitoring, Analysis, and Reporting Technology (SMART). The Standard edition of Helix enables automated appliance updates and is available with all licensed Unitrends appliances. Add the Premium edition to detect and remediate Windows VSS and HDD/SSD disk issues.

Helix appliance updates

Keeping your Unitrends appliances up to date is critical for optimal security and performance, and enables you to benefit from the latest features and fixes. Helix checks for appliance updates and automatically installs them as they become available, ensuring you have the latest enhancements at your disposal.

How it works

Start by opting in to automated appliance updates, as described in "To configure automated appliance updates". Helix then periodically checks for available updates. If an update is found, Helix creates a pending install task and runs the install as soon as there are no actively running backup or recovery jobs. (Note that the install terminates any running backup copy jobs.)

To configure automated appliance updates

1 Ensure that these prerequisites have been met:

- The appliance must be running release 10.4.3 or higher. If needed, install appliance updates.
- Outbound Helix communication and updates are performed over the following ports:

Port, Protocol, and Rule	Hostname or IP Address
5721: • TCP • UDP • Outbound	173.247.66.64
443: • HTTPS • Outbound	kaseyagroup-appliance-registry.jfrog.io

Port, Protocol, and Rule	Hostname or IP Address
443: • HTTPS • Outbound	repo.unitrends.com

2 Log in to the appliance UI.

You must log in directly to the appliance. You cannot configure automated updates for a managed appliance.

3 On the **Configure > Appliances** page, select the appliance and click **Edit**.

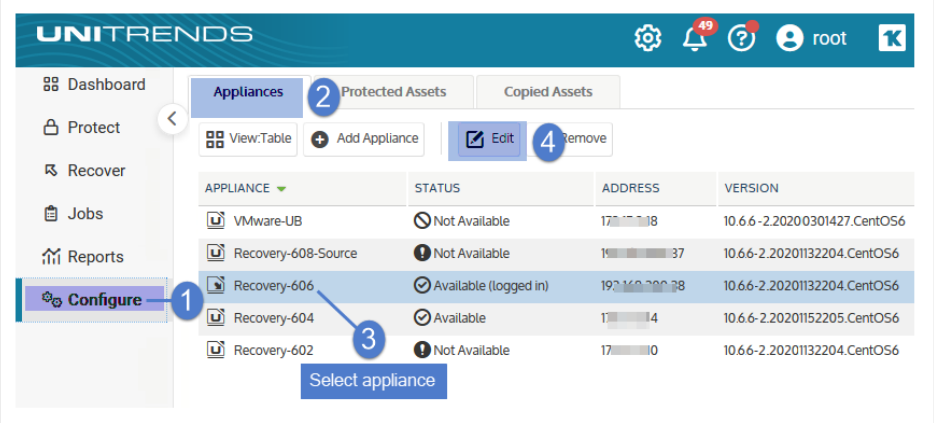

4 Select **Email** and specify the recipient where Helix notifications will be sent:

- Helix sends email notifications to the first recipient that has the Appliance box checked. Unitrends recommends that you use an email distribution list to ensure uninterrupted receipt of notifications as employee roles change in your company.
- SMTP configuration is not required because notifications are sent from Helix (rather than from the appliance itself). Helix notifications are sent regardless of whether SMTP has been configured on the appliance.
- Notifications are sent from this address: *Helix@unitrends.com*. Ensure that your SMTP server is configured to accept notifications from this address.

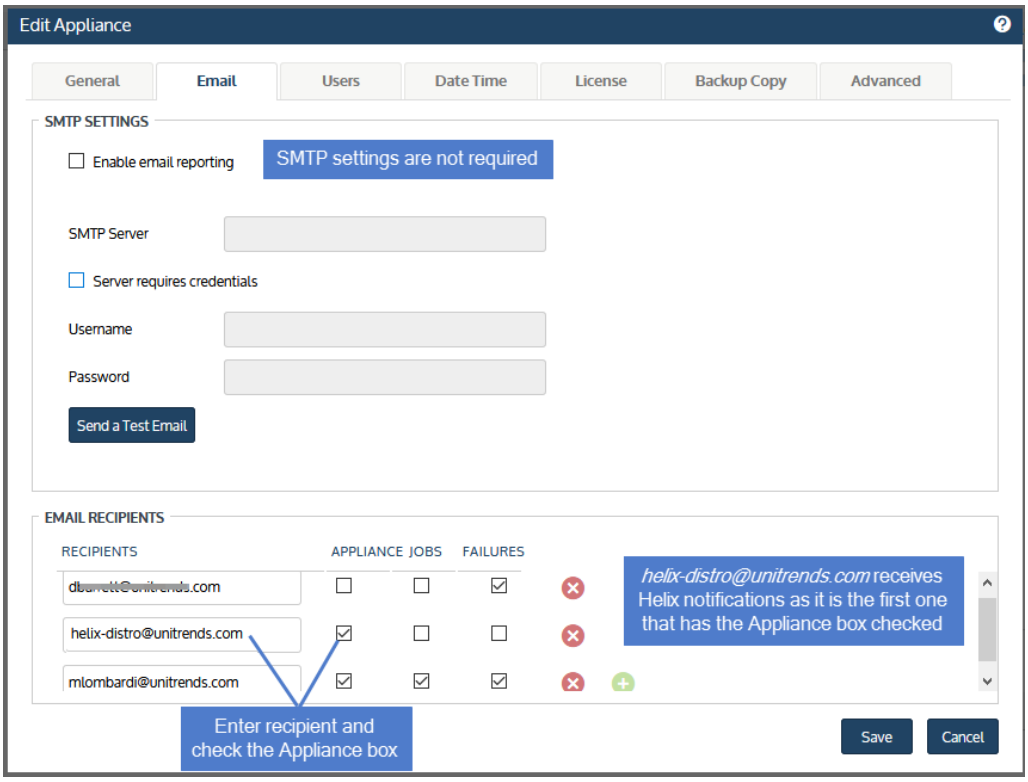

5 Click **Advanced** and select **Support Toolbox**.

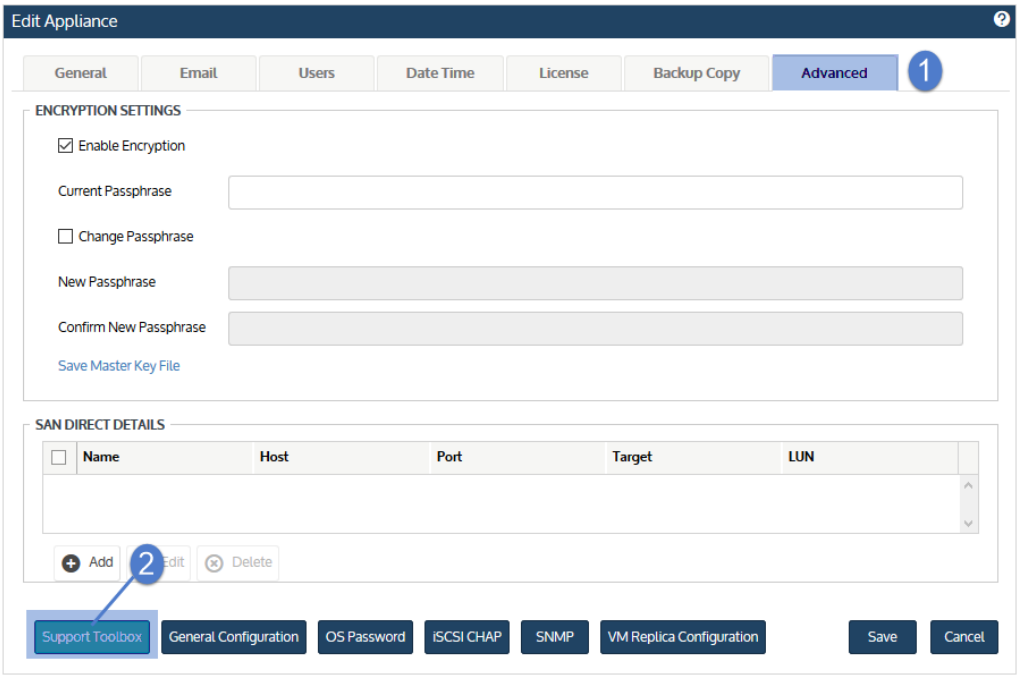

The Support Toolbox (Advanced) dialog displays.

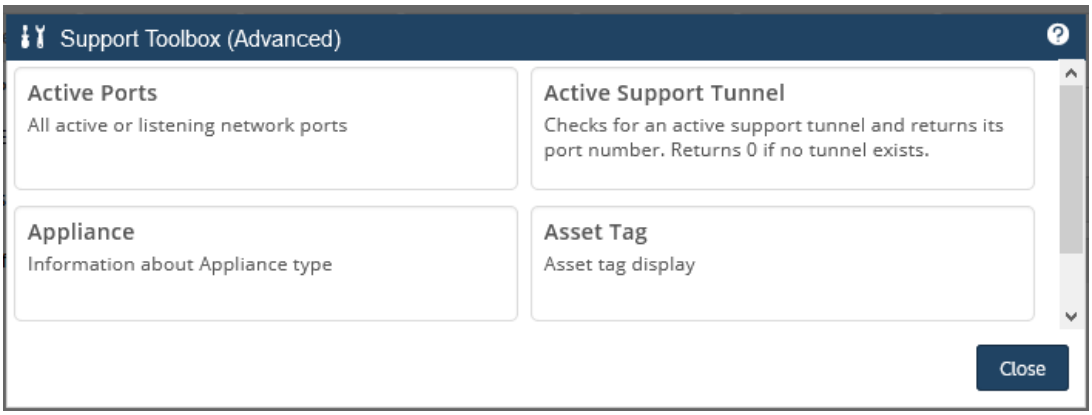

6 Scroll down and click **Helix On/Off** to enable Helix.

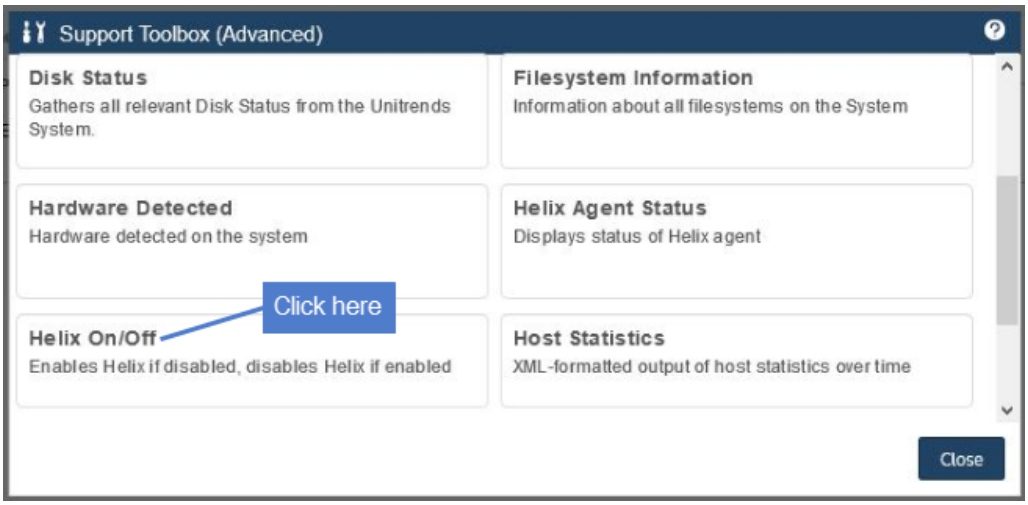

To verify that Helix has been enabled, click **Helix Agent Status**:

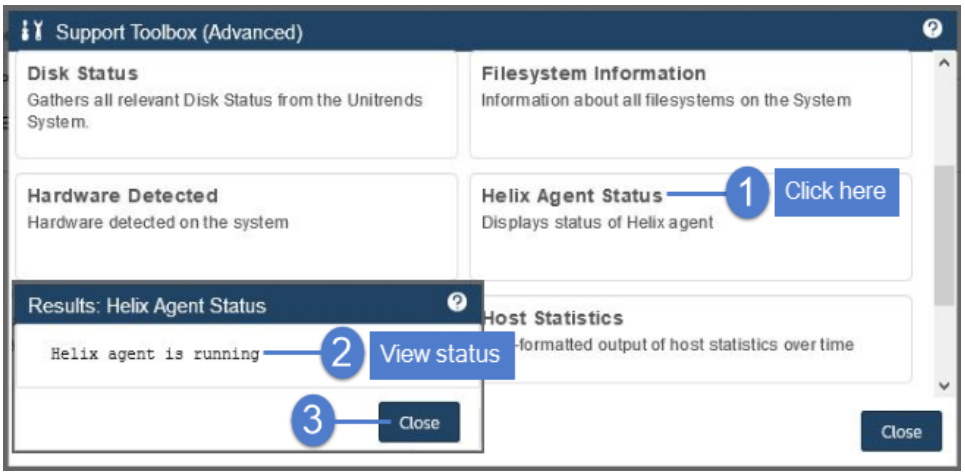

7 Click **Helix Auto-update On/Off** to enable automated updates:

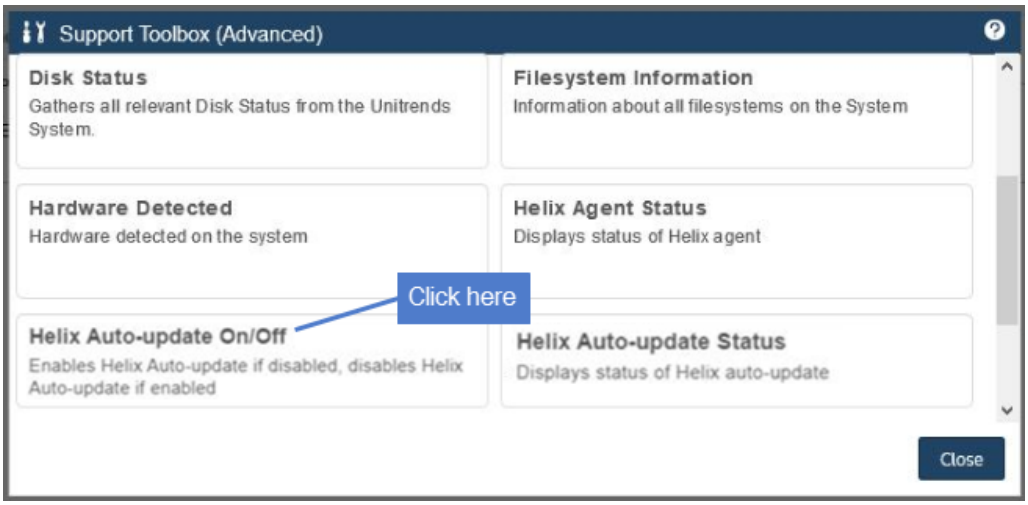

To verify that the auto update feature has been enabled, click **Helix Auto-update Status**:

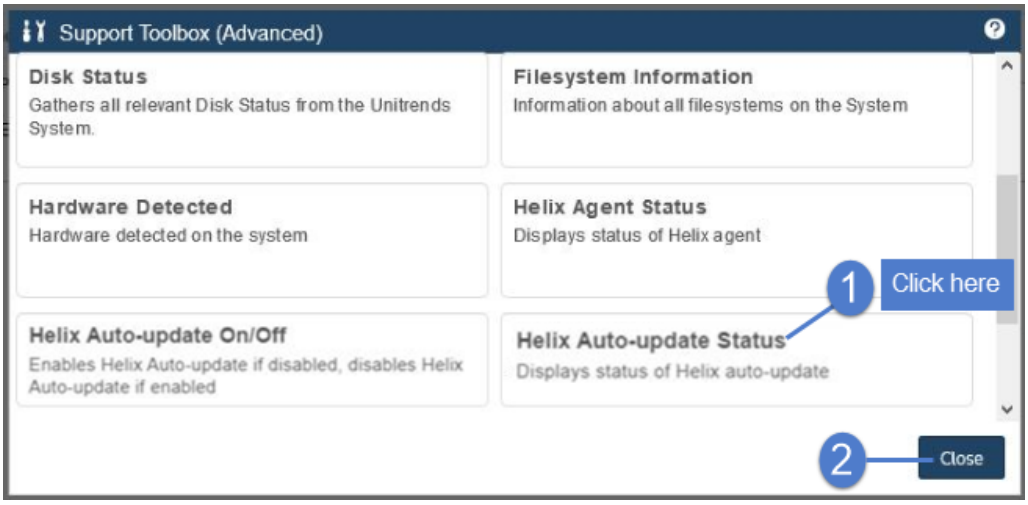

VSS monitoring and remediation with Helix

Backup is the last line of defense to a myriad of threats to your data. Successful backups cannot happen without your environment cooperating. It is your responsibility to monitor, detect, and remediate conditions that impact backups. Unitrends Helix is the intelligent engine that automates these tasks so you can focus on more important ones.

Troubleshooting environmental issues that impact backups wastes valuable time. Helix identifies and fixes the most common backup problems, without you lifting a finger. Volume Shadow Copy Services (VSS) are critical for data protection, but are the most commonly reported root cause for backup failures. Helix monitors your protected Windows assets and remediates VSS issues proactively— saving you from spending time analyzing logs and calling Support. But VSS remediation is just the beginning. Over time, Helix will be enhanced to intelligently increase Recovery Point Objectives (RPOs) and automate recoveries after certain failure conditions.

To use this feature, contact your Unitrends Sales Representative to purchase Helix Premium. Upon purchasing Helix Premium, you receive a Welcome email containing a link you can use to download the Helix agent. Then simply install the Helix agent on your protected Windows assets, either manually or to multiple assets by using your endpoint management tool. Once the agent is installed, Helix automatically starts detecting and fixing a variety of VSS issues. If an issue is detected, Helix sends an email notification about the error and fix.

Note: See "Installing the Helix agent" on page 1293 for agent requirements and installation procedures.

HDD/SSD disk monitoring and remediation with Helix

Helix leverages Self-Monitoring, Analysis, and Reporting Technology (SMART) to check the disk health of your protected Windows assets and notify you if a drive is in need of repair or is no longer reliable. To use this feature, contact your Unitrends Sales Representative to purchase Helix Premium. Upon purchasing Helix Premium, you receive a Welcome email containing a link you can use to download the Helix agent. Then simply install the Helix agent on your protected Windows assets, either manually or to multiple assets by using your endpoint management tool.

Note: See "Installing the Helix agent" on page 1293 for agent requirements and installation procedures.

Once the agent is installed, Helix automatically runs SMART disk status checks daily and sends an email notification when a drive has flagged a general SMART status other than *OK*. The email notification includes the status returned, as well as information about the Windows asset. The purpose of this alert is to help identify hard drives that may be in a pre-failure or error state in the Windows systems that Unitrends protects.

Next steps upon receiving a Helix SMART notification

The Helix SMART disk drive notification indicates that the Windows management interface on your machine has detected a SMART issue on one or more of its hard drives. This alert is typically, but not always, raised before a drive completely fails. Receipt of such an alert may indicate a system hardware health issue on the named client that should be investigated further. For remediation steps, see this KB article: Unitrends Helix: SMART disk drive problem on <machine name>.

IMPORTANT! Please DO NOT contact Unitrends Support about this issue. Unitrends Support is unable to provide assistance to customers who receive this alert as this alert relates only to 3rd party hardware that Unitrends does not sell or support. This alert is provided to notify you of potential hardware issues so you can avoid system crash and recovery. This alert does not imply Unitrends coverage for identified issues in 3rd party hardware.

Helix SMART disk monitoring limitations

The following limitations apply:

- Unitrends does not warrant or guarantee that your system will correctly identify individual drive failures. Please ensure you are also monitoring or running vendor provided diagnostic tools and are using RAID arrays in all production systems where possible to avoid data loss.
- SMART errors and alerts do not necessarily mean a drive is defective or that it needs to be replaced. SMART error conditions may be dismissed by the hardware vendor or may not be covered by pre-failure warranty. Unitrends is

not responsible for labor or costs associated with troubleshooting SMART errors reported/detected by wmic or cases that lead to unnecessary replacement.

- This tool is only able to identify that a drive has flagged a general SMART status other than *OK*, and displays this status. Detail information about drive health is not available via this tool. 1st party or 3rd party tools may be required to ascertain drive health details needed to resolve the issue.
- This tool reports all drives present, including the status of removable media. Unstable media or drives that enter sleep states may produce sporadic results.
- This tool only reports the status of local attached drives on systems that have the Helix Agent deployed. It cannot report the status of SAN or NAS drive systems and cannot be used to monitor VMWare, Xen, or Nutanix hypervisor storage status. This tool is currently limited to Windows operating systems that support the Unitrends Helix agent.
- This process is limited by the features and output of Microsoft wmic. Your system may have a failing drive that we cannot identify for any of the reasons below or for other undocumented reasons, and Helix should not be trusted as the sole source for managing your individual system health. Examples of drive issues that cannot be detected by SMART:
  - Drives are not seen on the bus at the time of the scan because they have hard failed.
  - Drives that are not returning queries because they are in a hard failed state.
  - Drives whose SMART status is obscured by a RAID controller or other firmware, drivers, or monitoring.
  - Failures of the wmic calls made through Windows.

Installing the Helix agent

If you have purchased Helix Premium, simply install the agent on your protected Windows assets to use the VSS and SMART disk remediation features. If an issue is detected, Helix sends an email notification about the error and fix.

Notes:

- Notifications are sent to the email address that was specified when you purchased Helix Premium. Unitrends recommends that you use an email distribution list to ensure uninterrupted receipt of notifications as employee roles change in your company.
- Notifications are sent from this address: *Helix@unitrends.com*. Ensure that your SMTP server is configured to accept notifications from this address.

Helix agent requirements

The Helix agent can be used to protect Windows assets that meet these requirements:

- Supported operating systems – Windows Server 2012/2012 R2 and higher, Windows 8 and higher.
- Port requirements – Port 5721 must be open outbound from the Windows asset to all IP addresses, for the TCP and UDP protocols.
- Free space and RAM – The Windows asset must have at least 100 MB of available space and 100 MB of RAM to install the Helix agent.

Note: The Helix agent can be installed on an asset that is running another Unitrends or Kaseya Windows agent. It is not necessary to uninstall these other agents.

Use one of these procedures to install the Helix agent:

- "To install the agent on multiple assets"
- "To install the agent on a single asset"

To install the agent on multiple assets

To deploy the agent to multiple assets, download the agent installer, *KcsSetup.exe*, by using the link in your Helix Welcome email. Use your endpoint management tool to run the installer on multiple assets.

To install the agent on a single asset

1. Download the agent installer, *KcsSetup.exe*, to the Windows asset.

   You can access the installer by using the link in your Helix Welcome email.

2. Log in to the Windows asset as administrator.
3. In Windows Explorer, browse to the download location and double-click the agent installer, **KcsSetup.exe**.

   The agent is installed.

4. Click **OK** to close the Agent Setup message.

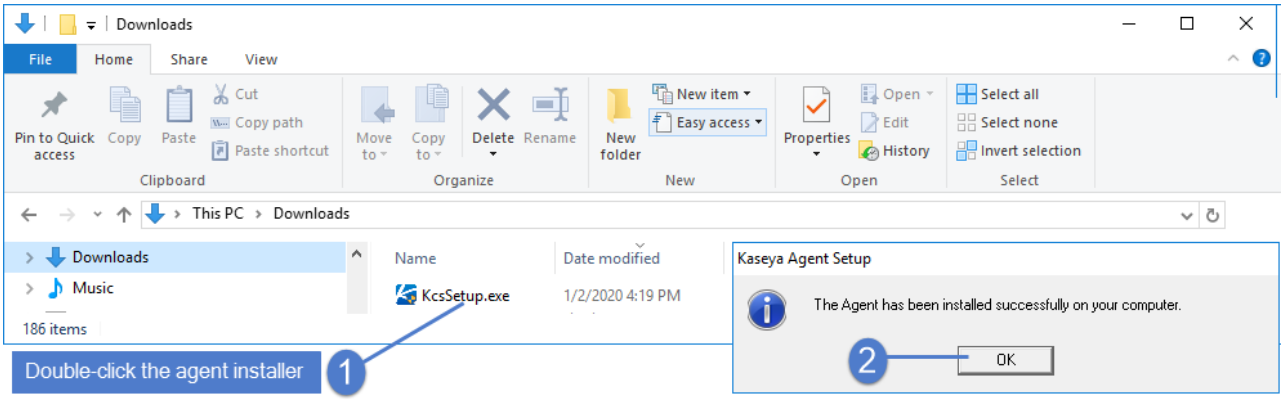

Chapter 23: Appliance Disaster Recovery

This chapter contains information for planning and implementing a disaster recovery (DR) strategy, and provides detailed instructions for recovering your Unitrends appliance. The disaster recovery (DR) procedure restores the configuration settings of the failed appliance. Next, you have the option to import the last backups of each protected asset. This enables you to quickly spin up a new appliance that has the same settings as the original, without having to add assets, set up schedules, and reconfigure settings. See these topics for details:

- "Preparing for appliance DR"
- "Performing DR from a hot backup copy" on page 1299
- "Performing DR from a cold backup copy" on page 1302
- "Licensing the DR target appliance" on page 1304

Preparing for appliance DR

To ensure fast and successful appliance DR, it is important that you create a strategy and implement your plan long before a failure occurs. See these topics to create your plan and prepare for appliance DR:

- "Record information about the original appliance"
- "Run backup copies at regular intervals" on page 1296
- "Create a plan for obtaining a fresh DR target appliance" on page 1296
- Review the "Requirements and considerations for appliance DR" on page 1296

Record information about the original appliance

Use the following table to record appliance information that will be needed to perform DR. You can find this information in the appliance UI on the Edit Appliance dialog (select **Configure > Appliances > Edit** and view the General and License tabs).

Appliance information for DR	Value
Appliance Name	
Appliance IP	
Asset Tag	
Feature String	
License Key	

Appliance information for DR	Value
Encryption?	If backups are encrypted, you will need the encryption passphrase to recover last backups.

Run backup copies at regular intervals

You can recover your appliance and its last backups from a hot or cold backup copy. Choosing which to use will be based on your particular needs, but the most effective disaster recovery strategy requires that these choices have been made well in advance.

Note: Disaster recovery from a tape backup copy is not supported.

Appliance metadata (its system state) is automatically included in each hot and cold backup copy. This metadata contains information such as protected assets, job schedules, storage configuration, and other appliance settings. Recovering this metadata restores an appliance in the event of a disaster. Be sure to run backup copy jobs at regular intervals to ensure you have a recent copy that you can use for appliance DR.

Create a plan for obtaining a fresh DR target appliance

A fresh appliance to recover to must be available. This DR target appliance can be a:

- Recovery Series or Recovery MAX appliance — Use a new or re-imaged appliance. If you need to re-image an appliance, see How to reimage a Unitrends appliance for instructions.
- Unitrends Backup virtual appliance — Deploy a new Unitrends Backup appliance.

  If the failed appliance is a Unitrends Backup virtual appliance and its backup storage is still available, you can opt to deploy by using this original backup storage to retain those backups.

Requirements and considerations for appliance DR

The following requirements must be met before you perform appliance DR:

- The DR target appliance must be freshly imaged or newly deployed and must be running the latest version of the Unitrends software.
- The DR target appliance must have local network access enabled:
  - If you can log in to the appliance UI by entering ***https://<applianceIPaddress>/ui/*** in a browser on the local network, access is enabled.
  - If you cannot log in and receive the message *Managed by UniView*, local network access is disabled. Enable local network access as described in "Disable or enable local network access to an appliance" on page 170.
- The operating system of the target appliance must be the same version as the original (failed) appliance or a higher version than the original (failed) appliance. For example, you cannot recover a failed CentOS 7 appliance to a CentOS 6 target appliance. To check the appliance OS version, select **? > About** in the appliance UI:

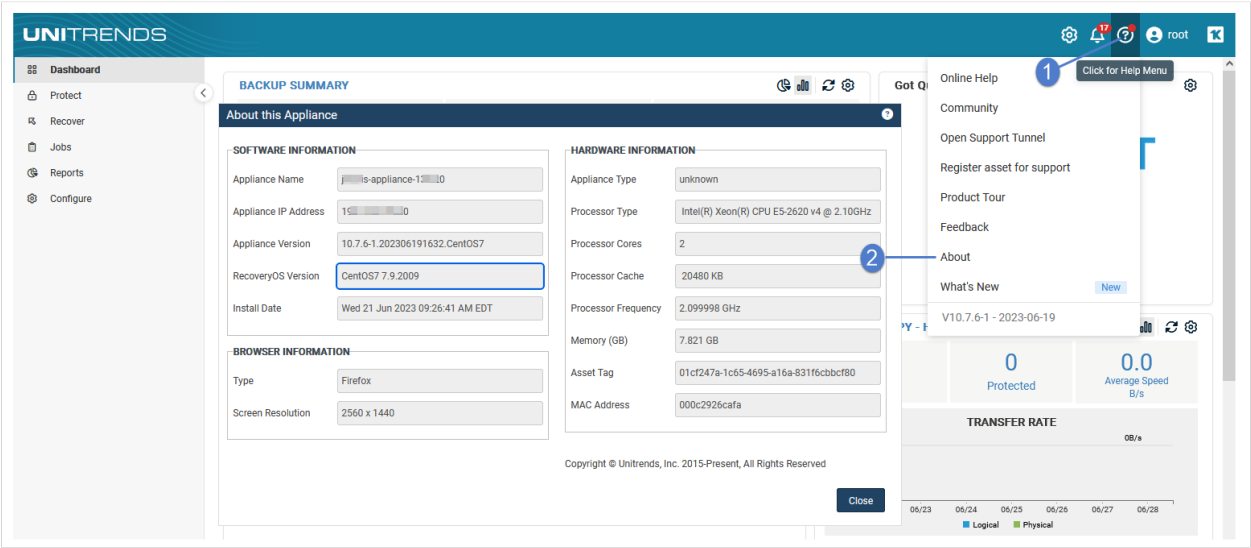

- The DR target appliance must be set up and configured onto the network. See the applicable guide below for instructions:
  - Quick Start Guide for Gen 10 Recovery Series Appliances
  - Quick Start Guide for Unitrends MSP Recovery MAX Appliances
  - Quick Start Guide for ION and ION+ Appliances
  - Deployment Guide for Unitrends Backup on VMware
  - Deployment Guide for Unitrends Backup on Hyper-V
  - Deployment Guide for Unitrends Backup on Nutanix AHV
  - Deployment Guide for Unitrends Backup on Citrix XenServer
  - Deployment Guide for Unitrends Backup in Microsoft Azure
  - Deployment Guide for Unitrends Free on VMware
  - Deployment Guide for Unitrends Free on Hyper-V
- The DR target appliance must have a minimum of 200GB of backup storage space available or as much storage as the original appliance, whichever is greater.
- You must have a hot or cold backup copy from which to recover. (DR from tape backup copy is not supported.)

Additional DR considerations are listed here. You will be asked to make choices during the DR process. It is best to consider your options before you start the DR procedure.

- During DR you will choose whether to retain the DR target appliance's storage settings or to recover the storage settings from the original (failed) appliance. For details on each option, see "Selecting storage configuration during DR".

- If encryption was configured on the original appliance, you must configure encryption on the DR target appliance before you start the DR procedure. You must configure encryption to use the encryption passphrase of the original appliance. For details, see "Encryption" on page 155.
- During DR you will choose a storage target where backups will be imported:
  - For Recovery Series and Recovery MAX physical appliances — Accept the default *Internal* storage target (this is the only storage target).
  - For Unitrends Backup virtual appliances — If you want to import backups to a different storage target, you must add the storage to the target appliance before you start the DR procedure. See "About adding backup storage to a Unitrends Backup appliance" on page 199 for details.
- The DR procedure does not recover any Windows replicas or VM replicas. You must create new replicas after running the DR procedure.
- The DR procedure does not recover any SLA policies. Any backup and backup copy schedules created by SLA policies are recovered, but you must recreate SLA policies after running the DR procedure.

Selecting storage configuration during DR

During the DR procedure, you will be asked whether to retain the storage settings on the DR target appliance or to recover storage settings from the original appliance. Differences are described here:

Storage	Description
DR using storage settings of the new target appliance	Storage configuration of the target appliance is retained while recovering system metadata. Schedules from the original appliance are updated to use the default backup storage on the new appliance. Example where DR target appliance storage settings are retained during DR: • Original (failed) appliance: two backup storage targets, *Internal* and *NewBackups*. Backups for Asset1 go to *NewBackups*. • New DR target appliance: two backup storage targets, *Internal* and *MoreBackups*. After DR, the new appliance still contains storage targets *Internal* and *MoreBackups*. Since the original *NewBackups* target does not exist, the schedule for Asset1 is updated so that its backups are written to the default storage, *Internal*.
DR using storage settings of the original appliance	Storage configuration of the original appliance is recovered. Schedules from the original appliance use the original configuration. Example where original appliance storage settings are retained during DR: • Original (failed) appliance: two backup storage targets, *Internal* and *NewBackups*. Backups for Asset1 go to *NewBackups*. • New DR target appliance: two backup storage targets, *Internal* and *MoreBackups*. After DR, backups for Asset1 continue to be written to the *NewBackups* target.

Performing DR from a hot backup copy

To recover the appliance by using a hot backup copy, verify that the requirements in "Preparing for appliance DR from a hot backup copy" have been met, then run the "To perform appliance DR from a hot backup copy" procedure.

IMPORTANT! This procedure configures the DR target appliance to match the original (failed) appliance. The procedure overwrites the appliance's database and hosts file and removes any existing backups.

Preparing for appliance DR from a hot backup copy

These prerequisites must be met before you perform DR from a hot backup copy:

- The original (failed) appliance must be running Unitrends version 9.1 or higher.
- The DR target appliance must be set up and configured to meet the general "Requirements and considerations for appliance DR" on page 1296.
- The DR target appliance must be given a unique, temporary hostname (this name will be replaced with the hostname of the failed appliance during the DR process). Do not use a hostname that has ever been known to the failed appliance or to the hot backup copy target appliance.
- Add the DR target appliance to the hot backup copy target appliance that houses the copy you will use for the DR. For details, see "To add an appliance" on page 419.
- If recovering encrypted backups, you must configure encryption on the DR target appliance. While configuring encryption, you must enter the encryption passphrase of the original (failed) appliance. For details, see "Encryption" on page 155.

To perform appliance DR from a hot backup copy

Use this procedure to recover the original appliance from a hot backup copy. DR is performed using a command line DR tool that you run from the backup copy target appliance (the appliance that received backup copies from the original appliance). You can quit the recovery process within the DR tool at any time by entering **E** (exit) at the `Selection:` prompt.

1. Using a terminal emulator, such as PuTTY, connect to the backup copy target appliance using the following:
   - Appliance IP address
   - Port 22
   - SSH connection type
2. Log in as user *root*.
3. Issue this command to launch the DR tool:

   ```
   # /usr/bp/bin/disaster_recovery
   ```
4. At the `Selection:` prompt, enter **1** to recover from a hot backup copy target.
5. Follow the prompts to perform the DR. See "Selection details" for a description of each selection.

6 While copies are being imported to the new appliance, you can exit the DR tool and close the terminal session. To monitor progress of the import, log in to the DR target appliance UI and select **Jobs > Active Jobs** (or **Jobs > Recent Jobs** to see completed imports).

7 (Optional) If you have copied backups of additional assets from the original (failed) appliance to another backup copy target, use these steps to recover those backup copies:

- Using a terminal emulator, such as PuTTY, connect to the second backup copy target appliance using the following:
  - Appliance IP address
  - Port 22
  - SSH connection type
- Log in as user *root*.
- Issue this command to launch the asset-only DR tool:

```
# /usr/bp/bin/disaster_recovery clientsOnly
```

- Follow the prompts to select the original (failed) appliance that you are recovering, the target (new) appliance where backup copies will be recovered, and the assets whose last backups will be imported to the DR target appliance.
- While copies are being imported to the new appliance, you can exit the DR tool and close the terminal session. To monitor progress of the import, log in to the DR target appliance UI and select **Jobs > Active Jobs** (or **Jobs > Recent Jobs** to see completed imports).
- Repeat these steps as needed to import copies from additional target appliances.

8 License the appliance as described in "Licensing the DR target appliance" on page 1304.

9 Disaster recovery sets the appliance root operating system password to this default value: *unitrends1*. Change the appliance root password as described in "Change the appliance operating system password" on page 146.

Selection details

DR tool prompt	Description
Select appliance to restore	Select the original (failed) appliance that you will recover.
Choose target for recovery	Select the new appliance where the failed appliance will be recovered. If you do not see your DR target appliance in the list, add its hostname and IP (by entering the number for the `Add new host to hosts file` option).
Please go to *DRtargetAppliance* and run 'disaster_ recovery metadata'	When you see this prompt, leave the DR tool open and running. Open a second terminal emulator session and do these steps: 1 Connect to the new DR target appliance using the following: – Appliance IP address – Port 22 – SSH connection type 2 Log in as user *root*. 3 Issue this command to recover metadata to the DR target appliance: `# /usr/bp/bin/disaster_recovery metadata` 4 Enter **y** to continue. 5 Choose a storage configuration by entering one of the following: • **y** to use the storage settings of the new DR target appliance • **n** to use the storage settings of the original (failed) appliance 6 Enter **y** to continue. 7 When you see `metadata restore completed`, return to the DR tool and press any key to continue.
Select Assets	Specify the assets whose last backups will be imported to the DR target appliance. This step is optional. • Enter **a** to import last backups of all assets. • Enter a number or numbers to import last backups of a subset of assets (examples: **2** or **1,2,4**). • Enter **e** to exit without importing backups.

DR tool prompt	Description
	Note: If you deployed a Unitrends Backup appliance by using storage from the original (failed) appliance, all backups on that storage are already present on the DR target appliance.
Backup copies from the selected assets to the new appliance have started	Last backups are being imported to the DR target appliance. You can now exit the DR tool and close the terminal session. To monitor progress of the import, log in to the DR target appliance UI and select **Jobs > Active Jobs** (or **Jobs > Recent Jobs** to see completed imports).

Performing DR from a cold backup copy

To recover the appliance by using a cold backup copy, verify that the requirements in "Preparing for appliance DR from a cold backup copy" have been met, then run the "To perform appliance DR from a cold backup copy" procedure.

IMPORTANT! This procedure configures the DR target appliance to match the original (failed) appliance. The procedure overwrites the appliance's database and hosts file and removes any existing backups.

Preparing for appliance DR from a cold backup copy

These prerequisites must be met before you perform DR from a cold backup copy:

- The DR target appliance must be set up and configured to meet the general "Requirements and considerations for appliance DR" on page 1296.
- If recovering encrypted backups, you must configure encryption on the DR target appliance. While configuring encryption, you must enter the encryption passphrase of the original appliance. For details, see "Encryption" on page 155.
- The DR target appliance must have access to the cold copy you will use for the DR procedure. Do the following to enable access to the cold copy:
  - Add the cold backup copy target to the DR target appliance as described in "Backup copy targets" on page 214.
  - If the cold copy target is a USB or eSATA device, load the disk(s) containing the cold copy into the device.

To perform appliance DR from a cold backup copy

Use this procedure to recover the original appliance from a cold backup copy. DR is performed using a command line tool that you run from the DR target appliance (the appliance to which you will recover the original failed appliance). You can quit the recovery process within the DR tool at any time by entering **E** (exit) at the `Selection:` prompt.

1. Using a terminal emulator, such as PuTTY, connect to the DR target appliance using the following:
   - Appliance IP address

- Port 22
- SSH connection type

2 Log in as user *root*.

3 Issue this command to launch the DR tool:

```
# /usr/bp/bin/disaster_recovery
```

4 At the `Selection:` prompt, enter **2** to recover from a cold backup copy target.

5 Follow the prompts to perform the DR. See "Selection details" for a description of each selection.

6 After you perform the recovery, license the appliance as described in "Licensing the DR target appliance" on page 1304.

7 Disaster recovery sets the appliance root operating system password to this default value: *unitrends1*. Change the appliance root password as described in "Change the appliance operating system password" on page 146.

Selection details

DR tool prompt	Description
Select Target	Select the cold backup copy target.
Would you like to use the storage configuration of the new or the original appliance?	Choose a storage configuration by entering one of the following: • **y** to use the storage settings of the new DR target appliance • **n** to use the storage settings of the original (failed) appliance
Would you like to restore appliance metadata?	Enter **y** to continue.
Encryption is enabled (only displays if the cold copy is encrypted and you have not configured encryption on the DR target appliance)	If you see this prompt, leave the DR tool open and running. Log in to the DR target appliance UI and configure encryption. Enter the encryption passphrase of the original (failed) appliance. For details, see "Encryption" on page 155. Once you have configured encryption, return to the DR tool and continue with the DR procedure.
Select Assets	Specify the assets whose last backups will be imported to the DR target appliance. This step is optional. • Enter **a** to import last backups of all assets. • Enter a number or numbers to import last backups of a subset of assets (examples: **2** or **1,2,4**). • Enter **e** to exit without importing backups.

DR tool prompt	Description
	Note: If you deployed a Unitrends Backup appliance by using storage from the original (failed) appliance, all backups on that storage are already present on the DR target appliance.
Select Storage for *DRtargetAppliance*	Select the storage where imported backups will be written. To use the default (Internal) storage, enter `1`. If your DR target is a Unitrends Backup appliance and you have added more storage targets, additional options display.
This will overwrite backups.	Any existing backups stored on the DR target appliance will be overwritten. Enter `y` to continue.
Beginning recovery, this will take some time.	Last backups are being imported to the DR target appliance. You can now exit the DR tool and close the terminal session. To monitor progress of the import, log in to the DR target appliance UI and select **Jobs > Active Jobs** (or **Jobs > Recent Jobs** to see completed imports).

Licensing the DR target appliance

After performing DR, use one of the following procedures to license the DR target appliance:

- "To license a Recovery Series or Recovery MAX appliance"
- "To license a Unitrends Backup appliance" on page 1305

To license a Recovery Series or Recovery MAX appliance

1. Log in to the DR target appliance UI.
2. On the **Configure > Appliances** page, select the appliance and click **Edit**.
3. On the Edit Appliance dialog, click the **License** tab.
4. Click **Add License Info** and enter the following information from the original (failed) appliance:
   - License Key
   - Feature String
5. Click **Save** to apply the license information.

To license a Unitrends Backup appliance

Because the newly deployed DR target appliance has a different MAC address, you must apply a new license. Do one of the following:

- If you have the email containing the license key of the original (failed) appliance:

  1. From the UI, select **Configure** > *ApplianceName* > **Edit** > **License** > **Upgrade**. The registration form displays.
  2. Select **I need to activate my purchase**.
  3. Enter the email address where you want to receive the license key.
  4. In the Unitrends Backup Registration Center, enter the activation code from the license email and submit the form.

  A new license key will be sent to the email you specified. Once you receive it, apply the license as described in "To license a Recovery Series or Recovery MAX appliance" on page 1304 above.

- If you no longer have the email containing the license key of the original (failed) appliance, contact Untirends Support (as described in "Support for Unitrends appliances" on page 25) to request a new license.

  Once you receive the new license, apply it as described in "To license a Recovery Series or Recovery MAX appliance" above.

This page is intentionally left blank.

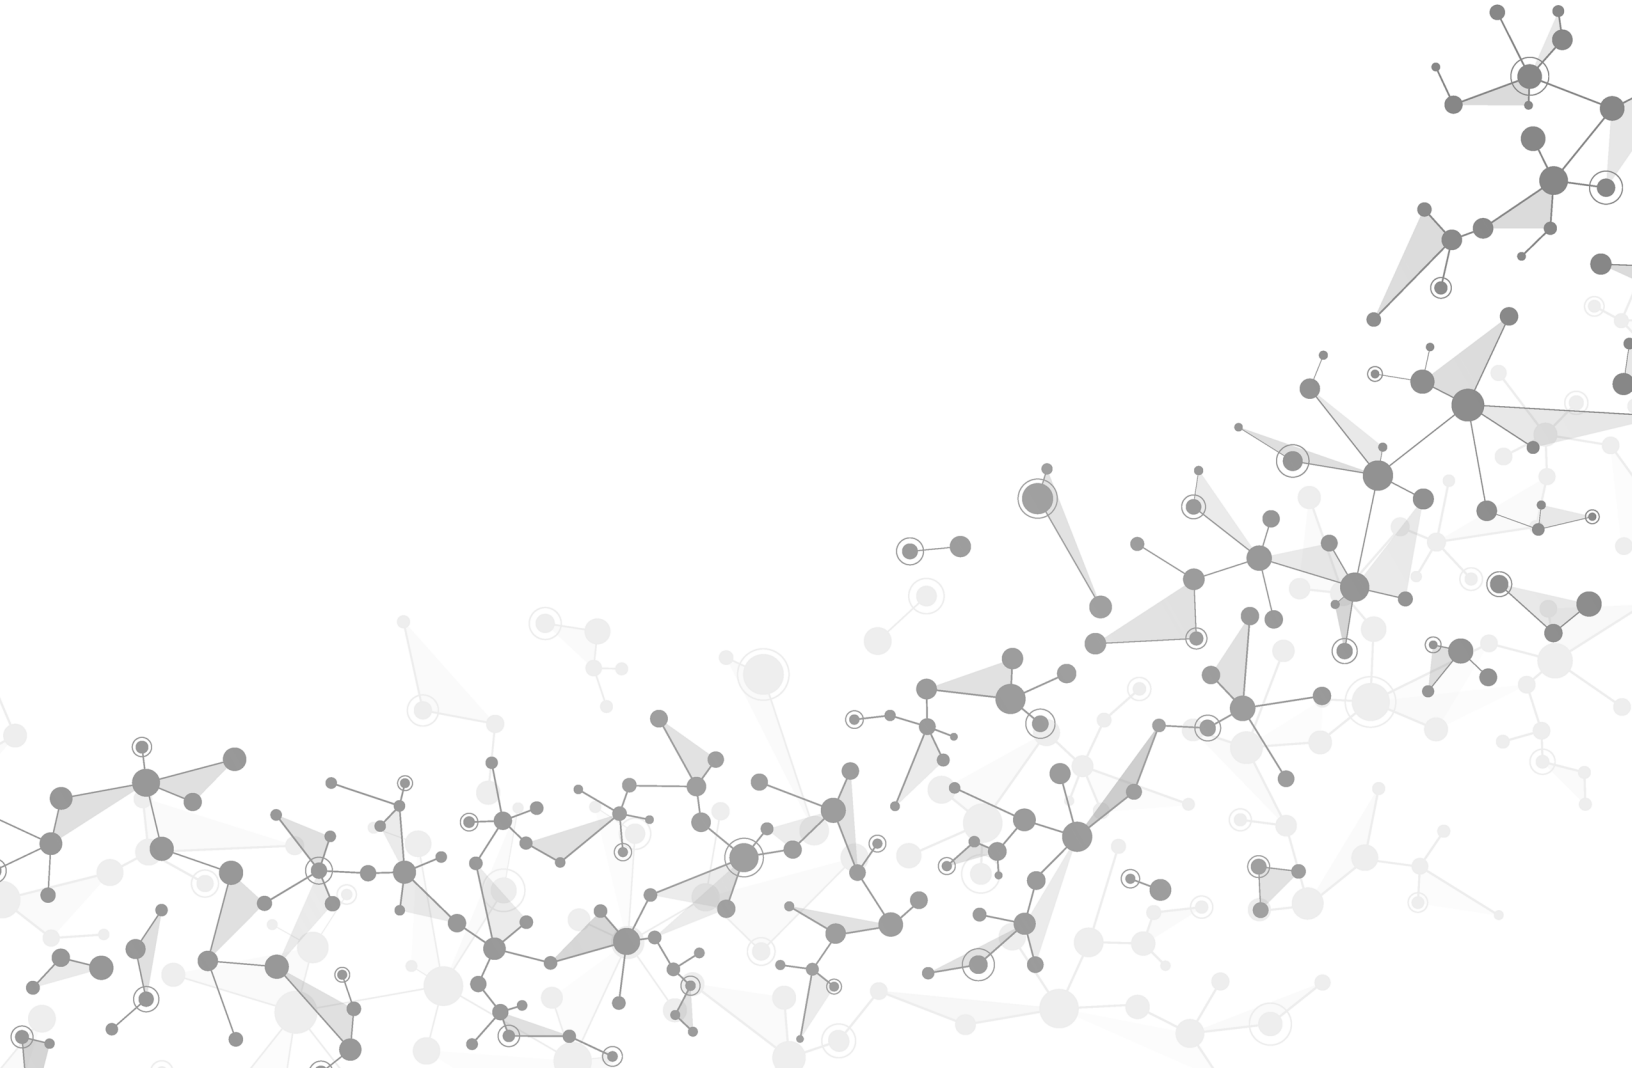

Chapter 24: Reports

Your Unitrends data protection solution provides the ability to create, generate, edit, schedule, and distribute a variety of reports.

You can generate reports to display information about backup jobs, recovery operations, backup copies, as well as your appliance and its storage use. When generating reports, you can configure the number of rows that display per page, filter results to display only entries that meet specified parameters, and alter the sort order for values.

Once generated, you can export all reports as a PDF or CSV (Excel) file.

See the following topics for more information:

- "Working with reports" on page 1307
- "Types of reports" on page 1313

Working with reports

Use the following procedures to generate, export, and distribute reports:

- "To generate a report" on page 1307
- "To customize a report" on page 1308
- "To export a report" on page 1309
- "Emailing reports" on page 1310
- "Modifying the From address for Appliance reports" on page 1311

To generate a report

1. Select **Reports** and click a **Category**.
2. Expand the report in the Available Reports list.
3. Select the appliance.
4. Specify the date range by selecting a From date and a To date.
5. Click **Generate Report**.

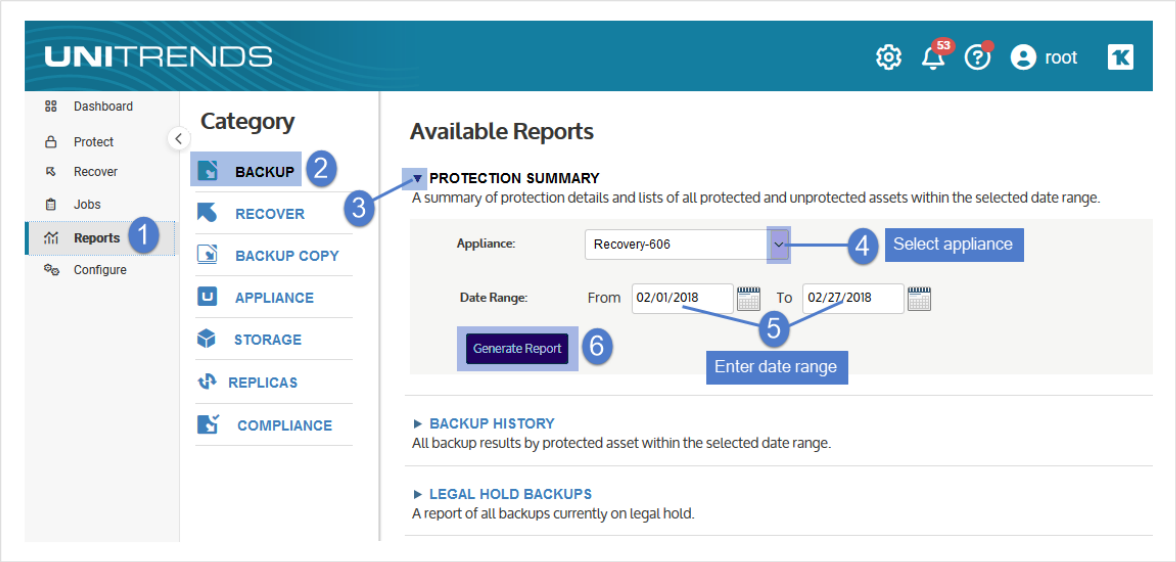

6 The report displays. For details on customizing the display, see "To customize a report". When you are finished viewing the report, click **Report Categories** to return to the Available Reports list.

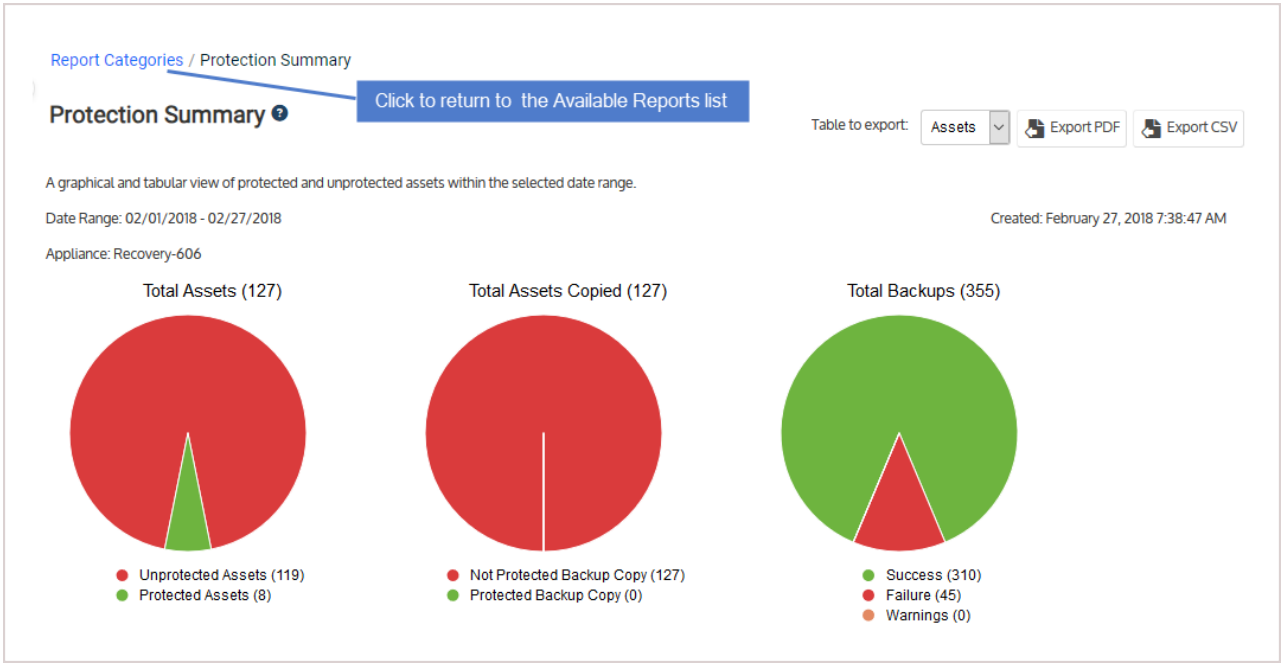

To customize a report

After generating a report, you can customize the display as follows:

- Click the icon to add or remove columns. Click **Reset column defaults** to restore the default display. Click **Clear all filters** to clear any column filters you have applied.
- Click the arrow next to a column name to sort values in ascending or descending order.
- Hover over a column border and drag to resize a column.
- Enter text in a column's filter field to display only rows that contain the string you entered.
- Modify the number of rows displayed on each page by selecting a new value in the **rows per page** field.

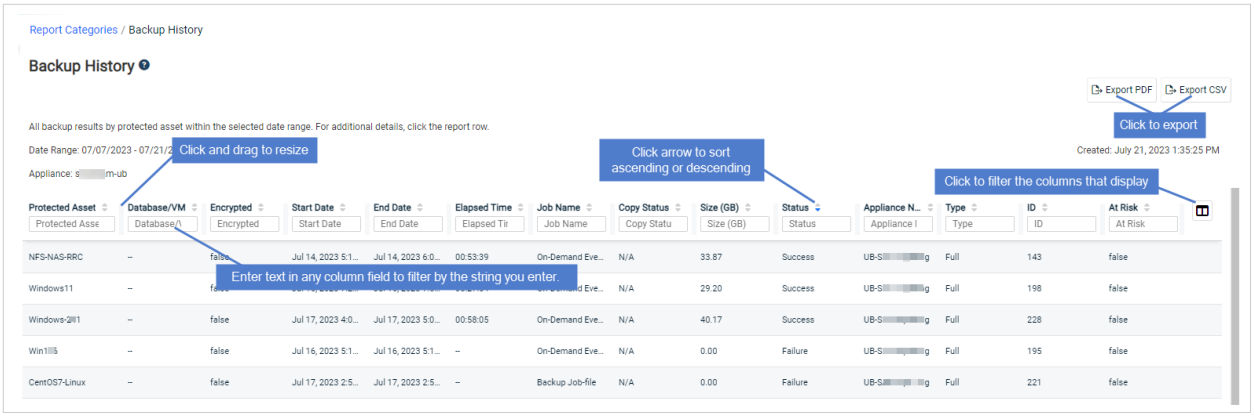

To export a report

Use this procedure to export a report as a PDF or CSV file.

1 Generate the report as described in "To generate a report" on page 1307.

2 From the generated report, do one of the following:

- To export to PDF, click **Export PDF**, select a row option, and click **Export**. The PDF report is downloaded to your computer.

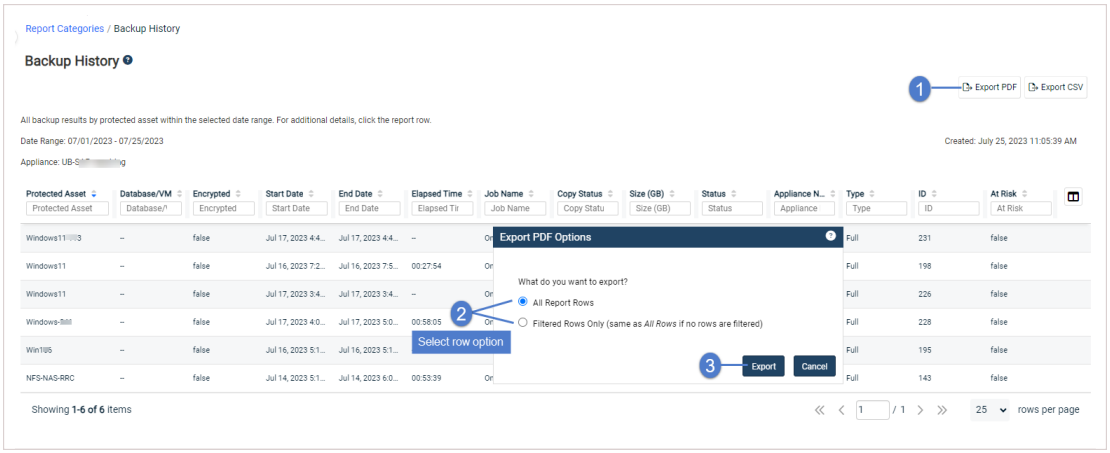

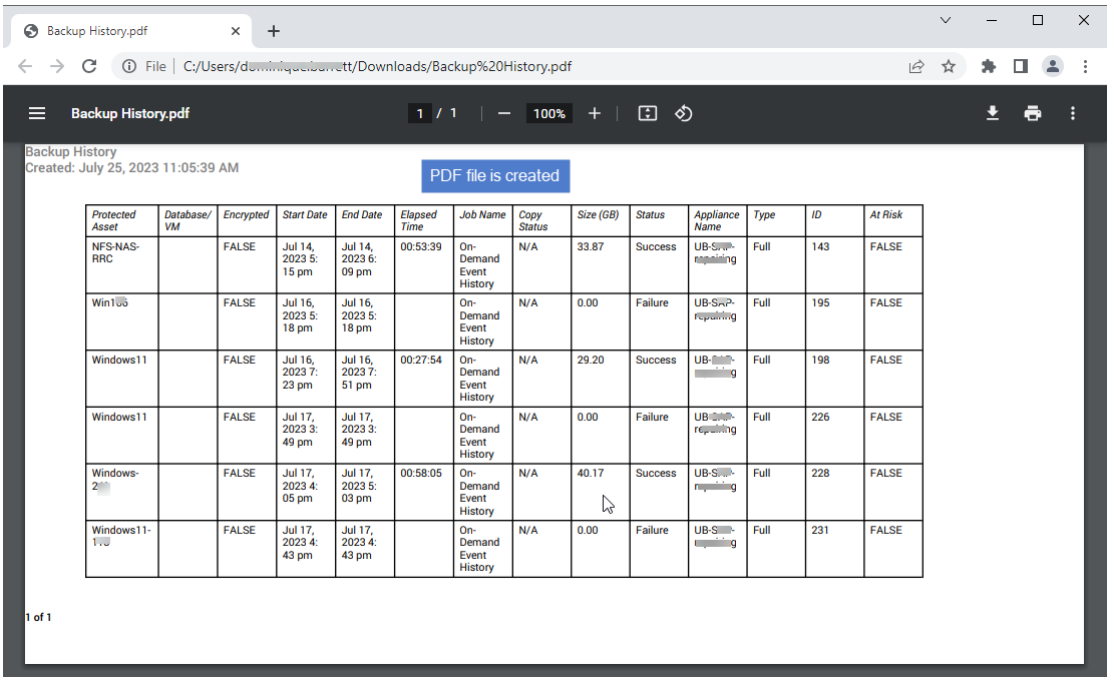

- To export to CSV, click **Export CSV**. The CSV report is downloaded to your computer.

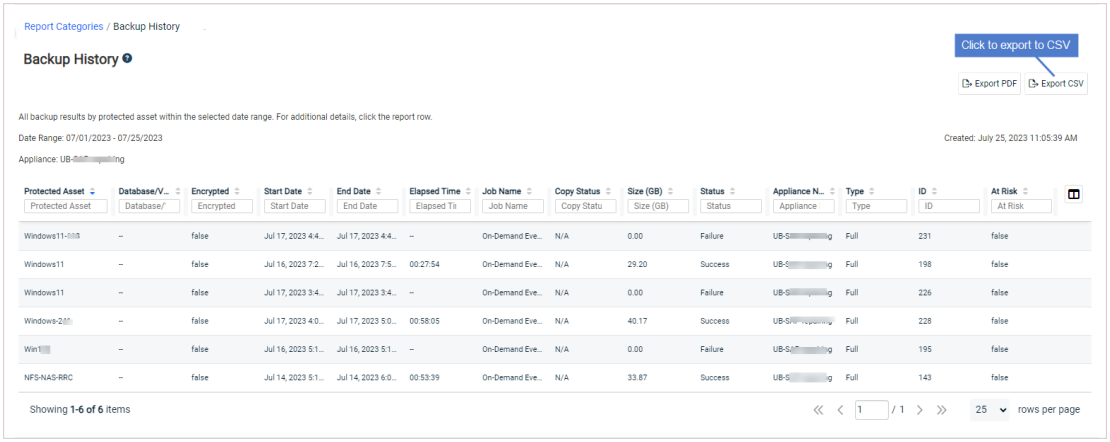

3 When finished, click **Report Categories** to return to the Available Reports list.

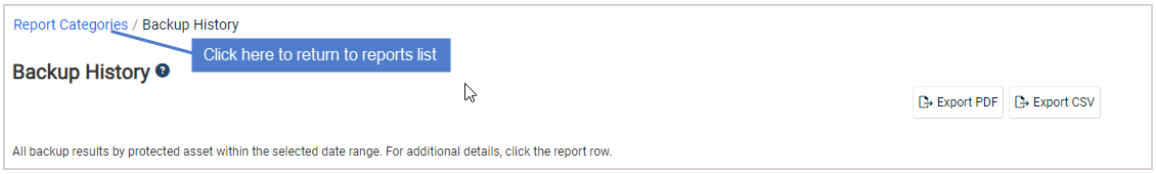

Emailing reports

You can configure an appliance to automatically email reports to designated recipients.

For additional information on configuring an appliance to send reports automatically by email, and on creating recipient lists, see "Email reporting" on page 117.

Modifying the *From* address for Appliance reports

If you have configured email recipients to receive System reports, the appliance sends these reports each day: Appliance Status report and Management Status report. The reports are emailed from this address: reports@*ApplianceHostname*. If desired, you can change this address to reports@*ApplianceDomainName*.

Note: Some relay servers are configured to modify the *From* address to a specified sender. If this is the case in your environment, this procedure will not update the *From* address used for Unitrends System reports.

To change the *From* address to reports@*ApplianceDomainName*

1 On the **Configure > Appliances** page, select the appliance and click **Edit**.

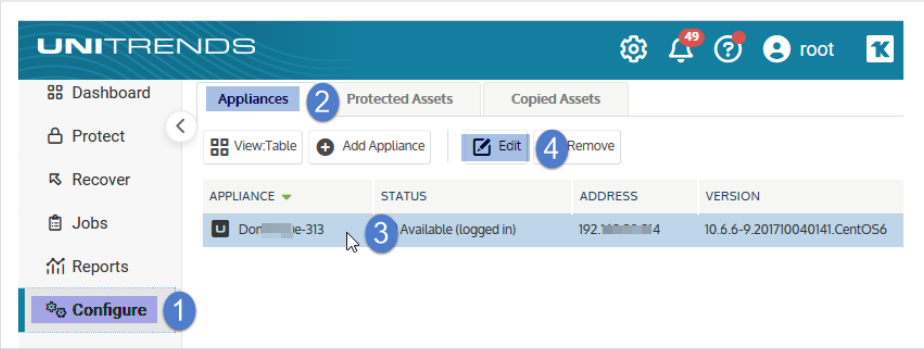

2 Click **Advanced** and select **General Configuration**.

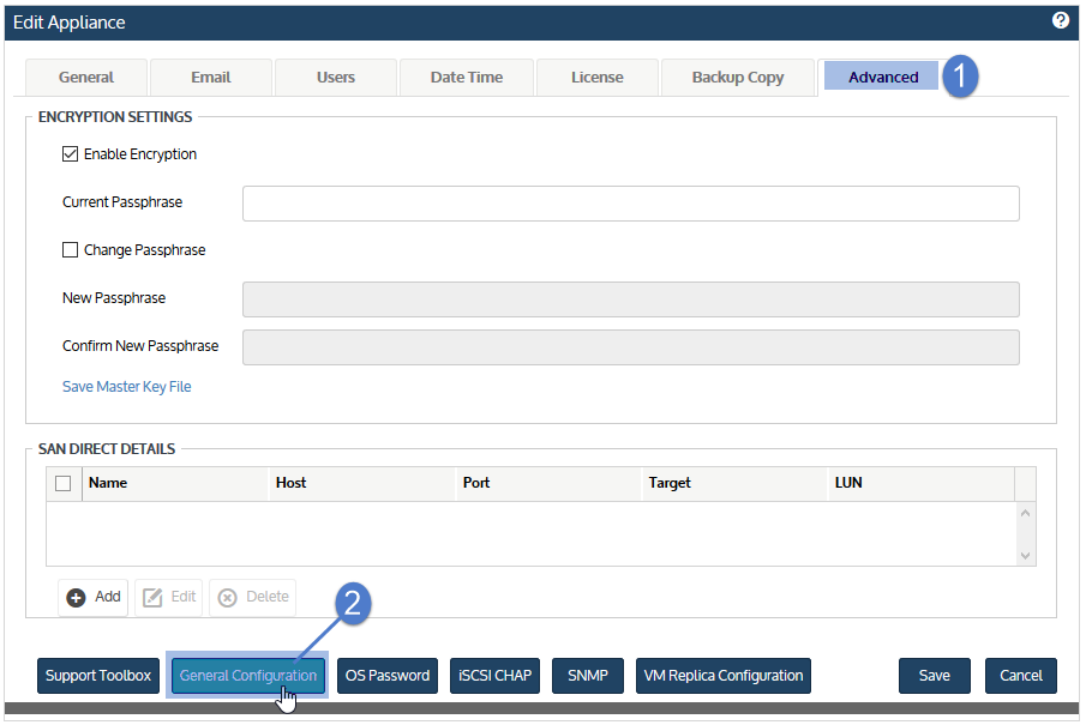

The General Configuration (Advanced) dialog displays.

3 Enter **system** in the Name field, then click the **SystemReportFromDomainNotHostname** row.

4 Enter **YES** in the Value field, then click **Save**.

5 Click **Close** to exit.

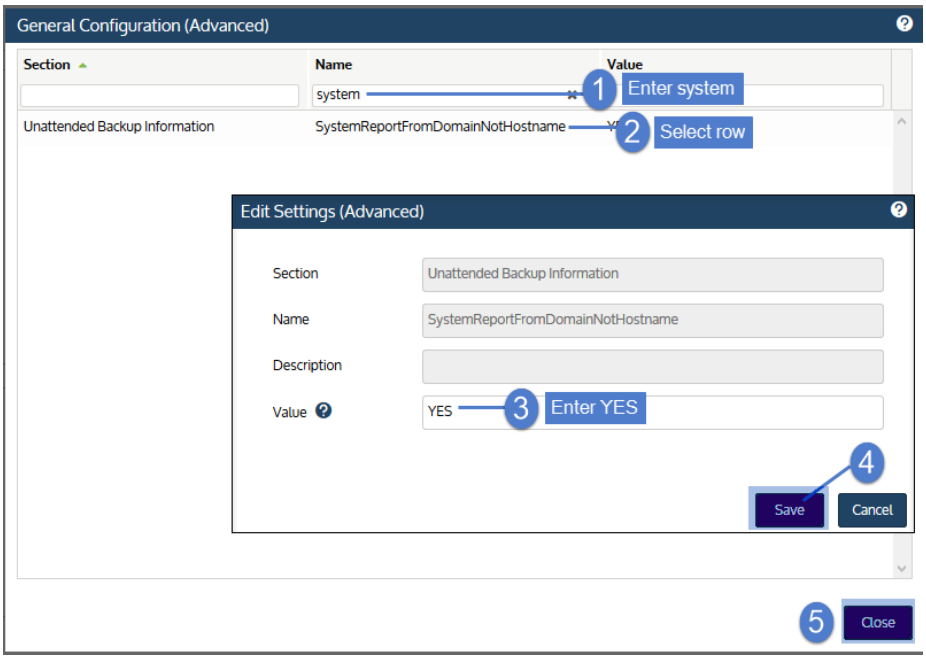

Types of reports

Reports are grouped by category. Select a category to view the available reports. See the following topics for details:

Category	Available Reports
"Backup reports" on page 1314	"Protection Summary report" on page 1331 "Backup History report" on page 1320 "Backup Failures report" on page 1322 "Weekly Status report" on page 1344 "Protection Policies report" on page 1345
"Recover reports" on page 1327	"Recovery History report" on page 1327 "Recovery Assurance report" on page 1329
"Backup Copy reports" on page 1331	"Protection Summary report" on page 1331 "Backup Copy Capacity report" on page 1336 "Backup Copy - Hot Targets report" on page 1337 "Backup Copies - Past 24 Hours report" on page 1340 "Storage Footprint report" on page 1341 "Backup Copy - Cold Targets report" on page 1342 "Weekly Status report" on page 1344 "Protection Policies report" on page 1345

Category	Available Reports
"Appliance reports" on page 1347	"Update History report" on page 1347 "Capacity report" on page 1348 "Load report" on page 1351 "Alerts report" on page 1352 "Trap History report" on page 1353 "Notifications report" on page 1355
"Storage reports" on page 1356	"Storage report " on page 1356 "Data Reduction report" on page 1359
"Replicas History report" on page 1360	"Replicas History report" on page 1360
"Retention Reports" on page 1363	"Legal Hold Backups report" on page 1363 "Long-Term Retention report " on page 1365 "Min-Max Retention report" on page 1368
"Compliance report " on page 1370	"Compliance report" on page 1370
"Email reports " on page 1372	"Appliance Status report" on page 1372 "Backup Copy Hot Targets report" on page 1374 "Backup Copy Job Notifications report" on page 1376 "Compliance report" on page 1377 "DCA Job Notifications report" on page 1379 "Management Status report" on page 1381 "Schedule report" on page 1384 "Schedule Success report" on page 1387 "Schedule Failure report" on page 1389

Backup reports

The Backup category contains reports that summarize information about your backup jobs. To view Backup reports, click on a report type under **Category**. Under **Available Reports**, click on the name of the report, and click **Generate Report**. (For detailed procedures, see "Working with reports" on page 1307.)

For details on each Backup report, see the following topics:

- "Protection Summary report" on page 1331
- "Backup History report" on page 1320
- "Legal Hold Backups report" on page 1363
- "Backup Failures report" on page 1322

- "Weekly Status report" on page 1344
- "Protection Policies report" on page 1345

Protection Summary report

The Protection Summary report summarizes your appliance's current protection details. The Protection Summary report:

- Prepares a summary log of all backups and backup copies.
- Displays the status of those backups and backup copies.
- Lists all protected and unprotected assets.

After generating a report, the Protection Summary report displays both a graphical and tabular view of protected and unprotected assets within a selected date range.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

Notes:

- The report includes backups that are currently stored on the selected appliance(s). Once a backup has been removed, it is not included in reports.
- The number of protected and unprotected assets does not include Unitrends appliances.

An example report is given here. Refer to the table below for a description of the charts and available report columns. Not all columns display by default. To customize your output, click the  icon to the right of the column header.

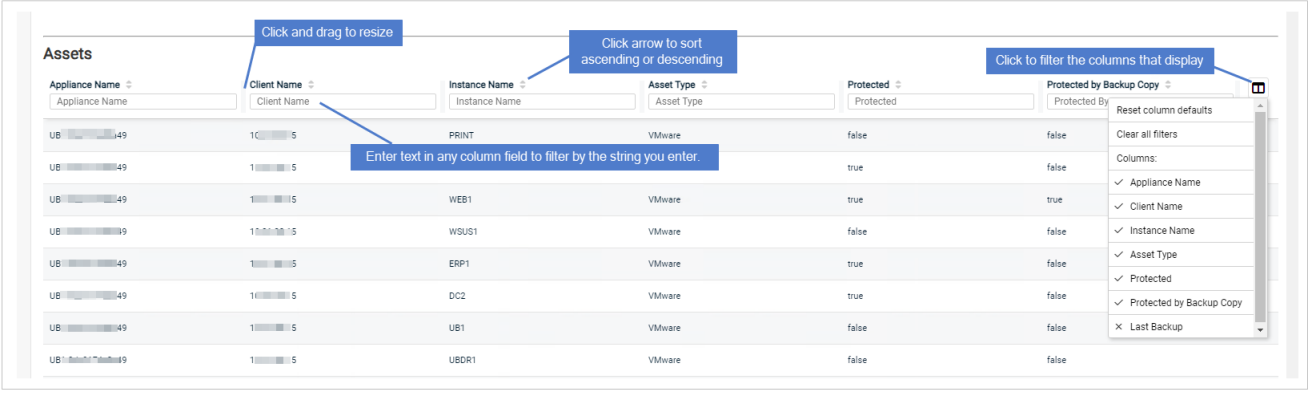

Item	Description
Protection Summary Pie Charts	Graphical view of protected and unprotected assets within the selected date range.

Item	Description
Total Assets	The total number of assets configured on the selected appliance(s). A pie chart shows the number of protected assets versus unprotected assets. Asset count does not include Unitrends appliances.
Total Assets Copied	The total number of assets and a pie chart that shows the number of assets protected by backup copies versus the number not protected by backup copies. Asset count does not include Unitrends appliances.
Total Backups	The total number of backups currently stored on the selected appliance(s) that ran within the specified date range. A pie chart shows the number of successful backups, the number of backups that completed with warnings, and the number of backups that failed.
Total Backups Copied	The total number of backups copies that ran within the specified date range. A pie chart shows the number of successful backups copies and the number of failed backup copies.
Backup Copies (*Target*)	Shows number of successful and failed backup copies by target.
Assets Table Column	Lists all assets configured on the selected appliance(s), regardless of whether backups or backup copies have run within the selected date range. See the Last Backup and Protected By Backup Copy columns to determine whether the asset has been protected within the specified date range.
Appliance Name	The name of the backup appliance.
Client Name	The name of the asset.
Instance Name	The name of the protected instance. Examples: • Agent-based asset - Contains *image-level* or is empty for file-level. • Virtual machine - Contains the VM name. • Database - Contains the database name. • Storage group - Contains the storage group name.
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, Linux, Oracle, Hyper-V, or VMware.
Protected	Indicates whether there is a successful backup of this asset that ran within the

Item	Description
	specified date range: True (yes) or False (no).
Protected by Backup Copy	Indicates whether a successful backup copy ran for the asset within the specified date range: True (yes) or False (no).
Last Backup	Start date and time of this asset's last successful backup that ran within the specified date range.
Backup Table Column	Lists all assets that have been protected by backups within the selected date range and provides summary backup information.
Appliance Name	The name of the backup appliance.
Client Name	The name of the asset.
Instance Name	The name of the protected instance. Examples: • Agent-based asset - Contains *image-level* or is empty for file-level. • Virtual machine - Contains the VM name. • Database - Contains the database name. • Storage group - Contains the storage group name.
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, Linux, Oracle, Hyper-V, or VMware.
Backup Types	Backup modes of this asset's backups that ran within the specified date range and are stored on the appliance. Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, and/or Transaction (SQL only).
Successes	The number of backup jobs that completed successfully within the specified date range.
Warnings	The number of backup jobs that completed with warnings within the specified date range.
Failures	The number of backup jobs that failed within the specified date range.
Last Backup	Start date and time of this asset's last successful backup that ran within the specified date range.

Item	Description
Next Backup	Date and time that this asset's next scheduled backup will run.
Next Backup Job	Name of the next backup job that will run for this asset.
Backup Copies to *Target* Table Column	Lists all assets that have been protected by backup copies on the given target, within the selected date range. The table provides summary information about these backup copies.
Appliance Name	The name of the backup appliance.
Client Name	The name of the asset.
Instance Name	The name of the protected instance. Examples: • Agent-based asset - Contains *image-level* or is empty for file-level. • Virtual machine - Contains the VM name. • Database - Contains the database name. • Storage group - Contains the storage group name.
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, Linux, Oracle, Hyper-V, or VMware.
Backup Copy Types	Backup modes of this asset's backups that were copied to the target within the specified date range. Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, and/or Transaction (SQL only).
Backup Copy Successes	The number of backup copy jobs that completed successfully within the specified date range.
Backup Copy Warnings	The number of backup copy jobs that completed with warnings within the specified date range.
Backup Copy Failures	The number of backup copy jobs that failed within the specified date range.
Last Backup Copy	Start date and time of this asset's last successful backup copy that ran within the specified date range.

Backup History report

The Backup History report contains backup results of all protected and unprotected assets within the selected date range for the selected appliance(s).

After generating the report, you can view the details of a specific job by clicking it in the grid. This opens the Backup Status: Report Entry dialog. For file-level backups, you can click **Details** in this dialog to view and/or export a list of files contained in the backup.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

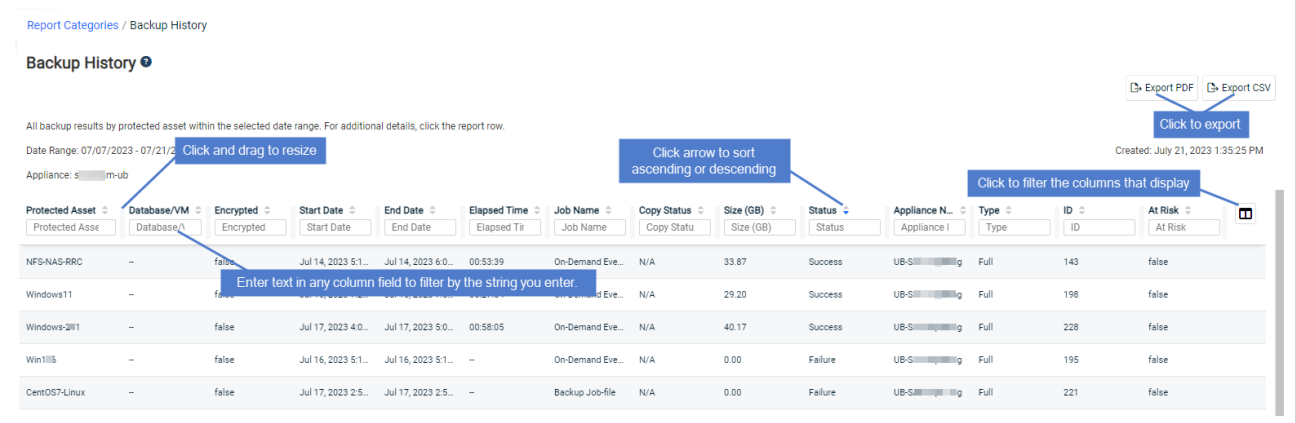

Report Field	Field Description
App Name	Application type: • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013. • For agent-based backups, contains *Image Level* or is empty for file-level. • For System Metadata contains system information, such as appliance configuration and settings. System Metadata is copied to the target during backup copy jobs if changes to this system information are detected.
Protected Asset	The name of the asset.
Complete	Indicates whether the job completed: *True* (yes) or *False* (no).

Report Field	Field Description
Database Name	For application backups, contains the name of the protected database. For host-level backups, contains the name of the protected virtual machine.
Encrypted	Indicates whether the backup is encrypted: *True* (yes) or *False* (no).
Start Date	The date and time at which the backup job started.
End Date	The date and time at which the backup job completed.
Elapsed Time	The amount of time it took for the job to complete, in *hh*:*mm*:*ss* format.
Files	For file-level backups, indicates the number of files in the backup.
Instance Name	For application backups, contains the database instance name. For host-level backups, contains one of the following: • The name of the virtual machine. • The name of the ESXi host if that host is being managed by a vCenter.
Job Name	Name of the backup job. *Job No Longer Exists* displays if this backup is no longer stored on the appliance.
Copy Status	Status that indicates whether the backup has been copied to the Unitrends Cloud or copied to another Unitrends appliance: • Completed - The backup copy job is complete. • Needed - Cannot determine the status of this backup's copy job. • In progress - The backup copy job is currently running. • Waiting - The backup copy job is queued but has not yet started. • N/A - This asset's backups are not configured for hot backup copy.
Size (GB)	The size of the backup, in gigabytes.
Status	Status of the backup job: *Success*, *Warning*, or *Failure*.
Appliance Name	The name of the backup appliance.

Report Field	Field Description
Type	Backup mode: Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, or Transaction (SQL only).
Certify Status	Indicates whether this Recovery Point has been certified by Copy Data Management or by the Unitrends ReliableDR product: *Success* (has been certified), *Failure* (has not been certified), or *None* (this job is not configured for certification).
ID	The system-generated job ID.
Verify Status	Indicates whether the appliance verified the backup and the verification status: *Success* (verified), *Failure* (verify failed), or *N/A* (this job is not configured to verify the backup).
Speed (MB/S)	Average transfer rate of the job, in megabytes per second.
At Risk	Indicates if a potential ransomware infection was discovered: *True* (potentially infected) or *False* (not infected).

Backup Failures report

The Backup Failures report summarizes information about all backup job failures with a specified date range for the selected appliance(s).

After generating the report, you can view the details of a specific job by clicking it in the grid.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

> **Note:** The report includes backups that are currently stored on the selected appliance(s). Once a backup has been removed, it is not included in reports.

An example report is given here. Refer to the table below for a description of the the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

Report Field	Field Description
App Name	Application type: • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013. • For Windows image-level backups, contains *Image Level*. • For file-level backups, contains "--".
Protected Asset	The name of the asset.
Complete	Indicates whether the job completed: *True* (yes) or *False* (no).
Database/VM	For application backups, contains the name of the protected database, instance, or storage group. For host-level backups, contains the name of the protected virtual machine.
Elapsed Time	The amount of time it took for the job to complete, in *hh:mm:ss* format.
Encrypted	Indicates whether the backup is encrypted: *True* (yes) or *False* (no).
End Time	The date and time at which the backup job completed.
Job Name	Name of the backup job. *Job No Longer Exists* displays if this backup is no longer stored on the appliance.

Report Field	Field Description
Instance Name	For SQL and SharePoint application backups, contains the database instance name. For VMware host-level backups, contains one of the following: • The name of the virtual machine. • The name of the ESXi host if that host is being managed by a vCenter.
Copy Status	N/A, which indicates that the backup has not been copied to the Unitrends Cloud or copied to another Unitrends appliance. (Failed backups are not copied.)
Size (GB)	The size of the backup, expressed in gigabytes.
Date	The date and time that the backup job started.
Status	Status of the backup job (*Failure*).
Appliance Name	The name of the backup appliance.
Type	Backup mode: Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, or Transaction (SQL only).
Certify Status	*None*, which indicates that this Recovery Point has not been certified by Copy Data Management or by the Unitrends ReliableDR product. (Failed backups are not certified.)
ID	The system-generated job ID.
Verify Status	Indicates whether the appliance verified the backup and the verification status: *Success* (verified), *Failure* (verify failed), or *N/A* (this job is not configured to verify the backup).

Weekly Status report

The Weekly Status report provides a status of the backups and backup copies that have run on the appliance over the last 7 days. Select an appliance and click **Generate Report** to view the report.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the report columns. To customize your display, click the icon to the right of the column header.

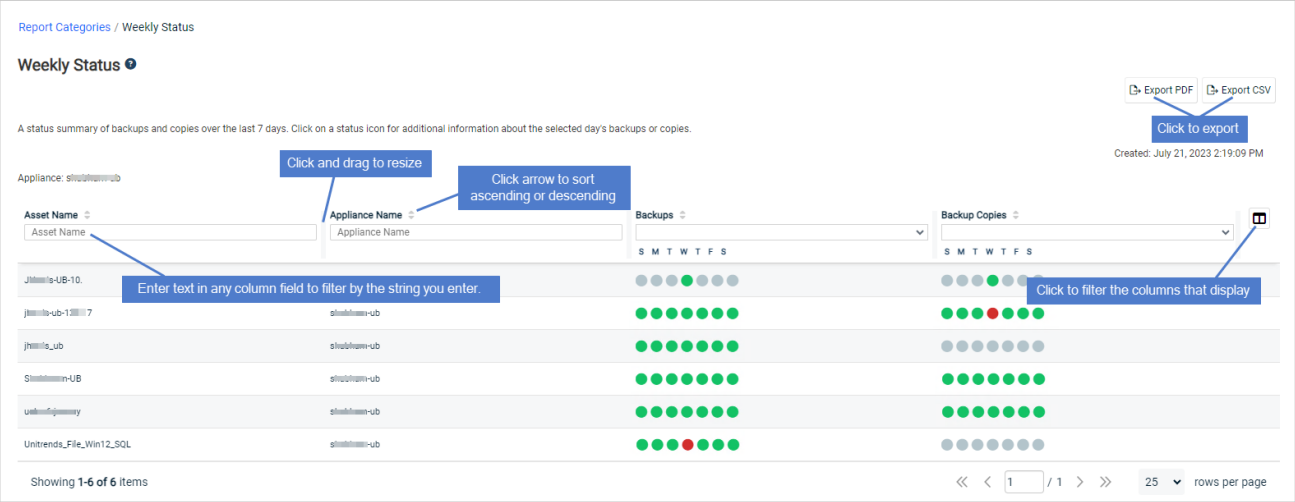

Item	Description
Asset Name	The name of the asset.
Appliance Name	The name of the appliance where the jobs ran.
Backups	Provides a high-level view of an asset's backups that ran over the last 7 days. Icons display for each day of the week: • A green checkmark indicates one or more successful backups ran that day. • A yellow icon indicates one or more backups completed with warnings. • A red icon indicates that one or more backups failed. • A gray circle indicates that no backups ran that day. Click a colored icon for a summary of the day's successes, warnings, failures, and backups in progress.
Backup Copies	Provides a high-level view of an asset's backup copies that ran over the last 7 days. Icons display for each day of the week: • A green checkmark indicates one or more successful backup copies ran that day. • A yellow icon indicates one or more backup copies completed with warnings. • A red icon indicates that one or more backup copies failed. • A gray circle indicates that no backup copies ran that day. Click a colored icon to see the backup copy target and a summary of the day's successes and failures.

Protection Policies report

The Protection Policies report shows protection policies by asset.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. For more detail, click a row to display the asset's Protection Policy Report Entry dialog.

Not all columns display by default. To customize your output, click the icon to the right of the column header.

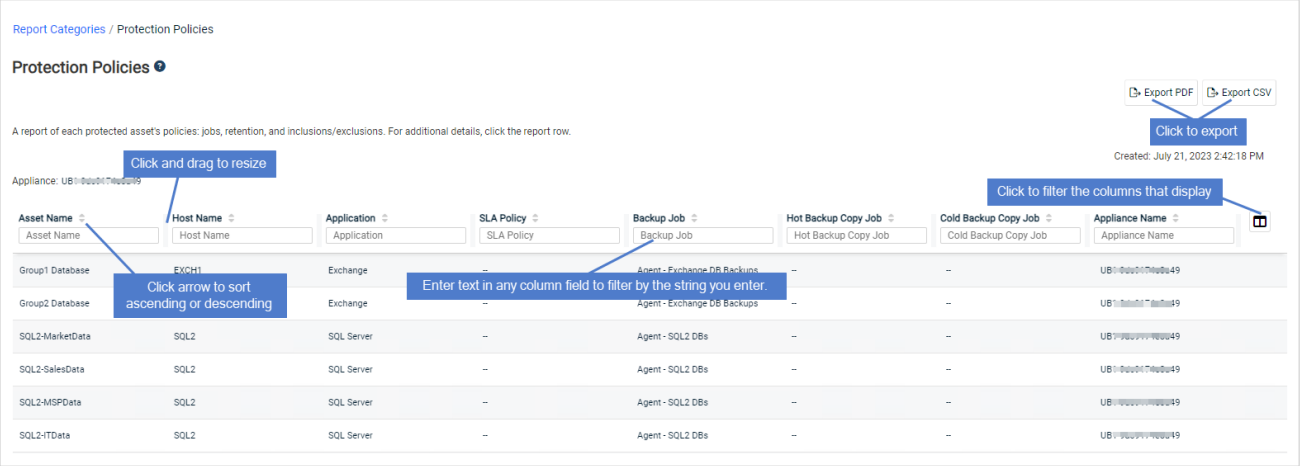

Report Field	Field Description
Asset Name	The name of the asset.
Host Name	Name of the machine hosting the protected asset.
Application	Application type: • For file-level backups, contains *file-level*. • For Windows image-level backups, contains *image-level*. • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V AHV. • For application backups, contains the application type. Examples: SQL Server, Oracle, Exchange Server.
SLA Policy	Name of the SLA policy that is protecting the asset. • This field is empty if the asset is not protected by an SLA policy.

Report Field	Field Description
	• Contains *Visible on local appliance only* if you are logged in to a manager appliance. Click the report row to view SLA policy details in the Protection Policy Report Entry dialog.
Backup Job	Name of the backup schedule that is protecting the asset. The name begins with the prefix *_SLA* if the schedule was created by an SLA policy.
Hot Backup Copy Job	Name of the schedule that is copying this asset's backups to a hot backup copy target. • This field is empty if the asset is not protected by a hot backup copy schedule. • The name begins with the prefix *_SLA* if the schedule was created by an SLA policy.
Cold Backup Copy Job	Name of the schedule that is copying this asset's backups to a cold backup copy target. • This field is empty if the asset is not protected by a cold backup copy schedule. • The name begins with the prefix *_SLA* if the schedule was created by an SLA policy.
Types	Lists the backup modes in the asset's backup schedule: Full, Incremental, Differential, Synthetic Full, Bare Metal, and/or Transaction (SQL only).
Appliance Name	The name of the backup appliance.

Recover reports

The Recover category contains reports that summarize information on Recovery and Recovery Assurance operations.

For details, see the following topics:

- "Recovery History report" on page 1327
- "Recovery Assurance report" on page 1329

Recovery History report

The Recovery History report provides information about recovery operations for a selected appliance(s) over a specified date range.

Notes:

The report does NOT include information about the following:

- VMware and Hyper-V instant recovery operations.

- VMware, Hyper-V, AHV, and XenServer file-level recovery jobs.
- Windows replica operations.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

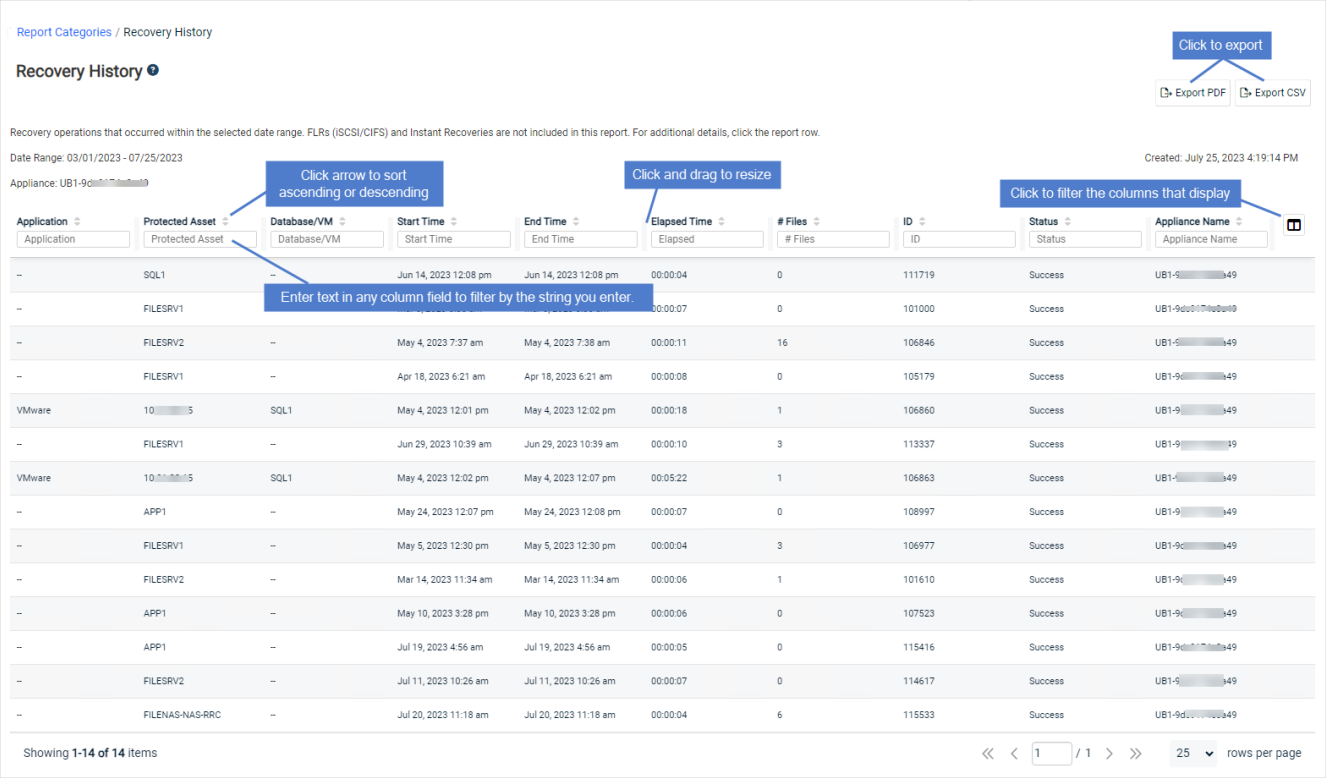

Report Field	Field Description
Application	Application type: • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013.
Protected Asset	The name of the asset.
Complete	Indicates whether the job completed: *True* (yes) or *False* (no).

Report Field	Field Description
Database/VM	For application backups, contains the name of the protected database or storage group. For host-level backups, contains the name of the protected virtual machine.
Start Time	The date and time that the recovery job started.
End Time	The date and time at which the recovery job completed.
Elapsed Time	The amount of time it took for the job to complete, in *hh:mm:ss* format.
# Files	The number of files recovered. (For recovery from host-level, image-level, and application backups, contains *1* as these backups do not contain individual files.)
ID	The system-generated job ID.
Instance Name	For application backups, contains the database instance name. For host-level backups, contains one of the following: • The name of the virtual machine. • The name of the ESXi host if that host is being managed by a vCenter.
Size (GB)	For recovery of Windows image-level backups, shows the amount of data, in gigabytes.
Status	Status of the recovery job: *Success*, *Warning*, or *Failure*.
Appliance Name	The name of the backup appliance.
Type	The type of recovery operation: • *Restore* for a standard backup recovery • *Image restore* for recovery of a Windows image-level backup • *Bare Metal Restore* for unified bare metal recovery.

Recovery Assurance report

The recovery assurance report lists all Recovery Assurance operations that occurred within a selected date range and their respective job results.

After generating the report, you can view the details of a specific job by clicking it in the grid.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

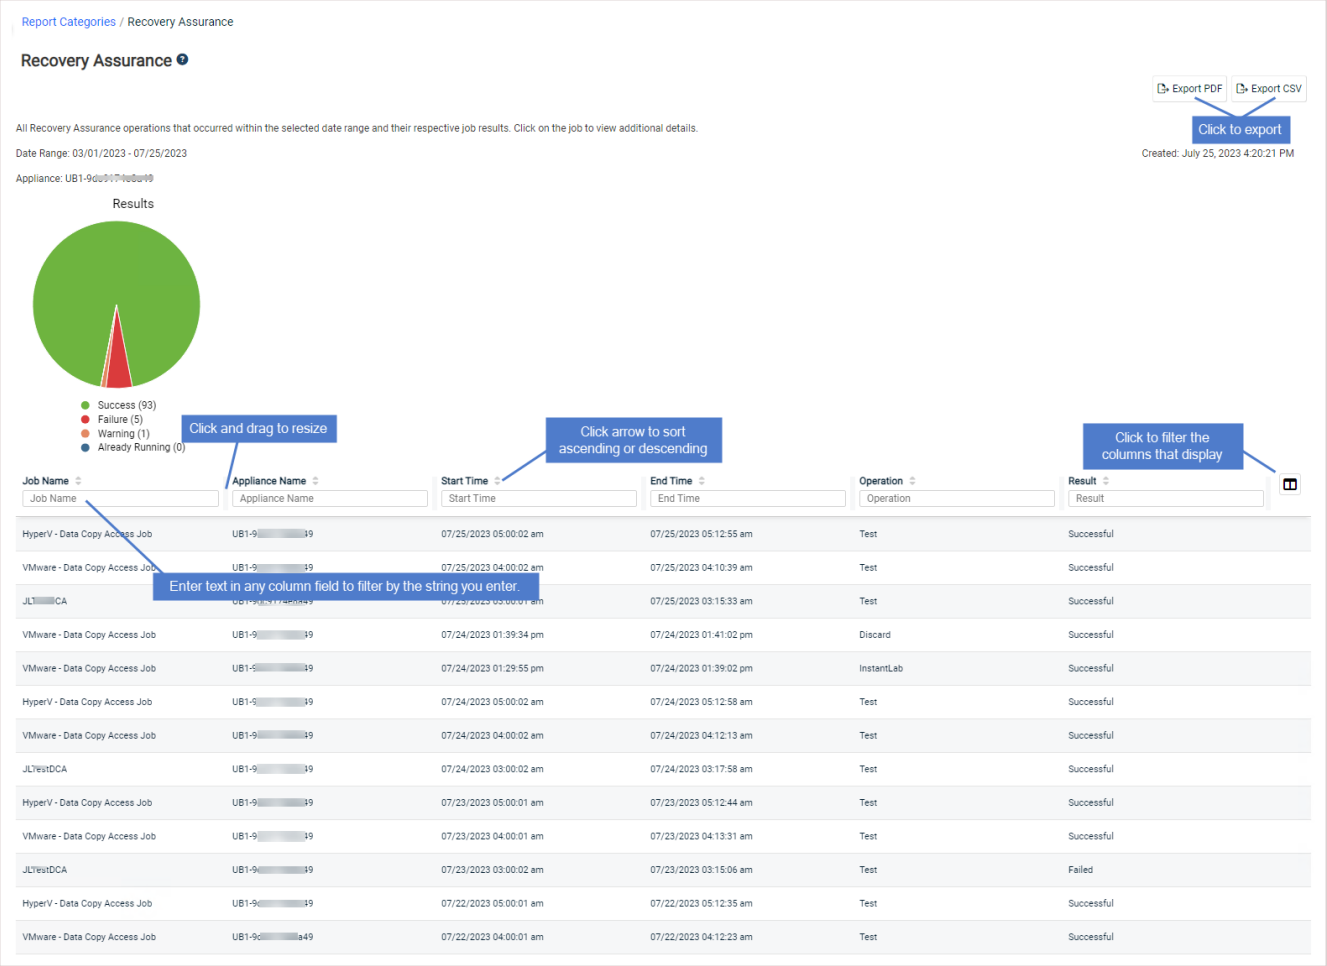

Report Field	Field Description
Job Name	The name of the data copy access job.
Appliance Name	The name of the backup appliance.
Start Time	The date and time that the data copy access job started.
End Time	The date and time that the data copy access job concluded.

Report Field	Field Description
Operation	The mode that the data copy access job ran in. Discarding a data copy access job is considered a distinct operation.
Result	Indicates whether the job failed or was successful.

Backup Copy reports

The Backup Copy category reports summarize information about your backup copy jobs.

Backup copies are duplicates of your backups stored on an off-site target. For an overview of how to use backup copies with your Unitrends appliance, see "Backup copies" on page 101.

For details, see the following topics:

Protection Summary report

The Protection Summary report summarizes your appliance's current protection details. The Protection Summary report:

- Prepares a summary log of all backups and backup copies.
- Displays the status of those backups and backup copies.
- Lists all protected and unprotected assets.

After generating a report, the Protection Summary report displays both a graphical and tabular view of protected and unprotected assets within a selected date range.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

Notes:

- The report includes backups that are currently stored on the selected appliance(s). Once a backup has been removed, it is not included in reports.
- The number of protected and unprotected assets does not include Unitrends appliances.

An example report is given here. Refer to the table below for a description of the charts and available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

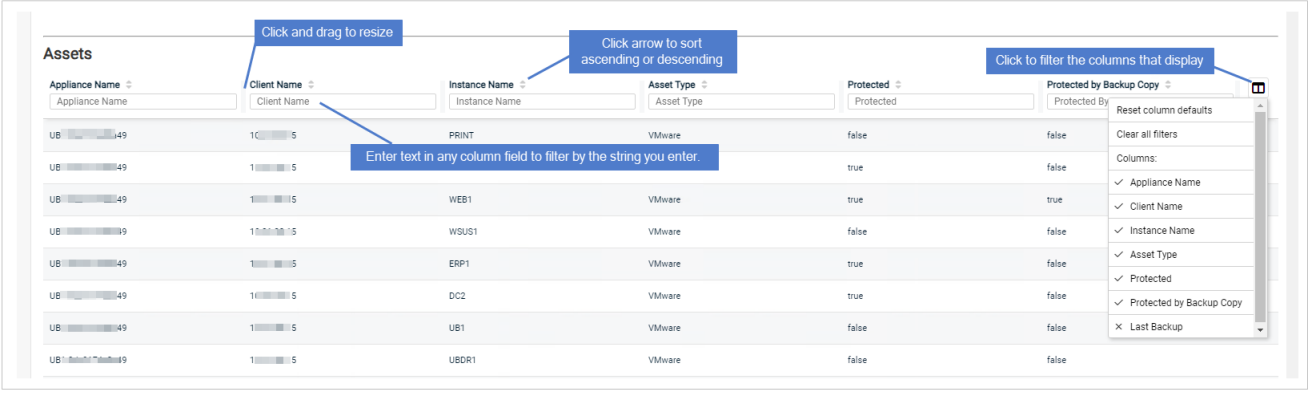

Item	Description
Protection Summary Pie Charts	Graphical view of protected and unprotected assets within the selected date range.
Total Assets	The total number of assets configured on the selected appliance(s). A pie chart shows the number of protected assets versus unprotected assets. Asset count does not include Unitrends appliances.
Total Assets Copied	The total number of assets and a pie chart that shows the number of assets protected by backup copies versus the number not protected by backup copies. Asset count does not include Unitrends appliances.
Total Backups	The total number of backups currently stored on the selected appliance(s) that ran within the specified date range. A pie chart shows the number of successful backups, the number of backups that completed with warnings, and the number of backups that failed.
Total Backups Copied	The total number of backups copies that ran within the specified date range. A pie chart shows the number of successful backups copies and the number of failed backup copies.
Backup Copies (*Target*)	Shows number of successful and failed backup copies by target.
Assets Table Column	Lists all assets configured on the selected appliance(s), regardless of whether backups or backup copies have run within the selected date range. See the Last Backup and Protected By Backup Copy columns to determine whether the asset has been protected within the specified date range.
Appliance Name	The name of the backup appliance.
Client Name	The name of the asset.
Instance Name	The name of the protected instance. Examples: • Agent-based asset - Contains *image-level* or is empty for file-level. • Virtual machine - Contains the VM name. • Database - Contains the database name. • Storage group - Contains the storage group name.

Item	Description
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, Linux, Oracle, Hyper-V, or VMware.
Protected	Indicates whether there is a successful backup of this asset that ran within the specified date range: True (yes) or False (no).
Protected by Backup Copy	Indicates whether a successful backup copy ran for the asset within the specified date range: True (yes) or False (no).
Last Backup	Start date and time of this asset's last successful backup that ran within the specified date range.
Backup Table Column	Lists all assets that have been protected by backups within the selected date range and provides summary backup information.
Appliance Name	The name of the backup appliance.
Client Name	The name of the asset.
Instance Name	The name of the protected instance. Examples: • Agent-based asset - Contains *image-level* or is empty for file-level. • Virtual machine - Contains the VM name. • Database - Contains the database name. • Storage group - Contains the storage group name.
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, Linux, Oracle, Hyper-V, or VMware.
Backup Types	Backup modes of this asset's backups that ran within the specified date range and are stored on the appliance. Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, and/or Transaction (SQL only).
Successes	The number of backup jobs that completed successfully within the specified date range.
Warnings	The number of backup jobs that completed with warnings within the specified date range.

Item	Description
Failures	The number of backup jobs that failed within the specified date range.
Last Backup	Start date and time of this asset's last successful backup that ran within the specified date range.
Next Backup	Date and time that this asset's next scheduled backup will run.
Next Backup Job	Name of the next backup job that will run for this asset.
Backup Copies to *Target* Table Column	Lists all assets that have been protected by backup copies on the given target, within the selected date range. The table provides summary information about these backup copies.
Appliance Name	The name of the backup appliance.
Client Name	The name of the asset.
Instance Name	The name of the protected instance. Examples: • Agent-based asset - Contains *image-level* or is empty for file-level. • Virtual machine - Contains the VM name. • Database - Contains the database name. • Storage group - Contains the storage group name.
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, Linux, Oracle, Hyper-V, or VMware.
Backup Copy Types	Backup modes of this asset's backups that were copied to the target within the specified date range. Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, and/or Transaction (SQL only).
Backup Copy Successes	The number of backup copy jobs that completed successfully within the specified date range.
Backup Copy Warnings	The number of backup copy jobs that completed with warnings within the specified date range.
Backup Copy	The number of backup copy jobs that failed within the specified date range.

Item	Description
Failures	
Last Backup Copy	Start date and time of this asset's last successful backup copy that ran within the specified date range.

Backup Copy Capacity report

The Backup Copy Capacity report summarizes information about the amount of data protected by hot backup copies on the selected Unitrends appliance or Unitrends Cloud backup copy target(s).

Run this report from the source backup appliance by selecting a backup copy target(s) and entering a date range. When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

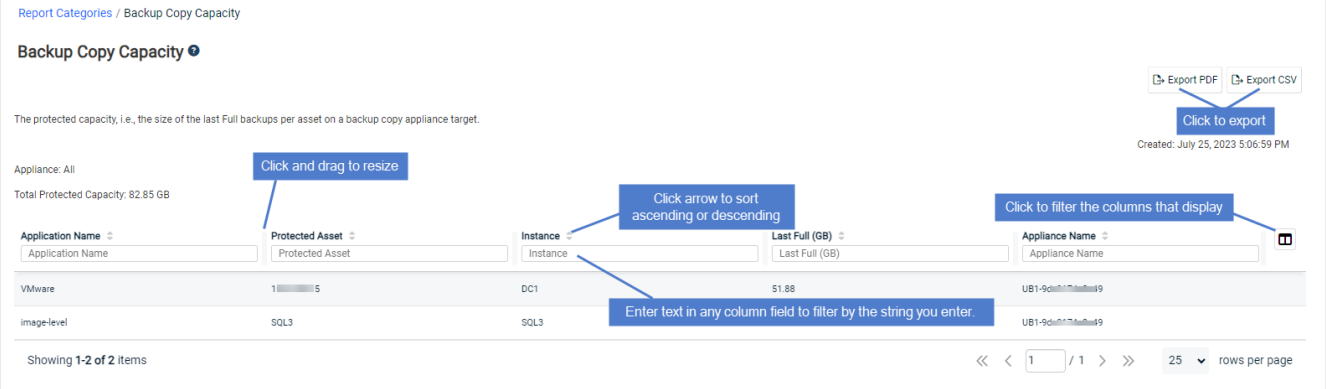

Report Field	Field Description
Application Name	Application type: • For file-level backups, contains *file-level*. • For Windows image-level backups, contains *image-level*. • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013. • For appliance configuration, contains *System Metadata*. System Metadata is automatically copied to the target during backup copy jobs and contains system

Report Field	Field Description
	information, such as appliance configuration and settings.
Protected Asset	The name of the asset.
Instance	For file-level and image-level backups, contains the name of the protected asset. For application backups, contains the name of the database, instance, or storage group. For host-level backups, contains one of the following: • The name of the virtual machine. • The name of the ESXi host if that host is being managed by a vCenter.
Last Full (GB)	The size of the last full backup copy, in gigabytes.
Appliance Name	The name of the source backup appliance that is sending copies to the hot target.

Backup Copy - Hot Targets report

The Backup Copy - Hot Targets report provides a summary of information about backups copied to the selected Unitrends Cloud or Unitrends appliance targets.

Run this report from the source or target backup appliance by selecting a backup copy target and entering a date range.

After generating the report, you can view the details of a specific job by clicking it in the grid. This opens the Backup Copy Status: Report Entry dialog. For file-level backups, you can click **Details** in this dialog to view a list of files contained in the backup copy.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

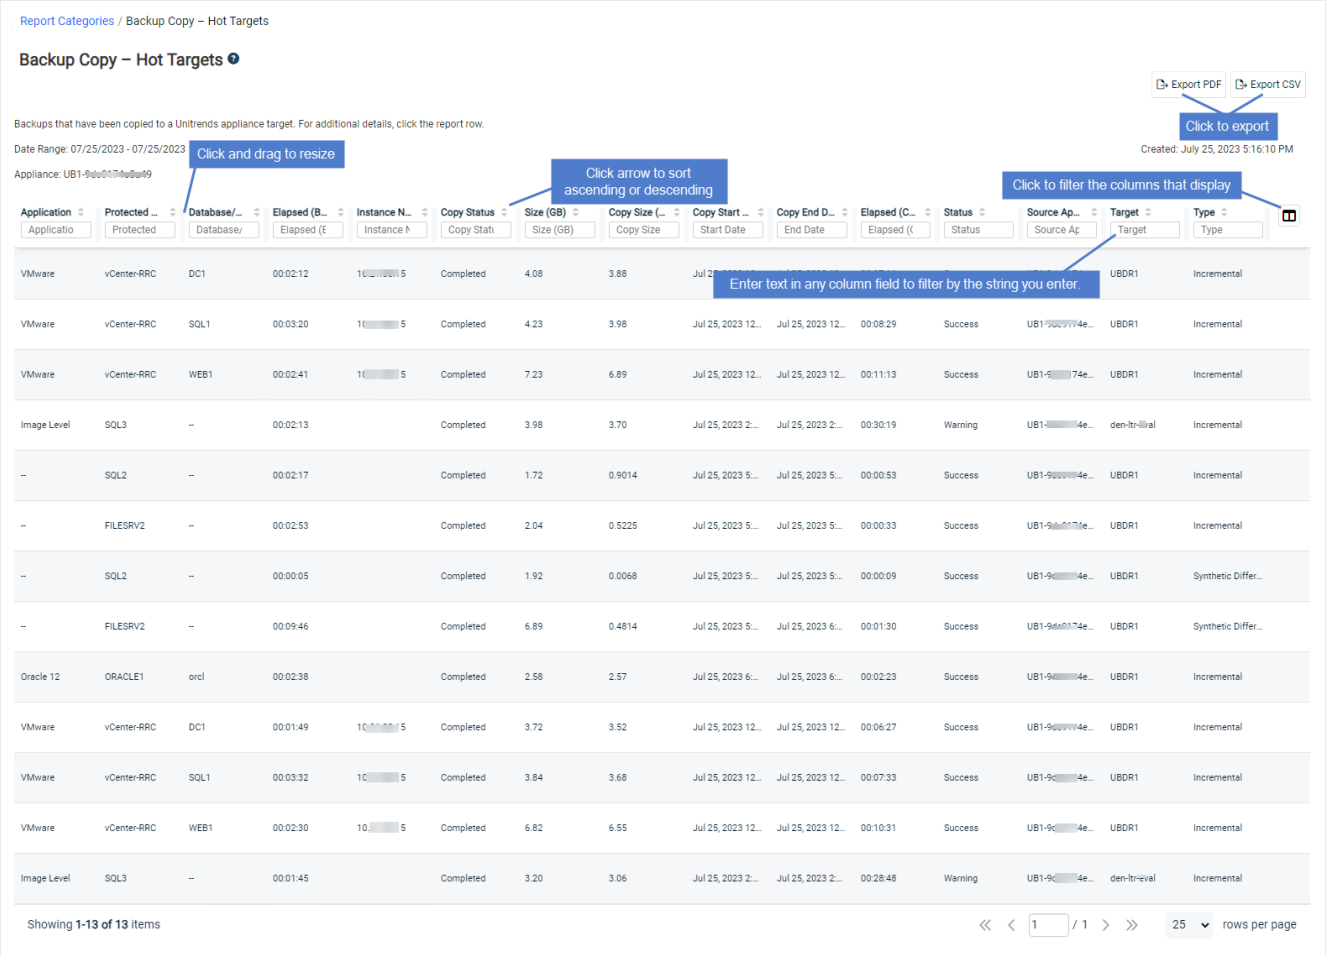

Report Field	Field Description
Application	Application type: • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013. • For file-level backups, contains "--". • For Windows image-level backups, contains *Image Level*.
Protected Asset	The name of the protected asset.
Complete	Indicates whether the job completed: *True* (yes) or *False* (no).

Report Field	Field Description
Database/VM	For SQL and Exchange application backups, contains the name of the protected database or storage group. For host-level backups, contains the name of the protected virtual machine.
Elapsed (Backup)	The amount of time it took for the original backup job to complete, in *hh:mm:ss* format.
Encrypted	Indicates whether the backup copy is encrypted. *True* (yes) or *False* (no).
End Date	The date and time at which the backup copy job completed.
Instance Name	For SQL and SharePoint application backups, contains the database instance name. For VMware host-level backups, contains one of the following: • The name of the virtual machine. • The name of the ESXi host if that host is being managed by a vCenter.
Copy Status	Status that indicates whether the backup has been copied to the Unitrends Cloud or copied to another Unitrends appliance: • *Completed* - The backup copy job is complete. • *Needed* - Cannot determine the status of this backup's copy job. • *In progress* - The backup copy job is currently running. • *Waiting* - The backup copy job is queued but has not yet started. • *Failure* - The backup copy job ran and failed. • *N/A* - This asset's backups are not configured for hot backup copy.
Size (GB)	The size of the original backup, in gigabytes. Displays when run from the source appliance only. Not applicable when run from the target appliance.
Copy Size (GB)	The size of the backup copy on the target, in gigabytes.
Start Date	The date and time that the original backup job started.
Copy Start Date	The date and time that the backup copy job started copying data to the target.
Copy End Date	The date and time that the backup copy job finished copying data to the target.
Elapsed (Copy)	The amount of time it took for the copy job to complete, in *hh:mm:ss* format.

Report Field	Field Description
Speed (MB/S)	Average transfer rate of the job, in megabytes per second.
Status	Status of the backup copy job: *Success*, *Warning*, or *Failure*.
Source Appliance	The name of the source backup appliance.
Type	Backup mode: Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, or Transaction (SQL only).
Certify Status	Indicates whether this Recovery Point has been certified by Copy Data Management or by the Unitrends ReliableDR product.
ID	The system-generated job ID.
Verify Status	Indicates whether the appliance verified the backup copy and the verification status: *Success* (verified), *Failure* (verify failed), or *N/A* (this job is not configured to verify the backup copy).

Backup Copies - Past 24 Hours report

The Backup Copies - Past 24 Hours report summarizes information about backups copied to the Unitrends Cloud or to Unitrends appliance targets within the last 24 hours. The report shows successful, active, and queued backup copy jobs.

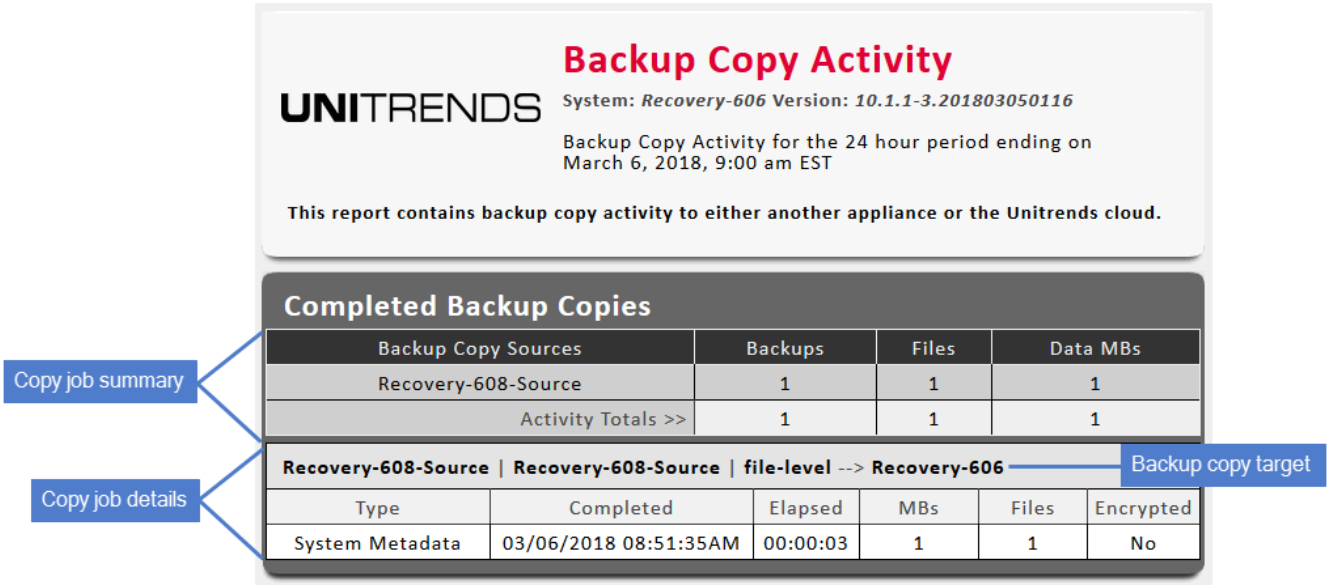

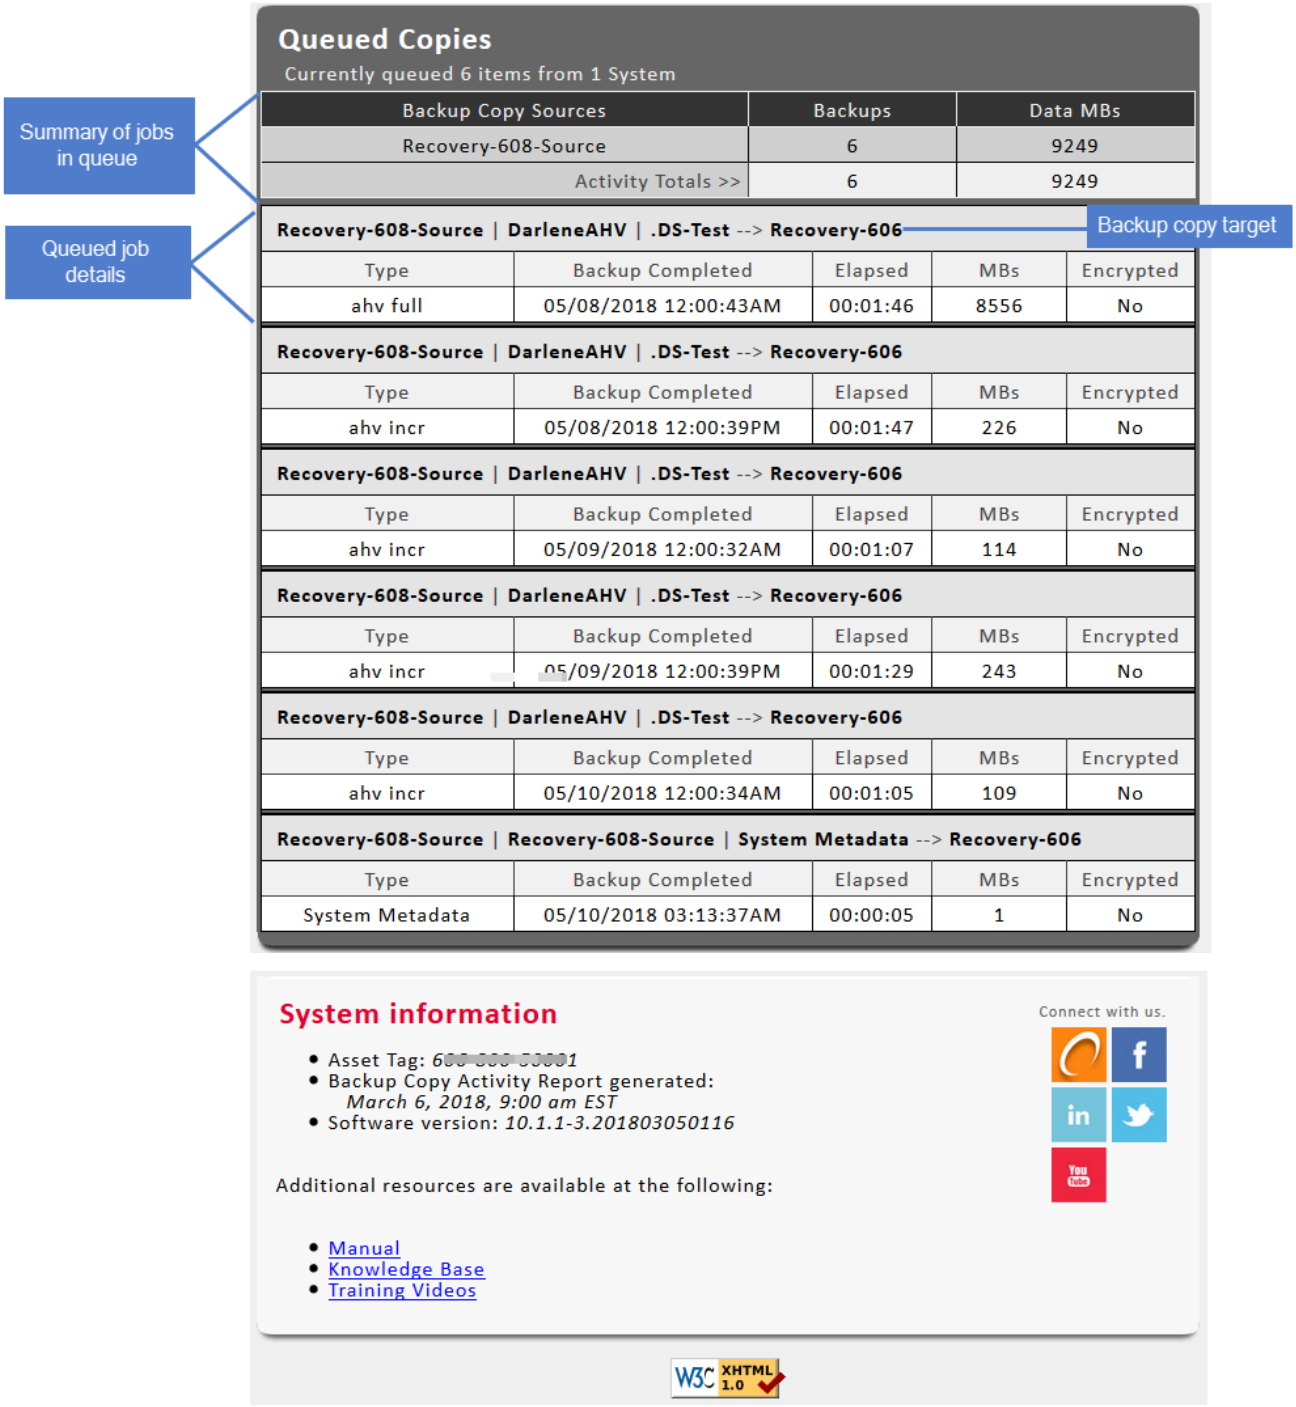

Storage Footprint report

The Storage Footprint report summarizes information about the amount of storage space sources are currently using on a target.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

Backup Copy - Cold Targets report

The Backup Copy - Cold Targets report summarizes information about each asset's backups that have been copied to cold storage within the selected date range for the selected appliance(s). Cold targets include: third-party cloud storage, attached disk storage, NAS and SAN storage, tape devices, eSATA devices, and USB devices.

After generating the report, you can view the details of a specific job by clicking it in the grid. This opens the Backup Copy Cold Targets Status Report Entry dialog.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

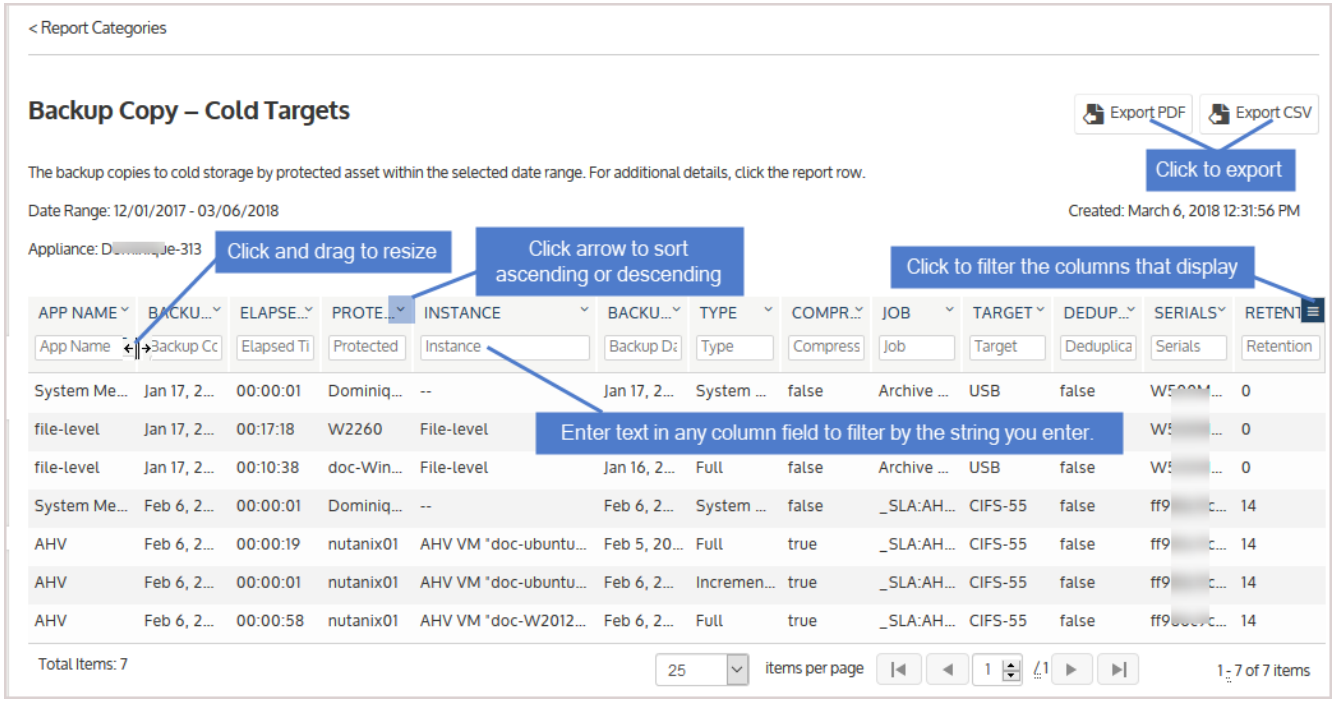

Report Field	Field Description
Appliance Name	The name of the backup appliance.
App Name	Application type: • For file-level backups, contains *File Level*. • For Windows image-level backups, contains *Image Level*. • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2.

Report Field	Field Description
	• For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013. • For system metadata backups, contains *System Metadata*. System metadata backups contain system information, such as job schedules and appliance configuration settings, that are used in the event of a disaster to recover the appliance. They are automatically created and copied to the target.
Backup Copy Time	The date and time when the backup copy job started.
Elapsed Time	The duration of the backup copy job, in hh:mm:ss format.
Protected Asset	The name of the asset protected by the backup.
Instance	Protected instance: • For file-level backups, contains *File-level*. • For Windows image-level backups, contains *Image-level*. • For application backups, contains the database instance name and database name. • For host-level backups, contains one of the following: – The virtual machine name and guid/uuid. – The ESXi host name if that host is being managed by a vCenter. • For system metadata backups, this column is empty.
Backup Date	The date and time the original backup job started.
Type	The type of backup that was copied: Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, Transaction (SQL only), or System Metadata.
Compressed?	Indicates whether the data is compressed: *True* (yes) or *False* (no).
Job	The name of the job.
Target	The name of the target on which the backup copy is stored.

Report Field	Field Description
Deduplicated?	Indicates whether the backup copy data has been deduplicated: *True* (yes) or *False* (no).
Serials	For tape, eSATA and USB devices, shows the serial numbers of the tapes or drives where the backup copy was written.
Label	For tape, eSATA and USB devices, shows the label assigned to the tapes or drives where the backup copy was written.
Barcodes	Barcodes of the tapes where the backup copy was written. (Applicable to tapes with barcode labels only.)
Retention Days	Minimum length of time the backup copy must be retained before it can be overwritten, in days.

Weekly Status report

The Weekly Status report provides a status of the backups and backup copies that have run on the appliance over the last 7 days. Select an appliance and click **Generate Report** to view the report.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the report columns. To customize your display, click the icon to the right of the column header.

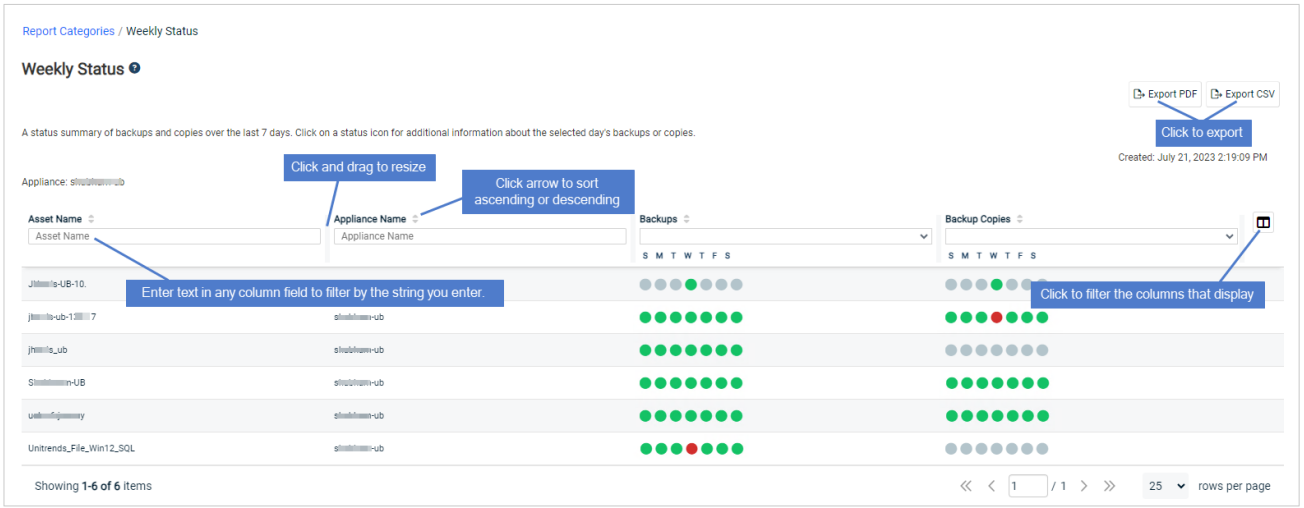

Item	Description
Asset Name	The name of the asset.
Appliance Name	The name of the appliance where the jobs ran.
Backups	Provides a high-level view of an asset's backups that ran over the last 7 days. Icons display for each day of the week: • A green checkmark indicates one or more successful backups ran that day. • A yellow icon indicates one or more backups completed with warnings. • A red icon indicates that one or more backups failed. • A gray circle indicates that no backups ran that day. Click a colored icon for a summary of the day's successes, warnings, failures, and backups in progress.
Backup Copies	Provides a high-level view of an asset's backup copies that ran over the last 7 days. Icons display for each day of the week: • A green checkmark indicates one or more successful backup copies ran that day. • A yellow icon indicates one or more backup copies completed with warnings. • A red icon indicates that one or more backup copies failed. • A gray circle indicates that no backup copies ran that day. Click a colored icon to see the backup copy target and a summary of the day's successes and failures.

Protection Policies report

The Protection Policies report shows protection policies by asset.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. For more detail, click a row to display the asset's Protection Policy Report Entry dialog.

Not all columns display by default. To customize your output, click the  icon to the right of the column header.

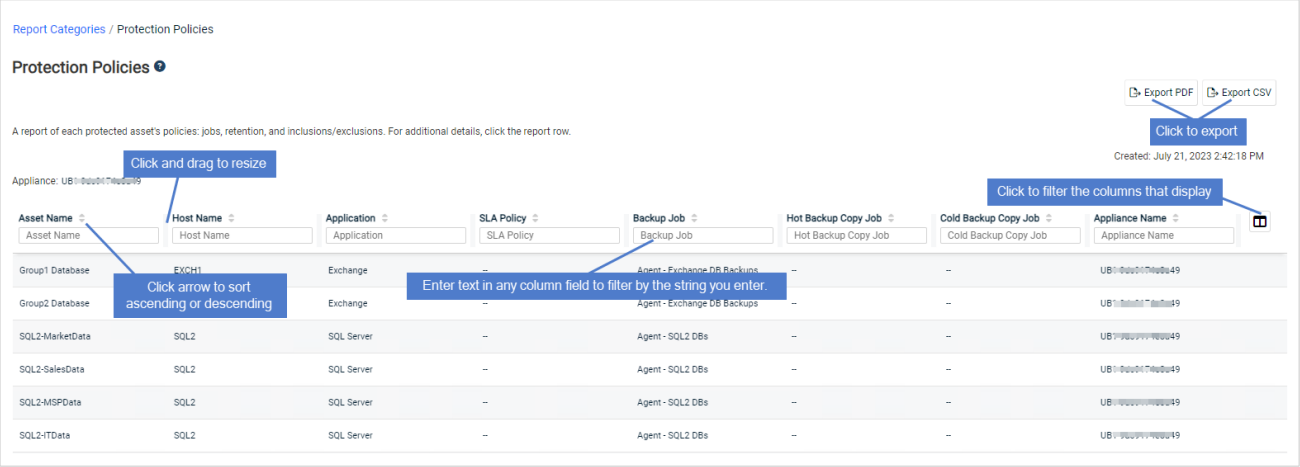

Report Field	Field Description
Asset Name	The name of the asset.
Host Name	Name of the machine hosting the protected asset.
Application	Application type: • For file-level backups, contains *file-level*. • For Windows image-level backups, contains *image-level*. • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V AHV. • For application backups, contains the application type. Examples: SQL Server, Oracle, Exchange Server.
SLA Policy	Name of the SLA policy that is protecting the asset. • This field is empty if the asset is not protected by an SLA policy. • Contains *Visible on local appliance only* if you are logged in to a manager appliance. Click the report row to view SLA policy details in the Protection Policy Report Entry dialog.
Backup Job	Name of the backup schedule that is protecting the asset. The name begins with the prefix *_SLA* if the schedule was created by an SLA policy.
Hot Backup	Name of the schedule that is copying this asset's backups to a hot backup copy target.

Report Field	Field Description
Copy Job	• This field is empty if the asset is not protected by a hot backup copy schedule. • The name begins with the prefix *_SLA* if the schedule was created by an SLA policy.
Cold Backup Copy Job	Name of the schedule that is copying this asset's backups to a cold backup copy target. • This field is empty if the asset is not protected by a cold backup copy schedule. • The name begins with the prefix *_SLA* if the schedule was created by an SLA policy.
Types	Lists the backup modes in the asset's backup schedule: Full, Incremental, Differential, Synthetic Full, Bare Metal, and/or Transaction (SQL only).
Appliance Name	The name of the backup appliance.

Appliance reports

The Appliance reports summarize and organize information about the performance of your Unitrends appliance.

For details, see the following topics:

- "Update History report" on page 1347
- "Capacity report" on page 1348
- "Load report" on page 1351
- "Alerts report" on page 1352
- "Trap History report" on page 1353
- "Notifications report" on page 1355

Update History report

The Update History report lists all software updates installed on the selected appliance within a specified date range.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

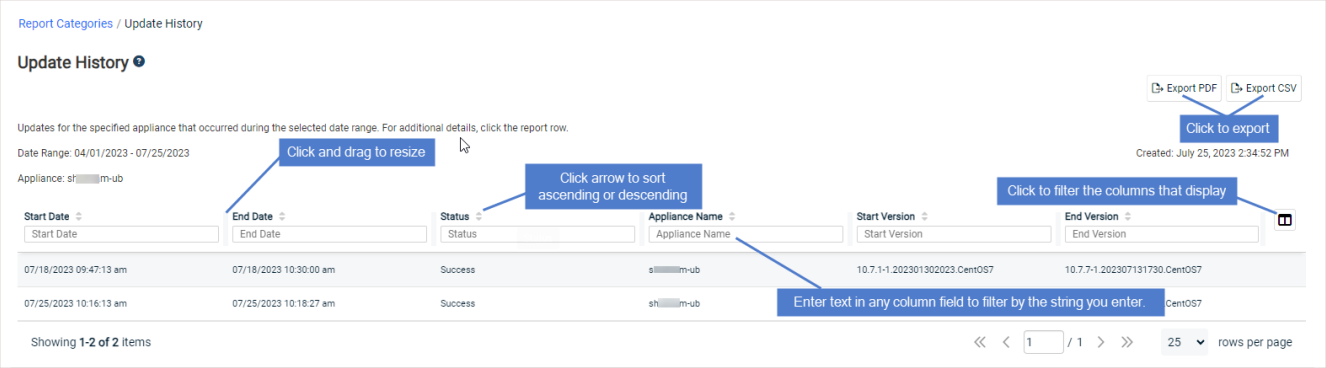

Report Field	Field Description
Start Date	The date and time when the update started.
End Date	The date and time when the update completed.
Status	The status of the update: Success or Failure.
Appliance Name	The name of the appliance to which the update was applied.
Start Version	The software version of the appliance before the update.
End Version	The software version of the appliance after the update ends.

Capacity report

The Capacity report provides information on storage capacity and how storage is currently being used. The top of the report provides summary information about the appliance's storage. Detail on each protected asset is given below.

To view the report, select an **Appliance** from the list and click **Generate Report**.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

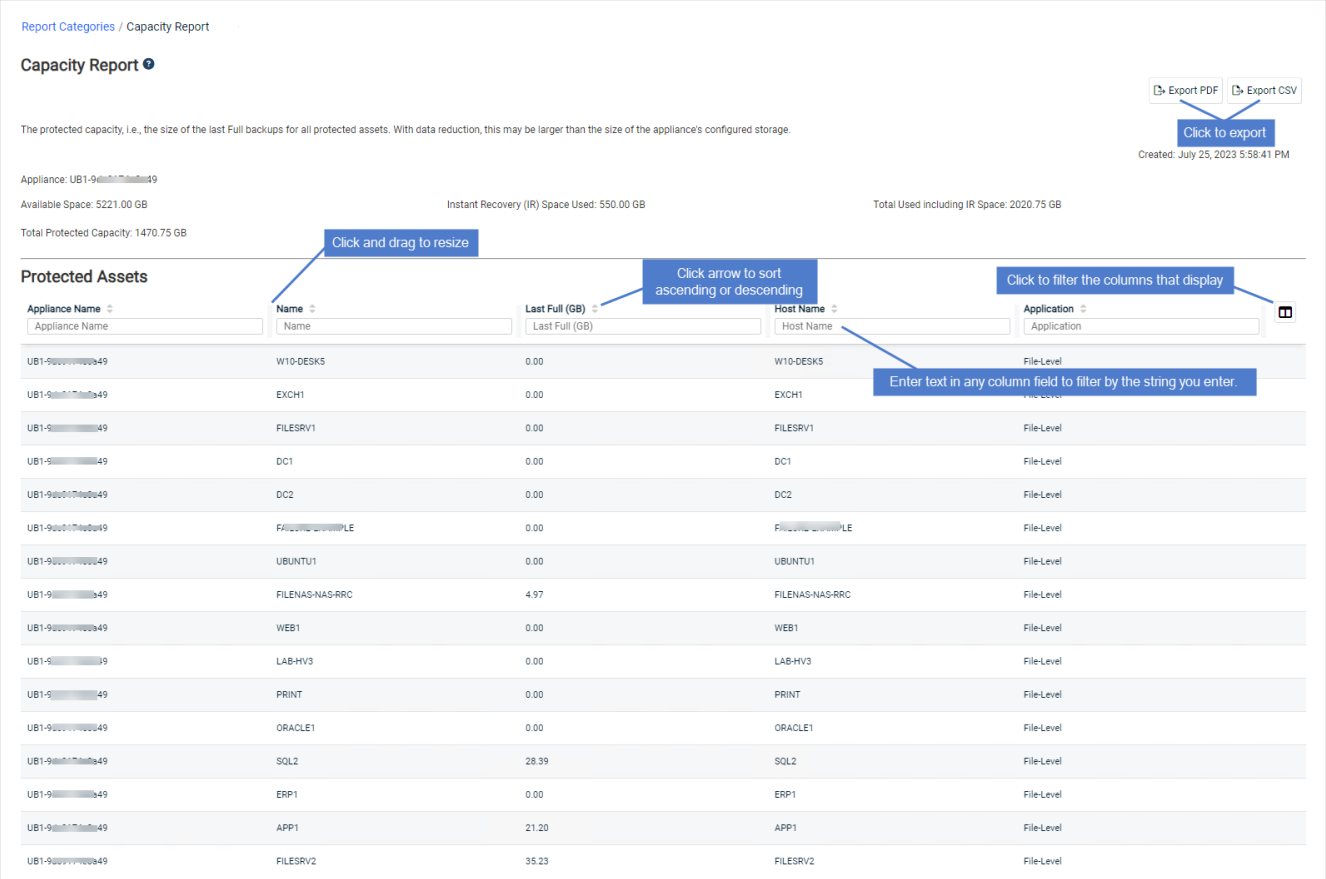

Report field	Field description
Appliance storage summary	The top of the report provides summary information about the appliance's storage.
Appliance	Name of the appliance.
Available Space	Total storage capacity of the appliance, in gigabytes. This is the total amount of space that can be used for all of the following: • Storing backups and imported backup copies • Storing hot backup copies (received from another appliance if this appliance is being used as a backup copy target) • Reserved for instant recovery

Report field	Field description
	Note: Available space does NOT indicate free space. Free space (space that can be used to store additional backups, copies, or instant recovery objects) is Available Space minus Total Used.
Instant Recovery (IR) Space Used	Total space reserved for instant recovery, in gigabytes.
Total Used including IR Space	Total space currently being used for all of the following: • Storing backups and imported backup copies • Storing hot backup copies (received from another appliance if this appliance is being used as a backup copy target) • Reserved for instant recovery
Total Protected Capacity	Total space currently being used to store the last full backups of all protected assets.
Protected Assets columns	Shows the amount of storage capacity used by all protected assets configured on the selected appliance(s).
Appliance Name	The name of the appliance.
Name	The name of the protected asset.
Last Full (GB)	The size of the asset's last full successful backup, in gigabytes.
# Fulls	The number of full backups stored on the appliance for this asset.
Host Name	Name of the machine hosting the protected asset.
Hypervisor	VMware, Hyper-V, XenServer, AHV, AWS or Azure (if applicable).
Application	Application type: • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V.

Report field	Field description
	• For application backups, contains the application type. Examples: SQL Server, Oracle, Exchange. • For file-level backups, contains *File-Level*. • For Windows image-level backups, contains *Image-Level*. • For System Metadata contains system information, such as appliance configuration and settings. System Metadata is copied to the target during backup copy jobs if changes to this system information are detected.

Load report

The Load report provides a graphical view of the load on the appliance over the selected date range.

Select the appliance and specify the date range to generate the report.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the report.

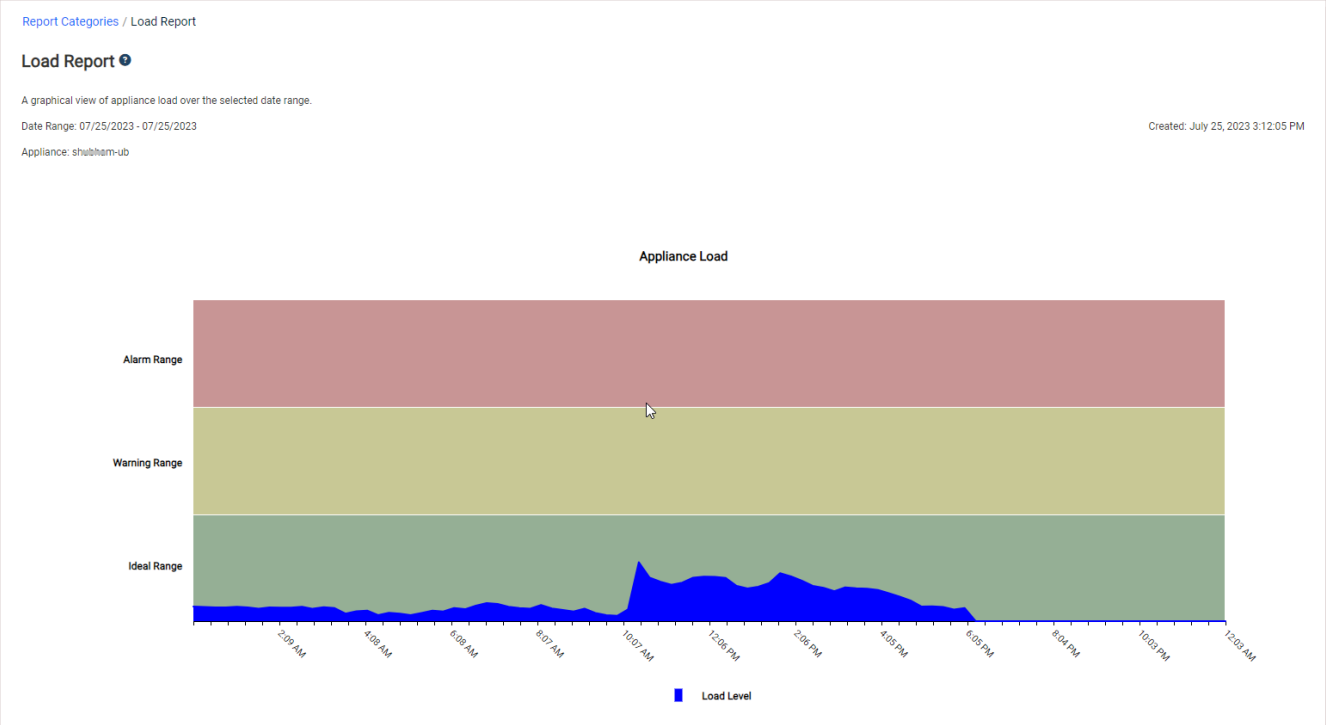

Report Field	Field Description
Date Range	The date range specified for the report.
Appliance	The selected appliance.
Appliance Load	The load level, depicted in blue, over the specified period of time. There are three load level ranges: *Ideal* (green), *Warning* (yellow), and *Alarm* (red).

Alerts report

The Alerts report lists all alerts generated for the selected appliance(s) within the specified date range.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

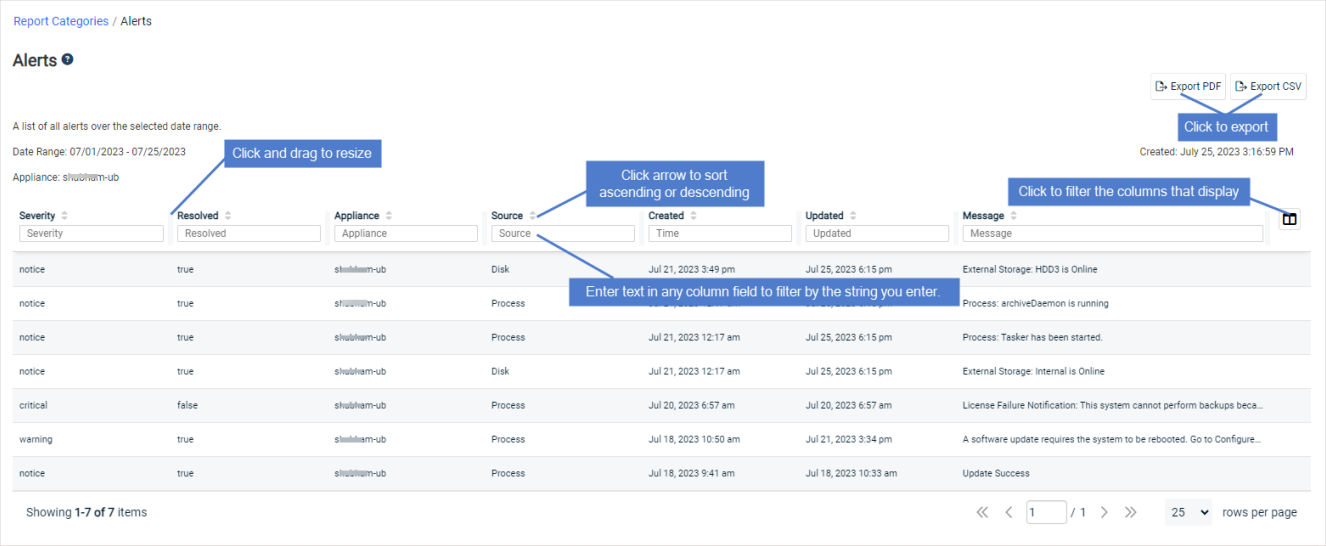

Report field	Field description
Severity	The level of severity assigned to the alert.
Resolved	Indicates whether the alert has been resolved: True (yes) or False (no).
Appliance	The name of the appliance that produced the alert.

Report field	Field description
Source	The component that generated the alert.
Created	The date and time the alert was generated.
Updated	The date and time the alert was last updated.
Message	Any system-generated message associated with the alert.

Trap History report

The Trap History report provides a list of all SNMP traps sent from the appliance over the selected date range. These traps are sent to Unitrends to enable proactive monitoring of the health of the appliance (if the necessary ports are open). You can also configure the appliance to send traps to your own network management server (for details, see "SNMP trap notifications" on page 178).

To view the report, select an appliance, enter a date range, and click **Generate Report**.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

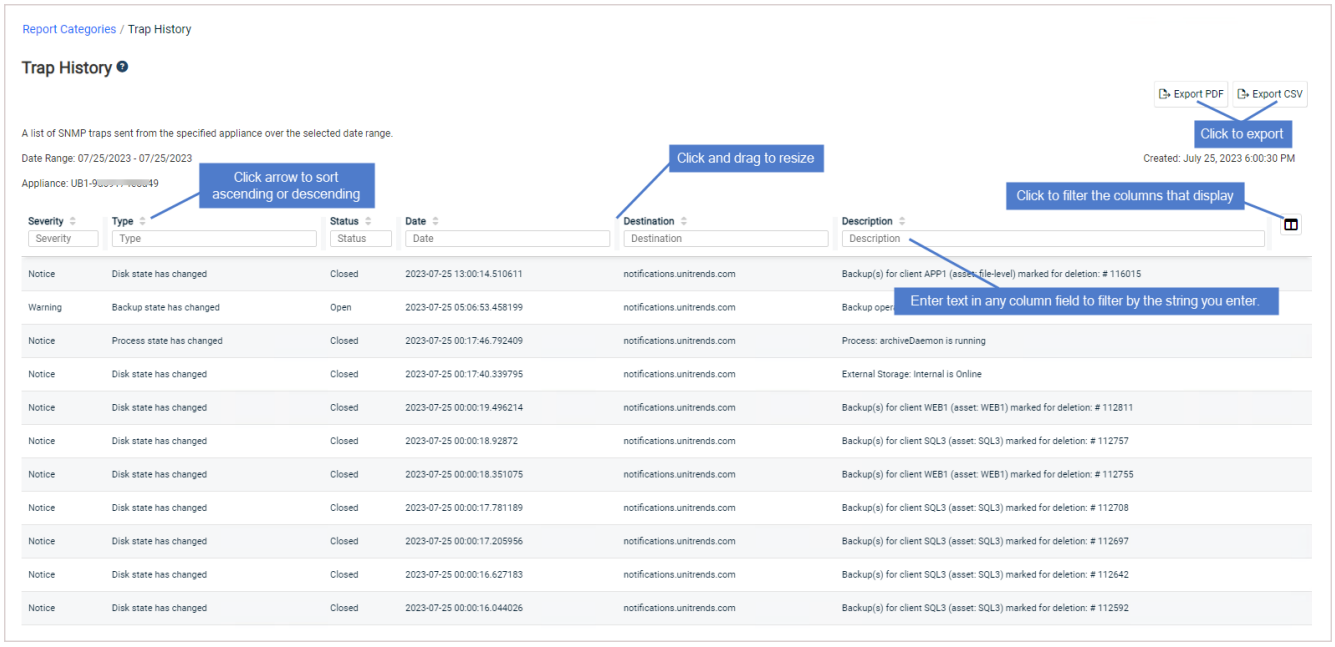

Item	Description
Appliance ID	ID of the appliance that ge the SNMP trap.
Appliance Name	Name of the appliance that sent the SNMP trap.
Severity	Severity of the issue: • Fatal - fix immediately, condition is causing failures or will cause failures • Warning - action needed to resolve, preventing optimal performance • Notice - notification, no action needed. Normal appliance operations will resolve the issue.
Type	Indicates the area impacted by the issue. Examples: • *Backup state has changed* indicates an issue with a backup operation (failed, did not run, etc.). • *Clients state has changed* indicates an issue with a protected asset (update is available, asset is not longer listening to connections, etc.). • *Disk state has changed* indicates an issue with disk storage (drive degraded, external storage is offline, etc.)
Status	Status of the issue - Open (not yet resolved) or Closed (resolved).
Date	Date and time the trap was generated.
Destination	Address where the trap was sent.
Description	Detailed description of the condition that triggered the trap.
OID	Object ID of the trap.
Object	Name of the appliance or protected asset impacted by the issue.
Community	Indicates the community to which the trap notification was sent. Public by default. You can set up custom trap notifications and specify a different community setting.

Notifications report

The Notifications report provides a list of all system notifications sent from the appliance over the selected date range. If your appliance is configured to email system notifications, these messages are also sent to the specified report recipients. (To check this, select **Configure > Appliances > Edit Appliance > Email** tab). Notifications that require action to resolve also display as alerts in the appliance UI.

To view the report, select an appliance, enter a date range, and click **Generate Report**.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

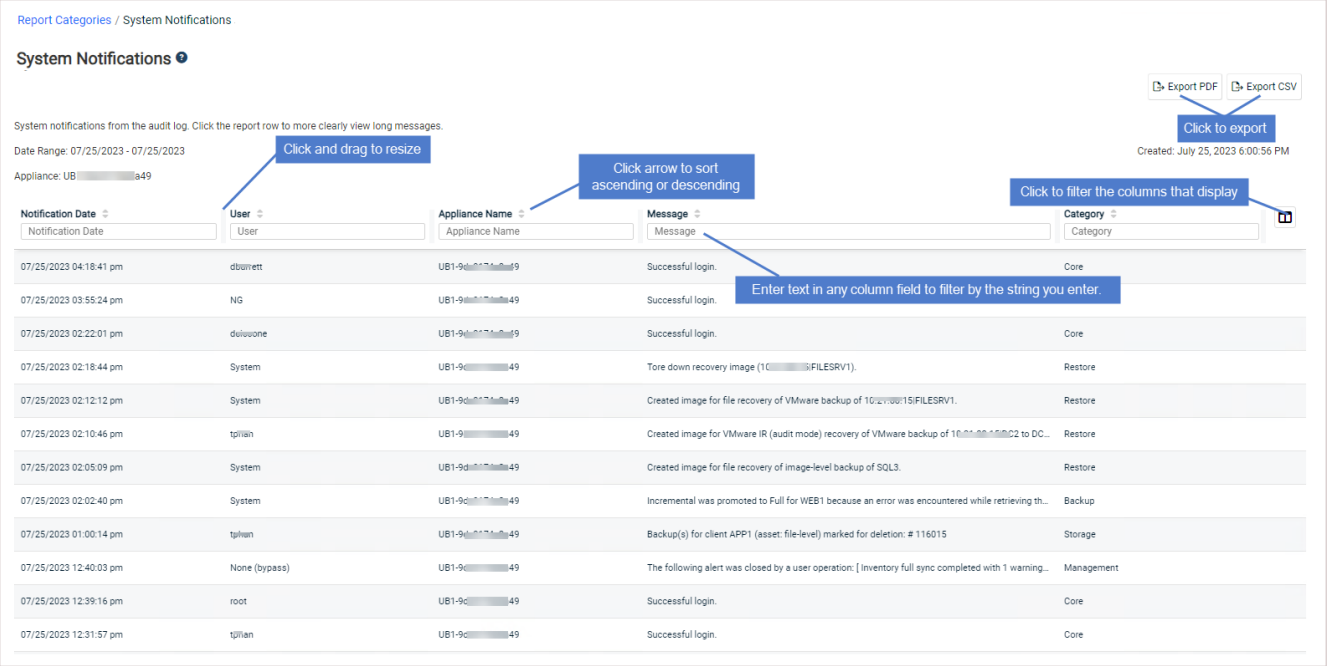

Item	Description
Notification ID	ID of the notification message template.
Notification Date	Date and time the notification was generated.
User	User that created the notification (*System*).
Appliance Name	Name of the appliance that generated the notification.

Item	Description
Message	Message text.
Category	Indicates the area impacted by the notification. For example, Backup, Storage, Management, etc.

Storage reports

The Storage reports provide information about your system's configured storage.

The Storage category features two available reports, Storage and Data Reduction.

For details, see the following topics:

- "Storage report " on page 1356
- "Data Reduction report" on page 1359

Storage report

The Storage report provides information about all configured backup storage, and displays a graphical and tabular view of this information. A pie chart displays a snapshot of the current information, and a bar graph displays information over the last thirty days.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the  icon to the right of the column header.

Report Field	Field description
Storage Totals Pie Chart	Graphical view of the appliance's backup storage.
Storage Totals	The appliance's total storage capacity. Includes all backup storage devices. Does not include backup copy target storage. **Note:** Reports and fields related to available storage on physical appliances report on the licensed storage instead of the total raw storage.
Free Storage	Amount of backup storage space available on the appliance.
Used Storage	Amount of backup storage space used on the appliance.
Backup Storage Usage Graph	Graph of storage used and available over the last 30 days.
Date	For a given day, provides a snapshot of the storage on the appliance (free versus used). Hover over a bar in the graph to see used and free capacity, in gigabytes.
Free Storage	Snapshot of the amount of backup storage space available on the appliance on a given day, in gray.
Used Storage	Snapshot of the amount of backup storage space used on the appliance on a given

Report Field	Field description
	day, in green.
Storage Table columns	Storage details by device.
Appliance ID	System-generated ID assigned to the appliance.
Appliance Name	Name of the backup appliance.
ID	System-generated ID assigned to the storage device.
Name	Storage device name. The initial backup storage device is *Internal*.
Type	Storage type: Internal, NAS, FC (Fibre Channel), AOE, and Direct-Attached Disk.
Protocol	For external storage, displays the connection protocol. For example, NFS or iSCSI.
Usage	How the storage is used by the appliance. For example, *stateless* for backup storage or *archive* for backup copy storage.
Online	Indicates whether the storage is currently online: True (yes) or False (no).
Status	Status of the storage: online or offline.
Allocated Storage (GB)	The amount of allocated storage for the storage device, in gigabytes.
MB Free	The amount of free space on the storage device, in megabytes.
Total Backup Data (GB)	The total amount of backup data on the storage device, in gigabytes.
Daily Growth	The average amount that used storage increased from one day to the next day. Averaged over the last 30 days.
Dedupe Ratio	For appliances that use deduplication, shows the amount of space saved on the storage device by not storing duplicate data blocks. Expressed as a ratio: amount of backup data before deduplication / storage space used after deduplication.

Report Field	Field description
Size History	Hover over this field to view the size history of the storage device over the last 30 days. Example: day:2018-04-12, mb_size:7511361, mb_used:2910551, mb_free:4600809 day:2018-04-13, mb_size:7511361, mb_used:2923749, mb_free:4587611 day:2018-04-14, mb_size:7511361, mb_used:2935473, mb_free:4575887 day:2018-04-15, mb_size:7511361, mb_used:3357588, mb_free:4153772 day:2018-04-16, mb_size:7511361, mb_used:3368479, mb_free:4142881 day:2018-04-17, mb_size:7511361, mb_used:3375857, mb_free:4135503 day:2018-04-18, mb_size:7511361, mb_used:3394316, mb_free:4117044 day:2018-04-19, mb_size:7511361, mb_used:3401852, mb_free:4109508 day:2018-04-20, mb_size:7511361, mb_used:3417450, mb_free:4093910 day:2018-04-21, mb_size:7511361, mb_used:3431555, mb_free:4079805 day:2018-04-22, mb_size:7511361, mb_used:3527589, mb_free:3983771 day:2018-04-23, mb_size:7511361, mb_used:3532847, mb_free:3978513 day:2018-04-24, mb_size:7511361, mb_used:3539796, mb_free:3971564 day:2018-04-25, mb_size:7511361, mb_used:3545368, mb_free:3965992 day:2018-04-26, mb_size:7511361, mb_used:3554935, mb_free:3956425 day:2018-04-27, mb_size:7511361, mb_used:3562747, mb_free:3948613 day:2018-04-28, mb_size:7511361, mb_used:3572589, mb_free:3938771 day:2018-04-29, mb_size:7511361, mb_used:3887922, mb_free:3623438 day:2018-04-30, mb_size:7511361, mb_used:3902066, mb_free:3609294 day:2018-05-01, mb_size:7511361, mb_used:3910491, mb_free:3600869 day:2018-05-02, mb_size:7511361, mb_used:3919974, mb_free:3591386 day:2018-05-03, mb_size:7511361, mb_used:3928374, mb_free:3582986 day:2018-05-04, mb_size:7511361, mb_used:3953364, mb_free:3557996 day:2018-05-05, mb_size:7511361, mb_used:3965509, mb_free:3545851 day:2018-05-06, mb_size:7511361, mb_used:4291814, mb_free:3219546 day:2018-05-07, mb_size:7511361, mb_used:4305007, mb_free:3206353 day:2018-05-08, mb_size:4294967295, mb_used:4294967295, mb_free:4294967295 day:2018-05-09, mb_size:4294967295, mb_used:4294967295, mb_free:4294967295 day:2018-05-10, mb_size:4294967295, mb_used:4294967295, mb_free:4294967295 day:2018-05-11, mb_size:4294967295, mb_used:4294967295, mb_free:4294967295

Data Reduction report

The Data Reduction report provides information (in graph form) about the amount of space saved by deduplication and compression for a selected appliance over a specified date range.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

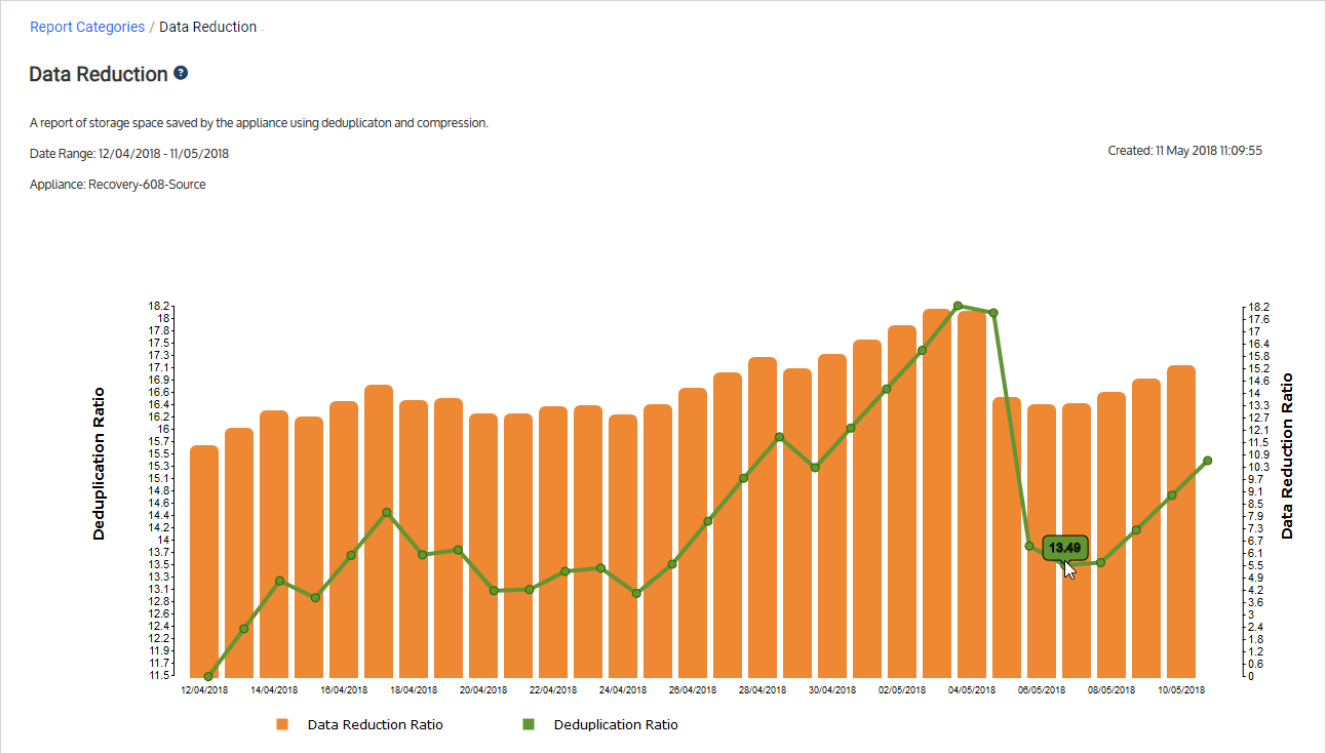

Replicas History report

The Replicas History report provides information about the backups that have been applied to replicas. To generate the report, select one or more appliances and specify a date range.

Replicas are set up by using the Create Windows Replica or Create Replica VMs dialog. The appliance then creates the replica from the most recent backup of the Windows or VM asset, and automatically applies all subsequent backups.

The appliance creates the replica by running a *create replica* operation and applies each subsequent backup by running one of the following:

- A *replica restore* to apply a host-level backup to a VM replica.
- A *virtual restore* to apply a file-level backup to a Windows replica.
- A *virtual restore* to apply a SQL or Exchange application backup to a Windows replica.

Each of these operations display as a row in the Replicas report. Click on a row to view details about the operation. See "Report details" on page 1361 for descriptions of each column in the report.

Report example:

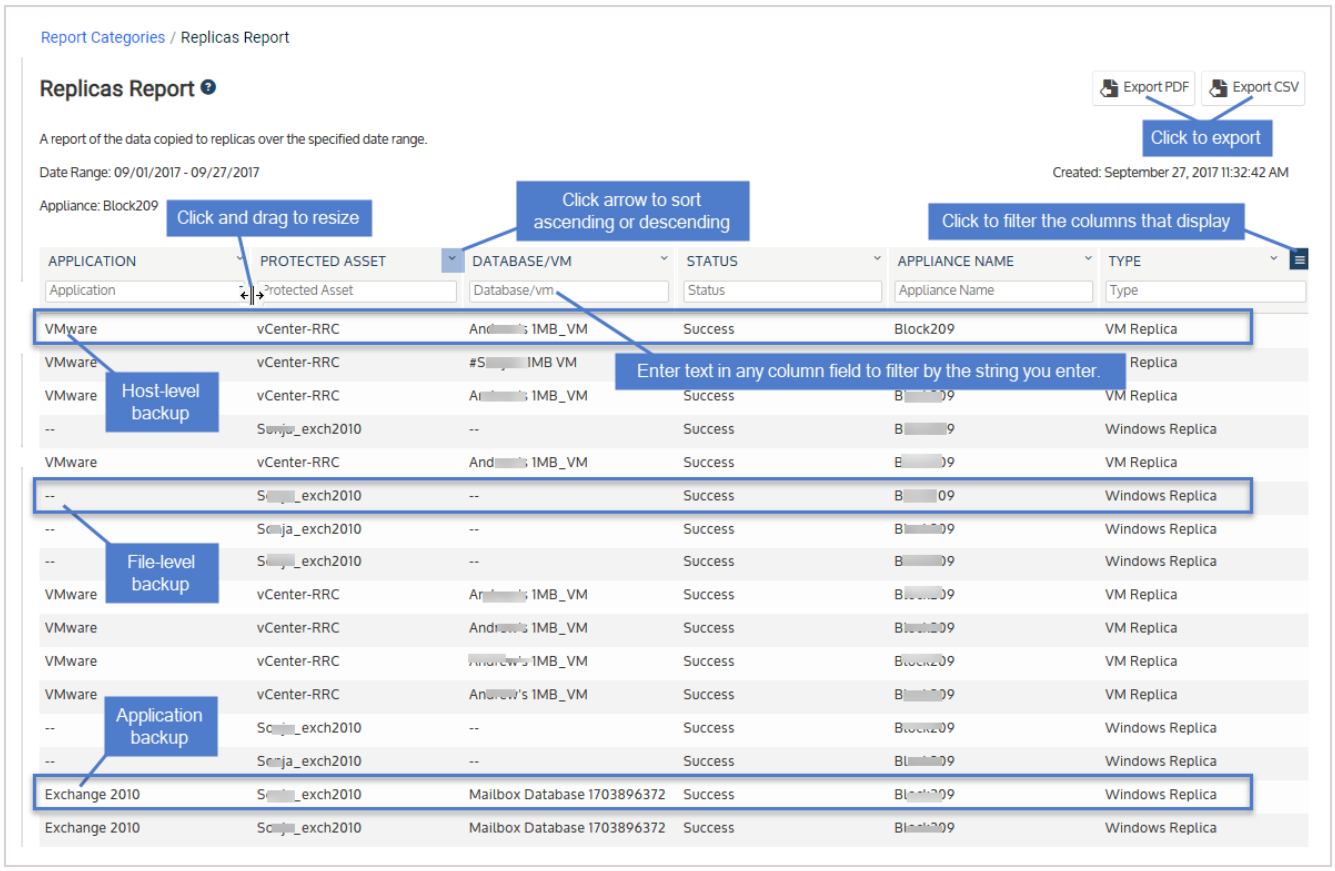

Report detail examples:

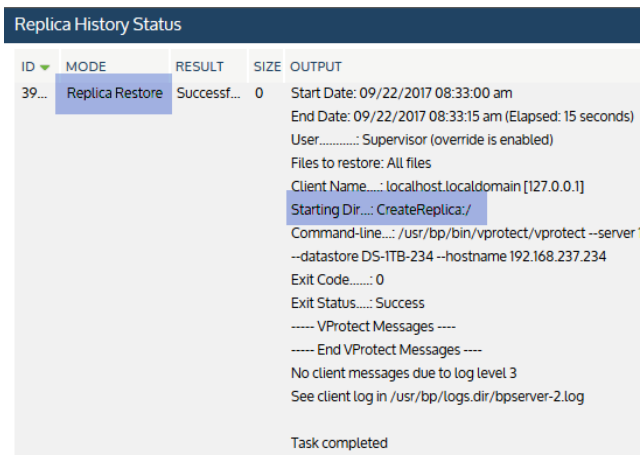

Report details

Refer to the following table for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

Report Field	Field Description
Application	Application type of the backup that was applied to the replica: • For host-level backups applied to VM replicas, contains *VMware*. • For application backups applied to Windows replicas, contains the application type. Examples: SQL Server 2012, Exchange 2013. • For file-level backups applied to Windows replicas, contains "--" (not applicable).
Protected Asset	Name of the original protected asset. • For host-level backups, contains one of the following: – The name of the virtual host where the original VM resides. – *vCenter-RRC* if the VM's virtual host is managed by a vCenter. • For file-level backups, contains the name of the original Windows asset. • For Exchange and SQL application backups, contains the name of the original Windows asset that hosts the application.
Complete	Indicates whether the job completed: True (yes) or False (no).
Database/VM	Database or VM name. • For host-level backups, contains the VM name. • For application backups, contains the database or storage group name. • For file-level backups, contains "--" (not applicable).
Elapsed Time	The amount of time it took to apply the backup to the replica, in hh:mm:ss format.
End Time	The date and time at which the job completed.
# Files	The number of files applied to the replica. (For host-level backups, contains *1* as these backups do not contain individual files.)
ID	The system-generated job ID.
Instance Name	Instance name. • For host-level backups, contains one of the following: – The name of the virtual host where the original VM resides.

Report Field	Field Description
	– *vCenter-RRC* if the VM's virtual host is managed by a vCenter. • For file-level backups, contains "--" (not applicable). • For Exchange and SQL application backups, contains "--" (not applicable).
Size (GB)	Not used.
Start Time	The date and time at which the job started.
Status	Status of the job: Success, Warning, or Failure.
Appliance Name	The name of the Unitrends appliance where the replica was created.
Type	Replica type: VM Replica or Windows Replica.

Retention Reports

The Retention category contains reports that summarize information about your backup retention. To view Retention reports, click on a report type under **Category**. Under **Available Reports**, click on the name of the report, and click **Generate Report**. (For detailed procedures, see "Working with reports" on page 1307.)

For details on each Retention report, see the following topics:

- "Legal Hold Backups report" on page 1363
- "Long-Term Retention report " on page 1365
- "Min-Max Retention report" on page 1368

Legal Hold Backups report

The Legal Hold Backups report summarizes all backups currently on legal hold for the selected appliance(s) within a specified date range. Legal holds are placed on assets in accordance with the "keep backups for *N* days" legacy asset-level retention setting. Legal holds are not used by the long-term retention scheme.

Note: Following the switch to long-term retention, backups continue to display in this report until the assets from which they originated are included in long-term retention policies.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the  icon to the right of the column header. (See "To customize a report" on page 1308 for details.)

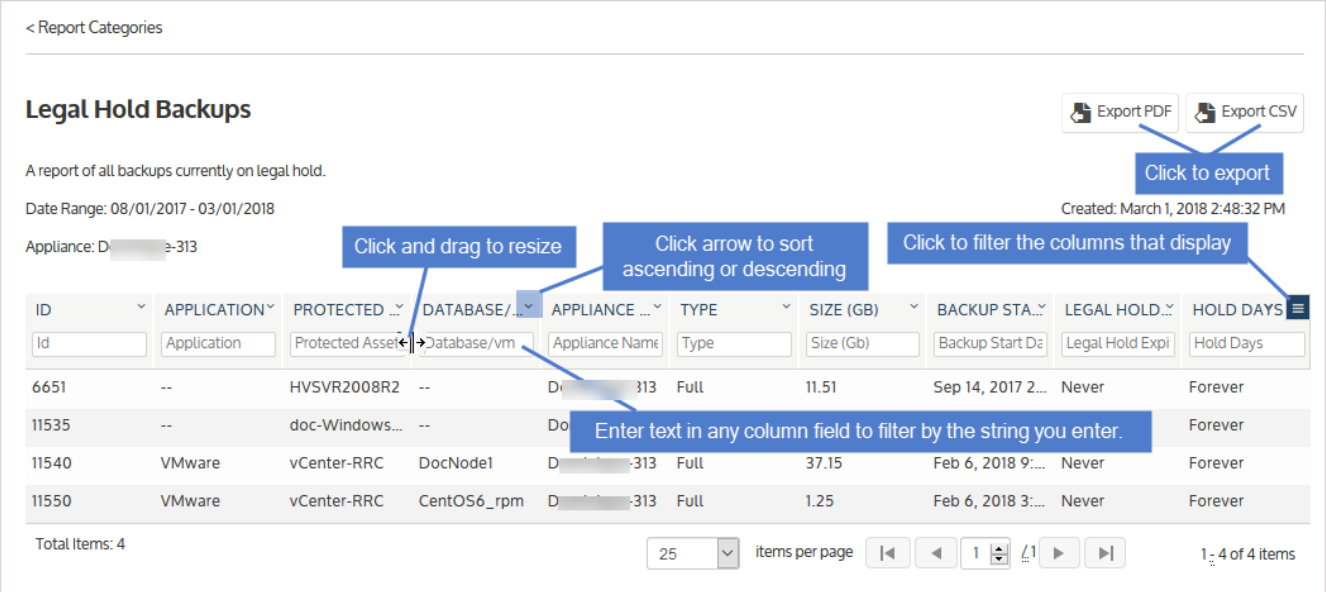

Report Field	Field Description
ID	The system-generated job ID.
Application	Application type: • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013.
Protected Asset	The name of the asset.
Database/VM	For SQL and Exchange application backups, contains the name of the protected database or storage group. For host-level backups, contains the name of the protected virtual machine.
Instance Name	For SQL and SharePoint application backups, contains the database instance name. For VMware host-level backups, contains one of the following: • The name of the virtual machine. • The name of the ESXi host if that host is being managed by a vCenter.
Appliance Name	The name of the backup appliance.

Report Field	Field Description
Type	Backup mode: Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, or Transaction (SQL only).
Size (GB)	The size of the backup, in gigabytes.
Backup Start Date	The date and time that the backup job started.
Legal Hold Expiration Date	The date on which legal hold settings expire and the backup follows standard retention policies.
Hold Days	The number of days the backup is under legal hold.

Long-Term Retention report

The Long-Term Retention report provides an overview of assets and their respective long-term retention policies.

Retention compliance for an individual asset can be viewed by clicking it in the table. The resulting dialog lists all retention points mandated by the policy and indicates if they are compliant. Compliant retention points have at least one successful backup available for each retention point.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

Notes:

- Source appliances may be selected when running this report; however, only the assets of managed sources display.
- A supplementary Retention Compliance Details report is also available. This report is not visible in the UI and must be manually retrieved from the appliance's file system. For instructions on generating and modifying the Retention Compliance Details report, see How to generate and modify the Retention Compliance Details report.

Refer to the following table for a description of the available report columns. Not all columns display by default. To customize your output, click the  icon to the right of the column header.

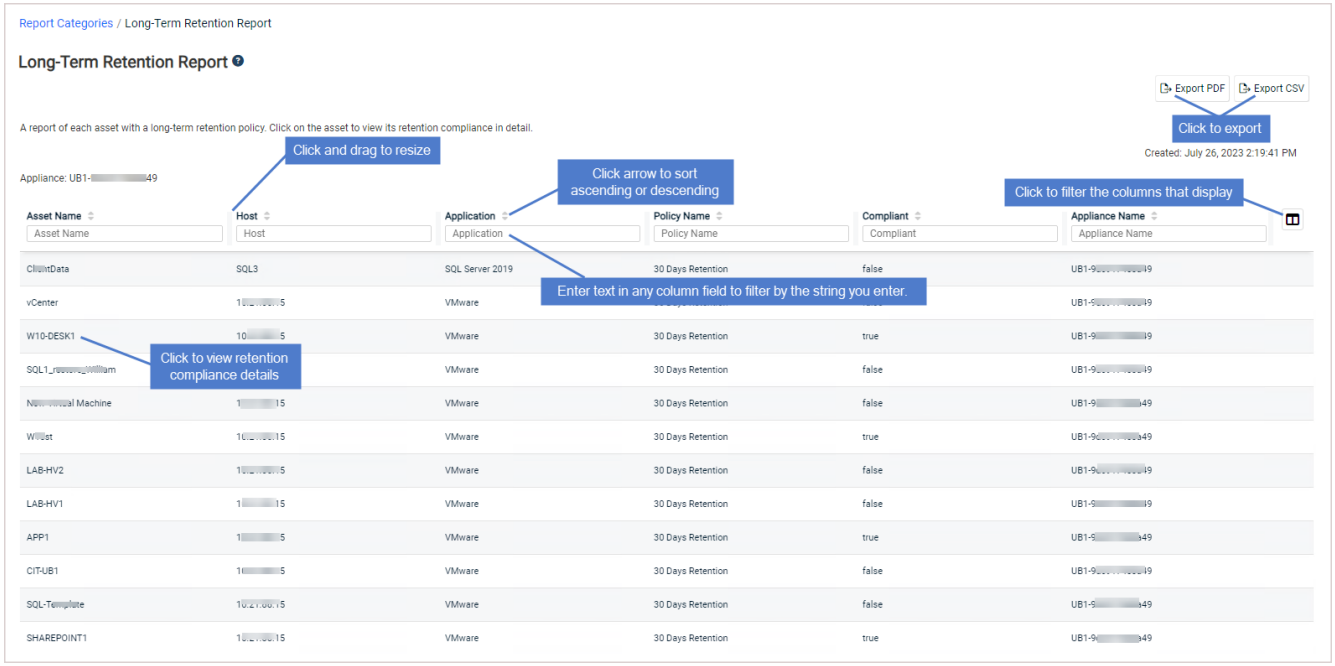

Retention compliance details:

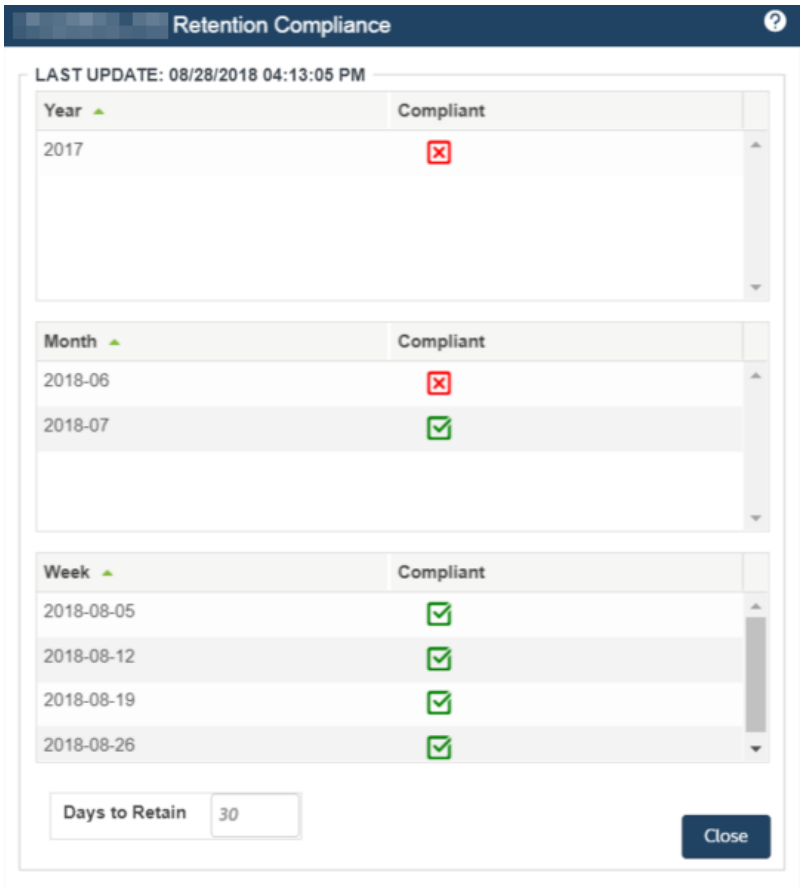

Report Field	Field Description
Asset Name	The name of the asset.
Host	Name of the machine hosting the asset.
ID	The instance ID.
Application	Application type: • For agent-based backups, contains *file-level* or *image-level*. • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013.

Report Field	Field Description
Policy ID	The ID of the asset's retention policy.
Policy Name	The name of the long-term retention policy applied to this asset.
Policy Description	The policy description that was entered when the policy was created.
Years	Years specified in the retention policy.
Months	Months specified in the retention policy.
Weeks	Weeks specified in the retention policy.
Days	Days specified in the retention policy.
System ID	The system ID of the appliance. The local appliance always has a value of *1*.
Appliance Name	The name of the appliance.
Compliant	Asset compliance with its associated retention policy: • *true* - The asset has at least one successful backup for each retention point. • *false* - The asset does not have at least one successful backup for each retention point.

Min-Max Retention report

The Min-Max Retention report provides an overview of assets and their respective retention settings.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

Refer to the following table for a description of the available report columns. Not all columns display by default. To customize your output, click the icon to the right of the column header.

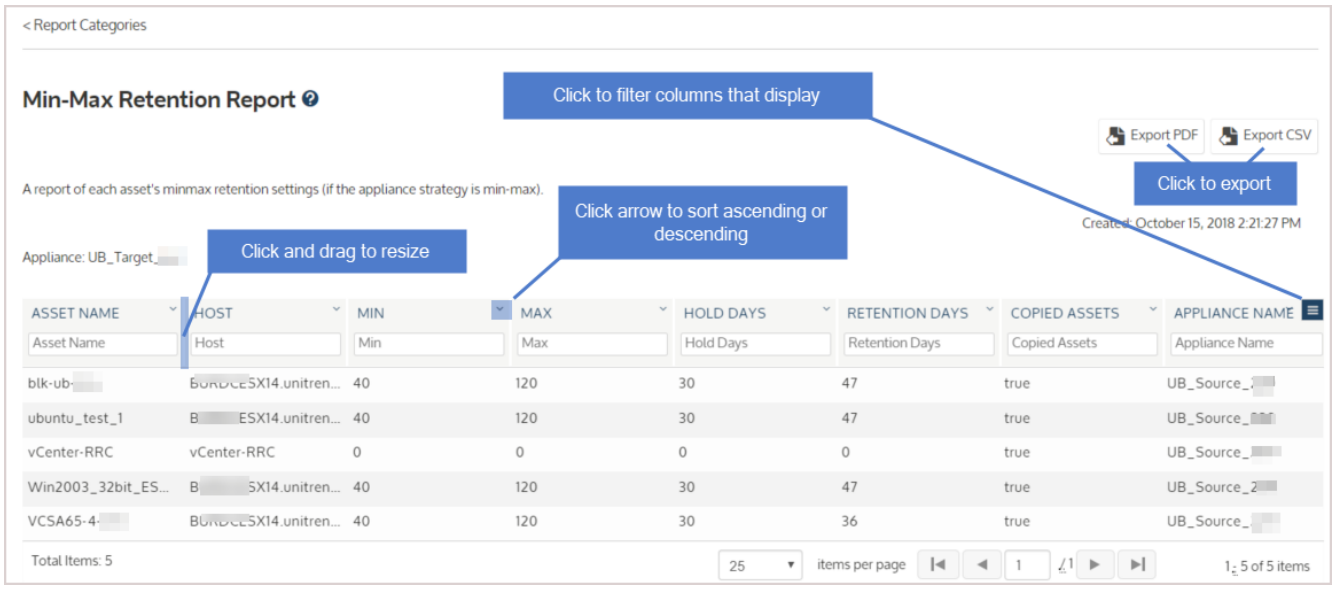

Report Field	Field Description
Asset Name	The name of the asset.
Host	Name of the machine hosting the asset.
ID	The instance ID.
Application	Application type: • For agent-based backups, contains *file-level* or *image-level*. • For host-level backups, contains the virtual host type. Examples: VMware, XenServer, Hyper-V 2012 R2. • For application backups, contains the application type. Examples: SQL Server 2012, Oracle 11, Exchange 2013.
Min	An email notification is sent if the asset has less than *N* days of backups stored on the appliance.
Max	Number of days after which the appliance will delete backups.
Hold Days	Number of days backups must be retained.
Retention Days	Backups are available for the preceding *N* number of days.

Report Field	Field Description
Copied Assets	*True* = The asset is a copied asset.
Appliance Name	The name of the appliance.

Compliance report

The compliance report provides information about RPO/RTO compliance.

For details, see the "Compliance report"

Compliance report

The Compliance report provides an overview of current RPO/RTO compliance based on the results of recent data copy access jobs.

After generating the report, you can view the details of a specific job by clicking it in the grid.

When finished, click **Report Categories** in the upper left to return to the Available Reports page.

Notes:

- The Compliance report automatically refreshes daily at 12:00 AM.
- Compliance reports are sent as emails when email reports are configured and the Appliance report type is selected. For more information on email reports, see "Email reporting" on page 117.

An example report is given here. Refer to the table below for a description of the available report columns. Not all columns display by default. To customize your output, click the  icon to the right of the column header.

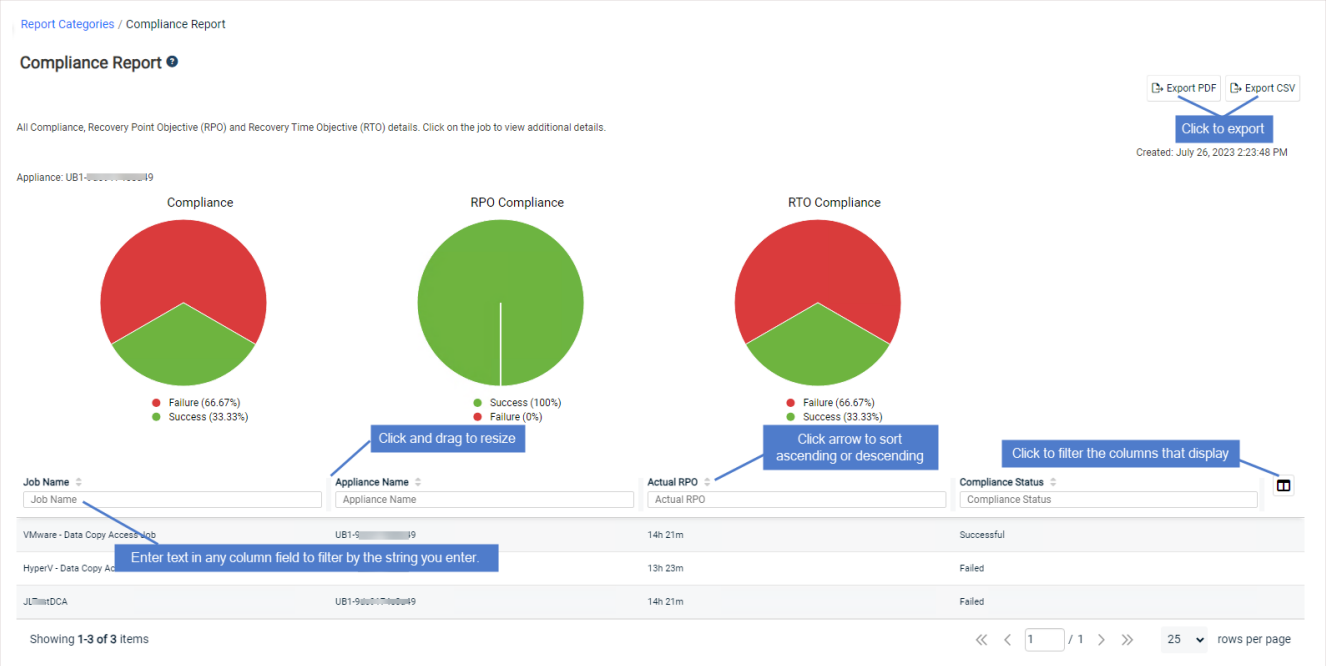

Report Field	Field Description
Job Name	The name of the data copy access job.
Appliance Name	The name of the backup appliance.
Profile Name	The name of the lab profile used in the data copy access job.
Profile RPO	The RPO specified in the lab profile you have selected for this data copy access job.
Profile RTO	The RTO specified in the lab profile you have selected for this data copy access job.
Actual RPO	The actual amount of data loss that would occur in the event of a disaster. Also known as RPA.
Actual RTO	The actual period of downtime that would occur in the event of a disaster. Also known as RTA.
Compliance Status	Indicates whether the data copy access job complies with RPO and RTO policies: • Failed - The job does not meet RPO and RTO objectives. • Passed - The job does meet RPO and RTO objectives.

Email reports

Email reports provide an additional avenue for you to receive regular updates on the status of your Unitrends Backup, Recovery MAX, and Recovery Series appliances. For information on configuring email reports, see "Email reporting" on page 117.

For details on specific email report types, see the following topics:

Appliance Status report

The Appliance Status report provides a summary of the status of all backup and hot backup copy jobs that occurred on the appliance in the last 24 hours. The first section of the report displays file-level and bare metal backups. The remaining sections show application and virtual machine backups for the following (as applicable): Windows image-level, Microsoft SQL Server, Microsoft Exchange Server, SharePoint, Oracle, UCS service profiles, NDMP, Hyper-V, VMware, XenServer, and AHV. To receive Appliance Status reports, you must select the Appliance report type when you configure email reporting for your appliance. (For details on configuring email reports, see "Email reporting" on page 117.)

A sample report is given here:

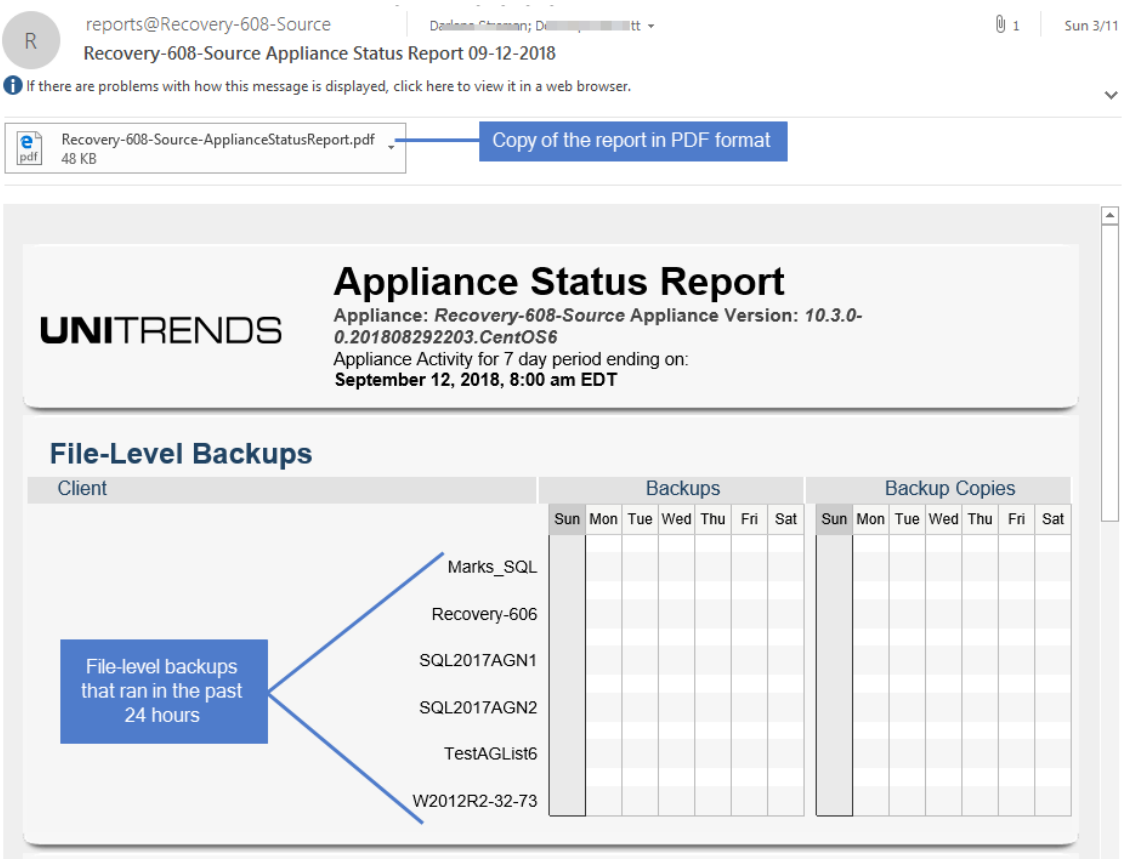
reports@Recovery-608-Source
Recovery-608-Source Appliance Status Report 09-12-2018
If there are problems with how this message is displayed, click here to view it in a web browser.
Recovery-608-Source-ApplianceStatusReport.pdf
48 KB
Copy of the report in PDF format
UNITRENDS
Appliance Status Report
Appliance: Recovery-608-Source Appliance Version: 10.3.0-0.201808292203.CentOS6
Appliance Activity for 7 day period ending on:
September 12, 2018, 8:00 am EDT
File-Level Backups
Client
Backups
Backup Copies
Sun Mon Tue Wed Thu Fri Sat
Marks_SQL
Recovery-606
SQL2017AGN1
SQL2017AGN2
TestAGList6
W2012R2-32-73
File-level backups that ran in the past 24 hours

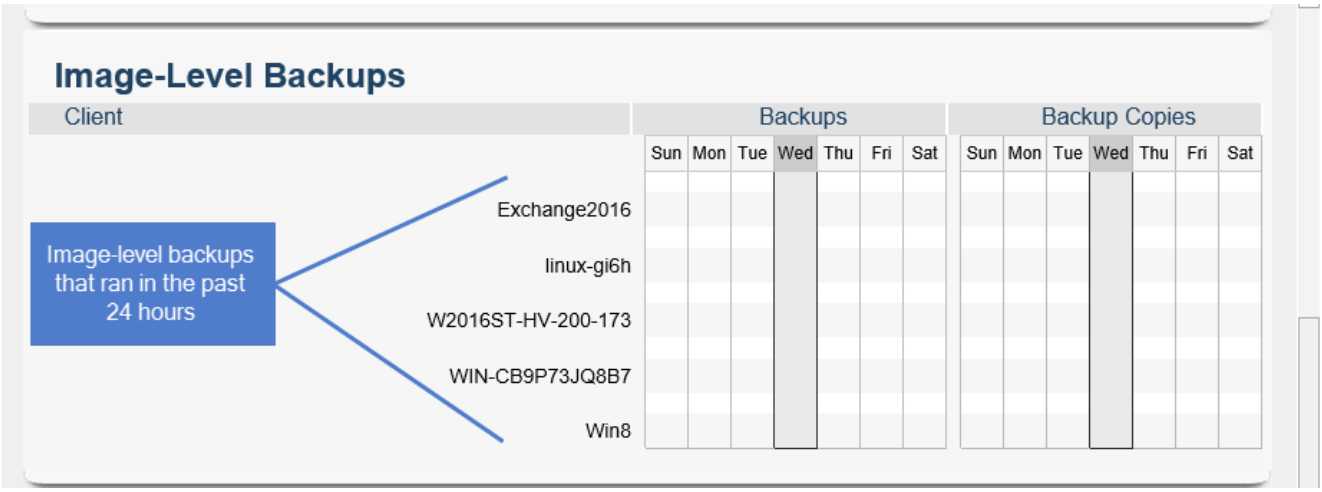
Image-Level Backups
Client
Backups
Backup Copies
Sun Mon Tue Wed Thu Fri Sat
Exchange2016
linux-gi6h
W2016ST-HV-200-173
WIN-CB9P73JQ8B7
Win8
Image-level backups that ran in the past 24 hours

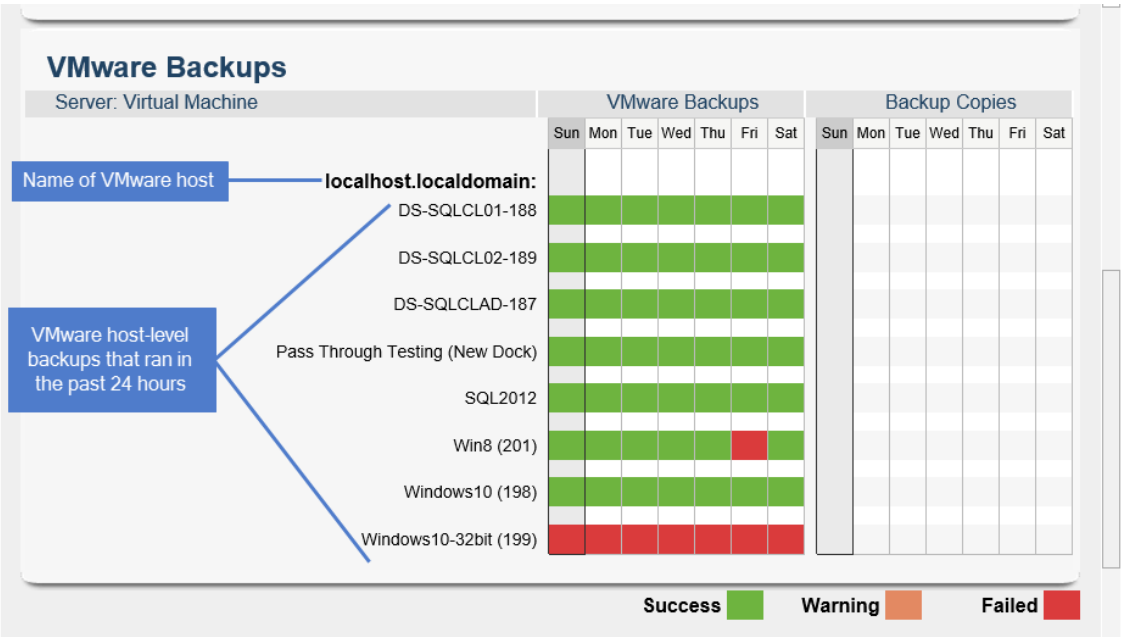

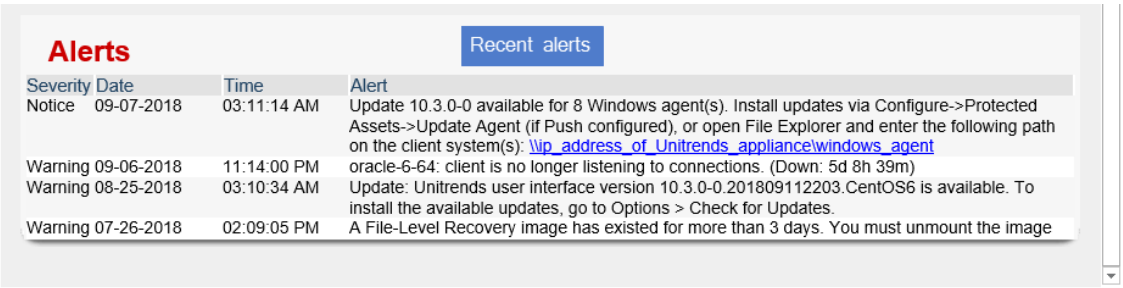

Backup Copy Hot Targets report

The Backup Copy Hot Targets report provides a summary of hot backup copy jobs that occurred on the appliance. This report is generated daily and documents activity that occurred in the preceding 24 hours. To receive Backup Copy Hot Targets reports, you must select the Appliance report type when you configure email reporting for your appliance. (For details on configuring email reports, see "Email reporting" on page 117.)

A sample report is given here:

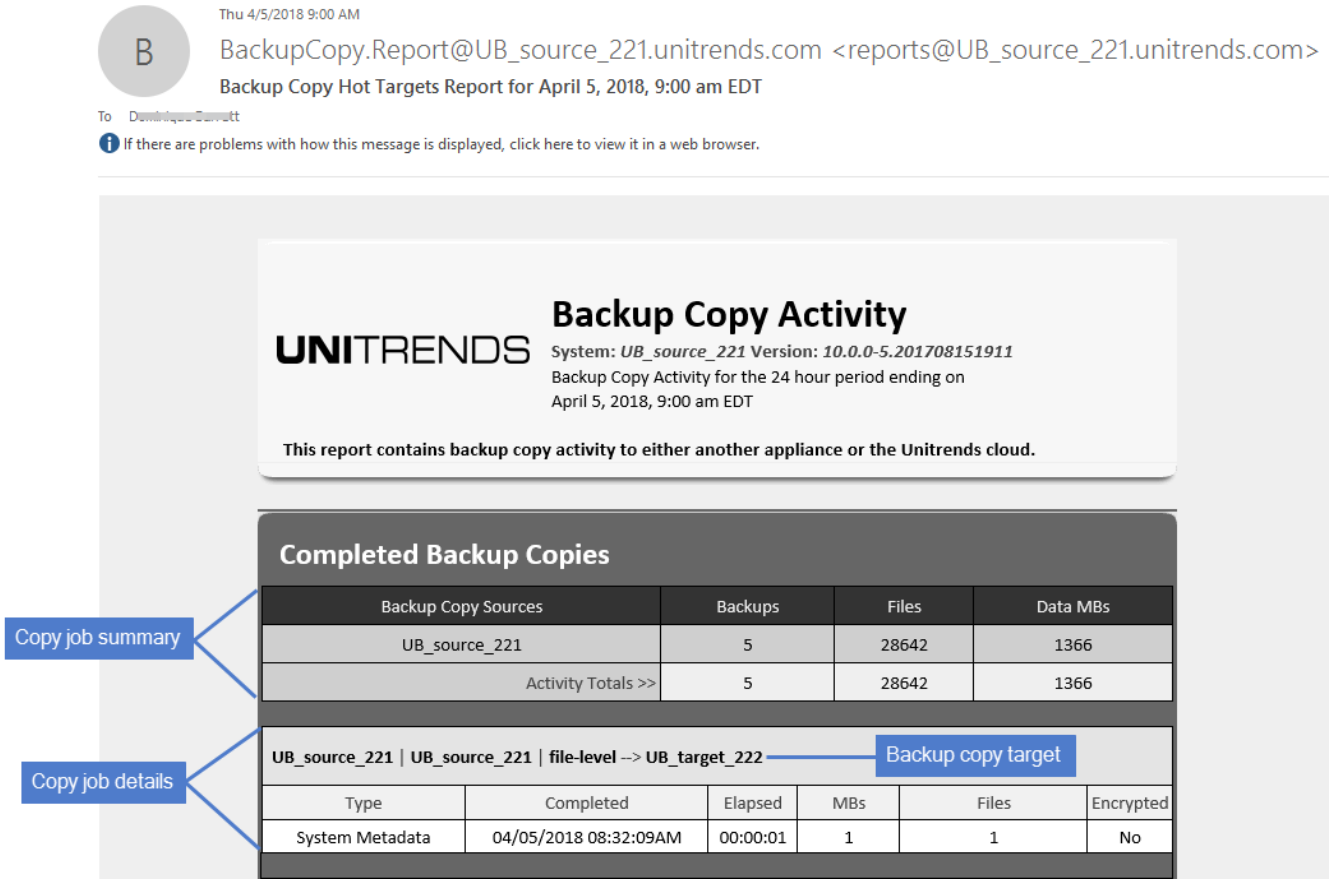

Thu 4/5/2018 9:00 AM

BackupCopy.Report@UB_source_221.unitrends.com <reports@UB_source_221.unitrends.com>

Backup Copy Hot Targets Report for April 5, 2018, 9:00 am EDT

To [redacted]

If there are problems with how this message is displayed, click here to view it in a web browser.

UNITRENDS

Backup Copy Activity

System: *UB_source_221* **Version:** *10.0.0-5.201708151911*

Backup Copy Activity for the 24 hour period ending on
April 5, 2018, 9:00 am EDT

This report contains backup copy activity to either another appliance or the Unitrends cloud.

Completed Backup Copies

Copy job summary

Backup Copy Sources	Backups	Files	Data MBs
UB_source_221	5	28642	1366
Activity Totals >>	5	28642	1366

Copy job details

UB_source_221 | UB_source_221 | file-level --> UB_target_222 — Backup copy target

Type	Completed	Elapsed	MBs	Files	Encrypted
System Metadata	04/05/2018 08:32:09AM	00:00:01	1	1	No

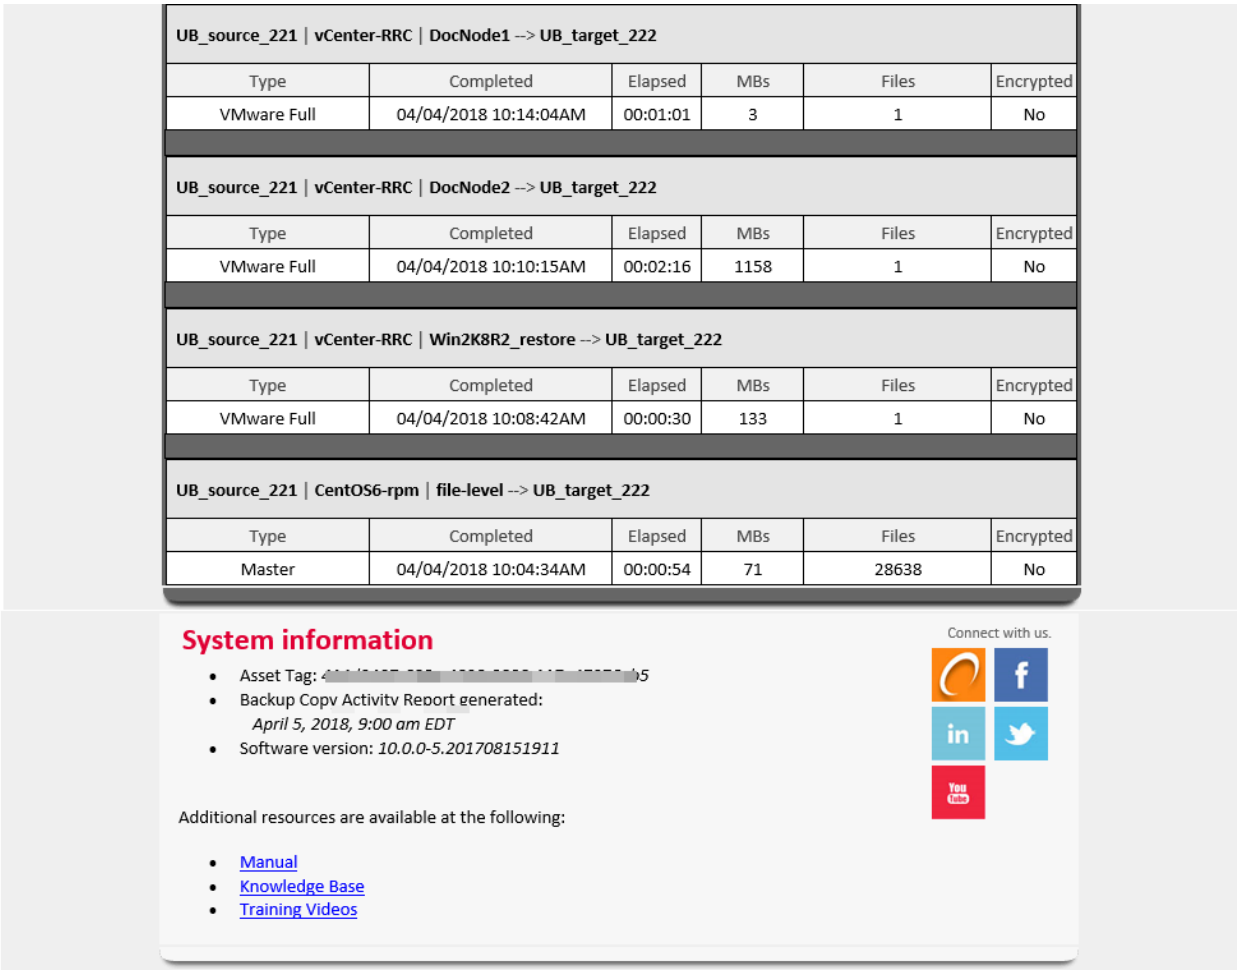

UB_source_221 | vCenter-RRC | DocNode1 --> UB_target_222

Type	Completed	Elapsed	MBs	Files	Encrypted
VMware Full	04/04/2018 10:14:04AM	00:01:01	3	1	No

UB_source_221 | vCenter-RRC | DocNode2 --> UB_target_222

Type	Completed	Elapsed	MBs	Files	Encrypted
VMware Full	04/04/2018 10:10:15AM	00:02:16	1158	1	No

UB_source_221 | vCenter-RRC | Win2K8R2_restore --> UB_target_222

Type	Completed	Elapsed	MBs	Files	Encrypted
VMware Full	04/04/2018 10:08:42AM	00:00:30	133	1	No

UB_source_221 | CentOS6-rpm | file-level --> UB_target_222

Type	Completed	Elapsed	MBs	Files	Encrypted
Master	04/04/2018 10:04:34AM	00:00:54	71	28638	No

System information

- Asset Tag: [redacted]
- Backup Copy Activity Report generated:
  April 5, 2018, 9:00 am EDT
- Software version: *10.0.0-5.201708151911*

Connect with us.

Additional resources are available at the following:

- Manual
- Knowledge Base
- Training Videos

Backup Copy Job Notifications report

The Backup Copy Job Notifications report is generated and sent each time a cold backup copy job finishes. To receive Backup Copy Job Notifications reports, you must select the Jobs report type when you configure email reporting for your appliance. (For details on configuring email reports, see "Email reporting" on page 117.)

A sample report is given here:

Compliance report

The Compliance email report provides an overview of current RPO/RTO compliance based on the results of data copy access (DCA) jobs that ran during the last week. The Compliance report automatically refreshes daily at 12:00 AM. To receive the report, you must select the Appliance report type when you configure email reporting for your appliance. (For details on configuring email reports, see "Email reporting" on page 117.)

Note: You can also generate the Compliance report from the appliance UI. For details, see "Compliance report" on page 1370.

An example report is given here. Refer to the table below for a description of the report columns.

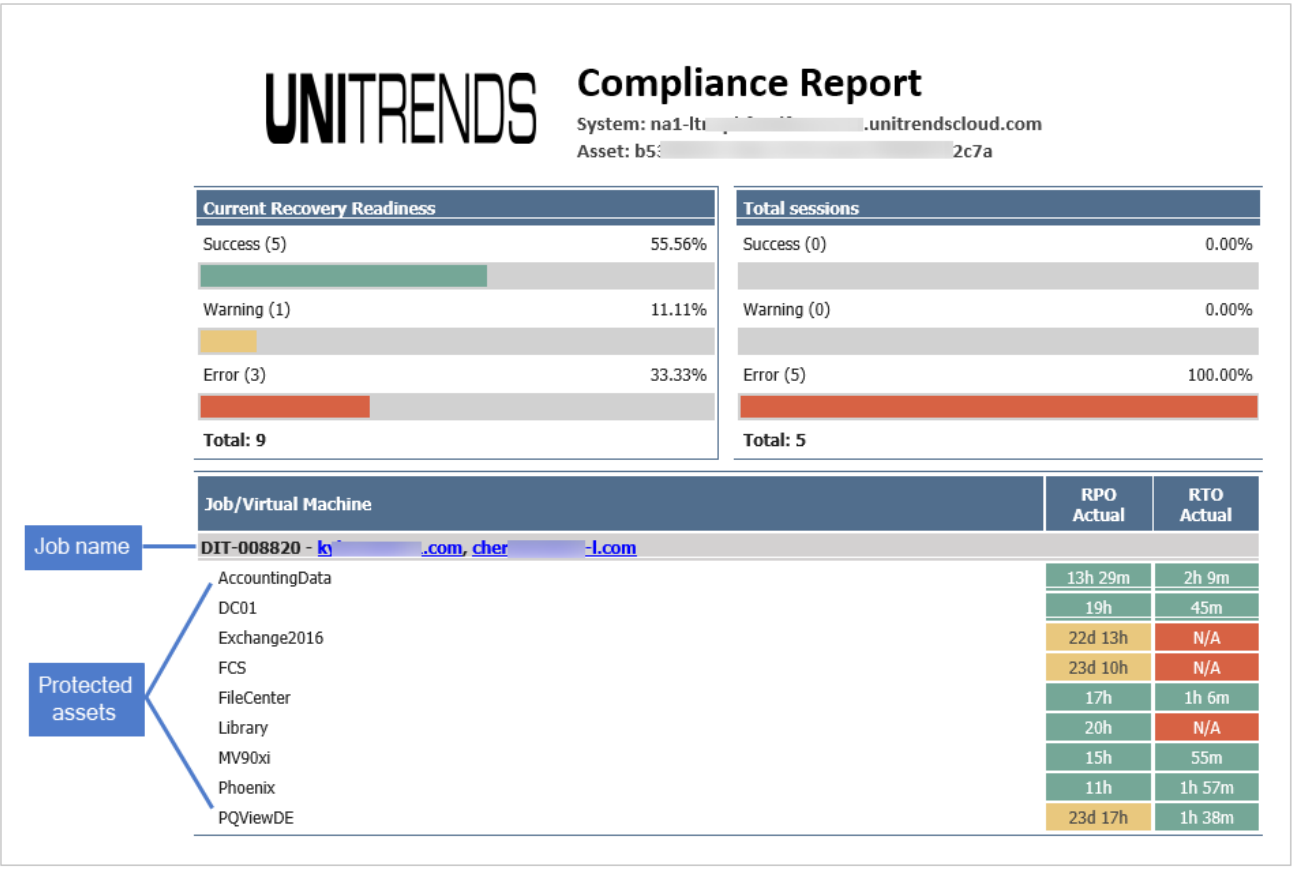

Report Field	Field Description
Current Recovery Readiness	Overview of the portion of assets that are in compliance versus assets that are not in compliance.
Success	Portion of assets that are in compliance with their job's RPO and RTO policies.
Warning	Portion of assets for which the DCA test job completed with warnings.
Error	Portion of assets for which the DCA test did not complete due to an error.
Total	The total number of assets protected by the job(s).
Total Sessions	Overview of the DCA job sessions that ran over the past week.
Success	Portion of sessions whose tests all completed successfully within the job's RPO

Report Field	Field Description
	and RTO window.
Warning	Portion of sessions where one or more tests completed with warnings.
Error	Portion of sessions where one or more tests were not completed due to an error.
Total	Total number of times the DCA job(s) ran over the last week.
Job/Virtual Machine	The name of the DCA job, followed by the names of each asset protected by the job.
RPO Actual	Recovery Point Objective Actual – Time since the asset's most recent recovery point: • Green indicates the asset is in compliance • Yellow indicates the asset's DCA test did not complete within the job's RPO policy • Red indicates the asset's DCA test did not complete due to an error
RTO Actual	Recovery Time Objective Actual – Amount of time it took to run the most recent DCA test of the asset: • Green indicates the asset is in compliance • Yellow indicates the asset's DCA test did not complete within the job's RTO policy • Red indicates the asset's DCA test did not complete due to an error

DCA Job Notifications report

The DCA Job Notifications report is generated and sent each time a data copy access test job finishes. The report includes screenshots of each DCA VM that was booted during the test. To receive DCA Job Notifications reports, you must select the Jobs report type when you configure email reporting for your appliance. (For details on configuring email reports, see "Email reporting" on page 117.)

A sample report is given here:

From: <reportuser@unitrends.com>
Date: Tue, Mar 2, 2021 at 2:21 PM
Subject: [recovery-8024s]Unitrends recovery-8024s Report (Test) - Success
To: <as[redacted]ail.com>

Test result

Screenshots of VMs tested

3 attachments (113 KB) Download all Save all to OneDrive - Kaseya

Job Report

System: recovery-8024s[redacted]ain
Asset: 8024S-10[redacted]

Test Report	
Job Name:	OnDemandJob
Session start time:	2021/03/02 05:18:52 PM
Total time:	2m44s

Recovery Step	Result	Progress %	Start Time	End Time	Description
OnDemandJob	Successful	100	2021/03/02 05:18:52 PM	2021/03/02 05:21:36 PM	Job Started
Get Job Properties	Successful	100	2021/03/02 05:18:52 PM	2021/03/02 05:18:52 PM	Finished
Boot Order: Group 1	Successful	100	2021/03/02 05:18:52 PM	2021/03/02 05:21:24 PM	Successful

PaulWin10	Successful	100	2021/03/02 05:18:52 PM	2021/03/02 05:21:04 PM	Finished
Start Instant Recovery	Successful	100	2021/03/02 05:18:52 PM	2021-03-02 05:18:56 PM	Finished
Setting Network	Successful	100	2021/03/02 05:18:56 PM	2021-03-02 05:20:56 PM	Finished
Capture Screen Shot	Successful	100	2021/03/02 05:20:56 PM	2021-03-02 05:20:59 PM	Finished
RPO/RTO Compliance	Successful	100	2021/03/02 05:20:59 PM	2021-03-02 05:20:59 PM	Not Configured
Tear down VM	Successful	100	2021/03/02 05:20:59 PM	2021-03-02 05:21:04 PM	Finished
Update Job Compliance (RPO/RTO)	Successful	100	2021/03/02 05:21:04 PM	2021-03-02 05:21:04 PM	Finished
PaulWin2008R2	Successful	100	2021/03/02 05:18:52 PM	2021/03/02 05:21:31 PM	Finished
Start Instant Recovery	Successful	100	2021/03/02 05:18:52 PM	2021-03-02 05:19:16 PM	Finished
Setting Network	Successful	100	2021/03/02 05:19:16 PM	2021-03-02 05:21:16 PM	Finished
Capture Screen Shot	Successful	100	2021/03/02 05:21:16 PM	2021-03-02 05:21:19 PM	Finished
RPO/RTO Compliance	Successful	100	2021/03/02 05:21:19 PM	2021-03-02 05:21:19 PM	Not Configured
Tear down VM	Successful	100	2021/03/02 05:21:19 PM	2021-03-02 05:21:31 PM	Finished
Update Job Compliance (RPO/RTO)	Successful	100	2021/03/02 05:21:31 PM	2021-03-02 05:21:31 PM	Finished
PaulWin2012R2efi	Successful	100	2021/03/02 05:18:52 PM	2021/03/02 05:21:36 PM	Finished
Start Instant Recovery	Successful	100	2021/03/02 05:18:52 PM	2021-03-02 05:19:20 PM	Finished
Setting Network	Successful	100	2021/03/02 05:19:20 PM	2021-03-02 05:21:20 PM	Finished

Management Status report

The Management Status report provides a summary of all backup jobs, backup copy jobs, and alerts that occurred on the managing appliance and its constituent managed appliances. Management Status reports are generated daily and document activity that occurred in the preceding 24 hour period. These reports are only generated by manager appliances. To receive Management Status reports from your managing appliance, you must select the Appliance

report type when you configure email reporting for your appliance. (For details on configuring email reports, see "Email reporting" on page 117.)

A sample report is given here:

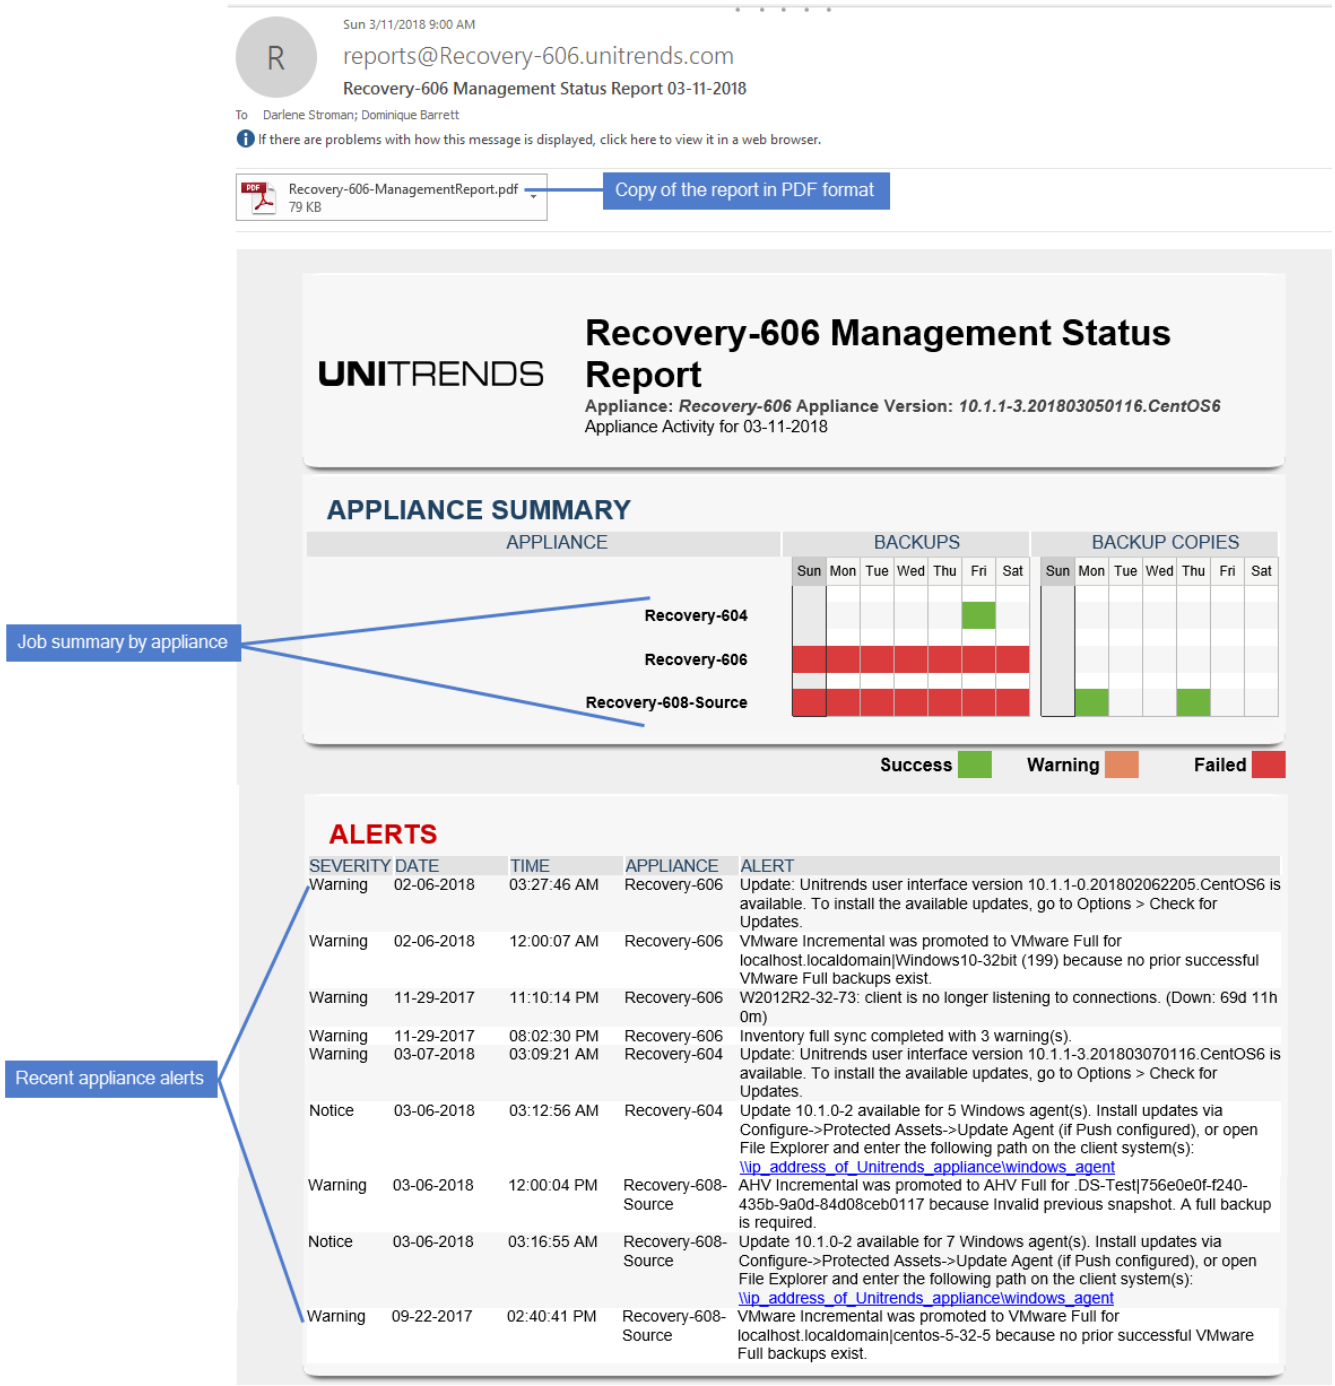

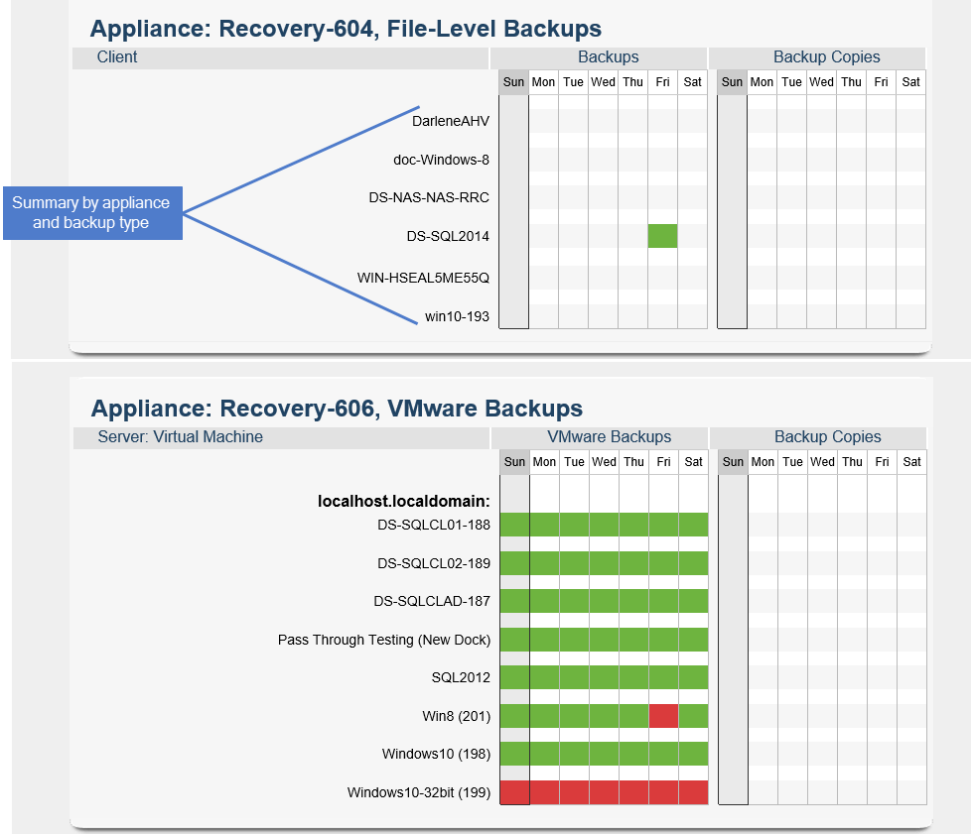
Appliance: Recovery-604, File-Level Backups
Client
Backups
Backup Copies
Sun Mon Tue Wed Thu Fri Sat
Summary by appliance and backup type
DarleneAHV
doc-Windows-8
DS-NAS-NAS-RRC
DS-SQL2014
WIN-HSEAL5ME55Q
win10-193
Appliance: Recovery-606, VMware Backups
Server: Virtual Machine
VMware Backups
Backup Copies
localhost.localdomain:
DS-SQLCL01-188
DS-SQLCL02-189
DS-SQLCLAD-187
Pass Through Testing (New Dock)
SQL2012
Win8 (201)
Windows10 (198)
Windows10-32bit (199)

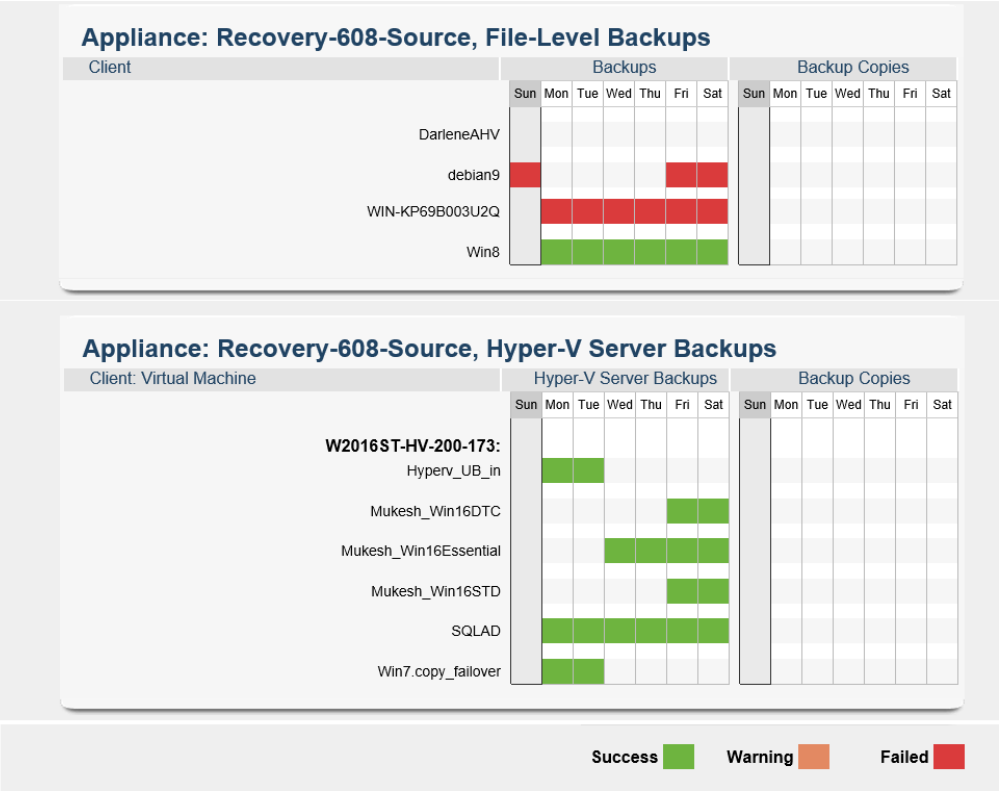

Schedule report

The Schedule report provides a summary of scheduled backup jobs that occurred on the appliance. The Schedule report is generated daily and documents activity that occurred in the preceding 24 hours. To receive Schedule reports, you must select the Jobs report type when you configure email reporting for your appliance. (For details on configuring email reports, see "Email reporting" on page 117.)

A sample report is given here:

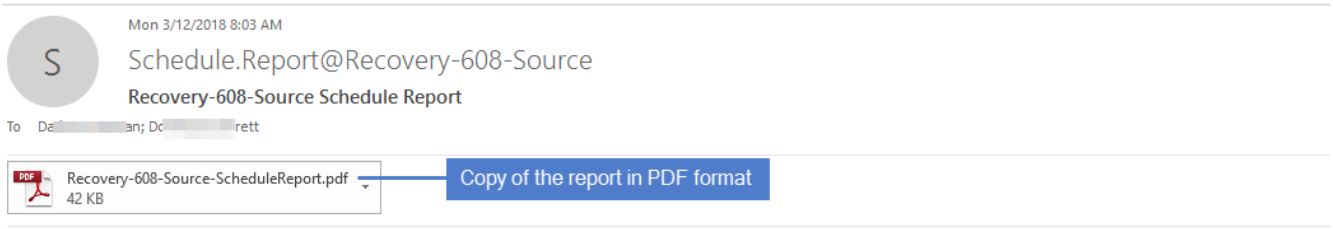

UNITRENDS

SCHEDULE REPORT

Schedule activity for 24 hours prior to: **Mon, 12 Mar 2018 08:01:55 -0400**

System name: **Recovery-608-Source**

There were 7 failures.

Verify status (V column): S=Success, F=Fail, A=Active, -=Not Applicable

Schedule Name

Schedule events for schedule 'Filter For DS in Name'

Client name: W2016ST-HV-200-173

Schedule description

Application name: Hyper-V 2016

Description: Filter For DS in Name

Failures: 0

3:00 AM Backups

Job details

Virtual Machine	Id	Backup Type	V	Files	MBs	Status
SQLAD	9742	Incremental	-	N/A	4813	Successful
Mukesh_Win16Essential	9743	Incremental	-	N/A	3665	Successful
Mukesh_Win16STD	9744	Incremental	-	N/A	827	Successful
Mukesh_Win16DTC	9745	Incremental	-	N/A	3	Successful
UB_ramp	9746	Full	-	N/A	10107	Successful

Schedule events for schedule 'Oracle-bkup'

Client name: orcl

Application name: Oracle 11

Description: Oracle-bkup

Failures: 1

3:00 AM Backups

Oracle Instance	Id	Backup Type	V	Files	MBs	Status
orcl	9739	Full	-	N/A	0	Failed

Schedule events for schedule 'Win Backup (old SLA job)'

Application name: file-level

Description: Job for SLA Policy "Win SLA Policy"; update by modifying the policy.

Failures: 2

3:00 PM Backups

Client Name	Id	Backup Type	V	Files	MBs	Status
Win8	9721	Incremental	S	1369	2547	Successful
WIN-KP69B003U2Q	9723	Incremental	F	0	0	Task failed: Could not contact client after 5 attempts.

3:00 AM Backups

Client Name	Id	Backup Type	V	Files	MBs	Status
Win8	9741	Incremental	S	1563	3112	Successful
WIN-KP69B003U2Q	9747	Incremental	F	0	0	Task failed: Could not contact client after 5 attempts.

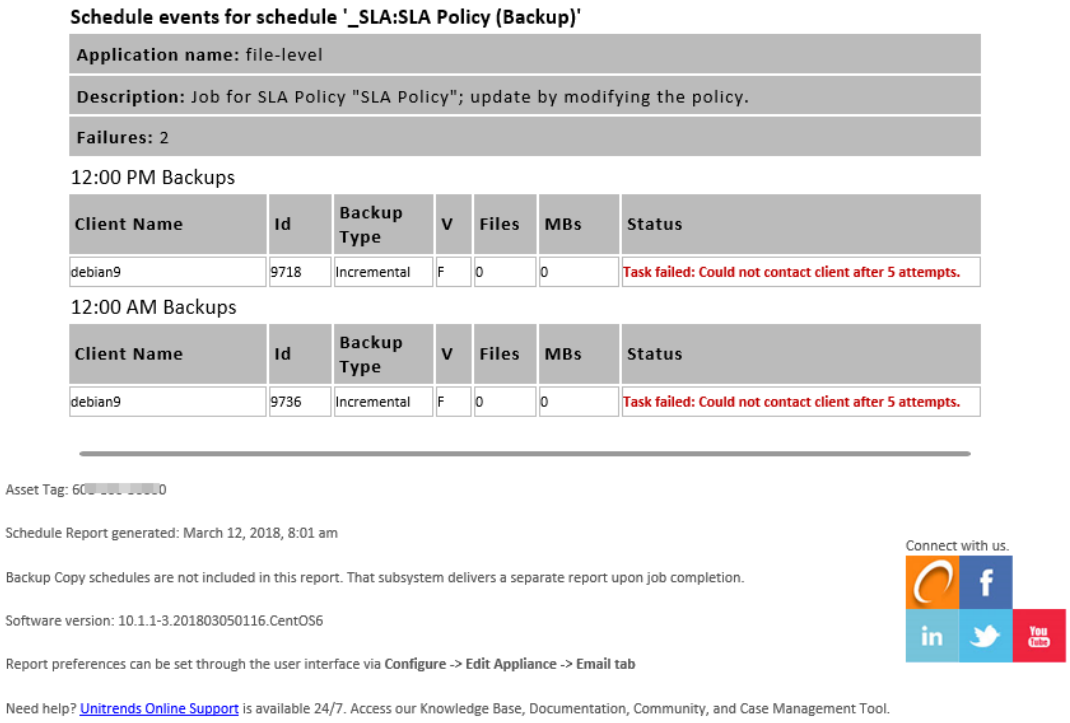

Schedule events for schedule '_SLA:SLA Policy (Backup)'

Application name: file-level

Description: Job for SLA Policy "SLA Policy"; update by modifying the policy.

Failures: 2

12:00 PM Backups

Client Name	Id	Backup Type	V	Files	MBs	Status
debian9	9718	Incremental	F	0	0	Task failed: Could not contact client after 5 attempts.

12:00 AM Backups

Client Name	Id	Backup Type	V	Files	MBs	Status
debian9	9736	Incremental	F	0	0	Task failed: Could not contact client after 5 attempts.

Asset Tag: 60[redacted]0

Schedule Report generated: March 12, 2018, 8:01 am

Backup Copy schedules are not included in this report. That subsystem delivers a separate report upon job completion.

Software version: 10.1.1-3.201803050116.CentOS6

Report preferences can be set through the user interface via **Configure -> Edit Appliance -> Email tab**

Need help? Unitrends Online Support is available 24/7. Access our Knowledge Base, Documentation, Community, and Case Management Tool.

Connect with us.

Schedule Success report

The Schedule Success report provides a summary of scheduled backup jobs that ran over the preceding hour and completed successfully. To receive Schedule Success reports, contact Unitrends Support for assistance (see "Support for Unitrends appliances" on page 25).

A sample report is given here:

UNITRENDS

SCHEDULE SUCCESS REPORT

Summary

Schedule events successful in the hour prior to: **July 3, 2019, 2:06 am**
System name: **Recovery-813**
There were 10 successful backups.

Schedule name

Successes for schedule 'BuzzWin7_8_FILE'

Schedule description

Application name: file-level
Description: BuzzWin7_FILE

Job details

Client Name	Backup Id	Backup Type	Status
BuzzWin7MBR	27824	Incremental	Successful
BuzzWin8ESX	27825	Incremental	Successful

Successes for schedule 'BuzzWin7_BLK'

Client name: BuzzWin7MBR
Application name: image-level
Description: BuzzWin7_BLK

Client Name	Backup Id	Backup Type	Status
Image Incremental	27816	Image Incremental	Successful
Image Incremental	27822	Image Incremental	Successful
Image Incremental	27830	Image Incremental	Successful

Successes for schedule 'BuzzWin8HVM_BLK'

Client name: BuzzWin8HVM
Application name: image-level
Description: Backup Job

Client Name	Backup Id	Backup Type	Status
Image Incremental	27814	Image Incremental	Successful

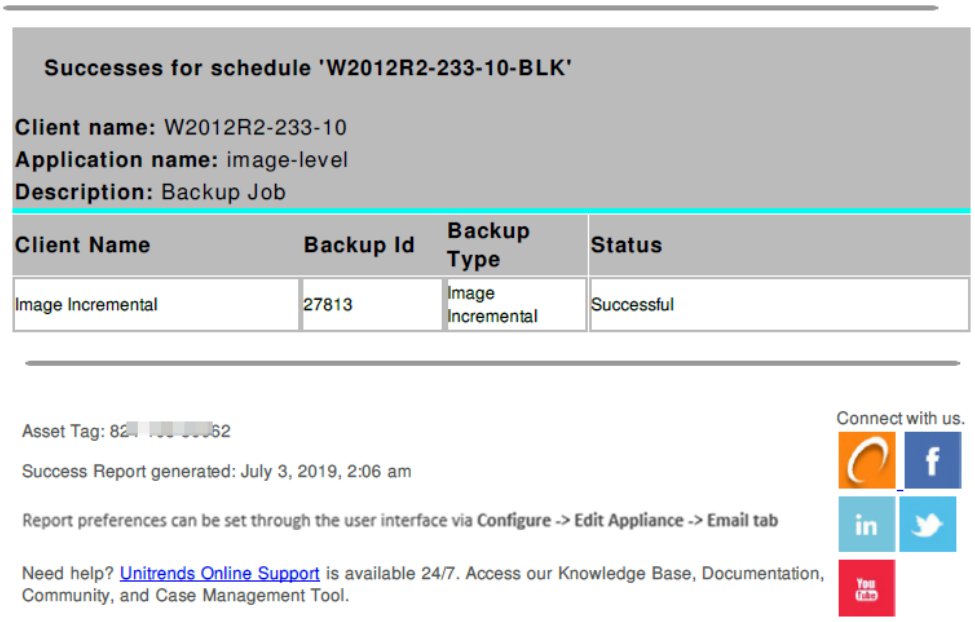

Successes for schedule 'W2012R2-233-10-BLK'

Client name: W2012R2-233-10
Application name: image-level
Description: Backup Job

Client Name	Backup Id	Backup Type	Status
Image Incremental	27813	Image Incremental	Successful

Asset Tag: 82 62

Success Report generated: July 3, 2019, 2:06 am

Report preferences can be set through the user interface via **Configure -> Edit Appliance -> Email tab**

Need help? Unitrends Online Support is available 24/7. Access our Knowledge Base, Documentation, Community, and Case Management Tool.

Connect with us.

Schedule Failure report

The Schedule Failure report provides a summary of scheduled backup jobs that ran over the preceding hour and completed in failed status. To receive Schedule Failure reports, you must select the Failures report type when you configure email reporting for your appliance. (For details on configuring email reports, see "Email reporting" on page 117.)

A sample report is given here:


UNITRENDS
SCHEDULE FAILURE REPORT
Summary
Schedule events that failed in the hour prior to: June 25, 2019, 4:00 pm
System name: Recovery-606
There were 2 failures.
Schedule name
Failures for schedule '_SLA:UEB10-14984 (Backup)'
Schedule description
Application name: file-level
Description: Job for SLA Policy "UEB10-14984"; update by modifying the policy.
Job details
Client Name
Backup Id
Backup Type
Status
SQL2017AGN1
1040467
Incremental
Task failed: Could not contact client after 5 attempts.
SQL2017AGN2
1040468
Incremental
Task failed: Could not contact client after 5 attempts.
Asset Tag: 60
01
Failure Report generated: June 25, 2019, 4:00 pm
Report preferences can be set through the user interface via Configure -> Edit Appliance -> Email tab
Need help? Unitrends Online Support is available 24/7. Access our Knowledge Base, Documentation, Community, and Case Management Tool.
Connect with us.