Administrator Guide for Recovery Series and Unitrends Backup

User Guide | January 2017

Administrator Guide for Recovery Series and Unitrends Backup

Release 9.1.1 | Version 2.01172017

Provide feedback

UNITRENDS

Copyright

Unitrends, Inc
200 Wheeler Road
North Tower, 2nd Floor
Burlington, MA 01803, USA
Phone: 1.866.359.5411

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, Citrix XenServer, NAS, SAN, Windows, Linux, SQL, SharePoint, Exchange, and Oracle.

Unitrends Recovery Series is a family of physical backup appliances for virtual and physical protection. This all-in-one solution integrates enterprise backup and recovery software with purpose-built hardware. It 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.

About this Guide

This guide describes how to administer Unitrends backup and recovery solutions using the current User Interface (UI). If you are using the legacy UI, see the Administrator Guide for Recovery Series and Unitrends Backup - Legacy Interface. For UI considerations, see *Which UI should I use?* in the Upgrade Guide for Recovery Series and Unitrends Backup.

Before getting started with this guide, you must deploy your appliance and configure network settings. For deployment instructions, see the following resources:

- Quick Start Guide for Recovery Series Appliances
- 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 in Microsoft Azure
- Deployment Guide for Unitrends Backup Installable Software
- Deployment Guide for Unitrends Free on VMware
- Deployment Guide for Unitrends Free on Hyper-V

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 appliances, Unitrends Backup editions, and Unitrends Free 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 Features
- Recovery Series Appliance Family Datasheet
- 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 appliance edition, environment, and type of data. 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.

- Overviews of features and instructions include considerations, requirements, and prerequisite information to assist in planning an effective protection strategy.
- Procedural information provides step-by-step instructions on performing backup and recovery operations. Instructions follow best practices for successful configuration and administration of your Unitrends data protection and recovery solution.
  See "Backup Administration and Procedures" on page 173 for instructions on performing, monitoring, and managing backup and backup copy jobs.
  For recovering backups, see the appropriate procedure for the backup type:
  - Host-level – See "Recovering Host-level Backups" on page 281.
  - Asset-level – See "Recovering Asset-level Backups" on page 305.
  - Applications – See "Recovering Application Backups" on page 341.
  - NAS – See "Recovering NAS Backups" on page 335.
  - iSeries – See "Recovering iSeries Backups" on page 365.
- 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.

Typographical convention	Description
`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 on VMware, Hyper-V, and Citrix XenServer 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 asset-level backups.

Term	Definition
Appliance	The Unitrends system that backs up and recovers data. Appliance can refer to: • A physical Recovery Series model. Consists of Unitrends hardware, Unitrends software, and additional configuration settings. • A Unitrends Backup system deployed as a virtual machine to a VMware, Hyper-V, Citrix XenServer, or Microsoft Azure environment. Consists of the Unitrends Backup VM, Unitrends software, attached storage, and additional configuration settings. • A Unitrends Backup system deployed to a CentOS or RHEL server. Also known as *Unitrends Backup Installable Software*. Consists of your CentOS or RHEL server, Unitrends software, and additional configuration settings.
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, 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. Equivalent to these legacy terms: • *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.

Term	Definition
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 and backup copies used to protect assets.
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 on VMware, Hyper-V, and Citrix XenServer appliances only. SAN or NAS storage that is connected directly to the Unitrends Backup VM over the iSCSI, CIFS, or NFS protocol. **Note:** SAN or NAS storage can also be used to store backup copies. This is supported on Recovery Series and Unitrends Backup appliances.
Global Options toolbar	Toolbar across the top of the user interface that includes several menus to quickly edit global options, perform administrative tasks, and access additional resources.
Initial backup storage	Applies to Unitrends Backup on VMware, Hyper-V, and Citrix XenServer 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 170GB - 64TB in size.
Initial disk	Applies to Unitrends Backup on VMware, Hyper-V, and Citrix XenServer appliances only. 100GB disk used to create the Unitrends Backup VM. While installing the EXE or OVA, you select a datastore on the ESXi host that the installer uses to create this disk.
Instant Recovery (IR)	Process that recovers a failed or corrupted virtual machine or Windows physical machine in minutes.
Job	Procedures performed to protect assets. Multiple job types exist, all of which can be monitored from the Active Jobs tile while in progress.

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. Created by performing instant recovery.
Resources	Amount of space, bandwith, disk space, memory, etc, consumed by the job or object.
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 on VMware, Hyper-V, and Citrix XenServer appliances only. Virtual machine created either by running the Unitrends Backup EXE installer or by deploying the Unitrends Backup OVA file.
Virtual Failover Client (VFC)	Virtual replica created during instant recovery of a failed Windows asset that can immediately assume the role of the original client in the event of a disaster.
Virtual Host	Host on which virtual machine assets reside. Also called a *hypervisor*.
Windows Instant Recovery (WIR)	Temporary solution for rapid recovery of a failed Windows asset. Creates a virtual replica of the asset that can immediately assume the role of the original in the event of a disaster.

Support for Unitrends Recovery Series and Unitrends Backup appliances

Support is provided through the following resources:

Unitrends support site

Access the Unitrends Support Site at http://www.unitrends.com/support, where you can:

- Download or upgrade your product
- Download latest agent releases
- Search Knowledge Base articles
- Connect with Community Forums
- Log a support case
- Access the Partner Service Portal

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.

Unitrends community forums

We encourage you to post questions in the forums and answer questions posted by other Recovery Series and Unitrends Backup users. Forums are moderated by Unitrends Customer Support. Access forum discussions at http://www.unitrends.com/support/connect-with-communities.

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. Access KB articles at http://www.unitrends.com/support/search-knowledge-base.

Unitrends product tour

After you complete the Quick Setup dialog, a short tour guides you through the necessary steps to begin protecting your environment. You can view this tour at any time by selecting **? > Product Tour**.

Navigating the User Interface

The user interface consists of a main dashboard and feature-specific tabs for easy navigation. Refer to the following topics for more information on navigating the user interface:

- "Dashboard" on page 22
- "Protect" on page 25

- "Recover" on page 26
- "Jobs" on page 28
- "Reports" on page 29
- "Configure" on page 30
- "Global Options toolbar" on page 33

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 Dashboard you can:

- Customize the Dashboard - To change the layout of the dashboard, click and hold the upper region of a tile, drag it to the desired location, and release. To specify which tiles to display, click the gear icon in the top right corner of the Backup summary tile, and select or clear tiles to display or hide them.
- Manually update the Dashboard tiles - The Dashboard tiles update hourly. To update the tiles at any time, click on the two-arrows icon in the top right corner of a tile. (Clicking the icon on the Backup Summary tile updates all other tiles.)
- Reset the Dashboard - To reset the Dashboard layout to its default mode for the current release, click the gear icon in the top right corner of the Backup Summary tile and click **Reset**.

See these topics for descriptions of each tile:

- "Active Jobs tile" on page 23
- "Backup Summary tile" on page 22
- "Got Questions? Ask the Community tile" on page 23
- "Storage tile" on page 23
- "Backup Copy - Hot Targets tile" on page 24
- "Backup Copy - Cold Targets tile" on page 24
- "Recover Summary tile" on page 24
- "Daily Feed tile" on page 25

Backup Summary tile

This tile displays the number of errors, assets not protected , and assets protected for all managed appliances. Bar graphs indicate backup performance.

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

- Errors - Displays all jobs, including canceled jobs, that ended in error within the last seven days. Clicking this number opens the **Backup Failures** report.
- Not protected - Displays the number of assets without a successful backup run within the last seven days. Clicking this number opens the **Protection Summary** report, which includes pie charts and tabular information on assests, backups, and backup copies. You can filter the table columns by selecting the drop-down in the upper right of a column.
- Protected - Displays the number of assets with a valid backup within the last seven days. Clicking this number opens the **Protection Summary** report, which includes pie charts and tabular information on assets, backups, and backup copies. You can filter the table columns by selecting the drop-down in the upper right of a column.

Got Questions? Ask the Community tile

This tile enables you to:

- Create a new forum account or add your email to an existing Unitrends community forum account by clicking the gear icon in the top right corner of the Got Questions? Ask the Community tile. Select **Create** or **Add**, fill in the field information, and click **Create Account** or **Add Account**.
- Specify the refresh rate for forum posts by clicking the gear icon in the top right corner of the Got Questions? Ask the Community tile, selecting a refresh rate, and clicking **Save**.

Storage tile

This tile shows details about available and used storage for all managed appliances. The tile displays the following information:

Field Name	Description
Type	Type of storage.
Name	Name of the storage.
Appliance	Name of the appliance associated with this storage.
Used/Allocated	The amount of available and used storage. Hovering over the bar in this column displays the amount of space used versus the amount still available.
Status	Hover over an 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).

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 links in the lower portion of the tile: All, Backup, Recover, Backup Copy.

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 only on appliances whose backups are being copied to the Unitrends Cloud or copied to another Unitrends appliance. This tile displays the number of errors, protected assets, and unprotected assets. Bar graphs indicate backup performance.

- 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. Clicking this number opens the "Backup Copy - Hot Targets report" on page 393.
- Protected - Displays the number of assets that have successful backup copies that ran within the last seven days. Clicking this number opens the "Protection Summary report" on page 389, which includes pie charts and tabular information on assets, backups, and backup copies. You can filter the table columns by selecting the drop-down in the upper right of a column.
- Average speed B/s - Displays the average backup copy speed of all jobs that ran in last seven days. Clicking this number opens the "Backup Copy - Hot Targets report" on page 393.
- 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).

Backup Copy – Cold Targets tile

Displays only on appliances whose backups are being copied to a cold backup copy target (eSATA, USB, tape, third-party cloud, attached disk, NAS, and SAN). 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. Use the legacy UI replication dashboard to view information from older sources.

- Errors - Displays the number of backup copy jobs with failures.
- 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.

Recover Summary tile

This tile displays details about recent recovery jobs on all managed appliances.

- Recent restores - Displays the number of recovery jobs in the last seven days. Clicking this number opens the Recovery History report.
- Active FLR - Displays the number of currently active file-level recovery objects. Clicking this number opens the File Level Recovery tab on the Recover page.

- Active IR - Displays the number of currently active instant recovery objects. Clicking this number opens the Instant Recovery tab on the Recover page.
- Avg speed - Displays the average recovery speed of all backup jobs within the last seven days (current day not included). The seven bars represent the seven days. Hover to the display the actual value and date for each bar.

Daily Feed tile

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

Protect

The Protect page provides status information about your protected assets. It shows an inventory of the assets and the status of backups and backup copies that ran over the last seven days. To adjust the refresh rate of the Protect page, click the configuration icon in its upper-right corner to display Tile Preferences. For more information about data protection, see the "Protection Overview" on page 35 chapter.

See the following for details on working with the Protect page:

- "Buttons" on page 25
- "Inventory tree" on page 25
- "Status table" on page 26
- "Filtering the Status table display" on page 26

Buttons

The Protect page contains the following buttons:

- Backup - Opens the Create Backup job dialog. For details on creating a backup job, see "Creating backup jobs" on page 173.
- Backup Copy - Opens the Create Backup Copy job dialog. For details on creating a backup copy job, see "Creating backup copy jobs" on page 186.

Inventory tree

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

- Appliances display as top-level nodes.
- Virtual hosts and physical assets 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 131 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.

Status table

Selecting an appliance or asset from the inventory tree populates the Status table, to the right, with details about backups and backup copies for each protected asset. The Status table can be filtered by the Name column. Backups run within the last seven days display status icons under the Backup and Backup Copies columns. Hovering over the status icons displays status information for a specific day. Clicking a status icon for an asset opens the Details screen. If the Backup and Backup Copies information areas are grayed-out, the selected asset has not been backed up within the past seven days.

Filtering the Status table display

Use the following to filter what is displayed in the Status table:

- Filter 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 backups that meet the condition you selected. For example, *Failed in the last 7 days* to see only the 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 backup copies that meet the condition you selected. For example, *Failed in the last 7 days* to see only the 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 copyin 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).

Recover

Use the Recover page to recover an entire asset, to recover individual files, or to perform instant recovery. 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" on page 26
- "File Level Recovery tab" on page 27
- "Instant Recovery tab" on page 27

Backup Catalog tab

This tab lists all backups, backup copies, and imported backups. Use the options in the Filter Backups area to customize the display.

Use these buttons while working with the backups or backup copies you have selected in the Backups 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 asset-level backups and copies only.

- Recover - Use to recover an entire asset from the selected 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 - A drop-down used to perform instant recovery of a Windows physical asset, a VMware virtual machine, or a Hyper-V virtual machine.
- 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.

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 more information, see "Recovering files from virtual machine backups" on page 283.

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:
  - The name of the recovery object
  - The creation date for the recovery object
  - The appliance name
  - Path to the recovery object (CIFS) if applicable (based on backup type)
  - iSCSI target for recovery object (if applicable)
  - Messages
- Browse/Download - Opens a File Browser for the selected object, where you can select files and/or folders to recover.
- Remove - Removes the selected file level recovery object .

Instant Recovery tab

Instant recovery enables you to recover a failed or corrupted VM or Windows physical asset and begin using it almost immediately. Performing VM instant recovery recreates the failed VM from a backup that you select. Performing Windows instant recovery creates a stand-by virtual replica of a Windows physical asset that is kept up to date as new backups of the original asset run. For more on instant recovery, see "Windows instant recovery" on page 316 and "Virtual machine instant recovery" on page 296.

Once you have performed instant recovery, use the Instant Recovery tab to view and work with the VM or Windows virtual replica object that the recovery created. Use these buttons while working with instant recovery objects:

- View Details - Show or hide the details pane.
- Edit - Displays the instant recovery configuration dialog for the selected instant recovery object.

- Go Live - Places a Windows instant recovery object into Live mode. Once in Live mode, instant recovery objects cannot be reset to another mode. All options other than Tear Down are unavailable.
- Audit - Places the selected Windows instant recovery object into Audit mode.
- Tear Down - Removes the instant recovery object from the appliance.

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" on page 28
- "Job Manager tab" on page 28
- "Recent Jobs tab" on page 28
- "Recent System Jobs tab" on page 29

Active Jobs tab

This tab displays all currently running jobs. Each asset protected in a backup job or backup is displayed as a separate running instance of that backup job. Use these buttons while working with active jobs:

- Create job - Select to open the Create Backup Job dialog or the Create Backup Copy Job dialog.
- View Details - Displays or hides the details pane for the selected job.
- Pause - Pauses the selected job.
- Cancel - Cancels 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. Use these buttons while working with the Job Manager:

- Create job - Select to open the Create Backup Job dialog or the Create Backup Copy Job dialog.
- View Details/Hide Details- Displays or hides the details pane for the selected job.
- Edit - Allows changes to the selected job.
- Disable/Enable - Disables or enables the selected job. An enabled job runs according to schedule and can be run on demand. A disabled job does not run.
- Delete - Removes the selected job.
- Run - Runs the selected job.

Recent Jobs tab

This tab displays backup and backup copy jobs that ran in the last seven days. Use these buttons while working with recent jobs:

- Create job - Select to open the Create Backup Job dialog or the Create Backup Copy Job dialog.
- Export CSV - Exports the job history as a CSV file.
- View log - Displays details for a selected job.

Recent System Jobs tab

This tab displays system jobs that ran in the last seven days. Use these buttons while working with recent system jobs:

- Create job - Select to open the Create Backup Job dialog or the Create Backup Copy Job dialog.
- Export CSV - Exports the job history as a CSV file.

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.

See the following for additional information:

Category	Available Reports
"Backup reports" on page 378	• "Protection Summary report" on page 389 • "Backup History report" on page 381 • "Legal Hold Backups report" on page 383 • "Backup Failures report" on page 384 • "Weekly Status report" on page 397
"Recover reports" on page 387	• "Recovery History report" on page 387
"Backup Copy reports" on page 389	• "Protection Summary report" on page 389 • "Backup Copy Capacity report" on page 392 • "Backup Copy - Hot Targets report" on page 393 • "Backup Copies - Past 24 Hours report" on page 395 • "Storage Footprint report" on page 395 • "Backup Copy - Cold Targets report" on page 396 • "Weekly Status report" on page 397

Category	Available Reports
"Appliance reports" on page 398	• "Update History report" on page 399 • "Capacity report" on page 399 • "Load report" on page 401 • "Alerts report" on page 401 • "Trap History report" on page 402 • "Notifications report" on page 403
"Storage reports" on page 404	• "Storage report " on page 404 • "Data Reduction report" on page 406

Configure

The Configure page enables you to manage appliances and assets. For procedures used to configure appliances and assets, see these topics in the Configuration chapter:

- "Appliance superuser account settings" on page 47
- "Appliance settings" on page 47
- "Protected assets" on page 111

The Configure page contains the following tabs:

- "Appliances tab" on page 30
- "Protected Assets tab" on page 31
- "Copied Assets tab" on page 32

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.

Appliance information

The Appliances tab displays the Unitrends appliance you are logged in to, 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.

Column	Description
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.
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 this appliance's UI.
- 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 about Storage, see "Backup storage" on page 69.
- Backup Copy Targets sub-tab - Use to add, edit, remove, disable/enable, and erase Backup Copy targets. For more information, see "Backup copy targets" on page 77.
- Network sub-tab - Use to view and edit network settings for each network adapter on the appliance. For more information, see "Networks" on page 65.
- 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 162.

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.

Asset information

The Protected Assets tab displays all assets that have been added to the appliance. 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.
Credentials	Indicates whether credentials have been assigned to the asset.
Retention	The asset's retention policy.
Agent Version	Unitrends agent version running on the asset.
Appliance	The name of the appliance that manages 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 appliance. For details, see "Protected assets" on page 111.
- 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 114.
- Manage Global VM Settings - Use to choose the quiesce setting that will be applied to all newly discovered VMware and XenServer VMs. You can also opt to apply this setting to current VMs. For details, see "Quiesce settings for host-level backups" on page 115.
- Update Agent - Use to install Windows agent updates on selected assets.
- Edit - Use to edit the selected asset. Modify various options, such as encryption, credentials, and retention. For details, see "Protected assets" on page 111.
- Remove - Use to remove the selected asset from the appliance. For details, see "Managing protected assets" on page 118.

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 by using the buttons across the top of the tab.

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 retention policy for this asset's backup copies.
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. For details, see "Managing retention settings" on page 130.
- Remove - Use to remove the selected copied asset from the appliance. For details, see "Managing protected assets" on page 118.

Global Options toolbar

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

- "Root (Avatar icon)" on page 33
- "Options (Gears icon)" on page 33
- "Help (question mark icon)" on page 34
- "Alert (Triangle '!' icon)" on page 34

Root (Avatar icon)

The default UI user account is *root*. From this menu you can view and edit your root user account details and log out.

Options (Gears icon)

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

- Inventory Sync - Use to update the inventory of protected virtual machines and databases.
- Check for updates - Use to check for appliance updates.
- Open the legacy interface - Use to open the appliance's legacy user interface.
- Deduplication Settings - Applies to Unitrends Backup appliances only. Use to modify the appliance deduplication level.

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.
- Product Tour – Select to open a tour that assists with registering a virtual host and creating a backup job, and shows you around the interface.
- 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.

Alert (Triangle '!' icon)

Click the Alert 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.

Chapter 2: Protection Overview

Unitrends Recovery Series 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 a protected *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 RTOs/RPOs. 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 the remainder of this chapter:

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. It is recommended to copy your local backups in order to recover from a disaster. See "Backup copies" on page 44.

Customize your backup strategy to meet your recovery point objectives (RPOs) and recovery time objectives (RTOs). 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 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:

- Run multiple jobs with multiple machines in each job.
- 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.

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.
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. • 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

Data is protected using these backup types:

- An asset-level backup protects an asset's file system and operating system. You must install a Unitrends agent on the asset for asset-level protection.

  Notes:

  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 asset-level backups are used to recover the machine (this is the recommended approach). But in the following cases a bare metal backup must be used instead:

  - To perform disaster recovery of a Windows 2003 asset to dissimilar hardware. (Supported for some distributions only. See the Compatibility and Interoperability Matrix for details.)
  - To perform disaster recovery of a Windows 2000 asset.
  - To perform disaster recovery in cases where the system state (boot and critical system volumes) has been excluded from asset-level backups. (All volumes and folders are included in asset-level backups by default. If you have opted to exclude this data, you cannot use asset-level backups for disaster recovery.)

  You must install the Windows bare metal agent to run bare metal backups. For details and requirements, see the Upgrade Guide for Recovery Series and Unitrends Backup.

- A host-level backup uses hypervisor snapshots to protect virtual machines. 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 server for application protection.
- A NAS backup protects data stored on a NAS device. You do not install an agent on the NAS asset.
- An iSeries backup protects 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 asset-level backups and hosted applications are protected with application backups.

Backing up virtual assets

For virtual assets, you can choose host-level or asset-level protection. Host-level backups capture files, application data, and virtual hardware. With asset-level protection, the appliance treats your VM as a physical asset to run asset-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 asset-level protection. For considerations specific to your environment, see "Protecting VMware virtual machines at the asset level" on page 216, "Protecting Hyper-V virtual machines at the asset level"

on page 207 , and "Best practices and requirements for XenServer protection" on page 220 to determine which approach to take.

Host-level protection	Asset-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 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.	You can recover Windows machines in minutes using the Windows instant recovery feature.
You can recover individual files from backups for VMs running Windows or Linux.	You can recovery 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 offsite target. These backups are duplicates (hence, "copies") 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 asset-level backups, host-level backups, application backups, NAS backups, and iSeries backups.

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 creating a backup job for a given asset, only supported modes are available for selection.

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 39
- "Incremental backup" on page 39

Full backup

A full backup captures all data on the asset:

- For asset-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 host-level backups, this includes VM metadata (configuration files) and blocks of all disks attached to the VM. For VMware, 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 229 or " NAS protection using NDMP" on page 230 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 261 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.

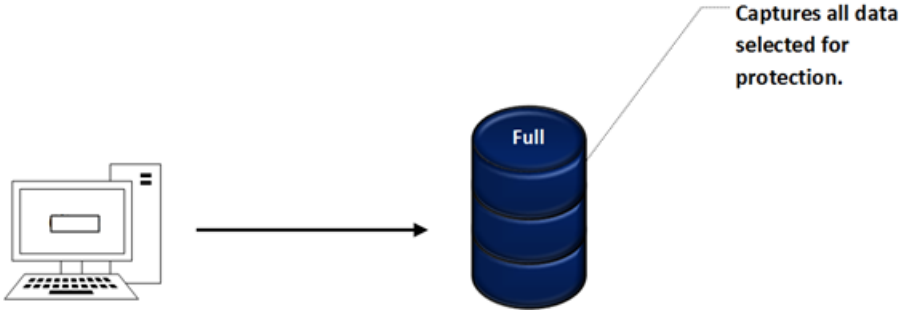

Incremental backup

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.

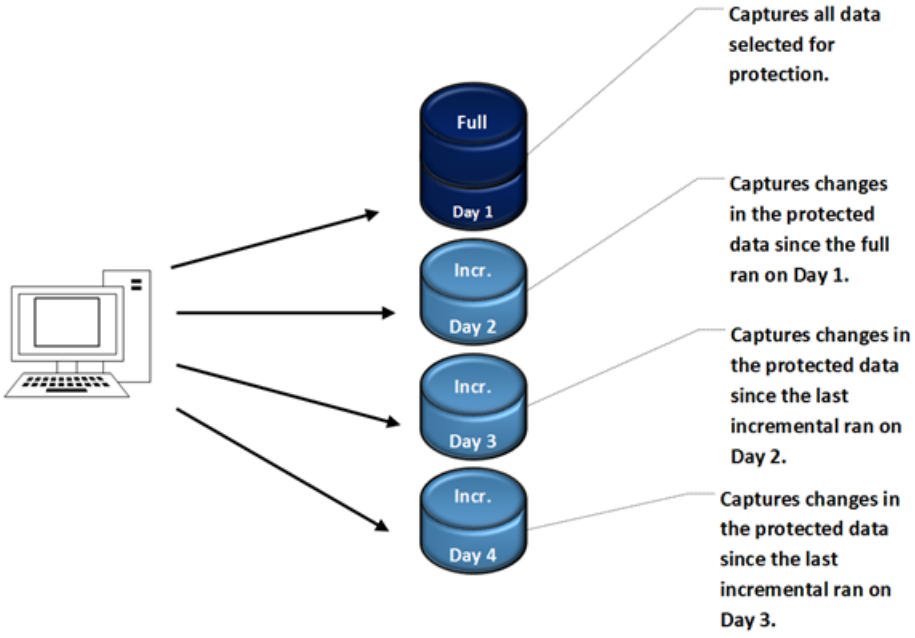

Differential backup

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.

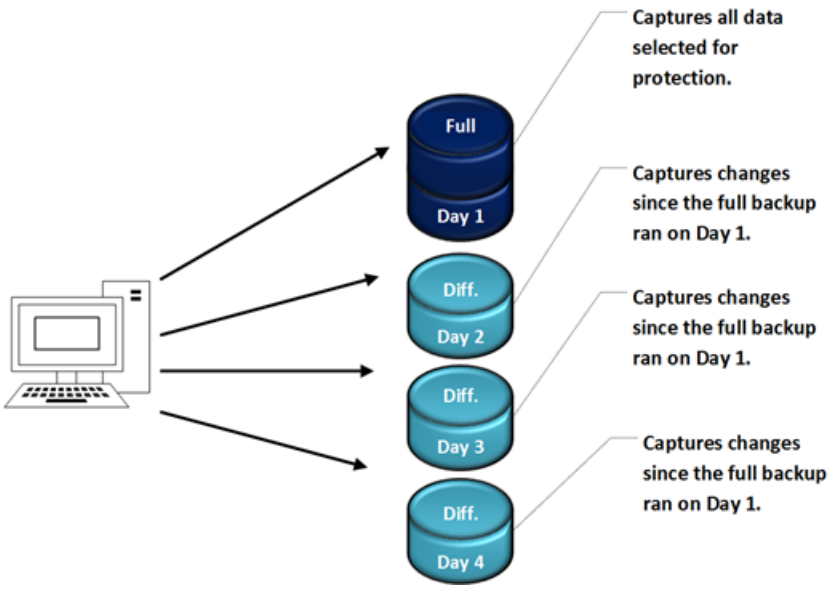

Selective backup

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 asset-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 asset-level backups are used to recover the machine (this is the recommended approach). But in the following cases a bare metal backup must be used instead:

- To perform disaster recovery of a Windows 2003 asset to dissimilar hardware. (Supported for some distributions only. See the Compatibility and Interoperability Matrix for details.)
- To perform disaster recovery of a Windows 2000 asset.
- To perform disaster recovery in cases where the system state (boot and critical system volumes) has been excluded from asset-level backups. (All volumes and folders are included in asset-level backups by default. If you have opted to exclude this data, you cannot use asset-level backups for disaster recovery.)

You must install the Windows bare metal agent to run bare metal backups. For details and requirements, see the Upgrade Guide for Recovery Series and Unitrends Backup.

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 for backup copy jobs as incrementals cannot be copied directly.

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 asset-level backups and host-level backups of VMware, Hyper-V, 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 KB 3560.

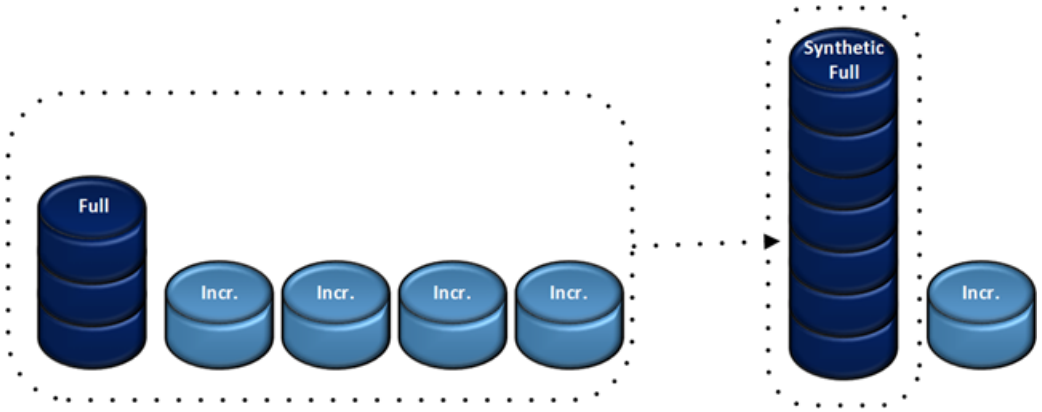

Backup groups

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 interdependencies 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.

The following diagrams illustrate backup groups:

- "Incremental forever backup groups" on page 42
- "Groups with full, differential, and incremental backups" on page 42
- "Selective backup in relation to a group" on page 43

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.

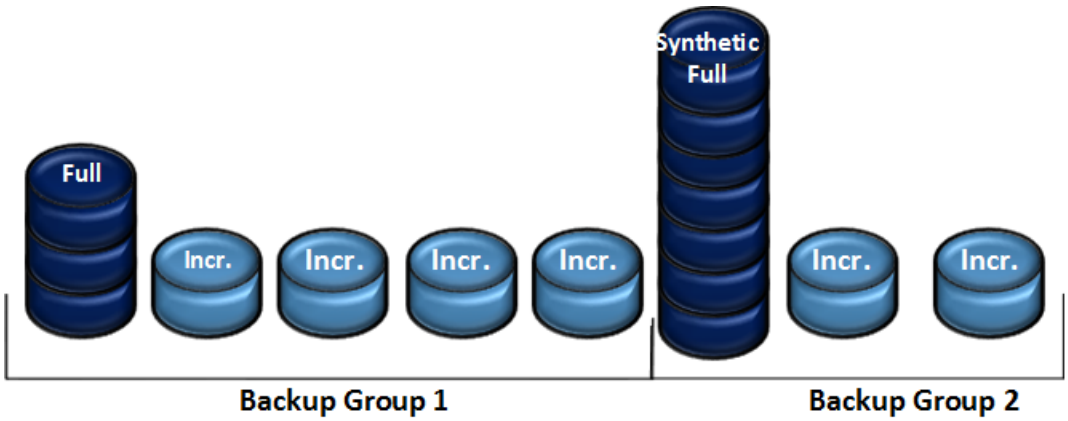

Groups with full, differential, and incremental backups

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

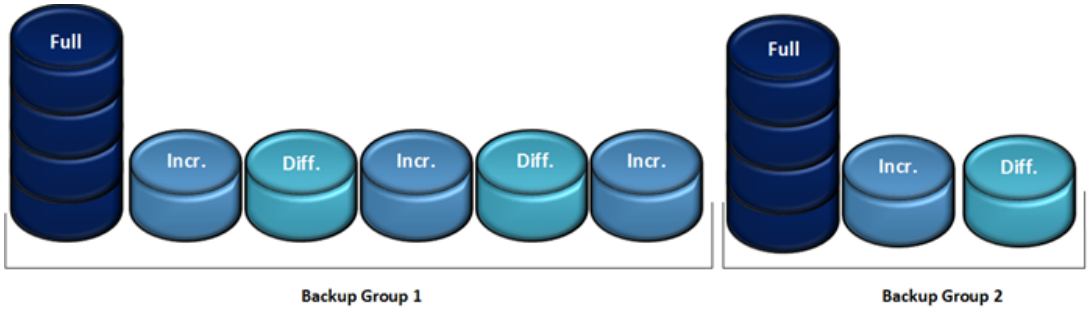

Selective backup in relation to a group

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 asset-level protection.

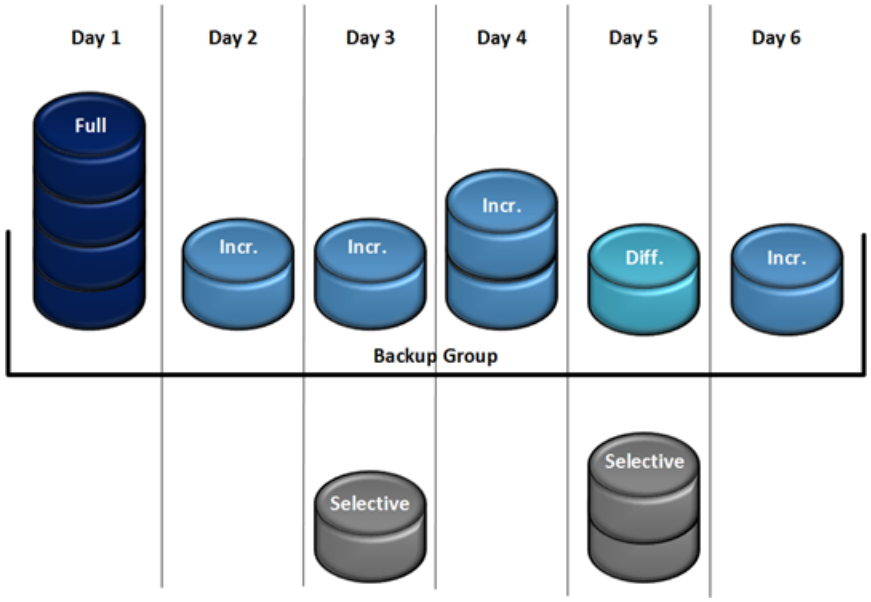

Backup strategies

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

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). 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. You can use retention policies to control how long backups remain on the appliance. Backups held by a policy are never purged. An asset's latest backups are also 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. You create retention policies that hold backups for a specified number of days and apply them to individual assets. You can create multiple policies and customize them to achieve individual RPOs and RTOs for each asset.

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

- Recovery Series 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 70 for details.

This storage capacity is used to store local backups and also for instant recovery. To use the instant recovery 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 69.

Backup copies

Backup copies are duplicates of your backups that are stored off-site. Unitrends recommends having a second copy of your backups on one of these targets 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, 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 that defines which backups to copy and other options. Backups are then copied according to the job 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 77 for procedures used to configure and add backup copy targets to your Unitrends appliance.
- "Backup Administration and Procedures" on page 173 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 instant recovery to spin up a virtual replica of the failed asset in minutes. This replica performs just like the original asset, so production downtime is reduced to just minutes. With Windows instant recovery, 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 281
- "Recovering Asset-level Backups" on page 305
- "Recovering Application Backups" on page 341
- "Recovering NAS Backups" on page 335
- "Recovering iSeries Backups" on page 365

Chapter 3: Configuration

This chapter provides instructions for performing administrative and configuration tasks for your appliances and protected assets.

See the following topics for details:

Appliance superuser account settings

By default, a superuser named *root* is created on the appliance. You cannot delete the root user. You can access and modify the user account settings by clicking the **Root** drop-down menu from the UI toolbar.

To view and edit account details

1. Click the **Root** drop-down menu and select **User Account** to open the Edit Account Details dialog.
2. (Optional) Edit the password.
3. Click **Save** to apply your changes or **Cancel** to exit the dialog without saving.

Appliance settings

From the **Configure > Appliances** page, you can manage your backup appliance and its storage, backup copy targets, and network. This section describes the appliance settings you can configure. Once you have reviewed this information, proceed to "Managing appliances" on page 55 for step-by-step configuration instructions.

Preparing to manage the appliance

Before you begin, see the following topics to determine which features to configure, and to gather required information:

Appliance network settings

During deployment, these network settings were configured for the appliance: IP address, subnet, gateway, primary DNS, hosts file, and port security:

- 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:
  - To connect the appliance to the Internet.
  - To add assets by 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 contains an entry for the appliance itself. Entries are automatically added when you add an asset to the appliance, or if you 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.
- The port security setting determines which ports are open and closed on the Unitrends appliance.

You can modify IP address, gateway, and DNS settings on the General tab of the **Configure > Appliances > Edit > Edit Appliance** page.

You can modify the hosts file and port security on the Network tab of the **Configure > Appliances** page.

Email reporting

You can configure the appliance to send system, job, and failure reports by email. To receive email reports, you must configure the appliance to use your SMTP server. You must also define email recipients. Supply the following on the **Email** tab of the **Configure > Appliances > Edit > Edit Appliance** page:

- The fully qualified SMTP server name or its IP address. (If a DNS record has not been configured, you must use the IP address of the SMTP server.)
- Username and password credentials if you have an externally-hosted SMTP server that requires authentication.

  Note: When using a non-local mail server or an internal SMTP relay configurations, 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.

- A valid test email address.
- Destination email addresses where reports will be sent.

Users

The User Interface (UI) manages and monitors Unitrends appliances. To access the UI requires a user account. 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 authentication. This section applies to creating users on the appliance. To use AD authentication, see "To set up Active Directory authentication" on page 56.

Each user account is assigned a role that defines the types of operations the user can perform on the appliance. Supported roles are:

Role	Description
Monitor	A user with this role is only able to view the status of operations, such as jobs, and to run reports. This user cannot create or start jobs or configure the system in any way.
Manage	A user with this role can view statuses and reports, start and view backup jobs, and perform other management tasks, such as adding or modifying assets and retention settings. However, this user cannot create or modify users other than modifying his or her own user account password.
Superuser	A user with this role, in addition to monitoring and managing systems, can add, edit, or delete users.
Administrator	This role is equivalent to the Superuser role. In the future this role will support the ability for a user to have different roles on different appliances.

Review these additional details before managing users:

- User accounts can only be used to access 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 system.
- To modify 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.
- See the "To edit an appliance" on page 55 procedure to add or modify users.
- Once you set up users, you can assign them to asset groups to control which assets they can access. For details, see "Grouping assets in custom folders" on page 131.

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.

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 appliances differ from those for virtual Unitrends Backup appliances:

- Recovery Series 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. After deploying the appliance, you must register and license it. Register the appliance as described in the applicable deployment guide, then apply the license you receive from Unitrends. See these deployment guides for details:
  - 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 in Microsoft Azure
  - Deployment Guide for Unitrends Free on VMware
  - Deployment Guide for Unitrends Free on Hyper-V
- Unitrends Backup Installable Software deploys without a license. After deploying, register the appliance as described in the Deployment Guide for Unitrends Backup Installable Software. Then apply the license you receive from Unitrends.
- See the "To edit an appliance" on page 55 procedure to add or modify the appliance license.

Configure the appliance as a backup copy target

You can configure the appliance as a backup copy target to store backups copied from another Unitrends appliance. 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. For details, see "Backup copy targets" on page 77.

Note: You can use Unitrends Backup on Hyper-V, Unitrends Backup on VMware, and Unitrends Backup on Citrix XenServer appliances either as a backup appliance or as a backup copy target. These appliances cannot perform both roles. You can use Recovery Series and Unitrends Backup installable software appliances as both a backup appliance and a backup copy target.

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

See the "To edit an appliance" on page 55 procedure to set up encryption.

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

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.

Access the Support Toolbox from the **Advanced** tab of the **Configure > Appliances > Edit > Edit Appliance** page. Scroll through the toolbox to find the information or task you are interested in. Hover over the options to see descriptions and helpful tips.

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 access these settings, click **General Configuration** on the **Advanced** tab of the **Configure > Appliances > Edit > Edit Appliance** page.

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. 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.

The appliance is deployed with these default operating system credentials:

- User *root*
- Password *unitrends1*

To change this password, click **OS Password** on the **Advanced** tab of the **Configure > Appliances > Edit > Edit Appliance** page.

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.
- For setup procedure, see "To configure iSCSI CHAP authentication" on page 62.

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.

For setup procedure, see "To set up SNMP trap notifications" on page 62.

Create a separate database partition on your Unitrends Backup appliance

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.
- You are using third-party deduplicated storage that is not designed to process database workloads.

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.

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:
  - *Determining your storage strategy* in the Deployment Guide for Unitrends Backup on VMware.
  - *Determining your storage strategy* in the Deployment Guide for Unitrends Backup on Hyper-V.
  - *Determining your storage strategy* in the Deployment Guide for Unitrends Backup on Citrix XenServer.
  - *Determining your storage strategy* in the Deployment Guide for Unitrends Free on VMware.
  - *Determining your storage strategy* in the Deployment Guide for Unitrends Free on Hyper-V.
  - *Disk and network configuration* in the Deployment Guide for Unitrends Backup Installable Software.

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, see "To create a separate database partition on your Unitrends Backup appliance" on page 64 to set up the partition and migrate the database.

Viewing 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.

Managing appliances

Use these procedures to manage existing appliances, add appliances you want to manage from this appliance's UI, and remove appliances you no longer want to manage:

To edit an appliance

Use this procedure to edit these appliance settings and to access the Support Toolbox:

- Appliance Name, IP address, and FQDN
- Email
- Users
- Date and time
- License

- Backup copy
- Encryption
- OS password
- CHAP authentication for iSCSI storage
- Advanced configuration settings, such as MaxConcurrentJobs and QuickSeek (for advanced users only)

For a description of each option, see "Appliance settings" on page 47. To edit DNS, port security, and the appliance hosts file, see "Networks" on page 65.

1. On the **Configure > Appliances** page, select the appliance and click **Edit**.
2. On the Edit Appliance dialog, modify information on the desired tab and click **Save**. Available tabs are:
   - General - Use to edit the appliance's name and IP Address. Make sure to remember the new IP address so you can continue to access the appliance's UI.
   - Email - Use to configure appliance email.
   - Users - Use to add, modify, or remove appliance users and to set up Active Directory (AD) authentication. For AD setup, see "To set up Active Directory authentication" below.
   - Date Time - Use to set the appliance date and time.
   - License - Use to view, modify, or upgrade the appliance license.
   - Backup Copy - Use to configure the appliance as a backup copy target. For details, see "Adding a Unitrends appliance backup copy target" on page 78.
   - Advanced - Use to set up or access these advanced features:
     - Encryption - Set up by providing a passphrase and saving the master key file. (For details, see "To configure encryption" on page 59 or "To change the encryption passphrase" on page 60.) Once you configure encryption, you can encrypt backups by asset (**Configure > Protected Assets > Edit Asset**).
     - Support Toolbox - Use to access additional information, such as log files, lists of running processes and services, and disk status, and to perform related administrative tasks. For details, see "To use the Support Toolbox" on page 61.
     - General Configuration - Edit advanced appliance configuration settings. For details, see "To configure advanced settings" on page 61.
     - OS Password - Use to change the appliance operating system password. (This is not the same as the UI password). For details, see "To change the appliance OS password" on page 61.
     - iSCSI CHAP - Use to set up CHAP authentication for use with iSCSI storage targets. For details, see "To configure iSCSI CHAP authentication" on page 62.
     - SNMP - Use send SNMP traps to your own Remote Monitoring and Management (RMM) software. For details, see "To set up SNMP trap notifications" on page 62.

To set up Active Directory authentication

You can now 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 group to which a user belongs determines which features that user can view and utilize. Users are granted one of the following privilege levels: *monitor*, *manage*, *administrator*, or *superuser*.

To authenticate using Active Directory

1 Create the following groups in your Active Directory domain:

Group	Description
Unitrends-Superuser	Members of this group are granted *superuser* privileges in the Unitrends UI.
Unitrends-Admin	Members of this group or domain administrators are granted *administrator* privileges in the Unitrends UI.
Unitrends-Manage	Members of this group are granted *manage* privileges 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.
Unitrends-Monitor	Members of this group are granted *monitor* privileges 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 groups to suit your environment. If you use your own names, be sure to enter these names when you configure AD authentication on the appliance. User group names in your AD domain must match the names you enter in step 8 on page 58.

2 Add users to the Unitrends domain groups as desired.

Users who are not domain 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 on page 58.

- Continue with step 4 on page 58 to add the AD server to the Unitrends appliance's hosts file.

4 Log in to the appliance UI.

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:

- On the **Configure > Appliances** page, select the appliance and click **Edit**.
- In the Edit Appliance dialog, click **Users**.
- Click **Add**, enter the following to add each Unitrends AD user group, then click **Save**:
  - Username - Name of the user group you have added to the AD domain. (See step 1 on page 57 for details).
  - Password - AD user group's password.
  - Confirm Password - Enter the password again.
  - Role - Select the role to apply to the AD user group.
- 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.

9 Click **Save**.

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 authenticate using Active Directory".

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 configure encryption

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

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

7 On the Edit Appliance dialog, click the **Advanced** tab.

8 Click **Save Master Key File**.

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

10 Log in to a Windows workstation as an administrator with full system access.

11 Launch File Explorer and enter the following path to access the master key file on the Unitrends appliance:

```
\\ApplianceIP\samba
```

12 Copy the master key file, called *crypt_image.iso*, to removable media and store it in a safe location.

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.

13 Once you have copied the key to removable media, delete *crypt_image.iso* from *\\ApplianceIP\samba* for increased security.

To change the encryption passphrase

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

2 On the Edit Appliance dialog, click 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 You receive a message indicating the Master Key File was saved to the appliance's samba share. Click **OK**.

7 Click **Save**.

8 Log in to a Windows workstation as an administrator will full system access.

9 Launch File Explorer and enter the following path to access the master key file on the Unitrends appliance:

```
\\ApplianceIP\samba
```

10 Copy the master key file, called crypt_image.iso, to removable media and store in a safe location.

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 Once you have copied the key to removable media, delete *crypt_image.iso* from *\\ApplianceIP\samba* for increased security.

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 In the Edit Appliance dialog, 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.

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 In the Edit Appliance dialog, click **Advanced** and select **General Configuration**.

4 Modify settings as desired and click **Save**

To change the appliance OS password

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 In the Edit Appliance dialog, click **Advanced** and select **OS Password**.

4 Enter the **Current OS Root Password**, the **New OS Root Password**, and confirm by entering it again in the **Confirm New OS Root Password** field.

The appliance is deployed with these default operating system credentials:

- User *root*
- Password *unitrends1*

5 Click **Save**.

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 In the Edit Appliance dialog, click **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.

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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)
- Enter the following command to configure the SNMP V3 username and password:

```
# /usr/bp/bin/cmc_snmpd script 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 In the Edit Appliance dialog, 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 In the Agent Configuration area, click **Download MIB**, and install it 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/*.

6 In the Trap Destinations area, click **Add**.

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

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

To add an appliance

1 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 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**.

To install appliance updates

1. Select the options icon in the upper-right.
2. Click **Check for Updates**.
3. Click **Apply Update**.

To restart an appliance

1. On the **Configure > Appliances** page, select the appliance and click **Edit**.
2. On the General tab, click **Shutdown/Restart**.

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 in "Create a separate database partition on your Unitrends Backup appliance" on page 53.

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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)
3. Issue this command to download the create database partition script:

   ```
   # wget ftp://ftp.unitrends.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 on page 65.

     ```
     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
```

To configure Unitrends Backup appliance deduplication settings

1 From the Global options at the top of the UI, select **Options > Deduplication Settings**.

2 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.

3 Click **Apply Settings**.

Networks

During deployment, these network settings were configured for the appliance: IP address, subnet, gateway, primary DNS, hosts file, and open ports.

- 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" on page 66 to modify IP address, subnet, gateway, or DNS settings.
- "To view or edit the hosts file" on page 66
- "To view or edit port security settings" on page 66

In addition to the standard appliance network settings, additional ports must be open if your appliance and its protected assets are separated by a firewall, to connect to the Internet, and to copy backups to a hot backup copy target. See "Additional ports" on page 67 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.
- 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 and click the **Network** tab.
2. Select the desired adapter and click **Edit**.
3. Enter the desired changes, and click **Save**.

To view or edit the hosts file

1. On the **Configure > Appliances** page, select the appliance and click the **Network** tab below.
2. On the Network tab, select the adapter (typically *eth0*) and click **Edit Hosts File**.
3. Do one of following:
   - To add an asset, click **Add**, enter all applicable infomation, and click **Save**.
   - To edit an asset, select it in the list, click **Edit**, modify infomation, 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**.

To view or edit port security settings

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 and click the **Network** tab below.
3. On the Network tab, click **Ports**.
4. 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.
5. Select a **Port Security** option and click **Save**.

Additional ports

Additional ports must be open if there is a firewall between your appliance and its protected assets, for connectivity to the Internet, and for connectivity to any hot backup copy target. See these tables for details:

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

Connectivity between the appliance and its protected assets

Task	Port, Protocol, and Rule	Hostname and IP Address	Notes
Protect assets that are separated from the appliance by a firewall.	1743: • TCP • Outbound from appliance • Inbound from protected asset 1745: • TCP • Outbound from protected asset • Inbound from appliance	Appliance hostname and IP Asset hostname and IP	If a firewall exists between the appliance and the assets (machines) you wish to protect, open these ports to enable communication and data transfer between the appliance and assets. You must also enter 1745 in the appliance master ini file. To do this: 1 In the appliance UI, select **Configure > Appliances > Edit > Advanced > General Configuration**. 2 In the Configuration Options Section, select **data** in the Name column. 3 In the Edit Settings dialog, enter **1745** in the Value field and click **Save**.

Connectivity between the appliance and the Internet

Task	Port, Protocol, and Rule	Hostname and IP Address	Notes
Product Updates	20 and 21: • FTP • Outbound 80: • HTTP • Outbound	ftp.unitrends.com 50.19.127.159	Used by the Unitrends appliance to perform software updates.

Task	Port, Protocol, and Rule	Hostname and IP Address	Notes
Remote Support	443: • HTTPS • Outbound	support-itivity.unitrends.com 74.202.224.68	Used for opening a remote tunnel to the Unitrends support team.
Proactive Monitoring	161 and 162: • TCP and UDP • Outbound	notifications.unitrends.com 104.130.228.89	Used for SNMP trap collection for all proactive monitoring.

Connectivity between the appliance and a hot backup copy target

Task	Port, Protocol, and Rule	Hostname and IP Address	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. • TCP and UDP • Outbound 443: • TCP • Outbound	Target appliance hostname and IP For Unitrends Cloud, the public-facing IP address provided by Unitrends.	Used for copying data to the Unitrends Cloud or your Unitrends target appliance.

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 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 70 for details.

- Unitrends Backup Installable Software is deployed to your CentOS or RHEL server. The amount of backup storage available is determined by the size of the */backups* partition. For instructions on adding more backup storage, see KB 5200. For details on initial configuration, see the Deployment Guide for Unitrends Backup Installable Software.

Instant recovery write space

When configuring backup storage, determine the percentage to use for backups and the percentage to reserve for instant recovery. Because your appliance is designed to retain as many local backups as possible, you should 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.

The storage reserved for instant recovery is used to store a disk image, known as a recovery object, based on the backup selected for recovery. Once the instant recovery completes, you must tear down the recovery object so the reserved space is made available to other instant recovery processes.

For more information, see "Virtual machine instant recovery" on page 296 and "Windows instant recovery" on page 316.

For both Recovery Series and Unitrends Backup appliances, you can view backup storage and modify the amount of storage allocated for backups and instant recovery.

To view backup storage

1. On the **Configure > Appliances** page, select the **Storage** tab.
2. 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 **Storage** tab.
2. 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:

- 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 71 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 71 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 71 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 74 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 KB 3350.

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 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" on page 72
- "To add a disk and expand the storage device" on page 72
- "To add a disk as a separate storage area" on page 73

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 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 the desired disk from the list of available attached disks and click **Add**.
   - To remove an attached disk, select it from the list of available attached disks and click **Remove**.
   - Clicking **Reset** reverts all disk settings to the original settings.
   - To refresh the list of available disks, click **Refresh**.
6. Adjust the storage allocation for Backups and IR as desired.
7. Click **Save**.
   On the **Configure > Appliances** page, the status initially displays as *Pending*. When the disk is finished attaching, the status converts to *Active* and the storage can be used.
8. If desired, 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 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 the desired 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**.
- Clicking **Reset** reverts all disk settings to the original settings.

7 Click **Save**.
On the **Configure > Appliances** page, the status initially displays as *Pending*. Once the disk attaches, the status converts to *Active* and the storage is available for use.

8 If desired, 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.

Before adding storage, be sure to review the recommendations in "About adding backup storage to a Unitrends Backup appliance" on page 70.

Procedures for adding external storage:

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 62.

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 173 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 173 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 173 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 173 for details.

NFS 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 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 105 for details on working with the target you added.
- See to "Creating backup copy jobs" on page 186 to start copying backups to the target.

Note: The types of backup copy targets supported vary by Unitrends appliance model. For example, attached disk is not supported on Recovery Series appliances. When adding backup copy targets to an appliance, only supported types display in the **Configure > Appliances > Backup Copy Targets > Add Target** list.

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

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 "Return to the source backup appliance and fine-tune settings by adjusting connection options" on page 81.
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 188.

Adding a Unitrends appliance backup copy target

The procedures for adding a Unitrends appliance backup copy target vary by whether you are setting up a single target appliance or setting up a pair of appliances that will each store copies of the others' backups (known as a *cross-copy* configuration). See one of the following for setup

procedures:

- "Adding one Unitrends appliance backup copy target" on page 79 to copy backups one way from a source appliance to a target appliance
- "Setting up cross-copy between two Unitrends appliances" on page 84 to copy backups in both directions, where each appliance is both a source appliance running local backups and a target appliance receiving backup copies from the other appliance

Adding one Unitrends appliance backup copy target

To add an appliance target, review the requirements in "Preparing to add one Unitrends appliance backup copy target", then set up the target as described in "To add one Unitrends backup copy target appliance" on page 80.

Note: Use these procedures to an appliance target that will only receive backup copies. To set up a pair of appliances that will each store copies of the others backups, see "Setting up cross-copy between two Unitrends appliances" on page 84.

Preparing to add one 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 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
    - Citrix XenServer
    - Microsoft Azure
  - Recovery Series appliances and Unitrends Backup Installable Software deployments 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 at least 200GB of available backup storage space.
- The configuration process creates a secure VPN tunnel. Be sure to use a network that is not in use in your environment. For example, if using netmask 255.255.255.0, the VPN tunnel must be on a dedicated subnet (default is 172.17.3.0).
- 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 64) before configuring the backup copy target.

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 one 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.

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 desired 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.
   - Secure Netmask IP - The netmask associated with the secure network above.
   - Secure Port - Port number to use for OpenVPN.
4. Click **Save**.

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**.
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.

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.
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:

   Note: By default, the backup copy target appliance does not manage backup copy source appliances. When the source appliance is not managed, backup copy jobs do not display on the **Jobs > Active Jobs** tab or on the **Dashboard > Active Jobs** tile. To view backup copy jobs from the target appliance, you must add management credentials for the source appliance.

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

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 options for optimal performance in your environment:

Option	Description	Procedure
Backup copy queue scheme	The 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 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 - This scheme is not supported for new targets. If you have updated your appliance from the legacy UI and were using the manual scheme, this still works, but you must switch back to the legacy UI to manually copy backups.	See "To configure connection options" on page 83.
Backup copy concurrency	The Max Concurrent Backup Copies setting 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 replications.	See "To configure connection options" on page 83.
Suspending backup copies	Use 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 83.

Option	Description	Procedure
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 83.
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 the backup copy job can use. The amount of bandwidth available for the job 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 83.

To set bandwidth throttling

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 83 above.
7. Click **Save**.
8. Repeat as necessary to create additional throttling schedules.

To configure connection options

1. Select **Configure > Appliances**.
2. Select the source backup appliance and click **Edit**.
3. On the **Backup Copy** tab, modify settings.

4 Click **Save**.

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 187.

Setting up cross-copy between two Unitrends appliances

In cross-copy configurations, each appliance acts as both a backup copy source and a backup copy target. This enables you to copy each appliance's local backups to the other appliance for added protection.

To set up cross-copy between two appliances, review the "Cross-copy requirements", then configure the appliances as described in "To set up cross-copy" on page 84.

Note: To set up an appliance target that will receive backup copies only, see "Adding one Unitrends appliance backup copy target" on page 79.

Cross-copy requirements

Both appliances must meet these requirements to use the cross-copy feature:

- Must be running release 9.0.0-7 or higher. It is highly recommended that both appliances be running the same version.
- Must be a Unitrends Backup Installable Software appliance or a Recovery Series model listed in the Recovery Series Appliance Family Datasheet.

  Note: Unitrends Backup appliances deployed to the following environments cannot be used for cross-copy: VMware, Hyper-V, Citrix XenServer, or Microsoft Azure.
- The appliance must have at least 200GB of available backup storage space.
- The configuration process creates a secure VPN tunnel. Be sure to use a network that is not in use in your environment. For example, if using netmask 255.255.255.0, the VPN tunnel must be on a dedicated subnet (default is 172.17.3.0).
- If one appliance is already managing the other (you are accessing both from one UI), you must log in to the managing appliance and remove the managed appliance (as described in "To remove a managed appliance" on page 64). Cross-copy cannot be configured if either appliance is managing the other.

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 set up cross-copy

To set up cross-copy, follow the steps in this procedure carefully. You must complete the steps in order and be sure to run each from the correct appliance. This procedure uses the terms *appliance1* and *appliance2* to identify the appliance where each step must be run.

Log in to both appliance UIs

To set up cross-copy, you will switch between the appliance UIs. Log in to both appliance UIs. Make a note of which appliance you will use as *appliance1* and which appliance you will use as *appliance2*.

Do these steps on *appliance1*

1. On the **Configure > Appliances** page, select *appliance1* and click **Edit**.
2. 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.
   - Secure Netmask IP - The netmask associated with the secure network above.
   - Secure Port - Port number to use for OpenVPN.
3. Click **Save**.

Do these steps on *appliance2*

1. On the **Configure > Appliances** page, select *appliance2*.
2. Click the **Backup Copy Targets** tab.
3. Select **Add Target > Unitrends > Unitrends appliance**.
4. Enter the following information about *appliance1*, then click **Save**:
   - Host Name of *appliance1*
   - (Required) User Name of *appliance1*
   - (Required) Password of *appliance1*

   Note: Do NOT enter a value in the Host IP address field. Instead, you must supply root or superuser credentials for *appliance1* by entering a User Name and Password. This enables *appliance2* to manage appliance1, which is required for cross-copy.
5. You receive a *Connection could not be verified* warning (unless you have previously installed your own custom security certificates on the *appliance1*). As long as you trust your network configuration, you can safely click **Continue**.
6. After clicking **Continue**, it takes several minutes to configure the OpenVPN tunnel. Before continuing, wait for the confirmation message that *appliance1* has been added as a backup copy target.

Return to *appliance1* and do these steps

1. On the **Configure > Appliances** page, select *appliance2*. *appliance2* has an **!** icon and a Pending status.
2. Click the **!** icon and enter the following:
   - Select **Allow backup copies to be sent from *appliance2*** and click **Save**.
   - Select the storage device where backup copies will be stored on *appliance1*.
3. Click **Save**.

4 Final configuration steps take place in the background and can take up to 20 minutes to complete. Configuration is complete when the status of *appliance2* changes from Pending to Not Available.

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.
- Once configuration is complete, you can go to *appliance2* and see that the status of *appliance1* is now Online.

5 On the **Configure > Appliances** page, select *appliance1*.

6 Click the **Backup Copy Targets** tab.

7 Select **Add Target > Unitrends > Unitrends appliance**.

8 Enter the following information about *appliance2*, then click **Save**:

- Host Name of *appliance2*
- Host IP address of *appliance2*
- (Optional) User Name of *appliance2*
- (Optional) Password of *appliance2*

9 You receive a confirmation message that *appliance2* has been added as a backup copy target. Click **OK**.

Return to *appliance2* and fine-tune connection options

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

Option	Description	Procedure
Backup copy queue scheme	The 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 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 - This scheme is not supported for new targets. If you have updated your appliance from the legacy UI and were using the manual scheme, this still works, but you must switch back to the legacy UI to manually copy backups.	See "To configure connection options" on page 83.
Backup copy concurrency	The Max Concurrent Backup Copies setting 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 replications.	See "To configure connection options" on page 83.
Suspending backup copies	Use 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 83.

Option	Description	Procedure
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 83.
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 the backup copy job can use. The amount of bandwidth available for the job 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 83.

To set bandwidth throttling

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 83 above.
7. Click **Save**.
8. Repeat as necessary to create additional throttling schedules.

To configure connection options

1. Select **Configure > Appliances**.
2. Select the source backup appliance and click **Edit**.
3. On the **Backup Copy** tab, modify settings.

4 Click **Save**.

Create jobs to start copying your backups.

Create a backup copy job each appliance to start copying local backups. For details, see "To create a backup copy job for a Unitrends appliance target" on page 187.

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 191.

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.
- To remove drives and insert new media, be sure to:
  - Power down the dock or recovery archive unit.
  - Swap the drive(s).
  - Power on the dock or recovery archive unit.

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 92.

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 92.

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 select Unitrends Recovery Series appliances and Unitrends Backup installable software deployments. Tape is not supported with Unitrends Backup on Hyper-V, Unitrends Backup on VMware, Unitrends Backup on Citrix XenServer, and Unitrends Backup on Azure appliances. For supported Recovery Series appliances, see the Recovery Series Appliance Family Data Sheet.

The following requirements must be met before setting up tape targets on Recovery Series appliances and Unitrends Backup installable software deployments:

System	Requirement
Recovery Series or Unitrends Backup installable software appliance	The appliance must be licensed with the advanced archiving (ADX) feature. Check for *ADX* in the Feature String under **Configure > Appliances > Edit > License**.

System	Requirement
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 92. • 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. • 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 procedures configure both the changer and the tape drive.)

- "To configure a multi-drive autochanger" below
- "To configure a single-drive autochanger" on page 94
- "To configure a tape drive" on page 96 (use this procedure if your tape device does not have an autochanger)

To configure a multi-drive 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**.

   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.
6. Enable the tape drive and autochanger:
   - On the Backup Copy Targets tab, select the tape drive and click **Enable**. Its status changes to *Online*.
   - On the Backup Copy Targets tab, 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 in the list on the Backup Copy Targets tab. Tape drives display as Type *tape*.
- Click **Edit** and 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 in the list on the Backup Copy Targets tab. The autochanger displays as Type *changer*.
- 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 192.
- If you did NOT check the *Use Unlabeled Tapes* box above, proceed to the next step to prepare the tapes.

11 Select the autochanger on the Backup Copy Targets tab 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 192.

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 single-drive 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. Click **Yes**.

6 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.

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

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

8 Enable the tape drive and autochanger:

- On the Backup Copy Targets tab, select the tape drive and click **Enable**. Its status changes to *Online*.
- On the Backup Copy Targets tab, select the autochanger and click **Enable**. Its status changes to *Online*.

9 (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 in the list on the Backup Copy Targets tab. Tape drives display as Type *tape*.
- Click **Edit** and check the Use Unlabeled Tapes box.
- Click **Save**.

10 (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 in the list on the Backup Copy Targets tab. The autochanger displays as Type *changer*.
- 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**.

11 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 192.
- If you did NOT check the *Use Unlabeled Tapes* box above, proceed to the next step to prepare the tapes.

12 Select the autochanger on the Backup Copy Targets tab and click **Prepare**.

13 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

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

WARNING! Clicking **Prepare** permanently deletes any existing data and formats the tapes.

15 Create a job to start copying your backups, as described in "To create a backup copy job for a tape target" on page 192.

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 on the Backup Copy Targets tab 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 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 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 192.

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, Google Cloud Storage Nearline, Amazon S3, or Rackspace. For details about creating an account and purchasing storage, see KB 3649.

- For bucket 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.
- For Amazon S3, Google Cloud Storage Standard, and Google Cloud Storage Nearline, you can use existing buckets that follow the supported naming conventions identified above. However, we recommend that you create unique folders for your Unitrends data.
- Amazon S3's Reduced Redundancy Storage (RRS) option is not supported.

You must enter the following account information when adding a cloud backup copy target to the appliance:

- Credentials for the storage bucket 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 provider.
- Name of the bucket that you are adding to the appliance.

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.

You can monitor the amount of data in your cloud storage from the Storage tile on the **Dashboard** or 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.

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. 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, or Rackspace 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 provider in the **Cloud Storage** drop-down.
6. Enter the bucket's name in the **Storage Path** field.

   To create a sub-folder within the bucket, 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. Enter the required credentials. These vary by storage provider:

   Note: Be sure to enter the credentials you use to access the particular bucket 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 Standard or Google Cloud Storage Nearline	Access Key Secret
Amazon	Access Key ID Secret Access Key
Rackspace	Username API Key

8 (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 97.

9 Click **Save**.

10 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.

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 In the Edit Appliance dialog, click **Advanced** and select **General Configuration**.

3 Change the CloudHook *UnmountAfter* setting to **0** and click **Save**

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, 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 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 99.
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 desired 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 189.

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 101

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 "CIFS configuration details" on page 101 below.

7 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 NAS target" on page 190.
- 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 NAS target" on page 190.

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 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 "NFS configuration details" on page 102 below.

7 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 NAS target" on page 190.
- 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 NAS target" on page 190.

NFS 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 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.

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 104

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 62.

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 190.
    - 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 190.

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 190.
    - 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 190.

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

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** to change information in the fields, and click **Save**.

Note: For Google, Amazon, 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 97.

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, 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 to **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 options for optimal performance in your environment:

Option	Description	Procedure
Backup copy queue scheme	The 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 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 - This scheme is not supported for new targets. If you have updated your appliance from the legacy UI and were using the manual scheme, this still works, but you must switch back to the legacy UI to manually copy backups.	See "To tune connection options for a Unitrends Cloud or Unitrends appliance target" on page 106.

Option	Description	Procedure
Backup copy concurrency	The Max Concurrent Backup Copies setting 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 replications.	See "To tune connection options for a Unitrends Cloud or Unitrends appliance target" on page 106.
Suspending backup copies	Use 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 tune connection options for a Unitrends Cloud or Unitrends appliance target" on page 106.
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 tune connection options for a Unitrends Cloud or Unitrends appliance target" on page 106.
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 the backup copy job can use. The amount of bandwidth available for the job 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 107.

To set bandwidth throttling

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 107 above.

7 Click **Save**.

8 Repeat as necessary to create additional throttling schedules.

To configure connection options

1 Select **Configure > Appliances**.

2 Select the source backup appliance and click **Edit**.

3 On the **Backup Copy** tab, modify settings.

4 Click **Save**.

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 procedures 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 drives and insert new media, be sure to:

- Power down the dock or recovery archive unit.
- Swap the drive(s).
- Power on the dock or recovery archive unit.

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 Click **Erase**.

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 appliances generates one for you.

8 Click **Confirm**. 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.

Notes:

If you need to remove tapes and insert new media, be sure to:

- Power down the tape device.
- Swap the tape(s).
- Power on the tape device.

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 import reference information for cold copies that were run on a different appliance:

1 Connect the media containing the cold copies to the target appliance that will import the reference information.

Notes:

If you need to remove drives or tapes and insert new media, be sure to:

- Power down the dock, recovery archive unit, or tape device.
- Swap the drive(s) or tape(s).
- Power on the dock, recovery archive unit, or tape device.

1. Log in to the target appliance.
2. On the **Configure > Appliances** page, select the appliance.
3. Click the **Backup Copy Targets** tab below.
4. Click **Scan for Media**.
5. Select the offline target.
6. Click **Enable**.

   The appliance brings the target online and imports reference information for any backup copies that were not found on the appliance.

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 109 for tape autochanger targets.
- "To initialize and erase cold backup copy media" on page 108 for all other cold backup copy targets, including single tape drive devices.

To swap out drives in an eSATA or USB cold backup copy target

If you need to remove drives and insert new media, be sure to:

1. Power down the dock or recovery archive unit.
2. Swap the drive(s).
3. Power on the dock or recovery archive unit.

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**.

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 asset-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 261.

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 136
Linux	"Installing and updating the Linux agent" on page 144
CentOS	"Installing and updating the Linux agent" on page 144
Debian	"Installing and updating the Linux agent" on page 144
Fedora	"Installing and updating the Linux agent" on page 144
Red Hat	"Installing and updating the Linux agent" on page 144
SUSE	"Installing and updating the Linux agent" on page 144
Ubuntu	"Installing and updating the Linux agent" on page 144
Solaris	"Installing and updating the Solaris agent" on page 158
Novell Netware	"Installing and updating the Novell Netware agent" on page 153
Novell OES Linux	"Installing and updating the Novell OES Linux agent" on page 155
Mac	"Installing and updating the Mac agent" on page 152
AIX	"Installing and updating the AIX agent" on page 150
SCO	"Installing and updating the SCO OpenServer agent" on page 157
UnixWare	"Installing and updating the UnixWare agent" on page 159
HP-UX	"Installing and updating the HP-UX agent" on page 151

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 118.

Feature	Supported for protected asset type? Physical	Virtual	Application	Configured where?
"Asset credentials" on page 114	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 115	Yes	Yes	Yes	Apply using Edit Asset.
"Encrypt backups" on page 115	Yes	Yes	Yes	Configure on Edit Appliance, then apply using: • Add Asset • Edit Asset

Feature	Supported for protected asset type? Physical	Virtual	Application	Configured where?
"Quiesce settings for host-level backups" on page 115	No	Yes, VMware and XenServer only	No	Configure globally, by host, or by VM: **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.

Once you have created a credential (as described in "Managing NAS assets" on page 121) you can apply the credential when adding or editing the following asset types:

- Virtual Hosts - Credentials are required for each vCenter, ESXi, Hyper-V, 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).
- 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.
- 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 - When you add a virtual host, any hosted VMs are discovered and display in the asset inventory tree under their virtual host. Credentials are required to enable application-aware protection of VMware Windows VMs and are optional for other VMs.
- Hosted applications - When you add an asset, any hosted applications are discovered and display in the asset 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 237.

Retention settings

The Unitrends appliance ingests new backups and retains them until there is no more backup storage space available. When backup storage is full, the oldest backups are 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 settings assure that the necessary recovery points are available on your appliance. For long term retention, copy backups to an off-site target as described in "Backup copies" on page 44.

Use retention settings to control how long an asset's backups are retained and the order in which backups are purged. When you add an asset, no retention is set, and the asset's backups are retained as long as possible until the system runs out of space (at which point the oldest backups are deleted). By setting retention policies, you can choose to retain backups of certain assets longer than others. To apply retention settings to assets, see "Managing retention settings" on page 130.

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 51. To encrypt an asset's backups, use the applicable Edit Asset procedure in "Managing protected assets" on page 118.

Quiesce settings for host-level backups

For host-level backups of VMware 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, 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* 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 127.	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 127.	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 129.

Quiesce setting	Description	Apply globally to all VMs	Apply to one host's VMs	Apply to selected VMs
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 VM disk snapshot is taken.	Apply to newly discovered VMs by selecting **Application 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 127.	To apply to hosted VMs, select **Overwrite this hypervisor's VMs to Application Consistent** in the Edit Virtual Host dialog. For details, see "To apply a quiesce setting to one host's VMs" on page 127.	To apply to one or more selected VMs, select **Application Consistent** in the Edit Assets dialog. For details, see "To edit a virtual machine asset" on page 129.
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 118 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 129.

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 249 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 111 for details on these features before running these 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 physical asset 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 **View: List**. Each row in the list describes a single asset.
  - To view assets in a table, click **View: Table**. Assets display in tiles on the left. Click an asset to view its details.
- Display options: select to filter the display by asset type or appliance.

Managing physical assets

Use these procedures to manage physical machine assets and VMs protected using a Unitrends agent.

Notes:

- For VMs protected at the host-level, see the "Managing virtual machine assets" on page 128 procedures instead.
- For iSeries assets, see "iSeries Backups Overview and Procedures" on page 261 instead.

To add an asset

The Unitrends agent must be installed for asset-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 111.

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 Enter or select optional settings as desired.

6 Click **Save**.

To create aliases for agent-based backups

For agent-based backups, you can create aliases for a single asset and run separate, radically different backup schedules for each. 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 backup job for agent-based assets" on page 174.

To create an alias for a protected asset

Note: Add the asset to the appliance, as described in "To add an asset" on page 119, before running this procedure.

1. On the **Configure > Appliances** page, select the appliance and click the **Network** tab below.
2. On the Network tab, select the adapter (typically *eth0*) and click **Edit Hosts File**.
3. Select the asset in the list and click **Edit**.
4. Enter a second hostname in the Aliases field, and click **Save**.
5. Add the alias as a protected asset by using the "To add an asset" on page 119 procedure:
   - For Hostname, enter the Alias name you entered in the hosts file.
   - For IP Address, enter the IP of the asset whose entry you modified in the hosts file. (The original protected asset and any related aliased assets must have the same IP address and different hostnames.)

To edit an asset

1. Select **Configure > Protected Assets**.
2. Select the desired asset.
3. Click **Edit**.
4. Modify settings as desired.
5. Click **Save**.

Removing an 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 Instant Recovery 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 asset

1. Select **Configure > Protected Assets**.
2. Select the asset you want to remove.
3. Click **Remove > Confirm**.

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 121
- "To add a NAS NDMP asset" on page 122
- "To edit a NAS asset" on page 122
- "Removing a NAS asset" on page 122

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.

10 Click **Save**.

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, enter a Credential Name, Username, and Password.

9 Click **Save**.

To edit a NAS asset

1 Select **Configure > Protected Assets**.

2 Select the NAS and click **Edit**.

3 Modify settings as desired, and click **Save**.

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 asset you want to remove.

3 Click **Remove > Confirm**.

Managing application assets

Use these procedures to manage application assets:

- "To add an application" on page 123
- "To add a UCS manager asset" on page 124
- "To edit a UCS manager asset" on page 124
- "To remove a UCS manager asset" on page 125
- "To edit an application" on page 125
- "To remove an application" on page 125

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 173. 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 237
- "SQL backup requirements and considerations" on page 240
- "SharePoint backup requirements and considerations" on page 250
- "Oracle backup requirements and considerations" on page 253
- "Cisco UCS service profile backup requirements and considerations" on page 258

Application	Add host procedure
Exchange	Add the Exchange server using the "To add an asset" on page 119 procedure. All hosted databases or storage groups display under the asset you have added.
SQL	Add the SQL server using the "To add an asset" on page 119 procedure. All hosted databases display under the asset you have added.
SharePoint	Configure and add the SharePoint server using the instructions in "SharePoint configuration prerequisites" on page 252. The SharePoint application displays under the asset you have added.
Oracle	Add the Oracle server using the "To add an asset" on page 119 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 asset" on page 119 for physical servers or "Adding a virtual host" on page 126 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

1. Select **Configure > Protected Assets**.
2. Click to expand the application's host.
3. Select the desired application and click **Edit**.
4. Modify settings as desired, and click **Save**.

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. This can be a vCenter, an ESXi server, a Hyper-V server, or a Citrix XenServer. Supported virtual hosts vary by Unitrends appliance, as described in the following table. (Additional version requirements apply. See the chapter "Host-level Backups Overview" on page 203 for details.)

Unitrends appliance	Supported virtual hosts
Recovery Series	• VMware vCenter • VMware ESXi server • Hyper-V server
Unitrends Backup Installable Software deployment	• VMware vCenter • VMware ESXi server • Hyper-V server
Unitrends Backup on VMware	• VMware vCenter • VMware ESXi server • Hyper-V server
Unitrends Backup on Hyper-V	• VMware vCenter • VMware ESXi server • Hyper-V server

Unitrends appliance	Supported virtual hosts
Unitrends Backup on Citrix XenServer	• VMware vCenter • VMware ESXi server • Hyper-V server • Citrix XenServer

Use these procedures to manage virtual hosts. For additional requirements and considerations, see "Host-level Backups Overview" on page 203.

Adding a virtual host

Before you start, see "Preparing to add a VMware host" on page 126, "Preparing to add a Hyper-V host" on page 126 or "Preparing to add a XenServer host" on page 126.

Preparing to add a VMware host

For VMware virtual machines, you can add standalone ESX hosts and vCenter servers. Note that you must add the vCenter server to protect templates or clustered VMs.

- Working with vCenter servers - When you add a vCenter server, the appliance automatically detects all ESX hosts managed by the vCenter and all VMs and templates residing on these hosts. When a vCenter server is added, Unitrends can work with VMware's vMotion technology to contact the appropriate server when backing up clustered VMs.

  In addition to adding the vCenter, we recommend that you also individually add each ESX host, so if your vCenter server goes down, Unitrends can contact the host directly to perform jobs.

- About adding ESX hosts - If you are not adding a vCenter, you must add the ESX host to your Unitrends appliance to protect its virtual machines. If you want to use the instant recovery feature or protect templates or clustered VMs, you must add a vCenter server as well.

  Once a virtual host is added, all VMs on that host are automatically discovered and can be selected for protection.

Preparing to add a Hyper-V host

Before adding the host, you must install the Unitrends Windows agent on the host. See "Installing the Windows agent" on page 136 for details.

Preparing to add a XenServer host

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.
- A stand-alone XenServer host where the Unitrends Backup VM resides.

To add a virtual host asset

1 Select **Configure > Protected Assets**.

2 Click **Add > Virtual Host**.

3 Complete all fields on the Add Virtual Host page.

> **Note:** 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 125.

4 Click **Save**.

To edit a virtual host asset

1 Select **Configure > Protected Assets**.

2 Select the desired virtual host asset.

3 Click **Edit**.

4 Modify settings as desired, and click **Save**.

To manage global quiesce settings

The global quiesce setting applies to all VMware 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.

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 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.

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 127. 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 ESX host, a Hyper-V host, or a XenServer host 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 ESX 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 ESX 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 > Confirm**.

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 127. All hosted VMs display under the host you have added.

To edit a virtual machine asset

1. Select **Configure > Protected Assets**.
2. Click to expand the VM's virtual host to display its VMs.
3. Select the desired VM and click **Edit**.
4. Modify settings as desired, and click **Save**.

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.

Preparing for application-aware protection

Before you start, create administrative credentials for the Windows VM, as described in "To add a credential" on page 130.

To apply the application aware quiesce setting

1. Select **Configure > Protected Assets**.
2. Click to expand the VM's virtual host.
3. Select the desired 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

To encrypt an asset's backups

1. Set up encryption on the appliance, as described in "To edit an appliance" on page 55.
2. Select **Encrypt Backups** using the applicable Edit Asset procedure. (See "To edit an asset" on page 120 or "To edit a virtual machine asset" on page 129.)

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 114. 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 130
- "To view all credentials" on page 130
- "To edit a credential" on page 130
- "To remove a credential" on page 130
- "To apply a credential to an asset" on page 130

To add a credential

Before you can apply credentials to an asset, you must create the credential using this procedure.

1. Click **Configure > Protected Assets > Manage Credentials > Add**.
2. Enter credential information and click **Save**.

This credential can now be applied to your assets.

To view all credentials

Click **Configure > Protected Assets > Manage Credentials**.

All credentials defined for this appliance display.

To edit a credential

1. Click **Configure > Protected Assets > Manage Credentials**.
2. Select the desired credential and click **Edit**.
3. Modify information as desired and click **Save**.

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 desired credential.
3. Click **Remove > Confirm**.

To apply a credential to an asset

1. Select **Configure > Protected Assets**.
2. Click to select the desired asset.

   Note: To apply the credential to a VM or application, expand the host asset to view and select the VM or application instance.

3. Click **Edit**.
4. Select the desired credential and click **Save**.

Managing retention settings

Retention settings are used to control how long backups are retained. Retention settings are described below. See "To apply retention settings to an asset" on page 131 to apply settings to assets.

Retention setting	Description
Minimum Retention	Minimum retention settings.

Retention setting	Description
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 an asset

1. Select **Configure > Protected Assets**.
2. Select the desired asset.
3. Click **Edit**.
4. Click **Manage Retention**.
5. Define settings as desired click **Save**.

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.

Notes:

- Folder groups do not display in the Backup Catalog or on reports.

- After upgrading to release 9.0.0-13 or higher, any navigation groups that were created in the legacy UI display in the current UI in Show Groups view. Group names, colors, and user settings created in the legacy UI apply to these groups in the current UI.

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 machines
- VMware virtual machines

See the following topics for details:

- "Unitrends users and asset groups" on page 132
- "Working with asset groups" on page 132
- "To create a top-level asset group" on page 133
- "To create an inner group" on page 134
- "To view or hide asset groups" on page 135
- "To edit an asset group" on page 134
- "To delete an asset group" on page 135

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 the following:
  - "Users" on page 49 to create Unitrends users on the appliance.
  - "To set up Active Directory authentication" on page 56 to create Unitrends users that authenticate using AD.

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. On the **Protect** page, click the **Show Groups** icon located above the inventory tree.
3. Click the **Manage Groups** pencil icon.
4. Click **Add**.
5. Enter a unique Group Name.
6. (Optional) Click the color drop-down to select a display color for this group's folder.
7. (Optional) On the **Users** tab, click to select users to add to the group.

   Notes: You cannot create users here. You can add existing Unitrends users to the group. To create users, see one of the following:

   - "Users" on page 49 to create Unitrends users on the appliance.
   - "To set up Active Directory authentication" on page 56 to create Unitrends users that authenticate using AD.

   - 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.
8. On the **Not in Group** tab, click to select assets to add to the group.
9. Click **Save**.

   The new group displays in the Manage Groups dialog.
10. Click **Close**.

    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. 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 that will contain the new inner group.
5. Click **Add**.
6. (Optional) Modify the Group Name. This name must be unique.
7. (Optional) Click the color drop-down to select a display color for this group's folder.
8. (Optional) On the **Users** tab, click to select users to add to the group.

   Notes: You cannot create users here. You can add existing Unitrends users to the group. To create users, see one of the following:

   - "Users" on page 49 to create Unitrends users on the appliance.
   - "To set up Active Directory authentication" on page 56 to create Unitrends users that authenticate using AD.

   - 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.

   Assets are moved from the parent group to this inner group.
10. Click **Save**.

    The new group displays in the Manage Groups dialog.
11. 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. 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 edit.
5. Click **Edit**.

6 Modify options as desired:

- Modify the Group Name. This 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.

7 Click **Save**.

8 Click **Close**.

To view or hide asset groups

To view asset groups, go to the Protect page and click the **Show Groups** icon located above the inventory tree.

To hide asset groups, go to the Protect page and click the **Hide Groups** icon located above the 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 asset-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 261.

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 136
Linux	"Installing and updating the Linux agent" on page 144
CentOS	"Installing and updating the Linux agent" on page 144
Debian	"Installing and updating the Linux agent" on page 144
Fedora	"Installing and updating the Linux agent" on page 144
Red Hat	"Installing and updating the Linux agent" on page 144
SUSE	"Installing and updating the Linux agent" on page 144
Ubuntu	"Installing and updating the Linux agent" on page 144
Solaris	"Installing and updating the Solaris agent" on page 158
Novell Netware	"Installing and updating the Novell Netware agent" on page 153
Novell OES Linux	"Installing and updating the Novell OES Linux agent" on page 155
Mac	"Installing and updating the Mac agent" on page 152
AIX	"Installing and updating the AIX agent" on page 150
SCO	"Installing and updating the SCO OpenServer agent" on page 157
UnixWare	"Installing and updating the UnixWare agent" on page 159
HP-UX	"Installing and updating the HP-UX agent" on page 151

Installing the Windows agent

To protect a Windows asset, you must install a lightweight agent on the Windows machine. Depending on the asset's operating system and configuration, this Windows core 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.

In addition to the regular core agent, you may need 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 integrated bare metal protection feature, which does not require the bare metal agent.

Windows agent requirements

The following requirements must be met before installing the Windows core agent:

- Administrative privileges for the user installing the agent.
- Approximately 1100 MB of free space on the Windows system drive, usually volume *C:*.

  (Applicable even if installing on a volume other than the Windows system drive.)
- Single Instance Storage (SIS) on Windows Storage Server 2008 is not supported and must be disabled for the agent to properly perform backups.
- 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.

In addition to the general "Windows agent requirements" on page 137, the following prerequisites must be met to push-install the Windows agent:

Item	Description
Windows versions supported	Windows XP, 2003 and later versions listed in the Compatibility and Interoperability Matrix are supported (32- and 64-bit).
Credentials	Trust credentials must be defined for the Windows server asset on the backup appliance. See "Managing asset credentials" on page 129 for details.
Windows environment	The Windows machine must be configured as described in "Windows configuration requirements" on page 137.

Windows configuration requirements

The following Windows configuration settings are required for the agent push feature:

- *Workstation* and *Server* services must be running and set to automatic restart.
- For Windows Vista and later, *Network discovery* and *Printer and File Sharing* must be enabled for the current network profile (in **Control Panel > Network and Sharing Center**).
- For Windows XP Professional and 2003, *File and Printer Sharing for Windows Networks* must be enabled for the network adapter itself. Select **Control Panel > Network and Sharing Center > Change Adapter Settings**, right-click the adapter, select **Properties** and check **File and Printer Sharing for Microsoft Networks**.
- For Windows XP Professional, turn off *Simple sharing* in **Control Panel > Folder Options > View > Use simple file sharing**.

- Trust credentials entered on the Unitrends Add Asset page 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 administrator account. Being a different member of the Administrators group is insufficient. If the administrator account is disabled, enable it by executing the following in an elevated command prompt: **`net user administrator /active:yes`**
  - Registry key *LocalAccountTokenFilterPolicy* must exist and be set to *1*.
- 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.
- For 32-bit assets only, the Windows machine must be configured with English as the default language. Select non-English default languages are supported for 64-bit machines. For a list of supported languages, see the Unitrends Compatibility and Interoperability Matrix.

Push-installing the Windows agent

Add the asset as described in "To add an asset" on page 119. The agent is installed automatically (assuming all "Windows agent requirements" on page 137 have been met).

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://support.unitrends.com/go/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 138
- "Command-line installer for Windows agents" on page 140
- "Agent deployment using Group Policy" on page 142

Agent installer for Windows

The agent installer loads all components in *Unitrends_Agentx86.msi*, *Unitrends_Agentx64.msi*, or *Unitrends_BareMetal.msi* onto the Windows asset during installation.

Consider the following before installing the Windows agent:

- For Microsoft SQL and Exchange servers, the Windows core agent includes SQL and Exchange components to protect these applications. For more information, see "Exchange backup requirements and considerations" and "SQL backup requirements and considerations".
- 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 core or bare metal agent

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 (i.e. local administrator).
2. Download the agent *.msi* file from https://support.unitrends.com/go/downloads.
   - For the core agent, click the 32-bit or 64-bit link for the applicable Windows version.
   - For the bare metal agent, click the release link next to *Bare Metal media creation*.
3. Launch the downloaded *.msi* file.
4. Review the license agreement and select **I Agree** to accept the terms and conditions.
5. Select an installation directory. The default directory is *C:\PCBP*. To install in another location (folder or volume), click **Browse** or manually enter the directory path.
6. Click **Back** to review or modify data, or click **Next** to begin the installation process. The installation can be interrupted at any time by clicking **Cancel**.

   Note: If the asset has a firewall enabled, the installer opens port 1743 and creates firewall exceptions for the necessary processes.

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" on page 139
- "Preparing to install on Windows Server 2012/2008" on page 140
- "To install the bare metal agent on a Windows asset running User Account Control" on page 140

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://support.unitrends.com/go/downloads.
   - On the Downloads page, click the release link next to *Bare Metal media creation*.
   - Save the file.
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 Recovery Series and Unitrends Backup appliances" on page 21).

Command-line installer for Windows agents

With the command-line installer, you can install, remove, and repair the Windows agent.

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 (upper case)	Parameter value	Default value
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.
INSTALLDIR		"C:\PCBP" **Note:** Directory name must not contain spaces and must be enclosed in double quotes.
IP		127.0.0.1
FIREWALL	True | False	False

Following are Windows agent installer command-line examples:

- Example 1 – Install *Unitrends_Agentx64.msi* with default values, where*<Deploy>* is the path of the downloaded agent:

```
msiexec /i "C:\<Deploy>\Unitrends_Agentx64.msi"
```

- Example 2 – Install *Unitrends_Agentx64.msi* with default values and turn on logging, where*<Deploy>* is the path of the downloaded agent and *C:\temp\Unitrends.log* is the path of the log file named *Unitrends.log*:

```
msiexec /quiet /l* "C:\temp\Unitrends.log" /i "C:\<Deploy>\Unitrends_
Agentx64.msi"
```

- Example 3 – Uninstall *Unitrends_Agentx64.msi*, where*<Deploy>* is the path of the downloaded agent:

```
msiexec /quiet /uninstall "C:\<Deploy>\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 140 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.

See the following topics for details:

- "Push installing agent updates" on page 143
- "Manually updating and removing Windows agents" on page 143

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 options toolbar. Note that alerts display only for assets that meet the push install requirements described in "Windows agent requirements".

To push install agent updates

1 Select **Configure > Protected Assets**.

2 All protected assets display. 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 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).

Manually updating and removing Windows agents

Use these procedures to manually update, move, or remove the Windows agent:

- "To manually update the Windows agent" on page 144
- "To move the agent to another location" on page 144
- "To remove or repair the Windows agent" on page 144

To manually update the Windows agent

Install the latest agent version as described in "Manually installing the Windows agent" on page 138. It is not necessary to uninstall existing agent software.

To move the agent to another location

If you are running out of space on *C:* or need to move the agent for some other reason:

1. Uninstall the agent as described in "To remove or repair the Windows agent".

   Log files remain in the *C:\PCBP* directory.

2. Move or delete the *C:\PCBP* directory as desired.

3. Manually install the agent in the desired location. See "Manually installing the Windows agent" on page 138 for details.

   Note: 1100 MB of free space is required on the system drive (usually volume *C:*), even if you are installing to a different volume. This space is needed for the installer program to run.

To remove or repair 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 140 for details.

Do one of the following to remove or repair the agent:

- In the Windows Add/Remove Programs interface, select the **Unitrends Agent** in the list, then select **Repair** to repair the agent or **Uninstall** to remove the agent.
- Run the *.msi* agent installer. During the install, select **Repair** to repair the agent or **Remove** to remove the agent.

Installing and updating the Linux agent

Unitrends protects most Linux distributions, including CentOS, Debian, Fedora, 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 144
- "Installing the Linux agent" on page 146

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 the table "Linux distributions and agent installers" on page 145 to determine which installer to use for your Linux asset. You can download the agent installers from the Unitrends Downloads page (https://support.unitrends.com/go/downloads). 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 • Oracle Linux • Red Hat	RPM-based installers. Automatically installs dependencies.
• All 32-bit and 64-bit distributions listed in the Unitrends Compatibility and Interoperability Matrix	GZEXE installers
• Ubuntu	dpkg-based installers. Automatically installs dependencies.

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 Red Hat Linux environments. These dependencies replace XINETD, which was a dependency for earlier versions.

Operating System	Dependencies
Red Hat 5 i386	• vixie-cron • tcp_wrappers Packages are located on the installation media.
Red Hat 5 x86_64	• vixie-cron • tcp_wrappers Packages are located on the installation media.

Operating System	Dependencies
Red Hat 6 i386	• ed Packages are located on the installation media.
Red Hat 6 x86_64	• ed • 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.

Needed dependencies for Oracle on Linux

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.

For assistance with these packages, you can download the Oracle Dependency plug-in from the Unitrends Downloads page (https://support.unitrends.com/go/downloads). 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.

The table below provides details about the Samba packages needed for Oracle protection.

Dependency	Package name
mount.cifs	• samba-client (for Red Hat 5, Oracle Linux 5, and CentOS 5) • cifs-utils (for most other Linux distributions)

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 146
- "To install the Linux agent on CentOS, Oracle Linux, and Red Hat assets" on page 147
- "Installing the Linux agent on Ubuntu assets" on page 148

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 appropriate agent installer on the Linux machine that you want to add to the Unitrends appliance. You can download the installer from the Unitrends Downloads page (https://support.unitrends.com/go/downloads).
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 145.

   Run the appropriate command from step 4 on page 147 after installing the dependencies.
6. (Optional) To protect Oracle databases, install the Oracle Dependency from the Unitrends Downloads page (https://support.unitrends.com/go/downloads).
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 149. Add the Linux asset to your Unitrends backup appliance to begin protecting it. For instructions, see "To add an asset" on page 119.

To install the Linux agent on CentOS, Oracle Linux, and Red Hat assets

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 appropriate agent installer on the Linux machine that you want to add to the Unitrends appliance. You can download the installer from the Unitrends Downloads page (https://support.unitrends.com/go/downloads).

   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 145.

6 (Optional) To protect Oracle databases, install the Oracle Dependency from the Unitrends Downloads page (https://support.unitrends.com/go/downloads).

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 149.

8 Add the Linux asset to your Unitrends backup appliance to begin protecting it. For instructions, see "To add an asset" on page 119.

Installing the Linux agent on Ubuntu assets

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" on page 148
- "To install the Linux agent on Ubuntu using GDebi" on page 149

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 149.

1 Save the appropriate agent installer on the Linux machine that you want to add to the Unitrends appliance. You can download the installer from the Unitrends Downloads page (https://support.unitrends.com/go/downloads).

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 149.

6 Add the Linux asset to your Unitrends backup appliance to begin protecting it. For instructions, see "To add an asset" on page 119.

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 148.

1 Save the appropriate agent installer on the Linux machine that you want to add to the Unitrends appliance. You can download the installer from the Unitrends Downloads page (https://support.unitrends.com/go/downloads).

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 149.

5 Add the Linux asset to your Unitrends backup appliance to begin protecting it. For instructions, see "To add an asset" on page 119.

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 ports 1743 and 1745.
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
```

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://support.unitrends.com/go/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 asset" on page 119.

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://support.unitrends.com/go/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 asset" on page 119.

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://support.unitrends.com/go/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 asset" on page 119.

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 the 9.0.0 agent 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://support.unitrends.com/go/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 asset" on page 119.

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 Novell OES Linux agent

Before adding a Novell OES 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 these procedures to install or update the Novell OES Linux agent:

- "Preparing to install the Novell OES Linux agent" on page 155
- "Novell OES Linux agent restrictions and limitations" on page 156
- "To install or update the Novell OES Linux agent" on page 156

Preparing to install the Novell OES Linux agent

The Unitrends agent for Open Enterprise Server (OES) enables you to backup and restore OES

data.

Before installing a Novell OES Linux agent:

- Make sure your Novell OES Linux system and its applications are running supported versions 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 on a 64-bit OES system, the 32-bit runtime environment must be enabled (this is the default configuration).

Novell OES Linux agent restrictions and limitations

The following restrictions apply to the Novell OES Linux agent:

- Recovering individual files and folders from a backup copy is not supported. Only full backups can be recovered from a backup copy when restoring a TSA-based backup.
- Xen is supported only when it is running on OES 2 on SUSE Linux Enterprise 10.
- Hot bare metal is not supported for OES 2 on SUSE Linux Enterprise 10 and 11.
- Bare metal of Xen guest operating systems can be performed only if VT/AMD-V is supported by the host server's CPU and this support is enabled in BIOS.
- If a network mount is mounted on a directory with the same name as seen on the Novell OES Linux machine, then the backups can have difficulty traversing that file system. For example, if *server1:**data* is mounted to *data*, this presents a problem. The mount point should use a different name, such as *server1:**data* mounted to *netdata*. This is a known issue with TSAFS.

To install or update the Novell OES Linux agent

1. Log in to the Novell OES Linux machine as user root.
2. Verify that the novell-sms package is running on the OES system by entering the following command:

```
# service novell-smdrd status
```

3. If the service is not running, enter the following command:

```
# service novell-smdrd start
```

4. Place the agent installation file, *oes_cnt*, on the OES system. (Download the agent from the Unitrends Downloads page at https://support.unitrends.com/go/downloads.)
5. Grant execute permission to the file by running the following command:

```
# chmod +x oes_cnt
```

6. Begin the installation by executing the file:

```
# .\oes_cnt
```

7. Enter **y** to continue the installation and press **Enter** to continue.

8 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.

9 (Optional) Enter an email address to receive reports from the OES system.

10 Enter the hostname of the backup appliance.

11 You are asked if your server is behind a firewall. Answer **yes** or **no**. Answering **yes** forces communication over port 1745.

12 Select **Enter** to approve default port and autoexec settings.

13 After the connection is made to the TSA, enter the user name (root) and password as prompted. This enables SMS-TSA based backups.

Note: Backup and restore speeds are limited by the TSAFS performance. The TSAFS performance on an NSS file system is superior to performance on a non-NSS file system by as much as 300%. For more information on improving the TSAFS performance, refer to the following Novell document, Fine-Tuning SMS Performance.

The agent is installed and you are ready to add your Novell OES Linux asset to the Unitrends appliance. For details, see "To add an asset" on page 119.

Removing the Novell OES Linux agent

To uninstall the agent, log in to the Novell OES Linux 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 SCO OpenServer agent

Before adding a SCO OpenServer 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 SCO OpenServer agent:

- "Preparing to install the SCO OpenServer agent" on page 157
- "To install the SCO OpenServer agent" on page 158

Preparing to install the SCO OpenServer agent

The Unitrends agent for SCO OpenServer enables you to back up, verify, and restore SCO server data.

Before installing a SCO OpenServer agent:

- Make sure your SCO OpenServer 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 SCO OpenServer agent

1. Log in to the SCO OpenServer machine as user root.
2. Place agent file, *sco5_cnt*, on the SCO OpenServer system. (Download the agent from the Unitrends Downloads page at https://support.unitrends.com/go/downloads.)
3. Change to the directory containing the agent file:

```
# cd <download location>
```

4. Grant execute permission to the file by running the following command:

```
# chmod +x sco5_cnt
```

5. Begin the installation by executing the file:

```
# ./sco5_cnt
```

6. Enter **y** to continue the installation and press **Enter** to continue.
7. 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.
8. (Optional) Enter an email address to receive reports from the SCO system.
9. Enter the hostname of the backup appliance.
10. You are asked if your server is behind a firewall. Answer **yes** or **no**. Answering **yes** forces communication over port 1745.
11. Select **Enter** to approve default port and autoexec settings.

The agent is installed and you are ready to add your SCO OpenServer asset to the Unitrends appliance. For details, see "To add an asset" on page 119.

Removing the SCO OpenServer agent

To uninstall the agent, log in to the SCO 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 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://support.unitrends.com/go/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 asset" on page 119.

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

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://support.unitrends.com/go/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 asset" on page 119.

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:

- "To view all copied assets" on page 161
- "To edit retention of a copied asset" on page 161
- "To remove a copied asset" on page 162

To view all copied assets

1. On the **Configure > Appliances** page, select the appliance.
2. Click the **Copied Assets** tab.
3. Use these options to customize your view:
   - View options:
     - To view assets in a list, click **View: List**. Each row in the list describes a single asset.
     - To view assets in a table, click **View: Table**. Assets display in tiles on the left. Click an asset to view its details.
   - Display options: select to filter the display by asset type.

To edit retention of a copied asset

1. On the **Configure > Appliances** page, select the appliance that is protecting the asset.
2. Click **Copied Assets** and select the desired asset.
3. Click **Edit > Manage Retention**.
4. Modify retention settings as desired, and click **Save**.

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.

Retention setting	Description
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 remove a copied asset

1. On the **Configure > Appliances** page, select the appliance that is protecting the asset.
2. Click **Copied Assets** and select the asset you want to remove.
3. Click **Remove > Confirm**.

ConnectWise PSA Integration

This chapter provides information and procedures about integrating with the PSA tool, ConnectWise. See the following topics for details:

Introduction

If you use a Professional Services Automation (PSA) tool, this feature pertains to you. The Unitrends PSA feature automates the creation of tickets in the Managed Service Providers' (MSPs) PSA tools. Currently, ConnectWise is the only PSA tool that Unitrends supports. Unitrends supports all versions of ConnectWise. The following tickets are supported:

Ticket	Description
Service tickets	Tickets that are used to track issues in the system.
Billing tickets	Tickets that contain billing information for creating invoices. Please note that these show up as service tickets in the PSA tool and contain information required for billing.

Previously, when an issue was found, the information was entered manually into ConnectWise. This process was prone to errors and was time-consuming. The PSA Integration feature enables Unitrends software to automatically create a service ticket. The PSA Integration feature also automates the retrieval of billing information from the Unitrends system and creates a billing ticket.

Once set up, the Unitrends system sends service ticket and billing information from the company sites directly to the MSPs' PSA Integration tool.

Configuring the PSA tool

To enable the PSA Integration tool, you must:

1. Configure the ConnectWise tool. For details, see "Configuring settings in ConnectWise" on page 163.
2. Configure the Unitrends system. For details, see "Configuring the Unitrends PSA Integration feature" on page 164.
3. Configure PSA settings in the Unitrends system, if needed. For details, see "Configuring PSA settings in the Unitrends system" on page 166.

Configuring settings in ConnectWise

Follow these instructions to configure the settings in the ConnectWise application. This is the first step in configuring your PSA tool.

Prerequisites

Before you configure settings in ConnectWise, you must perform the following prerequisites:

Settings	Description
Set up the Integrator Login	This enables the Unitrends system to integrate with the PSA tool.
Select the APIs to use / ensure the Service Ticket API is enabled	Only the Service Ticket API is required. Ensure that it is enabled. When it is enabled, the Unitrends system can create service tickets in ConnectWise.
Ensure that the integrator company name is active	When the integrator company is active, ConnectWise is able to "communicate with" the Unitrends system.

To configure settings in ConnectWise

1. Ensure that the prerequisites are met. See "Prerequisites" on page 163 more information.
2. In the Integrator Login screen, enter and save the following information:

Field	Description
Username	Enter the user name.

Field	Description
Password	Enter a password.
Access Level	Select **All records** in the drop-down box.
Service Ticket API (checkbox)	Click the **Service Ticket API** checkbox.
Service Board	Select **Professional Services** from the Service Board drop-down box.
Callback URL	Enter a valid URL, such as www.connectwise.com or “localhost”.

Note: Prior to using the PSA feature, confirm the time zone entry is set up correctly. Please contact your ConnectWise administrator or consult the ConnectWise documentation.

3 Ensure that the ConnectWise company name is set to active in the ConnectWise application.

4 Continue to "Configuring the Unitrends PSA Integration feature" on page 164, which is the second step in configuring your PSA tool.

Configuring the Unitrends PSA Integration feature

On the PSA Configuration page, you can create, modify, delete, view, or save configuration and authentication information. You can also create a test service ticket to send to ConnectWise.

Note: Before you configure the Unitrends PSA integration feature, make sure you configure the Connectwise tool. See "Configuring settings in ConnectWise" on page 163 for steps.

To create the PSA configuration

When you create the PSA configuration, you are providing authentication information to connect with ConnectWise.

1 In the Unitrends UI, select **Configure > Appliances > Interactions > Add**.

2 Enter the following information in the Add PSA Configuration dialog:

Enter the following information in the Details box:

Field	Description
URL	Enter the URL that your company uses for ConnectWise. Enter the node name only. Do not include https:// or http://. Example of correct entry: *test.connectwise.com* Examples of incorrect entries: *www.connectwise.com, http://www.connectwise.com*
Company ID	Enter the integrator company ID from ConnectWise. Make sure the this is an exact match, including case. **Note:** Make sure the ConnectWise company configuration is active in the ConnectWise application. See "Configuring settings in ConnectWise" for more information.

Enter the following information in the Credentials box:

Field	Description
Credentials	Select New in the dropdown menu.
Credential Name	Enter a credential name if you want to set up a level of credentials for a group, such as Operators.
Username	Enter the Integrator Login user name from ConnectWise.
Password	Enter the Integrator Login password from ConnectWise.
Confirm Password	Re-enter your password.
Domain	Optional.

3 Click **Save**. ConnectWise will appear in the Tool column of the **Interactions** tab.

4 Click Connectwise in the in the Tool column to highlight it, then click the **Send Test Ticket** button to send a test ticket to ConnectWise. You can view the status of your test ticket by clicking the **View Test Ticket History** button.

Note: Sending the test ticket is crucial to determining that ConnectWise is receiving tickets.

5 Use the ticket number to confirm that the system sent a test ticket to ConnectWise.

You can go to ConnectWise to confirm that the test ticket was sent. Perform a ticket search, if necessary.

6 Continue to "Configuring PSA settings in the Unitrends system" on page 166, which is the next step in configuring your PSA tool, if needed.

Note: After you create the PSA configuration, you can modify or delete it, as necessary. See "Editing or removing a PSA configuration" on page 168.

Configuring PSA settings in the Unitrends system

There are PSA settings that allow you to configure data in the Unitrends system and data that the Unitrends system sends to Connectwise. This information resides in the Unitrends appliance's *master.ini* file.

To configure PSA settings in the Unitrends system

1. In the appliance UI, select **Configure > Appliances > Edit > Advanced > General Configuration**.
2. Locate items labeled PSA in the Section column.
3. Click the row you wish to modify.
4. Enter the new value in the Value field and click **Confirm**.

   This table lists the settings and descriptions you can change:

Settings	Descriptions
BaseDelay	Time, in seconds, to delay a retry attempt on a remote connection failure between the Unitrends system and Connectwise.
Board	The Connectwise service board where the ticket information displays. See the Service Board drop-down box in Connectwise for possible values. Verify the exact wording prior to making an entry in this field.
CWApiVersion	The version of the ConnectWise API used by your ConnectWise server.
DebugTrace	Whether or not to show the debug trace in the psa.log (1 = show the debug trace and 0 = do not show the debug trace).

Settings	Descriptions

ExclusionListFile	Use this setting to exclude tickets from the information that the Unitrends system sends to the Connectwise service board. To use this setting:
	1 Create a file in the appliance Samba share that contains the strings to be excluded, one per line. The values are from the Summary Description column in the Service Board List in the Connectwise PSA application. Wildcards are not supported. The appliance excludes any notifications that contain the strings you enter in this file.
	2 Enter the full path filename in the Value field. This field is case-sensitive.
	Sample exclusion list file text entries:
	`{noformat}`
	`Unitrends user interface version`
	`will expire in 30 days`
	`{noformat}`
	In the example above:
	• The first line excludes notifications that a new user interface is available.
	• The second line excludes the license expiration warning.

LoginCompanyId	If your Connectwise server is managing multiple customer destinations, a unique LoginCompanyId must be entered here.
Priority1, 2, 3, and 4	This field must match the Priority 1, Priority 2, Priority 3, and Priority 4 values that display in the Priority drop-down on the Service Board List in the Connectwise PSA application. Verify the exact wording prior to making an entry in this field.
ProcessNotifications	Whether or not to send an email notification (1 = send and 0 = do not send). The email address comes from the primary contact for the company in the Connectwise PSA application.
RetryCount	The number of times to retry a remote connection from the Unitrends system to Connectwise before the system times out.
Status	The initial status description of the ticket. Defaults to N for new, but you can change it to another value, such as *unassigned*.

Settings	Descriptions
WSDLLocation	To use this setting: 1 Locate your Connectwise URL under **Configure > Appliances > Interactions > URL**. 2 Enter the URL into the Value field using the following format: *https://<URL>/v4_6_release/apis/1.5/ServiceTicketApi.asmx?wsdl*

Editing or removing a PSA configuration

You can modify or delete a PSA configuration.

To modify or delete a PSA configuration

1. Select **Configure > Interactions**.
2. Under the Tool column, click ConnectWise to highlight it.
3. Click the **Edit** button to open the Edit PSA Configuration window.
4. To remove the PSA configuration, click the **Remove** button.

Viewing ticket history

The Ticket History page displays information about tickets that were created sucessfully and tickets that could not be created.

To view ticket history

1. From the **Configure** tab, select your appliance.
2. Click on the **Interactions** tab.
3. Select ConnectWise.
4. Click on **View Ticket History**.

Using the billingInvoker script

The billingInvoker script allows the following tasks to be performed from the command line interface:

- Generating a billing ticket.
- Creating a billing ticket schedule.
- Deleting a billing ticket schedule.

Billing tickets appear in the Unitrends Backup appliance UI under **Configure > Appliances > Interactions > View Ticket History**.

WARNING! It is recommended that all administration tasks are performed using the graphical user interface. The Unitrends operating system implementation is proprietary to Unitrends and should not be modified from the command line unless following an certified Unitrends procedure. Performing general administration tasks from the command line can result in undesirable outcomes.

To generate a billing ticket

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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)
3. Issue this command to change to the command directory:

```
# cd /usr/bp/bin
```

4. Run the script using this command:

```
# ./billingInvoker
```

5. You will see the following output:

```
PSA Billing
    Usage: billingInvoker [generateBill | changeBillDay | getBillDay |
    disableBilling]
    generateBill -- generates a billing ticket
    changeBillDay <day> -- changes the day when the billing ticket should
    be created. <day> should be between 1-31.
    disableBilling -- disable the billing invocation via cron
    getBillday -- shows the billing day of the month
```

6. Enter the following command:

```
# ./billingInvoker generateBill
```

To create a billing ticket schedule

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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)
3. Issue this command to change to the command directory:

```
# cd /usr/bp/bin
```

4 Run the script using this command:

```
# ./billingInvoker
```

5 You will see the following output:

```
PSA Billing
    Usage: billingInvoker [generateBill | changeBillDay | getBillDay |
    disableBilling]
    generateBill -- generates a billing ticket
    changeBillDay <day> -- changes the day when the billing ticket should
    be created. <day> should be between 1-31.
    disableBilling -- disable the billing invocation via cron
    getBillday -- shows the billing day of the month
```

6 Enter the following command:

```
# ./billingInvoker changeBillDay <day of the month you want the billing
ticket to be generated on>
```

Note: You can view the scheduled day by following steps 1-5 and entering the command `/billingInvoker getBillDay`.

To delete a billing ticket schedule

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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)

3 Issue this command to change to the command directory:

```
# cd /usr/bp/bin
```

4 Run the script using this command:

```
# ./billingInvoker
```

5 You will see the following output:

```
PSA Billing
    Usage: billingInvoker [generateBill | changeBillDay | getBillDay |
    disableBilling]
    generateBill -- generates a billing ticket
    changeBillDay <day> -- changes the day when the billing ticket should
    be created. <day> should be between 1-31.
    disableBilling -- disable the billing invocation via cron
```

```
    getBillday -- shows the billing day of the month
```

6 Enter the following command:

```
# ./billingInvoker disableBilling
```

Chapter 4: Backup Administration and Procedures

This chapter provides instructions for creating and managing backup jobs and backup copy jobs. The procedures in this chapter are initiated from the Jobs page of the Unitrends User Interface. See the following topics for details:

- "Creating backup jobs" on page 173
- "Creating backup copy jobs" on page 186
- "Managing scheduled jobs" on page 193
- "Managing active jobs" on page 197
- "Viewing recent jobs" on page 199
- "Viewing system jobs" on page 200
- "Deleting backups and backup copies" on page 201

For information on generating reports on backups, see "Backup reports" on page 378.

Creating backup jobs

See the following topics to create backup jobs:

- "Preparing to create backup jobs" on page 173
- "Selecting assets to back up" on page 174
- "Backup job procedures" on page 174

Note: 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 261.

Preparing to create backup jobs

Before creating jobs, add to the appliance the assets you want to protect, as described in "Managing protected assets" on page 118. Unitrends also recommends that you review the following information to develop the best protection strategy for your environment:

- "Protection Overview" on page 35, which covers backup modes, backup groups, and other general key concepts.
- The backups overview chapter for the asset type you want to protect, which covers requirements and other considerations:
  - "Host-level Backups Overview" on page 203
  - "Asset-level Backups Overview" on page 223
  - "NAS Backups Overview" on page 227
  - "Application Backups Overview" on page 237

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 asset type you want to protect. You can then select assets of this type in an inventory tree. You cannot select assets that do not match the asset type you picked. 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:

- Each physical asset and virtual host displays as a top-level node.
- Each application displays as a sub-node under its physical host asset. Individual databases display under their application node.
- Each VM displays as a sub-node under its virtual host.

Backup job procedures

Use these procedures to create backup jobs:

- "To create a backup job for agent-based assets" on page 174
- "To create a backup job for VMware assets" on page 176
- "To create a backup schedule for VMware assets by using regular expression filters" on page 177
- "To create a backup job for Hyper-V assets" on page 179
- "To create a backup job for XenServer assets" on page 180
- "To create a backup job for NAS CIFS or NFS assets" on page 180
- "To create a backup job for NAS NDMP assets" on page 181
- "To create an Exchange backup job" on page 182
- "To create an Oracle backup job" on page 183
- "To create a SQL backup job" on page 184
- "To create a SharePoint backup job" on page 184
- "To create a UCS service profile backup job" on page 185

To create a backup job for agent-based assets

1. Select **Jobs > Create Job > Backup**.
2. Select **Agent-Based Assets** 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.

   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 agent-based assets to the schedule.
5. (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:

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).

- 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.)
  - Run a new full backup upon creating or modifying included files.
- 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, do any of the following:
    - Type in the full path (e.g., *C:/Documents*).
    - **Browse** the asset.
    - Enter a regular expression. (Wildcards are not supported.)
  - Run a new full backup upon creating or modifying excluded files.

  Note: 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.
- Advanced - Click to exclude the system state from backup or to specify commands to run on the asset before or after a scheduled backup runs.

  Notes:

  - To perform bare metal recovery or Windows instant recovery the following must be included in the backup: system state and all boot and critical system (OS) volumes. If you need these features for the asset, do not specify data to include or exclude unless you are sure these volumes will be included.
  - 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*.

- 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*.
- 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 instant recovery features. Be sure one of the aliased assets contains all of these disks to use these features.

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, then click **Save**:

- In most cases, the standard backup modes can be used to create the schedule.
- If you need more granularity, choose the **Custom** mode and do these steps to create a custom backup calendar:
  - Click the calendar icon next to *Click to Edit*.
  - 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.)
  - In the Add Backup dialog, modify settings as desired, then click **Save**.
  - Repeat these steps to add other modes to the calendar.
  - Click **Save** to save the settings and close the Calendar dialog.
  - Click **Save** in the Create Backup Job dialog to save the calendar.
- 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 backup job for VMware assets

Note: 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 > Create Job > Backup**.
2. Select **VMware Assets** 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.
   - To locate an asset by name, use the **Search** field below.
   - To protect all VMs, select the virtual host.
4. (Optional) Check the **Auto-include new VMs** box to automatically add newly discovered VMs to the schedule.
5. (Optional) 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.
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, then click **Save**:
   - In most cases, the standard backup modes can be used to create the schedule.
   - If you need more granularity, choose the **Custom** mode and do these steps to create a custom backup calendar:
     - Click the calendar icon next to *Click to Edit*.
     - 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.)
     - In the Add Backup dialog, modify settings as desired, then click **Save**.
     - Repeat these steps to add other modes to the calendar.
     - Click **Save** to save the settings and close the Calendar dialog.
     - Click **Save** in the Create Backup Job dialog to save the calendar.
   - 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 backup schedule for VMware assets 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.

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.

Note: 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 > Create Job > Backup**.
2. Select **VMware Assets** 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.

   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:
   - Select a 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.
8. Click **Next**.
9. Enter a unique Job Name and select **Create a Schedule**.
10. 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, choose the **Custom** mode and do these steps to create a custom backup calendar:
  - Click the calendar icon next to *Click to Edit*.
  - 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.)
  - In the Add Backup dialog, modify settings as desired, then click **Save**.
  - Repeat these steps to add other modes to the calendar.
  - Click **Save** to save the settings and close the Calendar dialog.
  - Click **Save** in the Create Backup Job dialog to save the calendar.
- The schedule is created and runs at the dates and times you specified.

To create a backup job for Hyper-V assets

Note: 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 > Create Job > Backup**.
2. Select **Hyper-V Assets** in the **What do you want to backup?** 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. (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, then click **Save**:
   - In most cases, the standard backup modes can be used to create the schedule.
   - If you need more granularity, choose the **Custom** mode and do these steps to create a custom backup calendar:
     - Click the calendar icon next to *Click to Edit*.
     - 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.)
     - In the Add Backup dialog, modify settings as desired, then click **Save**.
     - Repeat these steps to add other modes to the calendar.
     - Click **Save** to save the settings and close the Calendar dialog.

- Click **Save** in the Create Backup Job dialog to save the calendar.

- 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 backup job for XenServer assets

Note: 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 > Create Job > Backup**.
2. Select **XenServer Assets** in the **What do you want to backup?** list.
3. In the Inventory tree, expand the asset and check boxes to select virtual machines to protect.

   To locate the asset by name, use the **Search** field below.
4. (Optional) 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. (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, then click **Save**:
   - In most cases, the standard backup modes can be used to create the schedule.
   - If you need more granularity, choose the **Custom** mode and do these steps to create a custom backup calendar:
     - Click the calendar icon next to *Click to Edit*.
     - 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.)
     - In the Add Backup dialog, modify settings as desired, then click **Save**.
     - Repeat these steps to add other modes to the calendar.
     - Click **Save** to save the settings and close the Calendar dialog.
     - Click **Save** in the Create Backup Job dialog to save the calendar.
   - 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 backup job for NAS CIFS or NFS assets

1. Select **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.

To locate the asset by name, use the **Search** field below.

4 (Optional) Click **Edit** to specify directories or files to include or exclude. Click **Save** to retain any changes.

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, choose the **Custom** mode and do these steps to create a custom backup calendar:
  - Click the calendar icon next to *Click to Edit*.
  - 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.)
  - In the Add Backup dialog, modify settings as desired, then click **Save**.
  - Repeat these steps to add other modes to the calendar.
  - Click **Save** to save the settings and close the Calendar dialog.
  - Click **Save** in the Create Backup Job dialog to save the calendar.
- 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 backup job for NAS NDMP assets

1 Select **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.

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, then click **Save**:

- Each selected volume is backed up in a separate job. Jobs are queued at the time you specified and run as NDMP connections become available.
- In most cases, the standard backup modes can be used to create the schedule.

- If you need more granularity, choose the **Custom** mode and do these steps to create a custom backup calendar:
  - Click the calendar icon next to *Click to Edit*.
  - 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.)
  - In the Add Backup dialog, modify settings as desired, then click **Save**.
  - Repeat these steps to add other modes to the calendar.
  - Click **Save** to save the settings and close the Calendar dialog.
  - Click **Save** in the Create Backup Job dialog to save the calendar.
- If you created a schedule, the job is queued to run at the date and times you specified.
- If you chose Now, the job queues immediately. Click **Active Jobs** to view the running jobs.

To create an Exchange backup job

Note: 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 > 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, then click **Save**:
   - In most cases, the standard backup modes can be used to create the schedule.
   - If you need more granularity, choose the **Custom** mode and do these steps to create a custom backup calendar:
     - Click the calendar icon next to *Click to Edit*.
     - 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.)
     - In the Add Backup dialog, modify settings as desired, then click **Save**.
     - Repeat these steps to add other modes to the calendar.
     - Click **Save** to save the settings and close the Calendar dialog.

- Click **Save** in the Create Backup Job dialog to save the calendar.

- If you created a schedule, the job will run 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

Note: 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 > 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 will show as available for backup. The Oracle Dependency from the latest agent release must be installed to protect Oracle data.

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, then click **Save**.

   Note: **For Oracle on Linux.** If you are running an incremental forever schedule, you must also exclude Oracle database directories from journal tracking. See KB 3358 for details.

   - In most cases, the standard backup modes can be used to create the schedule.
   - If you need more granularity, choose the **Custom** mode and do these steps to create a custom backup calendar:
     - Click the calendar icon next to *Click to Edit*.
     - 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.)
     - In the Add Backup dialog, modify settings as desired, then click **Save**.
     - Repeat these steps to add other modes to the calendar.
     - Click **Save** to save the settings and close the Calendar dialog.
     - Click **Save** in the Create Backup Job dialog to save the calendar.
   - 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

Note: 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 > Create Job > Backup**.
2. Select **SQL** in the **What do you want to backup?** list.
3. In the Inventory tree, expand the SQL server and application, then check boxes to select databases to protect.
   - To locate an asset by name, use the **Search** field below.
   - To protect all databases, select the SQL application.
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, then click **Save**:
   - In most cases, the standard backup modes can be used to create the schedule.
   - If you need more granularity, choose the **Custom** mode and do these steps to create a custom backup calendar:
     - Click the calendar icon next to *Click to Edit*.
     - 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.)
     - In the Add Backup dialog, modify settings as desired, then click **Save**.
     - Repeat these steps to add other modes to the calendar.
     - Click **Save** to save the settings and close the Calendar dialog.
     - Click **Save** in the Create Backup Job dialog to save the calendar.
   - 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

Note: 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 > 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, then click **Save**:
   - In most cases, the standard backup modes can be used to create the schedule.
   - If you need more granularity, choose the **Custom** mode and do these steps to create a custom backup calendar:
     - Click the calendar icon next to *Click to Edit*.
     - 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.)
     - In the Add Backup dialog, modify settings as desired, then click **Save**.
     - Repeat these steps to add other modes to the calendar.
     - Click **Save** to save the settings and close the Calendar dialog.
     - Click **Save** in the Create Backup Job dialog to save the calendar.
   - 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 > 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, then click **Save**:
   - In most cases, the standard backup modes can be used to create the schedule.
   - If you need more granularity, choose the **Custom** mode and do these steps to create a custom backup calendar:
     - Click the calendar icon next to *Click to Edit*.
     - 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.)
     - In the Add Backup dialog, modify settings as desired, then click **Save**.
     - Repeat these steps to add other modes to the calendar.
     - Click **Save** to save the settings and close the Calendar dialog.
     - Click **Save** in the Create Backup Job dialog to save the calendar.

- 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 an off-site target. See the following topics to create backup copy jobs:

- "Preparing to create a backup copy job" on page 186
- "Selecting assets for backup copy" on page 186
- "Backup copy job procedures" on page 187

Preparing to create a backup copy job

The target stores backup copies when you run a backup copy job. Before creating backup copy jobs, you must first add the target (see "Backup copy targets" on page 77). For an overview of the types of targets you can use, see "Backup copies" on page 44.

The following considerations apply to 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 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*.
- Only backups that completed successfully (green) or with warnings (yellow) are eligible for backup copy jobs.
- After creating the backup copy job, the last backups for the selected assets are copied to the target when the backup copy job runs.
- 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.
- 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 asset and virtual host displays as a top-level node.
- 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.

The number of clicks within a check box determines the selection. Review the following for additional information:

Check box	Description
Clear	Excludes this asset from backup.
Check mark	Includes this asset in backup.
Gray	Includes this sub-item in backup.
Red X	Excludes this asset from backup.
Red	Excludes this sub-item from backup.

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 187
- "To create a backup copy job for a Unitrends Cloud target" on page 188
- "To create a backup copy job for a third-party cloud target" on page 188
- "To create a backup copy job for an attached disk target" on page 189
- "To create a backup copy job for a NAS target" on page 190
- "To create a backup copy job for a SAN target" on page 190
- "To create a backup copy job for an eSATA or USB target" on page 191
- "To create a backup copy job for a tape target" on page 192

To create a backup copy job for a Unitrends appliance target

Note: Because scheduling is automatic, a particular asset can only be added to one backup copy job. It is a best practice to create a single backup copy job for all assets associated with a Unitrends appliance backup copy target. This will make monitoring easier.

1. Log in to the source backup appliance.
2. Click **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.
   - Only assets that have a completed backup are listed here.
   - Expand the tree as necessary to select VMs and applications.

- Select a virtual host to select all of its VM assets.
- Select a SQL application to select all of its databases.

5 Click **Next**.

6 Select the target Unitrends appliance 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 "Return to the source backup appliance and fine-tune settings by adjusting connection options" on page 81).

To create a backup copy job for a Unitrends Cloud target

Use this procedure to use the Unitrends Cloud service to store your backup copies.

Note: Because scheduling is automatic, a particular asset can only be added to one backup copy job. It is a best practice to create a single backup copy job for all assets associated with a Unitrends Cloud backup copy target. This will make monitoring easier.

1 Log in to the source backup appliance.

2 Click **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.

- Only assets that have a completed backup are listed here.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select a SQL application to select all of its databases.

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 "Return to the source backup appliance and fine-tune settings by adjusting connection options" on page 81).

To create a backup copy job for a third-party cloud target

Use this procedure to use a third-party cloud service to store your backup copies.

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 105.

1 Click **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.
   - Only assets that have a completed backup are listed here.
   - Expand the tree as necessary to select VMs and applications.
   - Select a virtual host to select all of its VM assets.
   - Select a SQL application to select all of its databases.
4. Click **Next**.
5. Select the cloud target in the **Backup Copy Target** list.
6. Set remaining Backup Copy Options and Schedule Details as desired.
   - In most cases, the Standard scheduling mode can be used to create the schedule.
   - If you need more granularity, choose the **Custom** mode and do these steps to create a custom calendar:
     - Click the calendar icon next to *Click to Edit*.
     - In the Calendar dialog, select Backup Copy in the Backup copy area and drag it to a day on the calendar. (You cannot drag to a day in the past.)
     - In the Add Backup Copy dialog, modify settings as desired, then click **Save**.
     - Click **Save** to save the settings and close the Calendar dialog.
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. The job copies the last backup of the mode you selected in step 6 on page 189 above (*Fulls only* or *All backup modes*).

To create a backup copy job for an attached disk target

1. Click **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.
   - Only assets that have a completed backup are listed here.
   - Expand the tree as necessary to select VMs and applications.
   - Select a virtual host to select all of its VM assets.
   - Select a SQL application to select all of its databases.
4. Click **Next**.
5. Select the disk target in the **Backup Copy Target** list.
6. Set remaining Backup Copy Options and Schedule Details as desired.
   - In most cases, the Standard scheduling mode can be used to create the schedule.
   - If you need more granularity, choose the **Custom** mode and do these steps to create a custom calendar:

- Click the calendar icon next to *Click to Edit*.
- In the Calendar dialog, select Backup Copy in the Backup copy area and drag it to a day on the calendar. (You cannot drag to a day in the past.)
- In the Add Backup Copy dialog, modify settings as desired, then click **Save**.
- Click **Save** to save the settings and close the Calendar dialog.

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. The job copies the last backup of the mode you selected in step 6 on page 189 above (*Fulls only* or *All backup modes*).

To create a backup copy job for a NAS target

1 Click **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.

- Only assets that have a completed backup are listed here.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select a SQL application to select all of its databases.

4 Click **Next**.

5 Select the NAS target in the **Backup Copy Target** list.

6 Set remaining Backup Copy Options and Schedule Details as desired.

- In most cases, the Standard scheduling mode can be used to create the schedule.
- If you need more granularity, choose the **Custom** mode and do these steps to create a custom calendar:
  - Click the calendar icon next to *Click to Edit*.
  - In the Calendar dialog, select Backup Copy in the Backup copy area and drag it to a day on the calendar. (You cannot drag to a day in the past.)
  - In the Add Backup Copy dialog, modify settings as desired, then click **Save**.
  - Click **Save** to save the settings and close the Calendar dialog.

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. The job copies the last backup of the mode you selected in step 6 on page 190 above (*Fulls only* or *All backup modes*).

To create a backup copy job for a SAN target

1 Click **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.

- Only assets that have a completed backup are listed here.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select a SQL application to select all of its databases.

4 Click **Next**.

5 Select the SAN target in the **Backup Copy Target** list.

6 Set remaining Backup Copy Options and Schedule Details as desired.

- In most cases, the Standard scheduling mode can be used to create the schedule.
- If you need more granularity, choose the **Custom** mode and do these steps to create a custom calendar:
  - Click the calendar icon next to *Click to Edit*.
  - In the Calendar dialog, select Backup Copy in the Backup copy area and drag it to a day on the calendar. (You cannot drag to a day in the past.)
  - In the Add Backup Copy dialog, modify settings as desired, then click **Save**.
  - Click **Save** to save the settings and close the Calendar dialog.

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. The job copies the last backup of the mode you selected in step 6 on page 191 above (*Fulls only* or *All backup modes*).

To create a backup copy job for an eSATA or USB target

1 Click **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.

- Only assets that have a completed backup are listed here.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select a SQL application to select all of its databases.

4 Click **Next**.

5 Select the disk target in the **Backup Copy Target** list.

6 Set remaining Backup Copy Options and Schedule Details as desired.

- In most cases, the Standard scheduling mode can be used to create the schedule.

- If you need more granularity, choose the **Custom** mode and do these steps to create a custom calendar:
  - Click the calendar icon next to *Click to Edit*.
  - In the Calendar dialog, select Backup Copy in the Backup copy area and drag it to a day on the calendar. (You cannot drag to a day in the past.)
  - In the Add Backup Copy dialog, modify settings as desired, then click **Save**.
  - Click **Save** to save the settings and close the Calendar dialog.

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. The job copies the last backup of the mode you selected in step 6 on page 191 above (*Fulls only* or *All backup modes*).

To create a backup copy job for a tape target

1 Click **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.

- Only assets that have a completed backup are listed here.
- Expand the tree as necessary to select VMs and applications.
- Select a virtual host to select all of its VM assets.
- Select a SQL application to select all of its databases.

4 Click **Next**.

5 Select the tape drive or changer in the **Backup Copy Target** list.

6 Set remaining Backup Copy Options and Schedule Details as desired.

- For descriptions of tape settings, see "Tape option details" on page 193 below.
- In most cases, the Standard scheduling mode can be used to create the schedule.
- If you need more granularity, choose the **Custom** mode and do these steps to create a custom calendar:
  - Click the calendar icon next to *Click to Edit*.
  - In the Calendar dialog, select Backup Copy in the Backup copy area and drag it to a day on the calendar. (You cannot drag to a day in the past.)
  - In the Add Backup Copy dialog, modify settings as desired, then click **Save**.
  - Click **Save** to save the settings and close the Calendar dialog.

7 Click **Save**.

The job is created and will run at the date and times specified. The job copies the last backup of the mode(s) you selected in step 6 (*Fulls only* or *All backup modes*) for each asset you selected in step 3.

Tape option details

Item	Description
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 • 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 successful backups of any mode.
Encrypt backup copies	Not used for tape devices. If the tape device is configured for encryption, copiees 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).
Retention Period	Use with the Overwrite option. Check the Retention Period box to specify the length of time a copy is retained before it is overwritten. To define the retention period, enter a number and select Days, Weeks, Months, or Years. Example: enter **2** and select **Weeks** to overwrite copies that are over 2 weeks old.

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:

Note: iSeries schedules are managed by using the dpuconfig console interface. See "iSeries Backups Overview and Procedures" on page 261 for details on working with iSeries schedules.

- "To delete a job" on page 197
- "To run a scheduled job on-demand" on page 197

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.
   - The following information displays for each scheduled job:

Column	Description
Name	The name of the scheduled job.
Asset	The name of the asset whose data is being backed up, copied, or recovered.
Status	The current status of the job: • Running - The job is running now. • 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 details of a job

1. Click **Jobs > Job Manager.**
2. Select the applicable job and click **View Details**.
3. Click **Hide Details** to return to the original page view.

To view or edit a backup job

1. Click **Jobs > Job Manager**.
2. Select the applicable 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) Select an asset in the Job Inventory Settings area and click **Edit** to modify 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:

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).

- 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.)
  - Run a new full backup upon creating or modifying included files.
- 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, do any of the following:
    - Type in the full path (e.g., *C:/Documents*).
    - **Browse** the asset.
    - Enter a regular expression. (Wildcards are not supported.)
  - Run a new full backup upon creating or modifying excluded files.

  Note: 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.
- Advanced - Click to exclude the system state from backup or to specify commands to run on the asset before or after a scheduled backup runs.

  Notes:

  - To perform bare metal recovery or Windows instant recovery the following must be included in the backup: system state and all boot and critical system

(OS) volumes. If you need these features for the asset, do not specify data to include or exclude unless you are sure these volumes will be included.

- 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*.
- 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*.
- 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 instant recovery features. Be sure one of the aliased assets contains all of these disks to use these features.

5 Click **Save** to save any changes you have made.

To view or edit a backup copy job

1 Click **Jobs > Manager**.

2 Select the applicable job and click **Edit**.

3 (Optional) Modify settings as desired and click **Save**.

Note: For Google, Amazon S3, 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 97.

To enable or disable a job

Note: Jobs are enabled by default. When a job is disabled, none of its scheduled backups run. Be aware that disabling jobs can leave assets unprotected.

1 Click **Jobs > Job Manager**.

2 Select the applicable job.

3 Click either **Enable** / **Disable**. Disabled jobs display as grayed-out and italicized.

To delete a job

1 Click **Jobs > Job Manager**.

2 Select the applicable 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 applicable job.

3 Click **Run**. The job queues immediately.

4 To monitor, pause, or cancel the job, go to **Jobs > Active Jobs**.

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 24 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 24 for an at-a-glance view of backup copy performance.
- For Unitrends appliance backup copy targets, see the "Copied Assets tab" on page 32 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 261 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.
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	Job type: Backup, Backup Copy, Import, or Recover.
Started	The date and time the job began.
Progress	A graphic bar representing the completed percentage of current job's progress.
Appliance	The appliance running the job.

To view job details

1. Select **Jobs > Active Jobs**.
2. Select a job in the list.
3. Click **View Details.** (To hide the details, click **Hide Details**.)
   These details display at the bottom of the screen.

Column	Description
Job ID	A system-generated ID number assigned to the job.
Application	The type of asset whose data is being backed up, copied, or recovered.

Column	Description
Host/Server	The name of the virtual host or physical server.
Message	Any system-generated message produced during the job.

To pause a job

Use this procedure to pause queued jobs. Pausing queued jobs can prove useful if you want to push a given job to the top of the queue.

1. Select **Jobs > Active Jobs**.
2. Select a queued job in the list.
3. Click **Pause.**

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 a job

1. Select **Jobs > Active Jobs**.
2. Select a job in the list.
3. Click **Cancel.**

Viewing recent jobs

To view recent jobs, select **Jobs > Recent Jobs**.

The Recent Jobs tab captures the results of job activity over the last 7 days:

- You can export and save job information as a CSV (Excel) file. To export this information, select a job in the table, and click **Export CSV**. To view and/or export a list of files in an asset-level backup, see "Backup History report" on page 381.
- To see additional information about a specific job, select the job and click **View Log** to see the associated log.
- You can order the information in the Recent Jobs table in different ways, according to values in any of the columns. To change the order of information in the Recent Jobs table, hover over the column title, and click through the options. For example, you can display the "Started" column values in descending or ascending dates.

The Recent jobs table contains the following information:

Column	Description
Name	The name of the job.
Status	The final status of the job.
Type/Mode	The job type (Backup, Backup Copy, or Recover) and mode (Full, Incremental, 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.

Backup Catalog

After a backup job completes, the resulting backup displays in the Backup Catalog on the **Recover > Backup Catalog** page. Each backup receives a color code, based on its status:

- Green backups completed successfully
- Yellow backups completed with a warning
- Red backups completed with an error.

For more information about the Backup Catalog and your recovery options, see the "Recovery Overview" chapter.

Job reports

Reports provide additional detail about the jobs that ran in the last 7 days, as well as information about older jobs. See the following topics for more information:

- "Backup reports" on page 378
- "Recover reports" on page 387
- "Backup Copy reports" on page 389

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.

You can export and save this information to a CSV (Excel) file. To export this information, select the job in the table and click **Export CSV**.

You can order the information in the Recent Jobs table in different ways, according to values in any of the columns. To change the order of information in the Recent Jobs table, hover over the column

title, and click through the options. For example, you can display the "Started" column values in descending or ascending dates.

Functions are enabled based on your user account privileges. If you do not have the appropriate credentials, consult your network administrator.

The Recent System jobs table contains the following information:

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.
Message	Displays system-generated messages related to recent system jobs.
Appliance	The appliance that ran the job.

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 42.

Notes:

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 109 for tape autochanger targets.
- "To initialize and erase cold backup copy media" on page 108 for all other cold backup copy targets, including tape drive devices that do not have an autochanger.

See the following procedures to delete backups and backup copies:

- "To delete backups and imported backup copies"
- "To delete hot backup copies"

To delete backups and imported backup copies

1. Log in to the backup appliance.
2. Select **Recover > Backup Catalog**.

   Protected assets display. If desired, use the Filter Backups fields to filter the display.
3. Click to expand the desired asset to view its backups and imported backup copies.
4. Check boxes to select the backups and imported backup copies to delete.

5 Click **Delete**.

6 In the Confirm Backup Deletion dialog, review the list of backups that will be deleted.

7 Click **Confirm** to delete the backups.

To delete hot backup copies

1 Log in to 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.

4 Enter other filter options as desired.

5 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, as shown here:

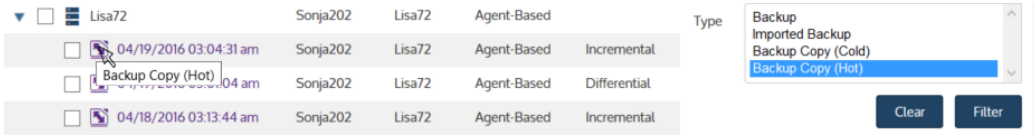

- If your target appliance is also being used as a backup appliance and its local backups are being copied to a hot backup copy target, the catalog lists both the hot backup copies stored on this appliance and any backups that were copied from this appliance to the hot backup copy target. (The hot backup copy target could be another appliance or the Unitrends Cloud).
  - You can delete the backup copies that are stored on this appliance only.
  - 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 deleted. If the backup copy is labeled *Backup Copy (Hot) on Target*, you must log in to its target appliance to delete the copy.

6 Check boxes to select the backup copies to delete.

7 Click **Delete**.

8 In the Confirm Backup Deletion dialog, review the list of backups that will be deleted.

9 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.

Chapter 5: Host-level Backups Overview

This chapter provides information for implementing host-level protection of VMware, Hyper-V, 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 281. For a general overview of Unitrends protection, see the "Protection Overview" on page 35 chapter.

Note: If you install a Unitrends agent on your VM and use asset-level protection, the Unitrends appliance treats the VM as a physical asset. See "Asset-level Backups Overview" on page 223 for information on protecting VMs with the agent.

Review the details in this chapter 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 111 to add your virtual host to the appliance
- "Backup Administration and Procedures" on page 173 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" on page 203.

Step 2: Install the Unitrends Windows agent on your Hyper-V host. See "Installing the Windows agent" on page 136.

Step 3: Add the Hyper-V host to your Unitrends appliance. See "Adding a virtual host" on page 126.

Step 4: Create backup jobs for your VMs. See "Creating backup jobs" on page 173.

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 203
- "General Hyper-V requirements" on page 204
- "Additional Hyper-V requirements" on page 205

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.
- To protect the file system and operating system of the Hyper-V host, you must run asset-level backups. For details, see "Asset-level Backups Overview" on page 223. Any files belonging to the Hyper-V application are automatically excluded from asset-level backups of the Hyper-V host.
- After making any configuration changes 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 VHD(X) format, you must run a new full backup to ensure the integrity of the VM's backup groups. After running a new full back up, you can continue protecting the VM with its existing schedule.
- 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 asset-level backups.
- In some cases, you may want or need to protect VMs using asset-level backups. For recommendations, see "Protecting Hyper-V virtual machines at the asset level" on page 207.
- To protect a VM with both host-level and asset-level (agent-based) backups, ensure that the VM's host-level and asset-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.

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 136. (It is not necessary to install agents on your virtual machines.) • For cluster configurations, be sure to install the same agent version on all hosts 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.

Item	Description
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: • 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/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 adhere to the following: • Generation 1 and 2 VMs are supported. Generation 2 VMs require Windows Server Hyper-V 2012 R2 as the recovery target. • The Hyper-V VM cannot be configured with pass-through disks. To protect a VM with pass-through disks, use asset-level protection instead.

Additional Hyper-V requirements

These additional requirements may apply to your environment.

Item	Description
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, it is recommended to install the agent on the VM and protect it as you would a physical server (leveraging Microsoft's VSS writers).

Item	Description
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, it is recommended to install the agent on the VM and protect it as you would a physical server (leveraging Microsoft's VSS writers).
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 KB 1334. • 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 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 KB 1335. • 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.

Item	Description
Faster incrementals on 2012 and 2012 R2	Unitrends leverages a Hyper-V changed-block-tracking (CBT) driver that greatly increases incremental backup performance on 2012/2012 R2 hosts. To use this driver, simply install Windows agent version 8.1.0-3 or higher on your Hyper-V hosts. The CBT driver is automatically installed with the Windows agent. To verify that the 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 136 for additional requirements.

Protecting Hyper-V virtual machines at the asset level

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 asset-level backups. Host- and asset-level backups provide you with different options.

Use the following topics to determine whether to run host- or asset-level backups of Hyper-V virtual machines:

- "General features of Hyper-V host-level and asset-level protection" on page 207
- "Asset-level protection examples" on page 208

General features of Hyper-V host-level and asset-level protection

General features of Hyper-V host-level and asset-level protection are given here:

Hyper-V protection strategy	Considerations
Host-level backups	• Quickest setup, do not need to add VMs individually or install the agent on each VM. • Automatically include new VMs in backup 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 296.

Hyper-V protection strategy	Considerations
Asset-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 restore. • For 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 backups for all recovery models other than *simple*. • Support backup of multi-node SharePoint farms. • Protect VMs configured with shared VHD(X)s. • Support Windows instant recovery (WIR) to quickly spin up a virtual replica of a failed Windows asset.

Asset-level protection examples

Specific instances when you might want to protect VMs at the asset level are described below. For instructions on setting up asset-level protection, see "Protected assets" on page 111.

Note: To protect a VM with both host-level and asset-level (agent-based) backups, ensure that the VM's host-level and asset-level jobs do not overlap. Running both simultaneously may lead to undesirable results.

VM configuration	Protection considerations
Hosted applications	
Hosted applications for which you need more granular control.	Use asset-level application backups to select individual databases to back up and recover.
Application versions that are not supported by Integration Services	Use asset-level application backups to protect the databases. Use asset-level protection for the VMs' file and operating systems.
Exchange	Use either host-level or asset-level protection.

VM configuration	Protection considerations
SQL	For *simple* recovery model databases, use either host-level or asset-level protection. For *full* or *bulk-logged* recovery model databases, use asset-level protection. (Host-level protection is not supported.)
SharePoint	Use asset-level protection. (Host-level protection is not supported.)
Oracle	Use asset-level protection. (Host-level protection is not supported.)
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 asset-level protection instead.
VMs in a cluster configuration with only one cluster shared volume	Unitrends recommends that you use asset-level protection.
VMs for which you would like to exclude volumes or large numbers of files when running backups	Use asset-level protections and exclude files from backups.
VMs functioning as large file servers for which you may need to frequently recover files.	Use asset-level protections so you can search for files to recover by name.
Windows VMs that you would like to protect with the Windows instant recovery (WIR) feature	Use asset-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 136. If the VMs you want to protect reside on Cluster Shared Volumes (CSVs), 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 (*DocCluster*) 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.
- Hover over the virtual machine name to see whether Integration Services are enabled on the VM. A message displays if they are not enabled. If they are enabled, hovering over the VM displays the VM name. (For VMs in a cluster setup, this may take a few minutes.) If Integration Services are not enabled, the VM is put in a saved state during the backup. See "General Hyper-V requirements" on page 204 for more information.
- For VMs hosted on Windows Server versions later than 2008 R2, backups are executed simultaneously. The number of jobs that run simultaneously 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:

- Clustered VMs display when you select their host node. However, to protect them, you must select the cluster itself (*DocCluster*). You cannot protect these VMs by selecting the owner node.
- Non-clustered VMs hosted on a cluster node do not display when you select the cluster (*DocCluster*). To protect these VMs, you must select the host node.
- To protect virtual machines hosted on a cluster node but that do not reside on CSVs, you must create a backup schedule for the node that hosts the VMs. You cannot protect them in the same schedule as the clustered VMs.
- If multiple virtual machines in a clustered environment (CSV) 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 ESX hosts directly or through a vCenter server. You can add ESX 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 214).

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" on page 211.

Step 2: Add the VMware host to your Unitrends appliance. See "Adding a virtual host" on page 126.

Step 3: Create backup jobs for your VMs. See "Creating backup jobs" on page 173.

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" on page 211
- "VMware virtual machine best practices and considerations" on page 212
- "General VMware requirements" on page 213
- "Additional VMware requirements" on page 214

VMware host best practices and considerations

You can add the following VMware servers to the Unitrends appliance to protect hosted virtual machines:

Item	Description
vCenter and managed ESX servers	If ESX servers belong to a vCenter and both are accessible on the network, Unitrends recommends that you add the vCenter and its ESX servers to the appliance. This enables the appliance to contact the vCenter for management operations (including vMotion support) and to directly contact the ESX servers for backup and recovery, potentially improving performance by reducing network traffic around the vCenter server.

Item	Description
vCenter only	If the ESX servers are accessible through a vCenter, adding the vCenter to the Unitrends appliance automatically detects all of the associated ESX 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 ESX servers. In this case, the appliance detects when VMs move between ESX servers in a cluster and contacts the appropriate server to perform backups.
ESX server only	If ESX servers are not accessible through a vCenter, or if only a subset of the VMs hosted on the vCenter's ESX servers are to be protected, you can add indvidual ESX servers. In this case, the appliance contacts the ESX 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 ESX 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 ESX 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 "General VMware requirements" on page 213 for details.
- In some cases, you may want or need to protect VMs by using asset-level backups. For recommendations, see "Protecting VMware virtual machines at the asset level" on page 216.
- To protect a VM with both host-level and asset-level (agent-based) backups, be sure to adhere to the following:
  - Ensure that the VM's host-level and asset-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 "Application-aware protection" on page 219 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 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 KB 3053.
- Host-level protection is not supported for the following (use asset-level backups instead):
  - VMs in a cluster configuration with a fault tolerant disk.
  - VMs with dynamic MAC addresses.
  - Independent and pass-though 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 hosted on Free ESXi versions.

General VMware requirements

The following requirements must be met for host-level protection of VMware virtual machines.

Item	Description
ESX host	Must be a licensed version listed in the Unitrends Compatibility and Interoperability Matrix. **Note:** Additional requirements and limitations apply to ESXi 6 and 6.5. See "Additional VMware requirements" on page 214 for details.
vCenter	Must be a licensed version listed in the Unitrends Compatibility and Interoperability Matrix. **Note:** Additional requirements and limitations apply to vCenter 6 and 6.5. See "Additional VMware requirements" on page 214 for details.
vCenter or ESX 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 ESX server to the backup appliance.

Item	Description
Virtual machine configuration	Verify the following VM configuration settings: • VM hardware version must be 4, 7, 8, 9,10, or 11. • VMware tools must be installed in the guest operating system to ensure file system and application consistency. • The VM must not be configured to use VM encryption. **Note:** **On Vsphere 5.5**. VMware introduced new SATA Virtual Hardware Controllers with vSphere 5.5 and VM Hardware Version 10. See KB 1102 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. **Note:** **On hardware version 4.** Only full backups are supported for VMs configured with hardware version 4.

Additional VMware requirements

These additional requirements may apply to your environment.

Item	Description
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.

Item	Description
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 ESX 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 appliances	For ESX 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 scheduling flexibility as the production network is not used during the backup. See the Unitrends Knowledge Base for requirements and setup procedures.
HotAdd backup for Unitrends Backup on VMware appliances	For ESX 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 the Unitrends Knowledge Base for requirements and setup procedures.
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).

Item	Description
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 at the asset level

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 asset-level backups. Host- and asset-level backups provide you with different options.

Use the following topics to determine whether to run host- or asset-level backups of VMware virtual machines:

- "General features of VMware host-level and asset-level protection" on page 216
- "Asset-level protection examples" on page 217
- "Application-aware protection" on page 219

General features of VMware host-level and asset-level protection

General features of VMware host-level and asset-level protection are given here:

Notes:

To protect a VM with both host-level and asset-level (agent-based) backups, be sure to adhere to the following:

- Ensure that the VM's host-level and asset-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 agent on each VM. • Automatically include new VMs in backup 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 Unitrends Recovery Series 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 asset-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.
Asset-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. • For 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 backups for all recovery models other than *simple*. • Support backup of multi-node SharePoint farms. • Support Windows instant recovery (WIR) to quickly spin up a virtual replica of a failed Windows asset.

Asset-level protection examples

Specific instances when you might want to protect VMs at the asset level are described below. For

instructions on setting up asset-level protection, see "Protected assets" on page 111.

VM configuration	Protection considerations
Hosted applications	
Hosted applications for which you need more granular control.	Use asset-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 "Application-aware protection" on page 219 for details). • Use asset-level protection for more granular control.
SQL	For simple recovery model databases, do one of the following: • Use host-level protection with the application-aware feature (see "Application-aware protection" on page 219 for details). • Use asset-level protection for more granular control. For full or bulk-logged recovery model databases, use asset-level protection. (Host-level protection is not supported)
SharePoint	Use asset-level protection. (Host-level protection is not supported.)
Oracle	Use asset-level protection. (Host-level protection is not supported.)
Disk configuration	
Cluster with fault tolerant disks	Use asset-level protection. (Host-level protection is not supported.)
Physical RDM disks	Use asset-level protection. (These disks are automatically excluded from host-level backups.)
Independent or pass-through disks	Use asset-level protection. (These disks are automatically excluded from host-level backups.)

VM configuration	Protection considerations
Sparse disks	Use asset-level protection. (Host-level protection is not supported.)
Dynamic MAC address	Use asset-level protection. (Host-level protection is not supported.)
VMs hosted on free ESXi versions	Use asset-level protection. (Host-level protection is not supported.)
Virtualized Active Directory (AD) servers for which you are not following Microsoft's best practices	Use asset-level protection.
VMs in Distributed File System environments for which you are not following Microsoft's best practices	Use asset-level protection.
VMs for which you would like to exclude volumes or large numbers of files when running backups	Use asset-level protections and exclude files from backups. (With host-level you can exclude virtual disks only. Asset-level provides more granular control.)
VMs functioning as large file servers for which you may need to frequently recover files	Use asset-level protections so you can search for files to recover by name.
Windows VMs that you would like to protect with the Windows instant recovery (WIR) feature	Use asset-level protection.

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. Once credentials have been applied, 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 truncated with VMware full and incremental backup.

Application	Log file truncation with VMware application-aware backup
SQL	Logs not truncated with VMware application-aware backup. Do the following: • Simple recovery model - no logs 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. See "Recommendations for bulk-logged recovery model" on page 249 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.

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 220.
2. Add the XenServer host to your Unitrends appliance. See "Adding a virtual host" on page 126.
3. Create backup jobs for your VMs. See "Creating backup jobs" on page 173.

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 220
- "XenServer virtual machine best practices and considerations" on page 221

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 "Asset-level Backups Overview" on page 223.

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 127.) 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 on the Add Virtual Host page. 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 on the Add Virtual Host page.

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 asset-level backups. If you choose to protect a VM with both host-level and asset-level (agent-based) backups, ensure that the VM's host-level and asset-level jobs do not overlap. Running both simultaneously may lead to undesirable results.

Chapter 6: Asset-level Backups Overview

This chapter provides details and requirements for asset-level backups. An asset-level backup protects an asset's file system and operating system. You can select to include or exclude files from asset-level backups.

Asset-level backups protect physical assets. For virtual assets, you can choose host-level or asset-level protection. Host-level backups capture files, application data, and virtual hardware. With asset-level protection, the appliance treats the VM as a physical asset to run asset-level and application backups. For more information on determining which backup type to use for a VM, see "Protecting Hyper-V virtual machines at the asset level" on page 207, "Protecting VMware virtual machines at the asset level" on page 216, or "Best practices and requirements for XenServer protection" on page 220.

Asset-level protection requires installing a Unitrends agent on the asset. Agent installation procedures vary by operating system. For detailed information on installing the agent on various operating systems, see "Installing the Unitrends agent" on page 111. After installing the required agent, add the asset to the appliance as described in "To add an asset" on page 119, then proceed to "Backup Administration and Procedures" on page 173 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 261 instead.

Considerations for asset-level backups

Consider the following information when preparing for asset-level backups:

- "Maximum file pathname lengths" on page 223
- "Default exclusions from asset-level backups of Windows servers" on page 224
- "Default exclusions from file-level backups of Linux servers " on page 224
- "Exclude active databases from asset-level backups" on page 225
- "Mac OS X Sleep Mode" on page 225

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 asset-level backups of Windows servers

By default, asset-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.

By default, 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 KB 2781.

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*

Exclude active databases from asset-level backups

Unitrends recommends excluding active databases from asset-level backups. Run application-level backups to protect active databases. For information on application backups, see "Application Backups Overview" on page 237.

Note: SQL files for system databases (such as master, model, and msdb) are always included to support the Windows instant recovery (WIR) feature. Do not exclude these if you will be using WIR for hosted SQL databases.

Some active databases are automatically excluded with asset-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 asset-level backups" on page 239.
- 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 asset-level backups" on page 249.

Mac OS X Sleep Mode

Either disable or configure the Mac OS X Sleep Mode to accommodate scheduled backups.

Chapter 7: NAS Backups Overview

This chapter 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 227 to determine which approach best suits your needs.
- " NAS protection using CIFS/NFS" on page 229 for additional CIFS and NFS requirements and considerations.
- " NAS protection using NDMP" on page 230 for additional NDMP requirements and considerations.
- "Start protecting the NAS asset" on page 235 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, differential, and selective backup modes.
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 asset-level or by volume.	Backup copy at the asset-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.) • 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.
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.
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. • Your NDMP NAS device must be configured with an MD5 password. Clear text passwords are not supported.
Network requirements	The following network requirements apply: • 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 66.

Feature	Description
Backup and recovery	See these topics for backup and recovery requirements and considerations: • "All jobs" on page 231 • "Backup jobs" on page 231 • "NetApp cluster protection" on page 232 • "Advanced configuration settings" on page 233 • "Recovery jobs" on page 234
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 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 234. • 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.

Feature	Description
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 KB 3662 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.

Feature	Description
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 *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 234. (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 234. • 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 234. • Maximum Running NDMP Jobs - The maximum number of running NDMP sessions per NAS NDMP asset defaults to 2. This value is accessible only from terminal. For more information, see KB 1313.

Feature	Description
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. For example, VNXe devices only support full volume-level recovery to the original location. 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 seleced 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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)
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 121 or "To add a NAS NDMP asset" on page 122.

Step 2: Run backup jobs as described in "To create a backup job for NAS CIFS or NFS assets" on page 180 or "To create a backup job for NAS NDMP assets" on page 181.

Chapter 8: Application Backups Overview

This chapter provides detailed information for protecting applications. Application backups only protect applications. To protect the application's host server, you must create separate asset-level backups to capture the server's file and operating systems.

Once you have reviewed the information in this chapter, install the required agent, add the application server to the Unitrends appliance, then proceed to "Backup Administration and Procedures" on page 173 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 2014, 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 239 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 asset-level backups" on page 239

Automatic exclusion of application data during asset-level backups

When you run asset-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

Additional requirements apply for Exchange CCR, SCR, and DAG environments. See the Unitrends Knowledge Base for details.

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 136.

Step 2: Add the Exchange server to the Unitrends appliance as described in "To add an asset" on page 119.

Step 3: Run backup jobs as described in "To create an Exchange backup job" on page 182.

SQL backup requirements and considerations

Review the following before implementing your SQL protection strategy:

Supported SQL features

Unitrends supports protection of the following SQL features:

- SQL system and user databases - To protect these databases, the only requirements are the ones described in "Agent prerequisites for Microsoft SQL" on page 241 and "SQL system requirements" on page 242. (Additional requirements apply if these databases are Always Encrypted, Stretch, or have disk storage on an SMB 3.0 share, as described below).

- SQL clusters - Unitrends supports protection of a variety of SQL cluster configurations, including cluster volumes, clustered shared volumes, AlwaysOn clusters, and failover clusters. To protect SQL clusters, see "Requirements and considerations" on page 241 for Windows agent, SQL system, and cluster requirements.
- Databases with disk storage on SMB 3.0 shares - To protect these databases, see "Requirements and considerations" on page 241 for Windows agent, SQL system, and SMB 3.0 share requirements.
- Always Encrypted databases - To protect these databases, see "Requirements and considerations" on page 241 for Windows agent, SQL system, and Always Encrypted database requirements.
- Stretch databases - Unitrends backups capture the data on the local SQL server only (and do not include any data in the Azure database). To protect these databases, see "Requirements and considerations" on page 241 for Windows agent, SQL system, and Stretch database requirements.

Requirements and considerations

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 and system requirements described below apply to all SQL protection. If you are protecting SQL clusters, data on SMB 3.0 shares, Always Encrypted databases, or Stretch databases, additional requirements apply. See the following topics for details:

Agent prerequisites 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 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.
- Additional agent version requirements apply to specific SQL features. For details, see the feature requirements sections.

SQL system requirements

In addition to the agent requirements:

- 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 operating system listed in the Unitrends Compatibility and Interoperability Matrix.
- The SQL application and server must be set up in a supported Microsoft deployment configuration.

Additional system requirements for SQL clusters and SMB 3.0

In addition to the agent and system requirements, the following are required to protect SQL clusters or data residing on SMB 3.0 shares:

- The Unitrends appliance and Windows agent must be running release 8.1 or higher.
- For clusters, all nodes in the SQL cluster must be running Windows agent release 8.1 or higher.
- For the applicable clustered or SMB 3.0 setup, SQL, and Windows versions must also meet the requirements in the following table:

Configuration	SQL	Windows
Cluster Volume	2005	2003 SP1, 2008, 2008 R2, 2012, 2012 R2
	2008	2008, 2008 R2, 2012, 2012 R2
	2012	2012, 2012 R2

Configuration	SQL	Windows
	2014	2012, 2012 R2
	2016	2012, 2012 R2
Clustered Shared Volume	2012	2012, 2012 R2
	2014	2012, 2012 R2
	2016	2012, 2012 R2
AlwaysOn Clusters	2012	2012, 2012 R2
	2014	2012, 2012 R2
	2016	2012, 2012 R2
SQL Failover Clusters	2014	2012, 2012 R2
	2016	2012, 2012 R2
SMB 3.0	2012	2012, 2012 R2
	2014	2012, 2012 R2
	2016	2012, 2012 R2

SQL cluster requirements and considerations

Consider the following before executing backups for databases hosted on servers configured in a cluster:

- You must add the cluster and each node in the cluster to the backup appliance, each as a separate asset. For details, see "Managing application assets" on page 123.
- When adding the cluster to the backup appliance, use the IP address of the clustered SQL server instance. This is the virtual IP address used to connect to the SQL server.
- There must only be one cluster IP address configured for each clustered SQL instance.
- To protect databases residing on cluster shared volumes (CSVs), you must select the cluster when creating backup jobs. You cannot protect these databases by selecting the owner node.
- To protect databases that are hosted on a cluster node but that do not reside on CSVs, you must create a backup schedule for the node that hosts those databases. You cannot protect them in the same schedule as the clustered databases.
- When backing up SQL nodes, include all local volumes, and exclude the system state.

- Additional considerations apply to SQL AlwaysOn Failover Cluster Instances. See "Considerations for SQL AlwaysOn Failover Cluster Instances (FCI)" for details.

Considerations for SQL AlwaysOn Failover Cluster Instances (FCI)

SQL Failover Clustering is a High Availability (HA) and Disaster Recovery solution. High Availability means that if one of the nodes in a SQL failover cluster fails, the secondary node is automatically promoted to the primary (active) node. There are some circumstances where a manual restart of the new primary database is required. See the following SQL documentation for details: Failover Policy for Failover Cluster Instances.

Unitrends provides seamless protection of your SQL environments in the event of a failover. Because the job schedule is attached to the cluster and not the individual nodes, the backups can continue as planned, providing uninterrupted protection of your SQL instance. For details about the SQL side of failover, see AlwaysOn Failover Cluster Instances (FCI). If your backups begin failing after a failover, see the following SQL documentation, Failover Cluster Troubleshooting.

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.

For Microsoft SQL recommendations, see the Recommendations section of AlwaysOn Failover Cluster Instances (FCI).

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 KB 1334.
- 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 244.
- 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 KB 1335.

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 173.

Requirements for SQL Always Encrypted databases

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.

Item	Always Encrypted database requirement or consideration
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 on the local SQL server with a paired database on 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 agent version	The SQL server must be running Windows agent release 9.0.0-13 or higher.
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.

Item	Stretch database requirement or consideration
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 to the Unitrends appliance, 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 Server recovery model considerations

The recovery model of your SQL databases 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 249 for details.
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 249 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 restoring 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 restoring 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 restored.
tempdb	A temporary workspace used by any session connected to the SQL Server instance and is used to hold intermediate or temporary data. For example, temporary tables, cursors, and data for sorting.	Every time SQL Server starts, this database is re-created. There is no reason to preserve this database by backing up or restoring.
distribution	Stores metadata and history data in support of SQL Server replication.	Present only if replication is configured.

Example SQL Server 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

Database	Backup Strategy
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 asset-level backups

When you run asset-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 the Windows instant recovery feature.

- Files in SQL *database/log* directories are excluded.

Start protecting SQL

After ensuring all requirements have been met, do the following to start protecting your SQL environment:

Step 1: Install the Windows agent on the SQL server as described in "Installing the Windows agent" on page 136.

Step 2: Add the SQL server to the Unitrends appliance as described in "To add an asset" on page 119.

Step 3: Run backup jobs as described in "To create a SQL backup job" on page 184.

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 asset-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 252.

Consider the following before implementing your SharePoint protection strategy:

- "SharePoint agent requirements" on page 251

- "SharePoint configuration prerequisites" on page 252

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 252.

- 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.
- Prerequisite configuration steps must be performed on the primary node, as described in "SharePoint configuration prerequisites" on page 252.
- 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 129. 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: KB 3061, KB 3067, KB 3066, KB 3058, KB 1147, and KB 3076.

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 253.

  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 253.
- 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 252.
- If you are unsure of the installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 252.

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 253. To check your SharePoint installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 252.

1. Install the Unitrends Windows agent on the SharePoint server as described in "Manually installing the Windows agent" on page 138.
2. Add the SharePoint server to the Unitrends appliance as described in "To add an asset" on page 119.
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 184 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 252. To check your SharePoint installation type, see "To determine the installation type for SharePoint 2013 and 2010 deployments" on page 252.

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 138.

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 asset" on page 119) and apply administrative trust credentials to the *Full Farm* database instance.

6 Proceed to "To create a SharePoint backup job" on page 184 to start protecting your SharePoint environment.

Oracle backup requirements and considerations

Use application backups to protect Oracle Database 11g and 12c on Windows, Linux, and Solaris platforms, and Oracle 10g on Windows platforms. Application backups ensure database consistency, whereas asset-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 254
- "Guidelines for creating Oracle credentials" on page 256
- "Start protecting Oracle" on page 257
- "Upgrading to newer Oracle versions" on page 257

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 111.
Oracle server	Add the Oracle server to the Unitrends appliance as described in "To add an asset" on page 119.
Credentials	Configure trust credentials for each application instance you wish to protect as described in "Managing asset credentials" on page 129. • 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 256.

Oracle Requirement	Description
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. • Each Oracle SID an Oracle server must be unique. • 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 12c, 11g, or 10g.
Oracle on Linux	• Version must be 12c or 11g. • Must install the Oracle Dependency as described in "Installing and updating the Linux agent" on page 144.
Oracle on Solaris	• Must be version 12c or 11g. • A Samba client for Solaris must be enabled. See KB 1303 for details. • Ensure the Solaris client has sufficient memory available. See KB 3169 for details. • Full pathname to each Solaris object cannot exceed 1024 characters. For details, see KB 3348.

Oracle Requirement	Description
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. • 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 KB 3358.

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 129 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 platform	Guidelines and requirements
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 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 136, "Installing and updating the Linux agent" on page 144, or "Installing and updating the Solaris agent" on page 158.

Step 2: Add the Oracle server to the Unitrends appliance as described in "To add an asset" on page 119.

Step 3: Run backup jobs as described in "To create an Oracle backup job" on page 183.

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 253.
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 173.

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 42 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

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 363. 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 asset-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 124.

Step 2: Run backup jobs as described in "To create a UCS service profile backup job" on page 185.

Chapter 9: iSeries Backups Overview and Procedures

This chapter 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 261 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 265.

Step 3: Run backup jobs as described in "Creating iSeries backup jobs" on page 267.

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.

Item	Requirement or consideration
Backup and recovery jobs	The following apply to backup and 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.
FTP	The FTP server must be configured and running on the iSeries. The FTP protocol is used by backup and recovery jobs.

Item	Requirement or consideration
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.
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.

Item	Requirement or consideration
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 always be the first object in the backup file. If it is recovered (which is via the **RSTUSRPRF** command), then a **RSTAUT** command will be executed after everything else is 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 will display before any other objects, except /Security Data, in the backup file. It is recovered using the **RSTCFG** command.

Item	Requirement or consideration
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)`

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" on page 265
- "To edit an iSeries asset" on page 266
- "To edit retention settings for an iSeries asset" on page 266
- "Removing an iSeries asset" on page 267

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 66.
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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)
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.
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 267.

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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)
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.
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 196 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 265. 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 267
- "To run an iSeries backup on-demand" on page 268

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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)

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.

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 77 to add a backup copy target to your backup appliance.
- "Creating backup copy jobs" on page 186 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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)

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.

Chapter 10: 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. To meet low RTOs, recover from local backups on the Unitrends appliance. For details on recovering from a given local backup, see the Recovery chapter for the applicable asset type:

- "Recovering Host-level Backups" on page 281
- "Recovering Asset-level Backups" on page 305
- "Recovering Application Backups" on page 341
- "Recovering NAS Backups" on page 335
- "Recovering iSeries Backups" on page 365

If a local backup is not available, you can recover from a backup copy as described in "Recovering Backup Copies" on page 271.

Chapter 11: Recovering Backup Copies

If a local backup is not available, you can recover from a backup copy. The same recovery operations that are used for local backups are supported for backup copies, but recovering from backup copies 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 276.
Third-party cloud (Amazon, 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 in the Cloud or on a target appliance?" on page 272
- "Requirements and limitations for recovering hot copies by using the source backup appliance" on page 272

Recovery procedures:

- "To import a hot backup copy" on page 275
- "To recover files from an asset-level backup copy by using the File Browser" on page 314
- "To recover files from an asset-level backup copy by using Search Files" on page 315
- "Recovering files from virtual machine backups" on page 283

How do I use my backup appliance to recover from hot copies that reside in the Cloud or on a target appliance?

To recover hot copies, you either import the backup copy to the source appliance or recover files directly from the hot copy in the Cloud or on the appliance target.

Supported recovery operations vary by backup type:

- For all backup types - 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 275. Once the backup copy has been imported, recover from it just like you would from any other local backup.
- For asset-level backup copies - Use the following additional options to recover files directly from the copy in the Cloud or on the appliance target. 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 the "To recover files from an asset-level backup copy by using the File Browser" on page 314 procedure.
  - Search for files in asset-level backup copies and download selected file(s) in a *.zip* file. For details, see the "To recover files from an asset-level backup copy by using Search Files" on page 315 procedure.
- For host-level backups of Windows or Linux VMs - 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 the "Recovering files from virtual machine backups" on page 283 procedure.

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. • Must be the source appliance that copied the backup to the 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.
Recovering files directly from the Cloud or appliance target	You can recover files directly from backup copies that reside in the 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" on page 273 • "Small downloads" on page 273 • "Large downloads" on page 274
Supported backup types	The backup that was copied must be one of the following: • An asset-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. • A VMware, Hyper-V, or XenServer host-level backup of a Windows or Linux VM.
Small downloads	For downloads that are 500MB or smaller, 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. **Note:** 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.

Requirement or limitation	Description
Large downloads	For downloads that are greater than 500MB, a *.zip* file is created in both of these locations: • 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. **Note:** 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, access the *.zip* file in the source appliance's */downloads* directory instead. • In the source appliance's */downloads* directory. Access the recovered files by entering *<SourceApplianceIP>/downloads* in an Internet browser. Do not download these files until you see the *Unitrends-Restore.zip* file. While the recovery is in progress, you see files in this directory, but the download is not complete until the *.zip* file has been created. (Recoveries are automatically removed from the */downloads* directory after 72 hours.) **Note:** The CIFS protocol is used to transfer files to the source appliance's */downloads* directory. If you have added your own target appliance and did NOT configure an OpenVPN tunnel connection, the outbound connection from the target to the source must be open for CIFS. If you are using OpenVPN, firewall protocol blocking does not apply.
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 Cloud or on a target appliance.
4. Enter other filter options as desired.
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 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 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 you wish to import.
7. Click **Import to Source**.

   The Import Backup Copies dialog displays. This 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 42.
8. Click **Import**.

   The message *Starting Import* displays.
9. Click **OK** to continue.

   Selected backup copies are imported to the appliance.
10. To verify that the import is complete, check for the backup copies in the Backup Catalog:
    - 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.

11 (Optional) Recover from the imported backup copy as you would from any regular backup.

Note: 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.

See these topics for recovery options and instructions:

- "Recovering Host-level Backups" on page 281
- "Recovering Asset-level Backups" on page 305
- "Recovering Application Backups" on page 341
- "Recovering NAS Backups" on page 335
- "Recovering iSeries Backups" on page 365

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 118 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 281
- "Recovering Asset-level Backups" on page 305
- "Recovering Application Backups" on page 341
- "Recovering NAS Backups" on page 335

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 275 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 Enter other filter options as desired.

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, as shown here:

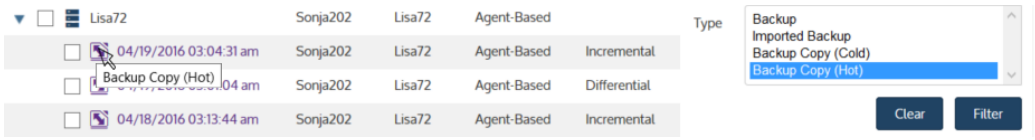

- If your target appliance is also being used as a backup appliance and its local backups are being copied to a hot backup copy target, the catalog lists both the hot backup copies stored on this appliance and any backups that were copied from this appliance to the hot backup copy target. (The hot backup 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 271.

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 asset-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 an asset-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 308
  - "Recover files from one cold backup copy by using the File Browser" on page 310.

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. Enter other filter options as desired.
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 displays. This 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 42.

9 Click **Import**.

To monitor the import's progress, click **View Jobs**.

10 Once the import completes, recover from the imported backup copy as you would from any regular backup.

> **Note:** 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.

To access the imported backup copy:

- In the Backup Catalog, be sure that the Imported Backup Type is selected in the filter options.
- Imported cold backup copies are light blue and the description *Imported Backup* displays when you hover over the backup copy icon, as shown here:

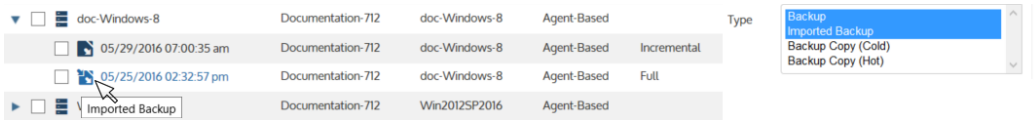

- See the following for recovery options and instructions:
  - "Recovering Host-level Backups" on page 281
  - "Recovering Asset-level Backups" on page 305
  - "Recovering Application Backups" on page 341
  - "Recovering NAS Backups" on page 335
  - "Recovering iSeries Backups" on page 365

Chapter 12: Recovering Host-level Backups

Host-level backups have the following recovery options:

- VM recovery – recover VMs that run any operating system.
- File-level recovery – recover files on VMs that run Windows or Linux.
- Instant recovery – recover VMware or Hyper-V VMs immediately.

The type of recovery depends on the type of data you want to recover. Review the following details to determine the best option:

- Use file-level recovery to recover only files from a backup or backup copy for VMs running Windows or Linux.
- For faster recovery of an entire VMware or Hyper-V VM, set up instant recovery.

  Instant recovery uses system resources that can impact the performance of other jobs. For non-critical VMs that you do not need recover immediately, Unitrends recommends a VM recovery.
- Recovery time depends on factors such as the amount of data on the VM and other tasks running on the appliance.
- If a local backup is not available, recover from a backup copy:
  - 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 278. 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 275. 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 283.
- For details on recovering a VMware template, see the Unitrends Knowledge Base.

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 imported backup copy to recreate the VM on the recovery target. The recovery target can be the original host or an alternate host running the same software version as the original or a later version.

Select a full backup, an incremental backup, or an imported backup copy for the recovery. With an incremental backup, the appliance uses all previous backups from the same backup group to recreate the VM.

Preparing to recover a virtual machine

Unitrends supports recovery of VMware, Hyper-V, 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 126 before recovering a VM.

About recovering VMware VMs

Review the following information on recovering VMware VMs:

- A recovered VM is configured with the latest hardware version supported by the target ESXi host.
- VMs with hardware version 10 can be recovered to ESXi 5.5 or to a higher version listed in the Compatibility and Interoperability Matrix.
- VMs with hardware version 11 can be recovered to ESXi 6 or to a higher version listed in the Compatibility and Interoperability Matrix.
- VMs with hardware version 13 can be recovered to ESXi 6.5 or to a higher version listed in the Compatibility and Interoperability Matrix.
- 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 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 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.
- A recovered VM is configured with the latest hardware version supported by the target Hyper-V host.
- 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.

About recovering XenServer VMs

Review the following information on recovering XenServer VMs:

- Unitrends recommends recovering to a XenServer host that is running the same version as the original host, or to a higher supported version listed in the Compatibility and Interoperability Matrix.
- 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

Use this procedure to recover an entire virtual machine.

1. Select **Recover** and click the **Backup Catalog** tab.

Use Filter Backups to the right to customize the backups that display.

2 Expand the desired asset and select a backup or imported backup copy to use for the recovery.

To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.

3 Click **Recover**.

4 Select from the following recovery options:

Recovery Options	Description
Target Location	Select the host where the VM will be recovered.
Target Resource	Select a resource pool. This field displays only if the ESXi host selected for the target location has resource pools. It does not display when recovering to a Hyper-V or XenServer host.
Target Storage	Select a datastore (ESXi host), a volume (Hyper-V host) or a Storage Repository (XenServer host).

5 Click **Next**.

A summary of the selected recovery options display.

6 (Optional) Modify the VM Name by clicking it in the Assets to Recover list and entering a new name. This is supported for VMware and XenServer VMs only.

7 Click **Save**. The job is queued immediately.

Notes:

- Recovery consists of two tasks. The first recovers the configuration files or metadata for the VM, and the second restores the data.
- To view job progress, click **Jobs > Active Jobs**.
- The recovered VM is created in a powered off state.
- After recovering a virtual machine, the next backup of the recovered VM is promoted to a full.

8 After the recovery job completes, go to the hypervisor and power on the recovered virtual machine.

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, 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:

- 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.
- 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.

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.

See these topics for more on supported recovery methods:

- "Recovery procedures overview" for a description of procedures by backup type and location.
- "Windows prerequisites and considerations" on page 285 to determine the file recovery options for your Windows VM.
- "Linux prerequisites and considerations" on page 286 to determine the file recovery options for your Linux VM.

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 281.

Recovery procedures overview

Use the "Windows file-level recovery" on page 287 and "Linux file-level recovery" on page 292 procedures to recover files. The procedure you use to create the recovery object is determined by 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 287 or "Linux file-level recovery" on page 292 procedure from the backup appliance to recover files from a local backup or from an imported backup copy. • For Windows - In "Step 1: 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 1: 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 278 or "To import a hot backup copy" on page 275.
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 287 or "Linux file-level recovery" on page 292 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 1: 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 1: 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 287 or "Linux file-level recovery" on page 292 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 1: 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 1: Create the recovery object", use this option: "To create the recovery object and recover from a hot backup copy by using the target appliance".

Windows prerequisites and considerations

The following requirements and considerations apply to recovering files from a host-level backup or

host-level backup copy of a Windows VM:

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 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 286.)
Recovery target requirements	If you are recovering by mounting the CIFS share or iSCSI LUN on a target machine, the target machine must meet these requirements: • The recovery target can be the original VM, a different VM running the same operating system, or a physical machine running the same operating system. • 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 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 must perform a VM recovery instead.

Linux prerequisites and considerations

The following requirements and considerations apply to recovering files from a host-level backup or host-level backup copy of a Linux VM:

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, select files to recover. • Mount the iSCSI LUN on a recovery target machine. From the target machine, select files to recover.
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 281. • 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".
Recovery target requirements	If you are recovering by mounting the CIFS share or iSCSI LUN on a target machine, the recovery target can be the original VM, a different VM running the same operating system, or a physical machine running the same operating system.
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 must recover the entire VM instead.

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. Before you start, be sure all requirements in "Windows prerequisites and considerations" on page 285 have been met.

- "Step 1: Create the recovery object"
- "Step 2: Recover files"
- "Step 3: Remove the recovery object from the appliance"

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 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 288, 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 288, 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 289, 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.
2. Select **Recover** and click the **Backup Catalog** tab.

   Use Filter Backups to the right to customize the backups that display.
3. Expand the desired 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 278 or "To import a hot backup copy" on page 275.)
4. Click **Recover Files**.
5. Follow the prompts to create the file recovery object.

   To view the recovery object, select **Recover** and click the **File Level Recovery** tab.

   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 using the host that manages it.

Proceed to "Step 2: Recover files".

To create the recovery object and recover from a hot backup copy by using the source backup 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)**.
- Select other filter options as desired.
- Click **Filter**.

4 Expand the desired asset and select the hot backup copy from which you want to recover files.

5 Click **Recover Files**.

6 Follow the prompts to create the file recovery object.

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.

- The recovery job starts and the recovery object is created in the Cloud or on the target appliance.
- In the Notice message, click **View FLR** to view the recovery object on the **File Level Recovery** tab. The recovery object Name is *AssetName (on the target)*.

Note: Recovery objects created in the Unitrends Cloud are automatically removed after 96 hours.

Proceed to "Step 2: Recover files" on page 290.

To create the recovery object and recover from a hot backup copy by using the 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)**.
- Select other filter options as desired.
- Click **Filter**.

4 Expand the desired asset and select the hot backup copy from which you want to recover files.

5 Click **Recover Files**.

6 Follow the prompts to create the file recovery object.

To view the recovery object, select **Recover** and click the **File Level Recovery** tab.

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 using the host that manages it.

Proceed to "Step 2: Recover files".

Step 2: 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 284.

- "To recover files by browsing and downloading to a .zip file" on page 290
- "To recover files by mounting the CIFS share" on page 290
- "To recover files by mounting the iSCSI LUN" on page 291

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.
4. Click **Download** then **Confirm** to download the selected files to a *.zip* file.
5. 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, do one of the following:

   - For downloads that are 500MB or smaller, you must run a new job.
   - For downloads that are greater then 500MB, access the recovered files in the source appliance's */downloads* directory by entering *<SourceApplianceIP>/downloads* in an Internet browser. Do not download these files until you see the *Unitrends-Restore.zip* file. While the recovery is in progress, you see files in this directory, but the download is not complete until the *.zip* file has been created. (Recoveries are automatically removed from the */downloads* directory after 72 hours.)

Proceed to "Step 3: Remove the recovery object from the appliance" on page 291.

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.

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.

7 Move selected files to the desired location on the recovery target.

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 291.

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 desired iSCSI target from the list.

The last part of the iSCSI identifier contains the VM name.

5 Click **Done**.

6 Use Windows Disk Management tools to assign drive letters and retrieve files.

> **Note:** 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 the desired location on the recovery target.

8 Use the iSCSI Initiator to disconnect from the LUN.

9 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**.

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. Before you start, be sure all requirements in "Linux prerequisites and considerations" on page 286 have been met.

- "Step 1: Create the recovery object"
- "Step 2: Recover files"
- "Step 3: Remove the recovery object from the appliance"

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 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 292, 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 293, 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 293, 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.
2. Select **Recover** and click the **Backup Catalog** tab.

   Use Filter Backups to the right to customize the backups that display.
3. Expand the desired 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 278 or "To import a hot backup copy" on page 275.)
4. Click **Recover Files**.
5. Follow the prompts to create the file recovery object.

   To view the recovery object, select **Recover** and click the **File Level Recovery** tab.

   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 using the host that manages it.

Proceed to "Step 2: Recover files" on page 294.

To create the recovery object and recover from a hot backup copy by using the source backup 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)**.
   - Select other filter options as desired.
   - Click **Filter**.
4. Expand the desired asset and select the hot backup copy from which you want to recover files.
5. Click **Recover Files**.
6. Follow the prompts to create the file recovery object.

   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.

   - The recovery job starts and the recovery object is created in the Cloud or on the target appliance.
   - In the Notice message, click **View FLR** to view the recovery object on the **File Level Recovery** tab. The recovery object Name is *AssetName (on the target)*.

   Note: Recovery objects created in the Unitrends Cloud are automatically removed after 96 hours.

Proceed to "Step 2: Recover files" on page 294.

To create the recovery object and recover from a hot backup copy by using the 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)**.
   - Select other filter options as desired.
   - Click **Filter**.
4. Expand the desired asset and select the hot backup copy from which you want to recover files.
5. Click **Recover Files**.
6. Follow the prompts to create the file recovery object.

   To view the recovery object, select **Recover** and click the **File Level Recovery** tab.

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 using the host that manages it.

Proceed to "Step 2: Recover files".

Step 2: Recover files

Use one of the following procedures to recover files. For a description of each method, see "Recovery procedures overview" on page 284.

- "To recover files by browsing and downloading to a .zip file" on page 294
- "To recover files to a Linux machine by mounting the iSCSI LUN" on page 294

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.
4. Click **Download** then **Confirm** to download the selected files to a *.zip* file.
5. 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, do one of the following:

- For downloads that are 500MB or smaller, you must run a new job.
- For downloads that are greater then 500MB, access the recovered files in the source appliance's */downloads* directory by entering *<SourceApplianceIP>/downloads* in an Internet browser. Do not download these files until you see the *Unitrends-Restore.zip* file. While the recovery is in progress, you see files in this directory, but the download is not complete until the *.zip* file has been created. (Recoveries are automatically removed from the */downloads* directory after 72 hours.)

Proceed to "Step 3: Remove the recovery object from the appliance" on page 295.

To recover files to a Linux machine by mounting the iSCSI LUN

Mount the iSCSI LUN to the target machine and copy the files.

1. Log in to the recovery target.
2. Enter the following command to change to the */tmp* directory:

```
# cd /tmp
```

3. Run the following command to copy the *iscsi_flr* script from the backup appliance:

```
# wget http://<appliance IP>/iscsi_flr
```

4. After the script downloads, add the execute permission:

```
# chmod +x iscsi_flr
```

5. Run the following command to mount the recovery object:

```
# ./iscsi_flr mount
```

6. Enter the appliance IP address.
7. Enter the full path of the mount point directory.

Notes:

- The full path is likely: /iscsi_flr.
- To view this path, select **Recover** and click **File Level Recovery** tab. Select the recovery target and click **Show Details**. The last part of the iSCSI details provides the path.

8. Move selected files to the desired location on the recovery target.
9. Run the following command from the */tmp* directory to disconnect from the LUN:

```
# ./iscsi_flr unmount
```

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 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**.

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

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, select a recovery point associated with a full, incremental, or differential backup. The appliance creates a new VM from the recovery point. The new VM can assume the role of the original and is available for use immediately.

Instant recovery uses a backup, or imported backup copy, to recreate the VM in the target location. The instant recovery process also creates a disk image, known as a recovery object, on the backup appliance. Data from this disk migrates to the new VM. The recovered VM remains fully operational during the migration.

The target location can be the original host or an alternate host running the same software version as the original or a higher supported version listed in the Compatibility and Interoperability Matrix. You must add the recovery target to the appliance that stores the backup, or imported backup copy, intended for the recovery.

When performing instant recovery, you can choose to perform the recovery in audit mode or instant recovery mode. Use audit mode to verify recovery points. Use instant recovery to replace a failed or corrupted VM.

Once the instant recovery completes, tear down the recovery object to make the reserved space available for other instant recovery processes.

See the following topics for details:

- "Audit mode"
- "Instant recovery mode" on page 297

Audit mode

Performing instant recovery in audit mode enables you to verify that a VM can be created from a backup or imported backup copy. When you select a backup or backup copy to audit, the appliance uses data from the backup or backup copy to create a disk image on the appliance and a new VM on the target host. This disk image is referred to as a recovery object. Although the VM resides on the host, it runs from the disk image created 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, delete it from the target host and tear down the recovery object from the appliance to free the system resources.

Instant recovery mode

Performing instant recovery enables you to replace a failed or corrupted VM. When you select a backup or imported backup copy, the appliance uses data from this backup to create a disk image or 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 migration completes, tear down the recovery object to free system resources.

Prerequisites and considerations

Unitrends supports instant recovery of VMware and Hyper-V virtual machines.

Prerequisites for VMware instant recovery

The following table describes the prerequisites and considerations for VMware instant recovery.

Prerequisite or consideration	Description
vCenter version and license	To perform instant recovery (in live mode), 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. **Note:** You can perform the audit process using a stand-alone ESXi server, but instant recovery in live 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 registered to the Unitrends appliance from which you are performing the instant recovery. • Must be running the same version or higher as the original ESXi server hosting the virtual machine. It is highly recommended that you recover to an ESXi version that matches the original. • Must have sufficient space 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 VMware backup of a virtual machine. (To run a backup, see "To create a backup job for VMware assets" on page 176.) • A local backup or backup copy.
VM hardware version	Since Unitrends does not create clones, the restored virtual machine 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.

Prerequisites for Hyper-V instant recovery

Consider the following as you plan for disaster recovery of Hyper-V virtual machines.

Unitrends software considerations and requirements

The following Unitrends software requirements must be met to perform Hyper-V instant recovery:

- The Unitrends Windows agent must be installed on the Hyper-V server hosting the VMs and on the Hyper-V server used as the restore target.
- Unitrends recommends reloading the list of VMs on the appliance after installing the agent on the Hyper-V server hosting your VMs. To reload the list of VMs, select **Options > Inventory Sync**.

Hyper-V and Windows Server considerations and requirements

The following Hyper-V server requirements must be met to perform Hyper-V instant recovery:

- The restore target must be:
  - 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.
- For the restore target, you must use a Hyper-V server that is running the same version or a higher version, as the Hyper-V server hosting the original VM. It is recommended that you restore to a Hyper-V server whose version matches that of the original, where possible.
- You can perform instant recovery using local backups or backup copies. The restore target can be the original Hyper-V server or an alternate Hyper-V server. The target server must be added to the Unitrends appliance from which you are performing the instant recovery.

  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 supported version.

General VM considerations and requirements

The following Hyper-V VM requirements must be met to perform Hyper-V instant recovery:

- All disks created on the Hyper-V VM must have unique names.
- Since Unitrends does not create clones, the restored VM is configured with the latest hardware version that is supported by the target Hyper-V server.
- Disks excluded during backup, including independent and physical RDM disks, are not restored with instant recovery for Hyper-V.
- Hyper-V Instant Recovery is not supported for Host component (AzMan Security Database for Hyper-V) backups of the original VM. These are not bootable backups.

Clustered VM considerations and requirements

The following requirements must be met to recover a clustered VM:

- Hyper-V instant recovery supports clustered and non-clustered virtual machines.
- 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.

Preparing for instant recovery

Unitrends recommends planning for instant recovery before a VM fails. Review the following information to assist in preparing for instant recovery. Included is a high-level process of tasks to complete before performing an instant recovery.

Step 1: "Back up assets"

Step 2: "Allocate storage for instant recovery"

Step 3: "Select and add a target host" on page 300

Step 4: "Perform instant recovery in audit mode" on page 301

Step 5: "Exit audit mode" on page 301

Step 1: Back up assets

Instant recovery can be performed at any time a long as there are backups or backup copies for your VMs and sufficient backup storage is allocated as instant recovery space. Run host-level backups for the virtual machines you wish to protect with instant recovery. For details, see "Host-level Backups Overview" on page 203.

Step 2: Allocate storage for instant recovery

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 disk for a recovered VM resides 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 disk. Storage allocated for instant recovery cannot be used for backups.

Supported storage procedures vary by Unitrends appliance type:

- Recovery Series 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.

If space is not available for instant recovery after a VM fails, adjust the storage allocation and use the newly configured storage for the instant recovery. For instructions on expanding your storage, see "Backup storage" on page 69 in the Configuration chapter.

Step 3: Select and add a target host

After verifying that there is enough space for the recovery, select a target host and add the host to the appliance. For instructions on adding virtual hosts and vCenter servers, see "Adding a virtual

host" on page 126.

Note: For VMware VMs, add the vCenter managing the target hosts.

Step 4: Perform instant recovery in audit mode

Perform instant recovery in audit mode to verify that backups and imported backup copies can be used to recover the VM. Repeat this procedure as needed to test new backups. The same dialog is used for audits and instant recovery. See "Performing instant recovery" on page 301 for instructions.

Access the virtual machine to ensure that it is operational and to verify the VM can be recovered in the event of a disaster. Applications on the VM requiring network access are not fully functional in audit mode.

Step 5: Exit audit mode

Changes to the VM in audit mode are lost upon exiting audit mode.

To exit audit mode

1. Select **Recover** and click the **Instant Recovery** tab.
2. Select the VM.
3. Click **Tear Down**.
4. Click **Confirm**.

Performing instant recovery

Perform the instant recovery after a VM fails.

A recovered VMware VM is created with the following default name: *<original_VM_name>_restore*. To edit the suffix, click **Edit Suffixes**.

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.

Note: To recover data from a cold backup copy, you must import the copy prior to performing the instant recovery.

To perform instant recovery

1. Select **Recover** and click the **Backup Catalog** tab.
   By default, the Backup Catalog displays in the assets view. Review the following for additional information:

Create Instant Recovery Job	Description
Protected asset icon	Click to view alphabetical listing of protected assets.

Create Instant Recovery Job	Description
Date icon	Click to view backups by date.
Appliance	Appliance on which the backup resides.
Host	Name of the virtual host.
Mode	Type of backup.

Use Filter Backups to filter the list of backups. Filter backups by Asset Name, Appliance, Mode, Date Range, Successes, and Type. Review the following for additional information:

Filter Backups	Description
Asset Name	Enter the name of the asset.
Appliance	Enter the name of the appliance.
Mode	Select the backup mode.
Held	Select to include backups on hold.
Successes	Select to include previous successful backups.
Type	Select the type of backup: • Backup • Imported Backup • Backup Copy

2 Select a host-level backup or imported host-level backup copy to use for the recovery.

3 Select **VM Instant Recovery** from the **Instant Recovery** list.

4 To perform the instant recovery in audit mode, select **Recover this VM in Audit Mode**. To recover a failed VM, clear this checkbox.

5 Select the following recovery options:

Recovery Options	Description
Target Location	Select the host to which to recover the VM.
Target Resource	Select a resource pool. This field displays only if the ESX host selected for the target location has resource pools. It does not display when recovering to a Hyper-V host.
Target Storage	Select a datastore (ESX host) or volume (Hyper-V host).

Recovery Options	Description
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	Select the target network switch. This field displays only when recovering a Hyper-V host.

6 Click **Next**.
A summary of the recovery options displays.

7 Click **Save**.

The new recovery object has been created and data migration begins. During this time, the VM is fully operational. Do not tear down the recovery object until the data migration completes.

Notes:

- About Windows server VM - 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.
- About 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 *Clt+Alt+F1* and log in as root. Then run *startx*.

Tearing down the instant recovery object

After performing an instant recovery, remove the recovery object from the appliance to free resources for the instant recovery. If the recovery was done in audit mode, this also deletes the VM from the host.

Note: Do not perform this procedure until the data migration completes and the VM is ready for use. Doing so requires having to restart the recovery process.

To tear down the instant recovery object

1 On the appliance used for the recovery, select **Recover** and click the **Instant Recovery** tab.

2 Select the applicable instant recovery object.

3 Click **Tear Down**.

Chapter 13: Recovering Asset-level Backups

This chapter describes recovery procedures for asset-level backups (backups that were run using a Unitrends agent). Assets must have an eligible backup or backup copy before running these procedures.

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. 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

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 asset" on page 119.
- 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 275 or "To import a cold backup copy" on page 278.

  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 314.

Run these procedures from the backup appliance to recover from backups or imported backup copies:

To recover from an asset-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. Select **Recover** and click the **Backup Catalog** tab.
3. Click **Search Files**.

4 In the **Type** list, select **Backup**.

5 Select the **Asset** whose backups and/or imported backup copies will be searched.

6 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.

7 Click **Search**.

The returned files display in the right panel.

8 Click to select files to recover:

- The files you select must be from a single backup.
- If you select files from multiple backups, the Save button becomes disabled.

9 Click **Next**.

10 Specify a Restore Target:

- Select an Asset. Choose from the agent-based assets that have been added to this appliance.
- (Optional) Enter a Directory path.
  - 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.

13 Click **Save**.

14 Click **OK** to close the Notice message.

Selected files are recovered to the target location. To view the running job, select **Jobs > Active Jobs**.

15 For Windows only, a reboot of the Windows asset may be required to reset locks on any files that were locked during the recovery.

To recover files from one asset-level backup by using the File Browser

Use this procedure to browse the contents of one backup or imported backup copy and recover selected files.

1. Log in to the backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.

   Use Filter Backups to the right to customize the backups that display.
3. Expand the desired asset and select a backup or imported backup copy to use for the recovery.
4. Click **Recover Files**.
5. Select or drag files and/or directories to recover.
6. Click **Next**.
7. Specify a Restore Target:
   - Select an Asset. Choose from the agent-based assets that have been added to this appliance.
   - (Optional) Enter a Directory path.
     - 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 Exclusions.
9. (Optional) Specify Advanced Options.
10. Click **Save**.
11. Click **OK** to close the Notice message.

    Selected files are recovered to the target location. To view the running job, select **Jobs > Active Jobs**.
12. For Windows only, a reboot of the Windows asset may be required to reset locks on any files that were locked during the recovery.

To recover an entire asset-level backup

Use this procedure to recover a backup or imported backup copy.

1. Log in to the backup appliance.
2. Select **Recover** and click the **Backup Catalog** tab.

   Use Filter Backups to the right to customize the backups that display.
3. Expand the desired asset and select a backup or imported backup copy to use for the recovery.
4. Click **Recover**.
5. Specify Recovery Options:
   - Select an Asset. Choose from the agent-based assets that have been added to this appliance.

- (Optional) Enter a Directory path.
  - 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.

7 Click **Next**.

A summary of the selected recovery asset and target location displays.

8 Click **Save**.

9 Click **OK** to close the Information message.

The selected backup is recovered to the target location. To view the running job, select **Jobs > Active Jobs**.

10 For Windows only, a reboot of the Windows asset may be required to reset locks on any files that were locked during the recovery.

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 310

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. Select **Recover** and click the **Backup Catalog** tab.
3. Click **Search Files**.
4. In the **Type** list, select **Backup copy (Cold)**.
5. Select the **Asset** whose backup copies will be searched.
6. 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.

7. Click **Search**.

   The returned files display in the right panel. The Serial # column shows the serial number of the disk(s) or tape(s) where the copy is stored.
8. Load the media that contains the files to import.
9. Click to select files to recover:
   - The files you select must be from a single backup copy.
   - If you select files from multiple backup copies, the Save button becomes disabled.
10. Click **Save**.
11. You see a message indicating that a selective backup will be created and imported to the appliance. Click **Save** to continue.
12. The recovery starts. Click **OK** to close the Notice message.

13 Select **Jobs > Active Jobs** to monitor progress of the import.

Once the import is complete, proceed to "Recover imported files".

Recover imported files

1 Select **Recover** and click the **Backup Catalog** tab.

2 Expand the desired 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 Specify Recovery Options:

- Select an Asset. Choose from the agent-based assets that have been added to this appliance.
- (Optional) Enter a Directory path.
  - 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.

5 (Optional) Specify Exclusions.

6 (Optional) Specify Advanced Options.

7 Click **Next**.

8 Review settings and click **Save**.

9 Click **OK** to close the Information message.

Selected files are recovered to the target location. To view the running job, select **Jobs > Active Jobs**.

10 For Windows only, a reboot of the Windows asset may be required to reset locks on any files that were locked during the recovery.

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.
  - 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.

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.
   - Click **Filter**.
4. Expand the desired 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. Select **Jobs > Active Jobs** to monitor progress of the import.

    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.
2. Expand the desired 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. Specify Recovery Options:
   - Select an Asset. Choose from the agent-based assets that have been added to this appliance.
   - (Optional) Enter a Directory path.
     - 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.
5. (Optional) Specify Exclusions.
6. (Optional) Specify Advanced Options.
7. Click **Next**.
8. Review settings and click **Save**.

9 Click **OK** to close the Information message.

Selected files are recovered to the target location. To view the running job, select **Jobs > Active Jobs**.

10 For Windows only, a reboot of the Windows asset may be required to reset locks on any files that were locked during the recovery.

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 asset" on page 119.

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 314.

- "To recover files from one asset-level hot backup copy by using the File Browser"
- "To recover an entire asset-level hot backup copy" on page 313

To recover files from one asset-level hot backup copy by using the File Browser

Use this procedure to browse the contents of one backup copy and recover selected files.

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.
- Click **Filter**.

4 Expand the desired 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 Specify a Restore Target:

- Select an Asset.

  The Asset list contains the agent-based assets that have been added to this target.

- (Optional) Enter a Directory path.
  - 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.

9 (Optional) Specify Exclusions.

10 (Optional) Specify Advanced Options.

11 Click **Save**.

12 Click **OK** to close the Notice message.

Selected files are recovered to the target location. To view the running job, select **Jobs > Active Jobs**.

13 For Windows only, a reboot of the Windows asset may be required to reset locks on any files that were locked during the recovery.

To recover an entire asset-level hot backup copy

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.
- Click **Filter**.

4 Expand the desired asset and select a backup copy to use for the recovery.

5 Click **Recover**.

6 Specify Recovery Options:

- Select an Asset.

  The Asset list contains the agent-based assets that have been added to this target.

- (Optional) Enter a Directory path.
  - 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.

8 Click **Next**.

A summary of the selected recovery asset and target location displays.

9 Click **Save**.

10 Click **OK** to close the Information message.

The selected backup is recovered to the target location. To view the running job, select **Jobs > Active Jobs**.

11 For Windows only, a reboot of the Windows asset may be required to reset locks on any files that were locked during the recovery.

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 305 procedures. For details on importing a backup copy, see "To import a hot backup copy" on page 275.

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 an asset-level backup copy by using the File Browser" on page 314
- "To recover files from an asset-level backup copy by using Search Files" on page 315

To recover files from an asset-level backup copy by using the File Browser

Use this procedure to browse an asset-level backup copy and recover selected files.

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.
   - Click **Filter**.
4. In the main Backup Catalog list, expand the asset whose files you wish to recover and select a *Backup Copy (Hot) on Target* to use for the recovery.

   To be certain you are viewing a backup copy that resides 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, select or drag files and/or directories to recover.
7. Click **Save**.
8. The message *Starting File Level Recovery on the Target* displays, indicating that the recovery has started.
9. To verify that the download starts, leave the message dialog open and view the status messages. Job status changes from *queued* to *active and downloading*.
10. Click **OK** to close the message.
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, do one of the following:

- For downloads that are 500MB or smaller, you must run a new job.
- For downloads that are greater then 500MB, access the recovered files in the source appliance's */downloads* directory by entering *<SourceApplianceIP>/downloads* in an Internet browser. Do not download these files until you see the *Unitrends-Restore.zip* file. While the recovery is in progress, you see files in this directory, but the download is not complete until the *.zip* file has been created. (Recoveries are automatically removed from the */downloads* directory after 72 hours.)

To recover files from an asset-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 Select **Recover** and click the **Backup Catalog** tab.

3 Click **Search Files**.

4 Select **Backup Copy (Hot)** from the **Type** list.

5 Select the **Asset** whose hot backup copies will be searched.

6 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.

7 Click **Search**.

The returned files display in the right panel.

8 Click to select files to recover:

- The files you select must be from a single backup copy.
- If you select files from multiple backup copies, the Save button becomes disabled.

9 Click **Save**.

10 The message *Starting File Level Recovery on the Target* displays, indicating that the recovery has started.

11 To verify that the download starts, leave the message dialog open and view the status messages. Job status changes from *queued* to *active and downloading*.

12 Click **OK** to close the message.

13 A *.zip* file containing the recovered files is placed in the default download location of the browser where the source appliance UI is running.

14 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, do one of the following:

- For downloads that are 500MB or smaller, you must run a new job.
- For downloads that are greater then 500MB, access the recovered files in the source appliance's */downloads* directory by entering *<SourceApplianceIP>/downloads* in an Internet browser. Do not download these files until you see the *Unitrends-Restore.zip* file. While the recovery is in progress, you see files in this directory, but the download is not complete until the *.zip* file has been created. (Recoveries are automatically removed from the */downloads* directory after 72 hours.)

Windows instant recovery

Unitrends Windows instant recovery (WIR) provides a way to recover a failed Windows physical asset. WIR creates a virtual failover client (VFC), a virtual replica of the client. A VFC can immediately assume the role of the original client in the event of a disaster.

When backups run on the original client, WIR restores them to the VFC, ensuring that the VFC always has all of the data from the original client.

The VFC can run on a:

- Recovery Series appliance (backup appliance or backup copy target appliance)

- ESXi host
- Hyper-V server

The VFC replaces the original windows asset until you can perform a bare-metal recovery to restore the asset to new physical hardware.

See the following topics for more details on WIR:

- "Windows instant recovery overview " on page 317
- "Windows instant recovery requirements and prerequisites" on page 318
- "Setting up a virtual failover client" on page 325
- "Windows instant recovery administration procedures" on page 326

Windows instant recovery overview

To set up Windows instant recovery (WIR), first back up the Windows asset and then create the virtual failover client (VFC).

Back up the Windows asset

WIR starts with an agent-based backup of the physical Windows server. The file-level full backup backs up:

- System files and folders.
- The system state, which includes the registry, IIS metabase, COM+ certificates, active directory information, and other key components necessary to restore the Windows machine.
- Disk metadata and layout, file system configurations, and other hardware-related information that enables WIR to reconstruct the standby VFC

IMPORTANT! Ensure that you have protected the Windows servers with the latest available release of the agent. With an agent older than release 8.0, the Windows server cannot acquire, from a full backup, the hardware metadata necessary to reconstruct the VFC. WIR can only create the VFC if it can find the original Windows protected asset.

If the Windows server acts as a host for mission-critical applications like SQL Server or Microsoft Exchange, use a scheduled backup strategy to protect the applications. With the application data and the operating system data protected, WIR can reconstruct the VFC to match the original Windows server.

Create the VFC

After configuring the Windows protected asset with scheduled backups, create the VFC as described in "Setting up a virtual failover client" on page 325. You can configure the VFC on the Unitrends appliance itself, or on a VMware or Hyper-V host.

You can use the following failover virtualization targets, whether the appliance protects onsite data or acts as a backup copy target:

VFC Target	Advantages
On-appliance failover virtualization	• No additional hardware required for failover virtualization. • Provides near-zero RTO without having to increase CapEx. • VFC is automatically protected while assuming the identity of the original Windows machine, thus providing continuous protection. • VFCs use the compute resources of the Recovery Series appliance. If the appliance is already under high load due to backups and backup copy (hot/cold) operations, there may not be sufficient resources to provide for VFCs. • Only Recover-Series appliances have this option.
VMware or Hyper-V host	• Leverage virtualization infrastructure for standby failover virtualization. • You can use failover virtualization as a migration strategy between the physical bare-metal Windows server to a virtual infrastructure. • In a failover scenario, provides the ability to dynamically scale resources in the virtual infrastructure to run the standby VFCs without any degradation of compute resources.

When you have created the VFC, all backups for the original protected Windows server automatically apply to the VFC. WIR keeps the VFC up-to-date with all backups. If the original Windows server then fails, the VFC can rapidly assume its identity.

Besides assuming the identity of the protected Windows asset, the VFC has two additional modes of operation:

Mode	Description
Live	The VFC assumes the identity of the original Windows server in a live production capacity.
Audit	The VFC can run on a private network, inaccessible from the production network, to verify the integrity of the VFC and the applications and data within it.

Windows instant recovery requirements and prerequisites

The following information covers both the requirements and prerequisites for Windows instant recovery (WIR).

- "VFC requirements" on page 319
- "Requirements for protected Windows server" on page 321
- "Storage allocation" on page 324
- "Configuring backups" on page 324

VFC requirements

The following topics cover the necessary requirements for different types of VFCs:

Requirements for running a VFC on a Recovery Series appliance

Requirement	Description
Unitrends system resources	VFC 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.
On-system retention	On-system retention is reduced because a portion of the appliance's storage is reserved for the VFC.
Use case for the VFC	Use temporarily until you can get new hardware and run bare-metal recovery.
UEFI-based assets	Cannot recover UEFI-based assets.
GPT-partitioned assets	Cannot recover GPT-partitioned assets.

Requirements for running a VFC under VMware ESXi

Requirement	Description
Hypervisor version	Must be running a paid version of ESXi 5.1, or a higher paid ESXi version listed in the Compatibility and Interoperability Matrix.
Compute	One VFC requires a minimum of 1024 MB of memory. This number must be a multiple of 4.
VFC client changeability	Once configured, do not change the VFC. Any alteration to the VFC (unless it is in Live mode) may lead the VFC to an inconsistent state.
Storage	Size of VFC is capped by what the hypervisor supports. If you configure the VFC with disks larger than 2 TB, the ESXi server must run version 5.5 or a higher paid ESXi version listed in the Compatibility and Interoperability Matrix.

Requirement	Description
Virtual hardware version	The VFC is configured with the highest hardware version that the hypervisor supports.
Virtual host asset	The ESXi host must be added to the appliance as an asset. See "Managing protected assets" on page 118.

Requirements for running a VFC under Hyper-V

Requirement	Description
Hypervisor version	The hypervisor must be one of the following: • A Windows Server with the Hyper-V role enabled, running Windows 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.
Host agent version	• The host must be running Unitrends agent version 8.0 or higher. • The host must be added to the Unitrends backup appliance as a protected asset.
Compute	One VFC requires a minimum of 1024 MB of memory (must be a multiple of 4).
VFC client changeability	Once configured, do not change the VFC. Any alteration to the VFC (unless it is in Live mode) may lead the VFC to an inconsistent state.
Storage	The size of the VFC is capped by what the hypervisor supports. The appliance creates a VFC with a disk of the same size as the disk on the original client. For Windows clients with disks larger than 2 TB, create the VFC on Hyper-V server version 2012 or higher. Create a VFC on Hyper-V server 2008 R2 with a VHD. Create a VFC on server version 2012/2012 R2 with a VHD(X).
VFC virtual machine generation	The client's firmware interface type determines the generation of the VFC VM. VFCs for BIOS-based clients are created as generation 1 VMs, and VFCs for UEFI-based clients are created as generation 2 VMs. A VFC for a UEFI-based client can run only on Hyper-V server version 2012 R2.
Passthrough disks	VFCs support pass-through disks. After booting the VFC in live mode and configuring the network settings, you must refresh and reconnect any existing iSCSI targets on the client.
Virtual host asset	The Hyper-V host must be added to the appliance as an asset. See "Managing protected assets" on page 118.

Additional requirements for running a VFC in a Hyper-V cluster environment

Requirement	Description
Registration	A VFC can run on a server in a cluster configuration. You must first install Unitrends Windows agent version 8.0 or higher on each node and then add each node and the cluster to the appliance from which you will create the VFC. Every node in the cluster must have the same agent version installed.
VFC location selection	To create a clustered VFC, select the cluster when specifying the location for the VFC. You cannot specify an owner node. If you select an individual node in the cluster, the VFC will not be clustered.
Network switch selection	For a clustered VFC, select the network switch common to all nodes in the cluster. If you do not select this switch, a VFC in live mode that fails over to another cluster will lose network connectivity.
2008 R2 clusters	To run the VFC on 2008 R2 servers in a cluster configuration, enable DCOM and WMI Virtualization access for all nodes in the cluster. For instructions, see KB 1140.
Live migration interoperability	During live migration of a clustered VFC, the Unitrends appliance cannot restore to the VFC, verify or audit the VFC, or boot it in live mode. If a restore or verify attempt takes place during a live migration, the appliance waits several minutes and then attempts the operation again. If you try to boot the VFC in audit or live mode during a live migration, the appliance notifies you that it cannot run the operation because of the migration and prompts you to attempt the operation again later.

Requirements for protected Windows server

The Windows server must meet the following requirements to use WIR:

Requirement	Description
Client Operating Systems	• 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** **Notes:** • *For 32-bit Windows XP clients, the VFC must reside on a Recovery Series appliance or a Hyper- V server. It cannot reside on an ESX server. • **A VFC running Windows 8 or higher cannot reside on Hyper-V server 2008 R2.
Server Operating Systems	• 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 • Windows 2012, 64-bit, all versions*** • Windows 2012 R2, 64-bit*** **Note:** ***A VFC running Windows Server 2012 or higher cannot reside on Hyper-V server 2008 R2.
Applications	• SQL Server 2005, 2008, 2012, 2014, and 2016. • Exchange 2003, 2007, 2010, 2013, and 2016. The following limitations apply: • WIR does not support Windows Cluster Server, and Hyper-V, Oracle, and SharePoint applications. You can use WIR to protect the Windows servers that host these applications, but the feature does not protect the applications. • For SQL applications - To recover SQL databases that are available and running when the VFC enters live mode, the master, model, and msdb system databases must be present in the asset-level full backup.

Requirement	Description
Firmware interface type	WIR supports BIOS- and UEFI-based assets. For UEFI-based assets, the VFC must be created on one of the following: • An ESXi server running paid ESXi version 5.1, or a higher paid ESXi version listed in the Compatibility and Interoperability Matrix. • A Hyper-V server running version 2012 R2. For a VFC running on a Hyper-V server, the original UEFI-based asset's operating system must be 64-bit and Windows 8 or higher. **Note:** A VFC for a UEFI-based asset cannot run on a Recovery Series appliance.
Disk configuration	WIR 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, WIR protects the data from all disks, but a maximum of four disks are accessible when booting. • For Windows 2003, the boot disk must be located on one of the first three disks if you will be running the VFC on an ESXi host.
Disk partition type	WIR protects Master Boot Record (MBR) partition types. GUID Partition Table (GPT) disks are supported with the following limitations: • Only the data volumes can be GPT. The boot and system volumes cannot be GPT. • A VFC for clients with GPT disks can run only on an external hypervisor. It cannot run on a Recovery Series appliance.
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 WIR recovers data in its non-deduplicated form, the volume must have enough capacity to house this non-deduplicated data.

Requirement	Description
Number of volumes	A client protected by WIR can have a maximum of 20 volumes, including the System Reserved volume and other unmounted volumes. A VFC created for a client with more than 20 volumes may fail to boot.
Separate boot and system partitions	For protected assets with boot and system partitions located on different disks, the system partition must reside on the first disk (Disk 0).
File System Configuration	WIR supports the following file systems: • NTFS • FAT/FAT32 • ReFS (Windows 2012 and later)
Active Directory	WIR supports an Active Directory database (NTDS) on the boot volume. (If it is not on the boot volume, the configuration is not supported, and you see an error message when you add the VFC.)

Storage allocation

Depending on the identity of the appliance (backup system or replication target), you can allocate storage among backups/replication, vaulting, and instant recovery. You cannot use resources allocated for a VFC for other purposes, such as backups, archives, or deduplication.

Note: This procedure is only required for on-host instant recovery on the Recovery Series appliance.

Before allocating storage, note the following:

- The system load and use.
- The amount of storage on the asset. To determine the storage amount to allocate for instant recovery on the appliance, figure the sum of space in use on all the original clients (for which you want to create VFCs). To determine the amount of space in use on an asset, select the computer window, and view the disks.

To allocate storage for the virtual failover client

1. On the **Configure > Appliances** page, select the **Storage** tab.
2. Select the **Internal storage** and click **Edit**.
3. Modify the percent of storage allocated for backups versus the percent of storage allocated for instant recovery, and click **Save**.

Configuring backups

Periodically back up clients and applications you want to protect with WIR. You do not have to create a special schedule for clients that WIR protects. If you implement WIR for a client that a backup schedule already protects, it does not produce a conflict. The existing schedule will continue

to function.

Unitrends recommends you use an incremental forever schedule, but any schedule with periodic full and differential backups, or periodic full and incremental backups, will work. For more details, see "Backups" on page 38.

For a SQL server protected by WIR, the recommended strategy uses a combination of SQL Full and Transaction Log backups. For more details, see "SQL backup requirements and considerations" on page 240.

For Microsoft Exchange application backups, the recommended strategy comprises a combination of full and incremental backups. For more details, see "Exchange backup requirements and considerations" on page 237.

Once you have created the VFC, WIR restores completed backups to that VFC, so the VFC stays current.

Setting up a virtual failover client

Use the following procedure to set up a VFC client. You can run the VFC on:

- A Recovery Series appliance (backup appliance or backup copy target appliance).
- An external hypervisor (Hyper-V or VMware).

To set up a virtual failover client

1. Select **Recover > Backup Catalog**.
2. Select **Instant Recovery > Windows Instant Recovery** (the dropdown button, not the tab).
3. In the Select Protected Asset list, select the applicable Windows asset and click **Next**. The list displays all protected assets eligible for WIR.
4. In the Define Location step, provide the location where you want to put the VFC. Depending on the selection, the UI displays options specific to that location. The list displays only locations for WIR.
   - If you select **Unitrends Appliance** as the location, you do not have to provide any more location configuration details. Click **Next** to continue.
   - If you select **VMware Host**, provide the following information:
     - Location: ESXi host to house the VFC.
     - (Optional) If your environment has configured resource pools, you can choose one to house the VFC. If your environment does not have any resource pools, skip this step.
     - Storage: Specific datastore location for the VFC disks.
     - Network: Selection of virtual networks discovered and available.
     - IP Address / Netmask / Gateway: Network configuration parameters to configure the VFC while restoring the backups and creating the stand-by virtual machine. When the VFC enters live or audit mode, it assumes the IP address of the original Windows server and not the configuration specified here.

       Click **Next** to proceed to the third step of the wizard.

- If you select **Hyper-V Host**, provide the following information:
  - Location: Hyper-V host to house the VFC.
  - Storage: Specific volume to house the VFC virtual disks.
  - Network: Selection of virtual networks discovered and available.
  - IP Address / Netmask / Gateway: Network configuration parameters to configure the VFC while restoring the backups and creating the stand-by virtual machine. When the VFC enters Live or Audit mode, it assumes the IP address of the original Windows server and not the configuration specified here. Click **Next** to proceed to the third step of the wizard.

5. In the Define VFC Options step, provide specific compute and exclusion options:
   - Name: If you create a VFC on a VMware or Hyper-V host, you can provide a custom name for the VFC.
   - Processor: Configurable number of processors connected to the VFC. Note that the VFC will be used in the event of a crash of the original Windows machine. The VFC will be used temporarily if it uses on-host failover virtualization, and the compute resources do not have to match the original physical Windows asset.
   - Memory: Configurable amount of memory attached to the VFC.
   - Volumes to include: Volumes marked as Critical are required to spin up the VFC. You can exclude other data volumes.
   - Email Verification Report: Selecting this option generates and emails a daily audit report with a screen shot of the VFC.
6. If applicable, provide application options and click **Next**.
7. Click **Save.** The UI creates a VFC for the selected Windows asset.

Windows instant recovery administration procedures

The following information provides details on administration procedures used in Windows instant recovery (WIR).

- "Accessing a VFC in audit mode" on page 326
- "Accessing a virtual failover client in live mode" on page 329
- "Tearing down a virtual failover client" on page 332
- "Monitoring and managing virtual failover clients" on page 332

Accessing a VFC in audit mode

Once you have created the VFC, you can check its status on the **Instant Recovery** tab on the Recovery page.

After setting up a VFC, audit it periodically to verify that it boots successfully. You can automate the audit process by enabling email verification reports, or you can perform manual audits.

You must complete at least one virtual restore before the VFC can boot in audit mode. No virtual restores take place while the VFC is in audit mode, but they resume when you take the VFC out of audit mode.

A VFC running in audit mode boots with no network interface. Auditing the VFC with the original client still online does not result in network conflicts or impact the original client in any way. However, applications on the VFC that require network access do not fully function in audit mode.

Automated audits for a virtual failover client

Note: WIR supports automated audits only for VFCs on a Recovery Series appliance or Hyper-V server. You must manually audit VFCs on an ESX server.

You can automate the audit process by enabling email verification reports for a VFC. If you enable the report, the appliance:

- Sends the VFC into audit mode after a restore.
- Takes a screenshot of the Windows screen after the VFC has had several minutes to boot.
- Sends the screenshot in an email to the addresses you entered when setting up reports for the appliance that manages the VFC.

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 VFC boots correctly.

The report runs once a day, but only after a restore. If the interval between restores lasts longer than 24 hours, you will not receive a report every day. If the VFC cannot boot, you will receive an email report indicating that the VFC cannot be verified.

You can enable verification reports when creating a VFC, or by modifying an existing VFC by using the instructions described in "Setting up a virtual failover client" on page 325

Manually auditing a virtual failover client

Manually auditing the VFC is a two-part process that:

- Sets the VFC to go into audit mode.
- Accesses the VFC in audit mode to verify that it boots successfully.

The procedures for accessing the VFC in audit mode vary depending on the location of the VFC.

Note: After you have finished auditing the VFC, you must take it out of audit mode, so virtual restores to the VFC can resume.

Audit mode procedures

Use these procedures while working in audit mode:

- "To set the virtual failover client to go into audit mode" on page 328
- "To access a virtual failover client in audit mode on a Recovery Series appliance" on page 328
- "To access a virtual failover client in audit mode on an external hypervisor" on page 328
- "To turn off audit mode for a virtual failover client" on page 329

To set the virtual failover client to go into audit mode

1. In the appliance that manages the virtual failover client (VFC), select the **Recover > Instant Recovery** tab.
2. Click on the VFC.
3. Check the **Audit** option.
   - You may see a message that the VFC will enter audit mode.
   - If, however, there is a restore in progress, the VFC does not go into audit mode until the restore completes.
4. Click **Confirm**.
5. You see that the VFC is now in audit mode: the audit column reads **Yes**, the State column **audit**, and the **Access** field displays the port number (for a VFC residing on a Recovery Series appliance) or IP address (for a VFC residing on an external hypervisor) that connects to the client.
6. To connect to the VFC so you can verify that it is functioning as expected, proceed to "To access a virtual failover client in audit mode on a Recovery Series appliance" on page 328 or "To access a virtual failover client in audit mode on an external hypervisor" on page 328.

To access a virtual failover client in audit mode on a Recovery Series appliance

Note: You must use a VNC viewer to access the VFC in audit mode on a Recovery Series appliance.

1. Set the VFC to go into audit mode. Upon entering audit mode, you can obtain a VNC port number by selecting the **Instant Recovery** job and clicking **View Details**.
2. Open a VNC viewer.
3. In the **Server** field, enter the IP address of the appliance, followed by a colon and the VNC port number. For example: 192.168.101.19:5905.
4. Click **Ok**. A Windows login screen displays, indicating the VFC is available.
5. Enter the credentials for the Windows client and press **Enter**.
6. After verifying that the VFC is running with its restored data, turn off audit mode.

To access a virtual failover client in audit mode on an external hypervisor

If you access the VFC before it has booted, you may see the first screen of the Windows Integrated Bare Metal Recovery Wizard. This screen displays because the instant recovery and integrated bare metal recovery features use the same ISO image to boot a recovered Windows machine. You should not attempt to complete the steps on the bare metal screen. After several seconds, the login screen for the original client displays.

1. Set the VFC to go into audit mode.
2. Connect to your hypervisor.

3 Locate the VFC in the list of virtual machines, and access it the same way you access all VMs on the hypervisor.

4 Enter the credentials for the Windows client and press **Enter**.

5 After verifying that the VFC is running with its restored data, turn off audit mode.

To turn off audit mode for a virtual failover client

1 On the Unitrends appliance that manages the VFC, select **Recover > Instant Recovery**.

2 Select the VFC in the list of WIR clients.

3 Select **End Audit** to stop the audit process. If backups successfully completed for the original client while the VFC was in audit mode, the VFC enters the Restore state when audit mode is turned off. If there are no backups to restore, its state is Idle.

Accessing a virtual failover client in live mode

If you are protecting a Windows asset with WIR and the asset fails, you can temporarily replace it with the VFC by booting the VFC into live mode. Because virtual restores constantly update the VFC with the original asset's data, the VFC can immediately assume the role of the original asset until you can recover it to new physical hardware. If the VFC resides on an external hypervisor, it can permanently replace the original client if the hypervisor has sufficient resources.

The original asset's backup and backup copy schedules protect the VFC in live mode, which allows any changes made to the VFC in live mode to be captured under the identity of the original assset, ensuring continuity of recovery points.

Live mode recommendations

Review these recommendations before going into live mode:

- The appliance begins sending alerts after a live VFC has run for 14 days.
- Recover the client to new hardware as soon as possible by using Unitrends bare metal recovery. (See the Bare Metal Protection and Recovery Guide for details.) The backup schedule for the original client protects data from the live VFC. Restore it after recovering the client to new hardware.
- A live VFC running on an external hypervisor does not use any of the appliance's resources. Instead, it uses the hypervisor's resources. The VFC can temporarily replace the original client, or, if the hypervisor has sufficient resources, the VFC can permanently replace the original client, as described here:
  - *Using the live VFC as a temporary for the original Windows asset.* If the VFC will replace the original asset only temporarily, recover the asset to new hardware as soon as possible, using Unitrends bare metal recovery, as described in the Bare Metal Protection and Recovery Guide. Data from the live VFC is protected by the backup schedule for the original asset. After recovering the Windows asset to new hardware, you will need to recover backups from the VFC to the new Windows asset. For details, see "To recover an entire asset-level backup" on page 307. After recovering the VFC's data to the new Windows asset, you should delete the VFC from the appliance and the hypervisor. For instructions, see "To tear down a virtual failover client" on page 332.

- *Using the live VFC as a permanent replacement for the original Windows asset.* If the VFC on a hypervisor will permanently replace the original asset, determine whether to continue protecting the VFC with the backup schedules of the original asset or whether to run Hyper-V or VMware backups for the VM. For details, see "Protecting Hyper-V virtual machines at the asset level" on page 207 or "Protecting VMware virtual machines at the asset level" on page 216.

  Note: It can take several minutes for a live VFC on a hypervisor to show up in the list of VMs to protect with VMware or Hyper-V backups. The VFC VM is not automatically added to a VM backup schedule.

  You should then delete the VFC's information from the appliance that was managing it. For instructions, see "Tearing down a virtual failover client" on page 332 . Be sure to delete the VFC from the appliance only, as you have the option to delete it from the hypervisor as well.

To boot a virtual failover client in live mode

This procedure provides instructions for booting a VFC in live mode, and recommendations for steps to take after the VFC is live. Shut down the original client before you boot the VFC in live mode.

1. In the appliance that manages the VFC, select the **Recovery > Instant Recovery** tab.
2. Click on the VFC.
3. Check the **Go Live** option.
4. Click **Confirm** to initiate the operation.

   Note: If a restore to the VFC is in progress, the VFC does not boot until the restore completes.

5. Complete one of the following, depending on where your VFC resides:

VFC location	Steps
Recovery Series appliance (backup system or replication target)	• Connect to the VFC by using VNC, then log in. To connect by using a VNC viewer, specify the IP address of the Unitrends appliance, followed by a colon and the VNC Port. • Click **View Details** under the **Recover > Instant Recovery** tab for the selected failover client to view the port number.
External hypervisor	Connect to the VFC by using the hypervisor manager.

6. If you see a message about reactivating Windows, you must activate the operating system by using your product key.
7. Reboot the VFC, if prompted. On first boot, Windows automatically performs some driver updates. When this process completes, the system prompts you to reboot.
8. 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.

9 Set the system clock. The asset may be running with the system clock time used by the latest backup. This issue may cause the client to boot with a past date or time.

10 From the Windows Control Panel, update the network properties for the adapter (the TCP/IPv4 address) by using one of these procedures:

Note: For a VFC that resides on an external hypervisor, the network settings you configured when creating the VFC are used only for virtual restores. You must assign new network settings after booting the VFC in live mode.

- For a VFC running on a Recovery Series appliance
  - If the original asset has a static IP address, assign the live VFC the same network settings as the original asset. This process ensures that the VFC functions as the original asset, and that the original asset's backup and scheduled jobs continue for the VFC.
  - 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 added appliance detects the live VFC after you connect it to your network. The appliance then treats the live VFC as if it is the original asset. No additional network configurations are necessary to ensure that scheduled backup and backup copy jobs continue for the client.
- For a VFC running on an external hypervisor
  - If the original asset has a static IP address, assign the live VFC the same network settings as the original client. This step ensures that the VFC functions as the original asset and that the original asset's scheduled backup and backup copy jobs continue for the VFC.
  - If the original asset has a static IP address and the hypervisor running the VFC does not have a network interface on the same subnet as the original asset, assign the VFC 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 and enter the new IP address. This process enables the appliance to treat the VFC as if it were the original asset.

11 In the Unitrends backup appliance that protects the original asset, perform the steps to finish the preparing the VFC and make the applications on the VFC available on the network.

12 (Optional) If you have created and updated the VFC by a backup copy on an appliance backup copy target, perform these steps from the backup appliance to which you added the original asset 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.

  The VFC can now perform the role of the original Windows asset.

13 Be sure to tear down the VFC as soon as possible. See "Live mode recommendations" on page 329 for options.

Tearing down a virtual failover client

This section provides instructions for deleting a VFC. Because the VFC uses appliance resources, you should delete a live VFC from the appliance after recovering the original asset to new physical hardware.

For a VFC running on a hypervisor, you have the option to delete its information from the appliance and to delete the VFC itself from the hypervisor. You can also delete only the information from the appliance without removing the VFC from the hypervisor.

When you delete a VFC, the appliance immediately removes it from the list of VFCs in the UI. However, it can take several minutes for the appliance to purge all information about the VFC. If you need to create a new VFC for the original asset, you must wait for this information purge. If it has not purged, the original asset will not display in the list of assets for which you can create a VFC.

To tear down a virtual failover client

1 Select the **Recover > Instant Recovery** tab.

2 Click to select the row containing the VFC you want to delete.

3 Click **Tear Down**. One of the following occurs depending on the location of the VFC:

- For a VFC residing on a Recovery Series appliance a box displays, asking you *Are you sure you want to stop this IR session?*. Click **Confirm** to delete the VFC.
- For a VFC residing on an external hypervisor, a box displays with options to delete the VFC from the hypervisor and the appliance or only from the hypervisor. Select the desired option and click **Tear Down**.

Monitoring and managing virtual failover clients

When you select the **Recover > Instant Recovery** tab, you can monitor a VFC by viewing details about its mode. This section explains the different modes.

The **Mode** column on the Instant Recovery screen indicates the current mode of the VFC. For example, whether it is a newly created VM, whether a restore is occurring, or whether it is in audit mode (see the table below for descriptions of all the possible modes).

The mode can change depending upon an action requested of the VFC. The action can be requested by the user or by the appliance managing the VFC.

Refer to the following table for explanations of the descriptions:

Mode	Description
New	State of a VFC for which no virtual restores have been performed. The VFC remains in this mode until a virtual restore has been performed.
Restore	A backup has completed, and the appliance has requested a restore. The VFC is in the Restore mode until the restore completes.
Idle	At least one backup has been performed to the VFC, but currently no action is occurring.
Halted	A backup has completed, and the appliance has requested a restore. The VFC goes into a Halted state if the restore cannot be performed. The following can occur when a VFC is in this mode: • If the restore could not be performed because the appliance could not reach the VFC, it tries again after several minutes, and the mode changes from Halted to Idle. After three failed attempts, the VFC's mode 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 client, the VFC's mode becomes Invalid, and it remains in Halted mode until a user deletes it.
Audit	A user has requested an audit, and the VFC has booted in Audit mode. For details about auditing a VFC, see "To set the virtual failover client to go into audit mode" on page 328.
Verify	The user has enabled verification reports. The appliance is taking a screenshot of the VFC's login screen in Audit mode. This screenshot is sent to the user in an email report to verify the VFC. After the verification completes, the VFC mode is Idle. For details about verification reports, see "Automated audits for a virtual failover client" on page 327.
Live	The user has requested the VFC to boot and replace the original asset. After the VFC mode changes to Live, the appliance no longer manages it, and virtual restores are no longer performed. Its mode is Invalid. For details about modes, see "Monitoring and managing virtual failover clients" on page 332. Once the mode of the VFC is Live, the only other mode it can enter is Off.
Off	The user has taken the VFC out of Live mode. A VFC in the Off mode can enter Live mode again, but its mode is Invalid, and it is no longer eligible for virtual restores.

Chapter 14: 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 335
- "Recovering NAS NDMP backups" on page 337

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 or files to recover.

See these topics for details:

- "Specifying the target recovery location" on page 335
- "To search multiple backups for files to recover" on page 335
- "To recover a NAS CIFS or NFS backup" on page 336
- "To recover files from one NAS CIFS or NFS backup" on page 337

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 CIFS or NFS NAS 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 Click **Recover > Search Files**.

2 For Type, select **Backup**.

3 For Asset, select the NAS whose files you want to recover.

4 Enter additional criteria as desired, and click **Search**.

All backups of this NAS stored on the appliance are searched for matching files.

5 In the results list, click to select files to recover.

6 Click **Next**.

7 Specify a recovery target:

- Select a NAS in the Asset list.
- In the Directory field, enter the full path where the backup will be recovered.

For details, see "Specifying the target recovery location" on page 335.

8 (Optional) Specify Exclusions.

9 (Optional) Modify Advanced Options.

10 Click **Save**.

The job is queued immediately.

11 Click **OK** to close the Notice message.

To view the running job, go to **Jobs > Active Jobs**.

To recover a NAS CIFS or NFS backup

Use this procedure to recover the entire backup.

1 Select **Recover**.

2 On the Backup Catalog tab, expand the NAS asset and select the desired backup.

- Use Filter Backups to the right to customize the backups that display.
- To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.

3 Click the **Recover** button.

4 Specify a recovery target:

- Select a NAS in the Asset list.
- In the Directory field, enter the full path where the backup will be recovered.

For details, see "Specifying the target recovery location" on page 335.

5 (Optional) Modify Advanced Options.

6 Click **Next**.

7 Review the recovery job settings, then click **Save**.

The job is queued immediately.

8 Click **OK** to close the Information message.

To view the running job, go to **Jobs > Active Jobs**.

To recover files from one NAS CIFS or NFS backup

1. Select **Recover**.
2. On the Backup Catalog tab, expand the NAS asset and select the desired backup.
   - Use Filter Backups to the right to customize the backups that display.
   - To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.
3. Click **Recover Files**.
4. In the File Browser, expand folders to view items in the backup.
5. Select or drag files or folders to recover.
6. Click **Next**.
7. Specify a recovery target:
   - Select a NAS in the Asset list.
   - In the Directory field, enter the full path where the backup will be recovered.

     For details, see "Specifying the target recovery location" on page 335.
8. (Optional) Specify Exclusions.
9. (Optional) Modify Advanced Options.
10. Click **Save**.

    The job is queued immediately.
11. Click **OK** to close the Notice message.

    To view the running job, go to **Jobs > Active Jobs**.

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. For example, VNXe devices only support full volume-level recovery to the original location.

  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 seleced files that contain non-UTF-8 compatible characters is not supported. Instead you must recover the entire backup.

See these procedures to recover from NDMP backups:

- "To recover an NDMP backup" on page 338
- "To recover files from an NDMP backup" on page 339

To recover an NDMP backup

Use this procedure to recover the entire backup.

1. Select **Recover** and click the **Backup Catalog** tab.

   Use Filter Backups to the right to customize the backups that display.

2. Expand the NDMP asset and select a backup or imported backup copy to use for the recovery.

   To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.

3. Click the **Recover** button.

4. 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.

5. Click **Next**.

6. Review the job details, then click **Save**.

   The recovery job is queued immediately.

7. 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

Use this procedure to recover selected files from an NDMP backup.

Notes:

- Recovering files is not supported for all NDMP filers. For example, VNXe devices only support recovering the entire backup to the original location. See the vendor documentation to determine whether recovering files is supported.
- Recovery of files that contain non-UTF-8 compatible characters is not supported.

1. Select **Recover** and click the **Backup Catalog** tab.

   Use Filter Backups to the right to customize the backups that display.

2. Expand the NDMP asset and select a backup or imported backup copy to use for the recovery.

   To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.

3. Click **Recover Files**.

4. In the File Browser, expand folders to view items in the backup.

5. Select or drag files or folders to recover.

6. Click **Next**.

7. 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.

8. 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 files that were recovered to an environment with different permissions, unmount the target volume and remount it with the appropriate permissions.

Chapter 15: Recovering Application Backups

This chapter 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 recover targets:

- "Recovering to the original Exchange server" on page 342
- "Recovering to a recovery area" on page 343
- "Recovering to an alternate location" on page 345

Recovering to the original Exchange server

Recovery to the Exchange server is the default recover 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 whether this is the problem: Exchange Database Is in a Dirty Shutdown State.

Condition	Explanation
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.

To recover a database or storage group to the original Exchange server

1. Verify all prerequisites have been met, as described in "Recovering to the original Exchange server" on page 342.
2. Select **Recover** and click the **Backup Catalog** tab.
3. Expand the Exchange server and select the applicable backup.
4. Click **Recover**.
5. Select **Original Exchange Server** as the Recovery Target.
6. (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**.

   All database and transaction log files are recovered directly to the original location.
9. 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 folder databases as the original storage group	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.

To recover a database or storage group to a recovery area

1. Verify all prerequisites have been met, as described in "Recovering to a recovery area" on page 343.
2. Select **Recover** and click the **Backup Catalog** tab.

   Use Filter Backups to the right to customize the backups that display.
3. Expand the Exchange server and select the desired backup or imported backup copy.

   To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.

4 Click **Recover**.

5 Select **Original Exchange Server** as the Recovery Target.

6 Select a Recovery Database from the drop-down.

7 (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**.

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 121.

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.
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. Kroll On-Track. 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.

To recover a database or storage group to an alternate location

1 Verify all prerequisites have been met, as described in "Recovering to an alternate location" on page 345.

2 Select **Recover** and click the **Backup Catalog** tab.

Use Filter Backups to the right to customize the backups that display.

3 Expand the Exchange server and select the desired backup or imported backup copy.

To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.

4 Click **Recover**.

5 Select **Alternate Location** as the Recovery Target.

6 Select an asset to which to recover the database or storage group.

7 Enter the desired Directory Path.

8 (Optional) Specify client-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**.

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 Kroll Ontrack Powercontrols to recover individual items from an Exchange backup.

Note: Kroll can be used with 32-bit versions of Outlook only. 64-bit Outlook versions are not supported. For a complete overview of Kroll Ontrack PowerControls for Exchange, including procedures and limitations, see Kroll's Exchange user guide available at http://www.krollontrack.com/support/user-guide-and-manuals/.

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 KOP 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 342), KOP or a third-party tool (e.g., Lucid8) can be used to search and recover individual Exchange items. Note that unlike KOP, third-party tools are certified and supported by third parties and not by 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.

To recover individual Exchange items directly from the Exchange backup

1. Select **Recover** and click the **Backup Catalog** tab.

   Use Filter Backups to the right to customize the backups that display.

2. Expand the Exchange server and select the desired backup or imported backup copy.

   To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.

3. Click **Recover Files**.

4. Click **Confirm** to create the recovery point.

Note: Creating the recovery point object can take some time. If you go to KOP and do not see any available items, check back later.

5. Use KOP to recover Exchange items. See "Recovering items with Kroll Ontrack PowerControls for Exchange" on page 348 for details.

6. Tear down the recovery object. For instructions, see "To view or tear down Exchange recovery objects" on page 347.

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 KOP 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, KOP or a third-party tool (e.g., Lucid8) may be used to search and recover individual Exchange items. Note that unlike KOP, 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 Kroll Ontrack PowerControls for Exchange

After completing step 1 on page 347 through step 4 on page 347 in "To recover individual Exchange items directly from the Exchange backup" on page 347, 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 Power Controls User Guide.

To recover items using Kroll Ontrack PowerControls for Exchange

Note: Creating the recovery object can take some time. If you do not see any available items in KOP, check back later.

1. Log in to your Windows machine with Kroll installed. Run Kroll 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 Kroll Ontrack PowerControls for Exchange.
9. Tear down the recovery object as described in "To view or tear down Exchange recovery objects" on page 347.

Recovering SQL backups

Unitrends supports recovery of full, differential, and transaction log backups of SQL databases. See the following topics for details:

- "Considerations for recovering SQL backups" on page 349
- "SQL recovery procedures" on page 350

Considerations for recovering SQL backups

Recovery procedures vary depending on what type of database and backup are being recovered. Consider the following before recovering SQL backups:

Recovery type	Requirements and considerations
All recovery operations	The following apply to all SQL recovery operations: • The entire database is recovered in a live state with each recovery operation. • Unitrends does not support granular recovery of SQL database records. • Databases must be recovered to a SQL application that is the same version or later than that of the original SQL application. Databases cannot be recovered to an older SQL version. • A SQL differential or transaction log backup can only be recovered to the original SQL instance. • When recovering a SQL differential or transaction backup, 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. • For 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.
User databases	A full user database backup can be recovered to the original location or to an alternate location. The alternate location can be any available SQL server that has been added as an asset to the Unitrends appliance (and is running a SQL version that is the same or later than that of the original SQL application). The database may also be renamed and recovered to a specified alternate path if desired.

Recovery type	Requirements and considerations
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. • 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.
Clusters	To recover a clustered database instance to an alternate location, you must select a path that is on a shared volume associated with that SQL instance.
Stretch databases	The backup must be recovered to the original database and instance. 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" on page 352 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: • 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

After you have reviewed the "Considerations for recovering SQL backups" on page 349, use one of the following procedures to recover a SQL backup or imported backup copy:

- "To recover one SQL full, differential, or transaction log backup" on page 350
- "To recover multiple SQL full, differential, or transaction log backups" on page 352
- "To recover a Stretch database backup" on page 352

To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.

To recover one SQL full, differential, or transaction log backup

1 Select **Recover**, and click the **Backup Catalog** tab.

Use Filter Backups to the right to customize the backups that display.

2 Expand the SQL server and database.

3 Select the backup or imported backup copy to recover.

4 Click **Recover** and select Recovery Options and Advanced Options, described in the following table:

Recovery Options	Description
Recovery Target	Select the asset where the database will be recovered.
Recovery Instance	Select the SQL instance to which you want to recover the database. **Notes:** • A full backup can be recovered to any SQL instance running the same or a later version than that of the original. • A differential or transaction log backup can only be recovered to the original instance.
Database	Enter a name for the recovered database. If the original SQL instance is selected and the database matches the original name, the original database files are overwritten during the recovery (regardless of whether you specify a new path). Otherwise a new database is created. For system databases, the database name cannot be changed.
Specify Path	Select **Specify Path** to enable **Browse**. Select a path on the restore target. **Notes:** • If this field is left blank, the files are recovered to their original locations. • If recovering to an alternate SQL instance on the same SQL server and no alternate target pathname is specified, the files are recovered to the default path of <Vol>:\UnitrendsRestore, where <Vol> is the volume on the client with the most free space. • If recovering to an alternate SQL instance, a path is required.
Point-in-time Recovery	(Optional) This option is available for transaction log backups only. For additional recovery points (down to the minute-level), check the **Point-in-time Recovery** box and select the desired recovery point by moving the Earliest/Latest slider.

Recovery Options	Description
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.

5 Click **Next**.

6 Click **Save**.

To recover multiple SQL full, differential, or transaction log backups

Use this procedure to recover multiple databases to their original locations. To recover to a different location, use the "To recover one SQL full, differential, or transaction log backup" procedure above instead.

1 Select **Recover**, and click the **Backup Catalog** tab.

Use Filter Backups to the right to customize the backups that display.

2 Expand the SQL server and database.

3 Select the backups or imported backup copies to recover.

- You can select multiple databases that reside on the same host server.
- 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.

4 Click **Save**.

To recover a Stretch database backup

A SQL backup of a Stretch database must be recovered to the original database and instance. 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.

1 Select **Recover**, and click the **Backup Catalog** tab.

Use Filter Backups to the right to customize the backups that display.

2 Expand the SQL server and database, then select the desired backup or imported backup copy.

3 Expand the database, and select the backup to recover. Use filter options as necessary to locate a specific database.

4 Click **Recover** and select these Recovery Options and Advanced Options:

Recovery Options	Description
Recovery Target	Select the original SQL server asset (where the database was backed up).
Recovery Instance	Select the original SQL instance that was backed up.
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.

5 Click **Next**.

6 Click **Save**.

The backup is recovered to the original location, creating a new SQL server database.

7 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 354

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 252.
- 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

Use these procedures to recover from a SharePoint backup:

- "To recover a SharePoint farm, site, service, or item" on page 355
- "To recover the entire farm on a standalone SharePoint 2013 or 2010 server" on page 356

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

1 Select **Recover** and click the **Backup Catalog** tab.

Use Filter Backups to the right to customize the backups that display.

2 Expand the SharePoint Farm and select the desired backup or imported backup copy.

To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.

3 Click **Recover Files**.

4 Click **Confirm** to create the file recovery object.

5 Click **OK**.

6 The system creates a share for this farm instance and starts the recovery. A row for this recovery displays in the grid.

7 Select the share and click **Show Details**.

8 The details form displays the full path of the share, \\<SourceSystemIP>\<ClientName>-Farm.

Note or copy the Network Path.

9 On the workstation used to recover items, map a network drive to the following location: *<SourceSystemIP>**<ClientName>*-*<Instance>*

10 Launch Explorer; right-click **Computer** and select **Map Network Drive**.

11 In the Folder field, enter the share displayed in the Network Path field.

12 Click **Finish**.

13 Recover the desired items using Windows or another third-party tool. For details on using Kroll, see "Recovering items with Kroll" on page 356.

Note: Creating the recovery object can take some time. If you do not see any available items, check back later.

14 Disconnect the network share once items have been recovered by right-clicking the share and selecting **Disconnect**.

15 On the backup appliance, tear down the recovery object as described in "To view or tear down SharePoint recovery objects" on page 354.

> **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

> **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 252.) 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" on page 355.

1 Select **Recover** and click the **Backup Catalog** tab.

Use Filter Backups to the right to customize the backups that display.

2 Expand the SharePoint Farm and select a backup or imported backup copy.

To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.

3 Click **Recover Files**.

4 Define the recovery options.

5 Click **Next**.

6 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 Kroll

Using Kroll 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 Kroll license must be applied. Speak with your Unitrends sales representative for information about obtaining a Kroll license.

The following procedure walks you through a typical item-level recovery. For a complete overview of Kroll Ontrack PowerControls for SharePoint, including procedures and limitations, see Kroll's SharePoint user guide available at http://www.krollontrack.com/support/user-guide-and-manuals/.

To recover items using Kroll

Note: Creating the recovery object can take some time. If you do not see any available items in KOP, check back later.

1. Follow step 1 on page 355 through step 12 on page 355 in "To recover a SharePoint farm, site, service, or item" on page 355.
2. From your Kroll 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 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 354.

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/<client_name>/ <instance>). In the second phase, the client 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 client. See the Bare Metal Protection and Recovery Guide 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 360 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 KB 3354.

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.

To recover an Oracle backup

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.

1. Select **Recover** and click the **Backup Catalog** tab.
2. Expand the Oracle asset and select the backup.

   To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.
3. Click **Recover**.
4. Select the desired recovery options.
5. 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 360.

To recover items from an Oracle backup

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.

1. Select **Recover** and click the **Backup Catalog** tab.
2. Expand the Oracle asset and select the backup.

   To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.
3. Click **Recover Files**.
4. Click **Confirm** to create the recovery object and start the recovery.
5. Click **OK** to clear the **Notice** dialog.
   A row for this recovery displays on the **File Level Recovery** tab.
6. 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.
7. On the workstation to be used to recover files, map a network drive to \\<*ApplianceIP*>\<*InstanceID*>.
8. Recover the desired items using an Oracle or third-party tool.
9. Disconnect the network share once files are recovered. To disconnect the share, select the share and click **Remove**.
10. 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 system. Because Oracle recovery to an alternate server is not supported, this procedure requires that you perform a disaster recovery (DR) of the Oracle client to a new client that is directly attached to the backup copy target. See the Bare Metal Protection and Recovery Guide for details.

Once the client has been recovered, you perform an Oracle granular recovery. This operation is only supported on Oracle on Windows platforms running Oracle 10g, 11g or 12c.

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 replicated backup of the source client taken after the database was deployed is required. To perform DR using Windows integrated recovery, the backup copy target must have been run

using version 7.3 or higher for BIOS-based clients, or version 7.4 or higher for UEFI-based clients. To perform DR using legacy Windows hot 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 the Bare Metal Protection and Recovery Guide 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. Click **Filter**.
6. Expand the list of backups under the Oracle DB instance and select a backup to recover.
7. Click **Recover Files**.
8. Click the **File Level Recovery** tab.
9. Select the share and click **Show Details**.

   Record or copy the CIFS path as it displays in the File Level Recovery Details window.
10. Log in to the target you created in step 1 on page 361.
11. Open Windows Explorer.
12. In Windows Explorer, navigate to the CIFs path recorded in Step 9.
13. Open the file `unitrends-<database>.env`. Note the following values:
    - Oracle SID
    - Oracle Home

- Backup #

14 Open a command-line prompt.

15 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>';
```

16 Run the following commands:

```
# restore database;
# recover database;
# alter database open resetlogs;
# quit;
```

The Oracle recovery is complete.

17 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" on page 362
- "To recover from a service profile backup" on page 363
- "Identifying files in UCS service profile backups" on page 363

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 363.
- 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.

To recover from a service profile backup

Use this procedure to recover the entire backup or selected service profiles, templates, pools, and policies.

1. Select **Recover** and click the **Backup Catalog** tab.

   Use Filter Backups to the right to customize the backups that display.
2. Select a backup or imported backup copy to use for the recovery.

   To import a backup copy, see "To import a cold backup copy" on page 278 or "To import a hot backup copy" on page 275.
3. Click **Recover**.
4. In the File Browser, expand folders to view items in the backup.

   To expand all folders, press * on your keyboard.
5. Select or drag files to recover. To recover the entire backup, select the top-level folder.
6. Click **Next**.
7. 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.
8. Click **Save**.
9. Click **OK** to close the Notice message.

   Selected items are recovered to the target location. To view the running job, select **Jobs > Active Job**.
10. 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-**

Chapter 16: 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 365 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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)
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.

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 365.

Chapter 17: 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 370
- "Performing DR from a cold backup copy" on page 372
- "Licensing the DR target appliance" on page 374

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 368
- "Create a plan for obtaining a fresh DR target appliance" on page 368
- Review the "Requirements and considerations for appliance DR" on page 368

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 backed up each time a hot or cold backup copy runs. 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 appliance – Use a new or re-imaged appliance. If you need to re-image an appliance, see KB 3207 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:

- Both the original (failed) appliance and the DR target appliance must be running Unitrends version 9.1 or higher.
- The DR target appliance must be set up and configured onto the network. See the applicable guide below for instructions:
  - Quick Start Guide for Recovery Series Appliances
  - 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 in Microsoft Azure
  - Deployment Guide for Unitrends Backup Installable Software
  - 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. See "To edit an appliance" on page 55 for details.
- During DR you will choose a storage target where backups will be imported:
  - For Recovery Series 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 70 for details.

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: 2 backup storage targets, *Internal* and *NewBackups*. Backups for Asset1 go to *NewBackups*. • New DR target appliance: 2 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*.

Storage	Description
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: 2 backup storage targets, *Internal* and *NewBackups*. Backups for Asset1 go to *NewBackups*. • New DR target appliance: 2 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 DR target appliance must be set up and configured to meet the general "Requirements and considerations for appliance DR" on page 368.
- If recovering encrypted backups, you must configure encryption on the DR target appliance. While configuring encyprtion, you must enter the encryption passphrase of the original appliance. For details, see "To edit an appliance" on page 55.

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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)

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. After you perform the recovery, license the appliance as described in "Licensing the DR target appliance" on page 374.

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, enter **7** to add its hostname and IP.
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*. (If you have not reset the OS root user password, the default password is *unitrends1*.) 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.

DR tool prompt	Description
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. **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 368.
- If recovering encrypted backups, you must configure encryption on the DR target appliance. While configuring encyprtion, you must enter the encryption passphrase of the original appliance. For details, see "To edit an appliance" on page 55.
- 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 77.
  - 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*. (If you have not reset the OS root user password, the default password is *unitrends1*.)
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. After you perform the recovery, license the appliance as described in "Licensing the DR target appliance" on page 374.

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. See "To edit an appliance" on page 55 for details. Once you have configured encryption, return to the DR tool and continue with the DR procedure.

DR tool prompt	Description
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. **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 appliance"
- "To license a Unitrends Backup appliance" on page 375

To license a Recovery Series 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 appliance" on page 374 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 Recovery Series and Unitrends Backup appliances" on page 21) to request a new license.

  Once you receive the new license, apply it as described in "To license a Recovery Series appliance" above.

Chapter 18: 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:

- "Types of reports" on page 377
- "Working with reports" on page 406

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 378	• "Protection Summary report" on page 389 • "Backup History report" on page 381 • "Legal Hold Backups report" on page 383 • "Backup Failures report" on page 384 • "Weekly Status report" on page 397
"Recover reports" on page 387	• "Recovery History report" on page 387
"Backup Copy reports" on page 389	• "Protection Summary report" on page 389 • "Backup Copy Capacity report" on page 392 • "Backup Copy - Hot Targets report" on page 393 • "Backup Copies - Past 24 Hours report" on page 395 • "Storage Footprint report" on page 395 • "Backup Copy - Cold Targets report" on page 396 • "Weekly Status report" on page 397

Category	Available Reports
"Appliance reports" on page 398	• "Update History report" on page 399 • "Capacity report" on page 399 • "Load report" on page 401 • "Alerts report" on page 401 • "Trap History report" on page 402 • "Notifications report" on page 403
"Storage reports" on page 404	• "Storage report " on page 404 • "Data Reduction report" on page 406

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 details, see the following topics:

- "Protection Summary report" on page 389
- "Backup History report" on page 381
- "Legal Hold Backups report" on page 383
- "Backup Failures report" on page 384
- "Weekly Status report" on page 397

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.

Click **>Report Categories** in the upper-left of the report 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, virtual hosts, and VM templates.

Refer to the following table for a description of the charts and available report columns. Not all columns display by default. To customize your output, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

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, virtual hosts, and VM templates.
Total Backup Copies	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, virtual hosts, and VM templates.
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 Copies	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.
System Name	The name of the backup appliance.
Client Name	The name of the asset.
Instance Name	The name of the protected database, storage group, or virtual machine (if applicable).

Item	Description
Last Backup	Date and time of this asset's last successful backup that completed within the specified date range.
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, 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).
Backup Table Column	Lists all assets that have been protected by backups within the selected date range and provides summary backup information.
System Name	The name of the appliance on which the asset is configured.
Client Name	The name of the asset.
Instance Name	The name of the protected database, storage group, or virtual machine (if applicable).
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, 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	Date and time when this asset's last successful backup job completed, within the specified date range.

Item	Description
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.
System Name	The name of the appliance on which the asset is configured.
Client Name	The name of the asset.
Instance Name	The name of the protected database, storage group, or virtual machine (if applicable).
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, Oracle, Hyper-V, or VMware.
Backup Copy 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).
Backup Copy Successes	The number of backup jobs that completed successfully within the specified date range.
Backup Copy Failures	The number of backup jobs that failed within the specified date range.
Last Backup Copy	Date and time when this asset's last successful backup job completed, 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 asset-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 of the report to return to the Backup category of 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 accordion icon to the right of the column header. See

"To customize a report" on page 406 for details.

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 asset-level backups, contains 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).
Database Name	For application backups, contains the name of the protected database. 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 Date	The date and time at which the backup job completed.
Files	For asset-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.

Report Field	Field Description
Replication 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.
Start Date	The date and time that the backup job started.
Status	Status of the backup job: Success, Warning, or Failure.
Appliance Name	The name of the backup appliance.
Type	Backup mode: Full, Incremental, Differential, Synthetic Full, Selective, Bare Metal, or Transaction (SQL only).
Job Name	Name of the backup job. *Job No Longer Exists* displays if this backup is no longer stored on the appliance.
Certify Status	Indicates whether this Recovery Point has been certified by the Unitrends ReliableDR product: Success (has been certified) or Failure (has not been 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).
Speed (MB/S)	Average transfer rate of the job, in megabytes per second.

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.

Click **< Report Categories** in the upper-left of the report 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 accordion icon to the right of the column header. See

"To customize a report" on page 406 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.
System Name	The name of the backup appliance.
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.

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 of the report to return to the Backup category of 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.

Refer to the following table for a description of the the available report columns. Not all columns display by default. To customize your output, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

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.
Protected Asset	The name of the asset.
Complete	Indicates whether the job completed: True (yes) or False (no).
VM Name	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	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.
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.
Replication Status	N/A, which indicates that the backup has not been copied to another Unitrends appliance. (Failed backups are not copied.)
Size (GB)	The size of the backup, expressed in gigabytes.

Report Field	Field Description
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 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.

Refer to the following table for a description of the report columns. To customize your display, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

Item	Description
Asset Name	The name of the asset.
Appliance Name	The name of the appliance where the jobs ran.

Item	Description
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.

Recover reports

The Recovery History report summarizes and organizes historical information about recovery jobs.

For details, see the "Recovery History report" on page 387.

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 and Hyper-V file-level recovery jobs.

The report does include information about Windows Instant Recovery operations.

When finished, click **< Report Categories** in the upper-left of the report to return to the Recovery category of the Available Reports page.

Refer to the following table for a description of the available available report columns. Not all columns display by default. To customize your output, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

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.
Certify Status	Indicates whether this Recovery Point has been certified by the Unitrends ReliableDR product.
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. For host-level backups, contains the name of the protected virtual machine.
Elapsed	The amount of time it took for the job to complete, in hh:mm:ss format.
Encrypted	Indicates whether the data is encrypted.
End Time	The date and time at which the recovery job completed.
# Files	The number of files recovered. (For recovery from host-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)	Amount of data, in gigabytes.
Start Time	The date and time that the recovery job started.
Status	Status of the recovery job: Success, Warning, or Failure.
Appliance Name	The name of the backup appliance.

Report Field	Field Description
Type	The type of recovery operation: • *Restore* for a standard backup recovery • *Virtual Restore* for Windows Instant Recovery or VM Instant Recovery. • *Bare Metal Restore* for integrated bare metal recovery.

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 44.

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.

Click **>Report Categories** in the upper-left of the report 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, virtual hosts, and VM templates.

Refer to the following table for a description of the charts and available report columns. Not all columns display by default. To customize your output, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

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, virtual hosts, and VM templates.
Total Backup Copies	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, virtual hosts, and VM templates.
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 Copies	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.
System Name	The name of the backup appliance.
Client Name	The name of the asset.
Instance Name	The name of the protected database, storage group, or virtual machine (if applicable).
Last Backup	Date and time of this asset's last successful backup that completed within the specified date range.
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, Oracle, Hyper-V, or VMware.

Item	Description
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).
Backup Table Column	Lists all assets that have been protected by backups within the selected date range and provides summary backup information.
System Name	The name of the appliance on which the asset is configured.
Client Name	The name of the asset.
Instance Name	The name of the protected database, storage group, or virtual machine (if applicable).
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, 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	Date and time when this asset's last successful backup job completed, within the specified date range.
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.
System Name	The name of the appliance on which the asset is configured.

Item	Description
Client Name	The name of the asset.
Instance Name	The name of the protected database, storage group, or virtual machine (if applicable).
Asset Type	The type of protected asset. Examples: SQL Server, Windows 2012, Oracle, Hyper-V, or VMware.
Backup Copy 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).
Backup Copy Successes	The number of backup jobs that completed successfully within the specified date range.
Backup Copy Failures	The number of backup jobs that failed within the specified date range.
Last Backup Copy	Date and time when this asset's last successful backup job completed, 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 and entering a date range. When finished, click **< Report Categories** in the upper-left of the report to return to the Backup Copy category of the Available Reports page.

Refer to the following table for a description of the available report columns. To customize your output, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

Report Field	Field Description
Application 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 appliance configuration, contains *System Metadata*. System Metadata is automatically copied to the target during backup copy jobs and contains system information, such as appliance configuration and settings.
Protected Asset	The name of the asset.
Instance	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.
Database Name	For application backups, contains the name of the protected database. For host-level backups, contains the name of the protected virtual machine.
Size (GB)	The size of the 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 asset-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 of the report to return to the Backup Copy category of 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 accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

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 appliance configuration, contains *System Metadata*. System Metadata is automatically copied to the target during backup copy jobs and contains system information, such as appliance configuration and settings.
Protected Asset	The name of the protected asset.
Complete	Indicates whether the job completed: True (yes) or False (no).
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 Time	The amount of time a backup copy job took to complete, in hh:mm:ss format.
Encrypted	Indicates whether the backup copy is encrypted.
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.

Report Field	Field Description
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	The amount of time it took for the job to complete, in hh:mm:ss format.
Status	Status of the backup copy job: Success, Warning, or Failure.
Source Appliance	The name of the appliance on which the asset is configured.
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 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).
Speed (MB/S)	Average transfer rate of the job, in megabytes per second.

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 of the report to return to the Backup Copy category of 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 of the report to return to the Backup Copy category of 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 accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

Report Field	Field Description
Appliance Name	The name of the backup appliance.
App Name	Application type: • For asset-level backups, contains *file-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 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.

Report Field	Field Description
Instance	Protected instance: • For asset-level backups, contains *file-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. – 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.
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.

Refer to the following table for a description of the report columns. To customize your display, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

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.

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 399
- "Capacity report" on page 399
- "Load report" on page 401
- "Alerts report" on page 401
- "Trap History report" on page 402
- "Notifications report" on page 403

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 of the report to return to the Appliance category of the Available Reports page.

Refer to the following table for a description of the available report columns. To customize your output, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

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 of the report to return to the Appliance category of the Available Reports page.

Refer to the following table for a description of the appliance summary and report columns. To customize your output, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

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.

Report field	Field description
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 **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.
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, or XenServer (if applicable).

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 asset-level backups, contains 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.

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 of the report to return to the Appliance category of the Available Reports page.

Refer to the following table 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 classifications: 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 of the report to return to the Appliance category of the Available Reports page.

Refer to the following table for a description of the available report columns. To customize your output, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

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.
Source	The component that generated the alert.
Time	The date and time the alert was generated.
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 53).

To view the report, select an appliance, enter a date range, and click **Generate Report**.

Click **>Report Categories** in the upper-left of the report to return to the Available Reports page.

Refer to the following table for a description of the report columns. Not all columns display by default. To customize your display, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

Item	Description
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).

Item	Description
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.
Appliance ID	ID of the appliance that generated the trap.
Appliance name	Name of the appliance that generated 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**.

Click **>Report Categories** in the upper-left of the report to return to the Available Reports page.

Refer to the following table for a description of the report columns. Not all columns display by default. To customize your display, click the accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

Item	Description
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.
Message	Message text.

Item	Description
Category	Indicates the area impacted by the notification. For example, Backup, Storage, Management, etc.
ID	ID of the notification message template.

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 404
- "Data Reduction report" on page 406

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 of the report to return to the Storage category of 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 accordion icon to the right of the column header. See "To customize a report" on page 406 for details.

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.
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.

Report Field	Field description
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 day, in green.
Storage Table	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 average amount that used storage increased from one day to the next day. Averaged over the last 30 days.
Daily Growth	The average amount that used storage increased from one day to the next day. Averaged over the last 30 days.

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 of the report to return to the Storage category of the Available Reports page.

Working with reports

Use the following procedures to generate, export, and distribute reports:

- "To customize a report" on page 406
- "To generate a report" on page 406
- "To export a report" on page 407
- "Emailing reports" on page 407
- "Modifying the From address for System reports" on page 408

To customize a report

When generating a report, you can customize the output as follows:

- Click the accordion icon to the right of the column header to add or remove columns:

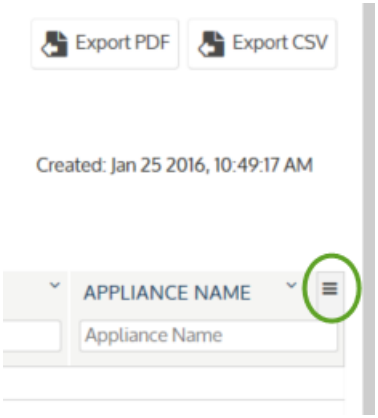

  - Click Reset column defaults to restore the default display.
  - Click Clear all filters to clear any column sorting you have applied.
- Click the arrow next to a column name to:
  - Sort values in ascending or descending order.
  - Adjust the column width.
  - Hide a column.
  - Pin the position of a column to the left or right.

To generate a report

1. Select the **Reports** page from the UI.
2. On the Reports page, click a **Category**.

> **Note:** The displayed Available Reports list varies depending on the Category you select. (When selecting the Reports page, the category defaults to Backup and displays the Available Reports list for that category.)

3. Click the applicable report from the Available Reports list.
4. If necessary, select the appliance.
5. Specify the date range by selecting a From date and a To date.
6. Click **Generate Report**.

> **Note:** After generating a report, you can configure the number of rows that display per page. Use the **Items per page** field below the bottom of the report to adjust this value. To modify the columns that a generated report displays, click the accordion icon (found to the right of the last column heading). To clear or modify the filters, select or deselect columns to display or hide.

7. When finished, click **< Report Categories** in the upper-left of the report to return to the Available Reports list

To export a report

1. To export a report as a PDF, or a CSV file, select the **Reports** page from the UI.
2. On the **Reports** page, click a **Category**.

> **Note:** The displayed Available Reports list varies depending on the Category selected. (When selecting the Reports page, the category defaults to Backup and displays the Available Reports list for that category.)

3. Click the applicable report from the Available Reports list.
4. If necessary, select the appliance.
5. Specify the date range by selecting a From date and a To date.
6. Click **Generate Report**.

> **Note:** After generating a report, you can configure the number of rows that display per page. Use the **Items per page** field below the bottom of the report to adjust this value. To modify the columns displayed in a generated report, click the accordion icon (found to the right of the last column header). To clear or modify the filters, select or deselect columns to display or hide.

7. Once a report generates, click **Export PDF** to export and save it as a PDF file, or click **Export CSV** to export and save it as a CSV (Excel) file.
8. When finished, click **< Report Categories** in the upper-left of the report 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 48.

Modifying the *From* address for System reports

If you have configured email recipients to receive System reports, the appliance sends these reports each day: System 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. On the Edit Appliance dialog, click the **Advanced** tab.
3. Click **General Configuration**.
4. In the General Configuration (Advanced) dialog, click the **Name** column to sort the names alphabetically.
5. Scroll down and select this Name: **SystemReportFromDomainNotHostname**.
6. Enter **YES** in the Value field and click **Save**.