FALCONSTOR

FalconStor HABANERO

USER GUIDE

FalconStor® Habanero™ User Guide

FalconStor Software, Inc.
701 Brazos Street, Suite 400
Austin, TX 78701 USA
Phone: 631-777-5188
Website: www.falconstor.com

12726

Contents

Introduction

FalconStor Habanero is a software-as-a-service offering designed to simplify secure offsite data protection for IBM Power customers. By integrating directly with existing IBM Power workloads, backup applications, and established operational processes, Habanero enables customers to establish cloud-based enterprise-grade offsite protection without deploying new infrastructure or changing how backups are run today.

Habanero provides enterprises secure offsite protection to meet recovery, compliance, and cyber-resilience requirements.

Web Portal

The portal is divided into five main sections, *Dashboard, Monitor, Manage, Analytics*, and *Administration*. Each section can be accessed from a tab on the menu bar that appears at the top of each portal page.

Dashboard

The initial landing page for the portal is the dashboard, a configurable display that provides a global view of your environment.

The dashboard is customizable and can have one or multiple tabs where users can add specific widgets to monitor resources. The widget information is updated every 10 seconds

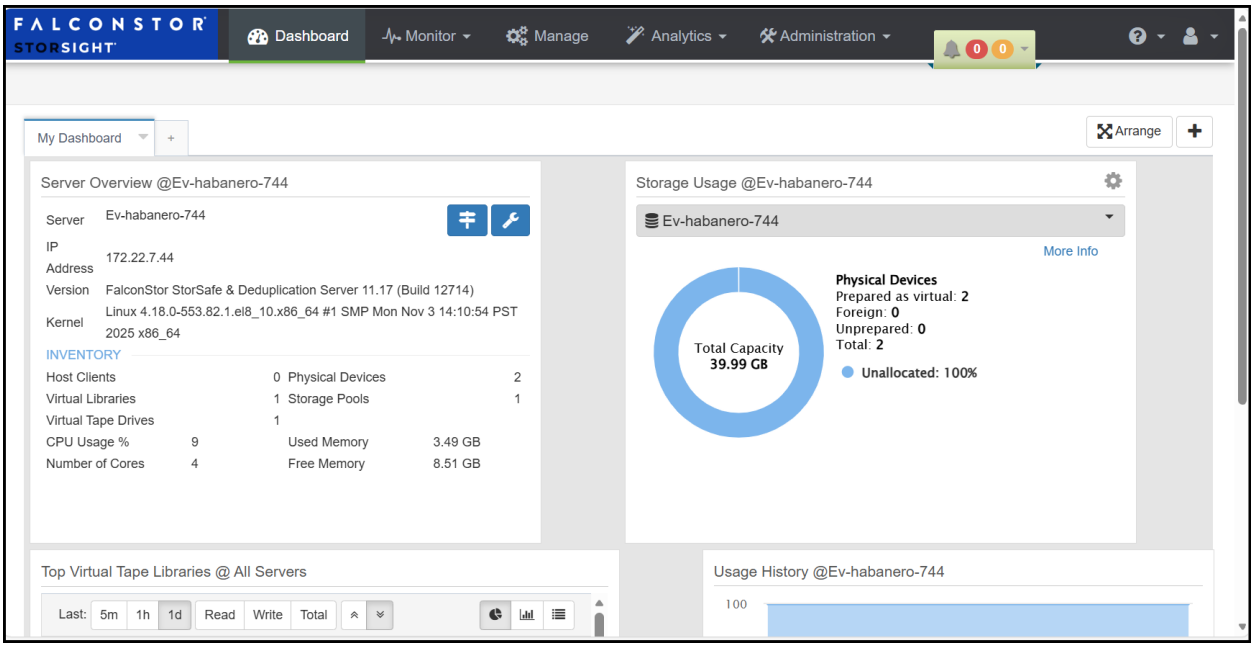

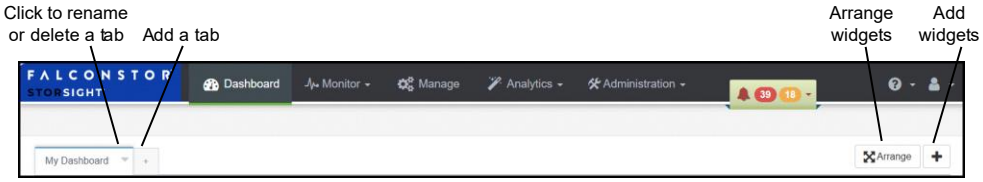

Widgets

Click the “+” icon on the far right to select from a list of available widgets. You can have up to eight widgets on a tab.

Click the *Arrange* icon to move widgets. Click the *Arrange* icon again to exit the move mode.

Click the *Manage* icon  on a widget to change the title and applicable parameters, or to remove the widget.

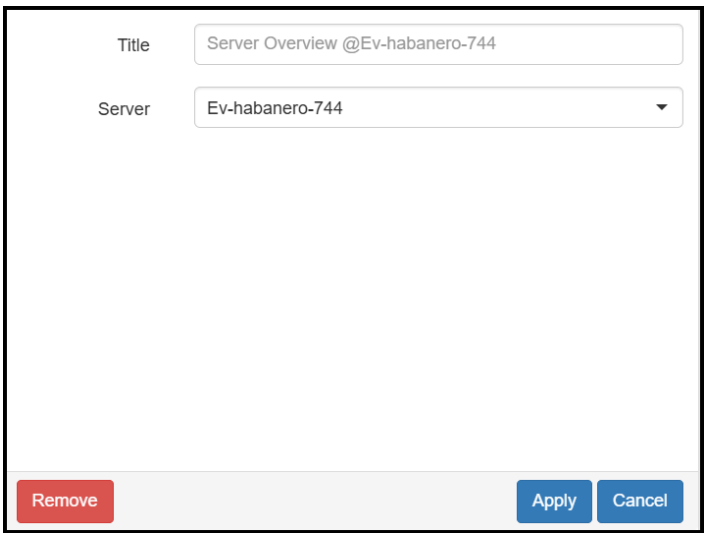

The following widgets are available:

Alerts widget displays alerts for all servers by default. You can select to display alerts for a specific server.

Server Overview widget displays general information, current storage usage, and usage history for the selected server. Click the *Server Ports* icon  to view the status of the server's ports and modules. Click the *Installed Patches* icon  to view patches on this server.

Storage Usage widget displays the server capacity graph and the number of devices. Click the *More Info* link to display details in a table.

Usage History widget displays a graph showing the amount of free and used storage capacity over the last week. You can highlight a point on the graph to see more detailed information.

Top Virtual Tape Libraries widget displays the virtual libraries with the highest read/write/total throughput for the past five minutes, one hour, or one day.

Global portal functions

Some functions are available on multiple pages in the portal.

Help 

The help icon appears on the top right corner of the portal and allows you to view StorSight version, build number, and end user license agreement.

User profile 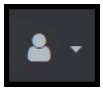

The user profile icon appears on the top right corner of the portal and allows you to set user settings, change your password, and log out of the server.

Use settings include the number of lines to display per page, which is valid for all pages with many potential items, like tapes and alerts. The default number of items per page is 50. The minimum is 25 and the maximum is 1000.

You can also set the automatic refresh of for reports deduplication, replication, and import/export jobs on the *Jobs* page. When enabled, the refresh rate can be set for every 10 seconds or every minute.

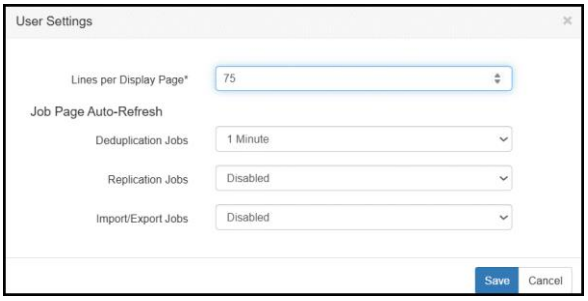

Column sort 

Click any column heading to sort the entire list based upon values in that column. You will see a triangle or inverted triangle on the column being sorted depending upon the direction of the sort.

Server sort 

Click the drop-down arrow to change the order of the server display.

Refresh 

Click the *Refresh* icon to update the page with the most recent information that is available. If you previously performed a search on the page, only the matching items will be displayed after refreshing. If you previously filtered information in a column, the unfiltered view will be displayed after refreshing.

Export/Send 

Click the *Export/Send* icon to export the information on the current page to a PDF or CSV file. If email and/or SMS have been configured by your administrator, you can also send a PDF or CSV file via email. When you export, you can convert size units of data in tables in order to have a unique unit for all values, for example, in a CSV, column.

Save/Reset Layout 

Click the *Save/Reset Layout* icon to save the current column layout or reset it to the system default.

Filter 

Click the *Filter* icon to filter the information displayed in a column. Click the three bars on a column, select items to display. You can search for items, from the *Search* box.

Arrange 

Click the *Arrange* icon to arrange columns in a table. Check or uncheck a column name to show or hide it. Click-hold a column name and drag it to another position to change the order of columns horizontally.

To permanently keep the new column layout, click the *Save/Reset Layout* icon and select *Save Layout*.

Search 

Type one or multiple strings, separated by space, in the *Search for* box to search all visible text in the current page. Clear the search field to return to the full view of the list.

If you need to search for text in a field that links to a pop-up window, use the *Filter* icon,

Configuration Overview

Perform the following steps to configure Habanero:

1. Open any HTML5-capable browser to connect to StorSight, the Habanero web portal, using the IP address, user account, and password that were provided to you by FalconStor.
2. Create your virtual tape libraries. Refer to Virtual Tape Libraries for more information.
3. Create virtual tapes for your newly created virtual tape libraries. Refer to Virtual Tape Libraries for more information.
4. Configure your host clients that are the backup application servers that are assigned to virtual tape libraries and drives for backup. Refer to Host Clients for more information.
5. Create a replication policy designating a server as the target server. Refer to Policies for more information.
6. If you want to use the Cloud Archive option to export tapes, add the object storage account using the credentials of the bucket that was provided to you by FalconStor. Refer to Object Storage for more information.

Virtual Tape Libraries

Tape Libraries

Virtual tape libraries emulate physical tape libraries and are used for backup by third-party tape or disk backup software, database backup utilities, and archiving applications.

Each virtual tape library consists of one or more virtual tape drives and can have one or more virtual tapes. Each virtual tape library and drive can be assigned to one or more backup application servers (clients).

Select *Manage* from the menu bar and click the *Virtual Libraries* tab, to display and manage the virtual tape libraries on the selected server.

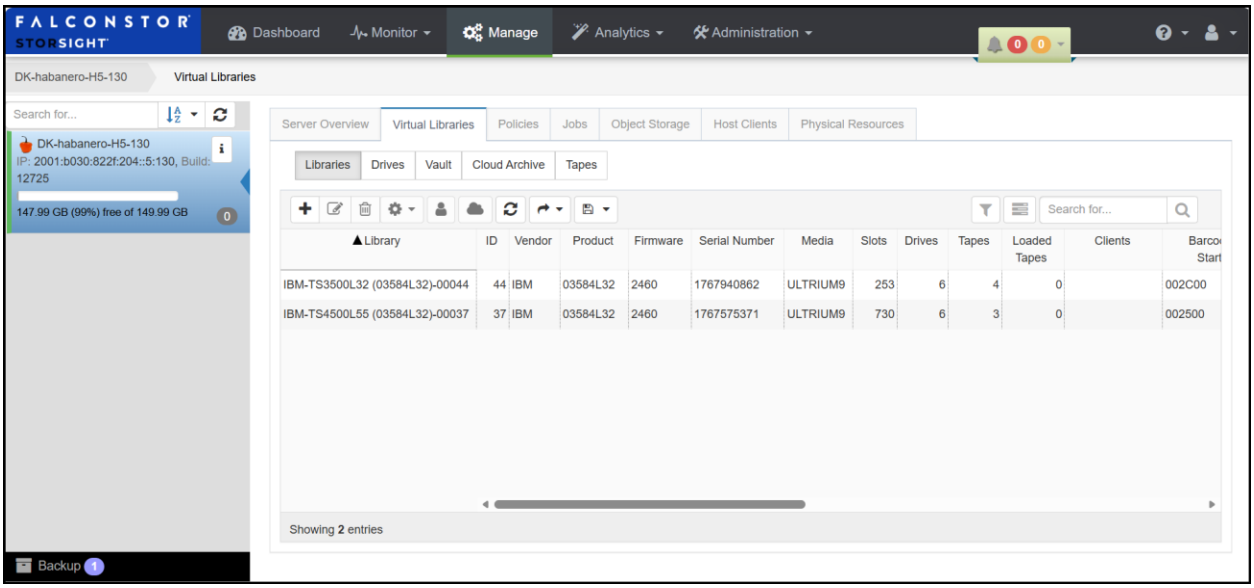

Create a virtual tape library

1. Click the “+” (*Create*) icon.
2. On the *General* tab, specify general library properties.

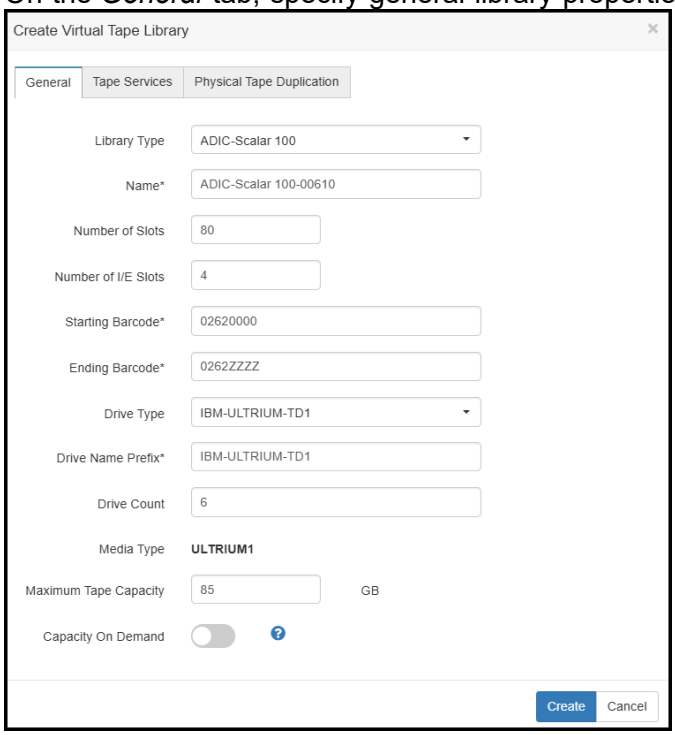

Library Type - Select the physical tape library you are emulating. If you have a physical tape library, you need to create a virtual tape library that resembles it in order for the virtual tapes to use the same format as the physical tapes. This is important for importing and exporting functions and guarantees that your backup application will accept the tapes.

For IBM i host clients, select the virtual tape library type as FalconStor FALCON TS3500L32 (03584L32) or FALCON TS3500L32 (03584L32) and the media type as ULTRIUM3 (LT03) or newer. By using Falcon library types, you get the 3584-403 device types to configure on IBM i host clients.

Name - Specify the library name, up to 64 characters. excluding < > & $ " / \’

Number of Slots - Maximum number of tape slots in your tape library.

Number of I/E Slots - Maximum number of slots used to take tapes in and out of the bin. Set the number of Import/Export slots to 1.

Starting/Ending Barcode - Indicate a range of barcodes that will be used when creating virtual tapes. By default, barcodes increment in an alphanumeric sequence; for example, XXX0009 to XXX000A. In order to set the barcode to increment in a numeric sequence (XXX0009 to XXX0010), you have to set the last three digits of the Ending Barcode field to 999; for example, XXX0999

Note that for IBM libraries, the default barcode range is set to six characters.

Barcode Suffix - Add the media type as a suffix to new tape barcodes. This is applicable when the length of barcode is a maximum of six characters.

Auto Loader - Support auto-loader mode on specific IBM libraries.

- *Leave Tapes in Library* - (Default mode) Tapes remain in the drive with an eject status. Once a tape is moved from the drive to the virtual vault by the user, the next available tape will be moved from the slot to a drive.
- *Move Tapes to Vault* - Move tapes from the tape drive to the virtual vault when a tape is unloaded.
- *Loop Tapes* - Unload the current tape to move it to a slot and load the next tape to the drive.

Drive Type - Select the drives used in your library.

Drive Name Prefix - The prefix is combined with a number to form the name of the virtual drive. The drive name can be up to 59 characters excluding < > & $ " / \'

Drive Count - Determines the number of virtual tape drives available. This translates into the number of concurrent backup jobs that can run. Backup software licensing considerations may affect the number of tape drives you wish to present to each client server. This number can exceed the standard number of drives for the library as long as the backup software supports it.

Media Type - Informational.

Maximum Tape Capacity - Indicate the maximum size for each tape, up to 8,192 GB.

- If you will not be exporting data to physical tape, you can enter any maximum capacity.
- If you will be exporting data to physical tape but you will not be using hardware or software compression, you can enter any maximum capacity, but if you enter a capacity that exceeds the native uncompressed capacity for the media, you may not be able to export to physical tape.
- If you will be exporting data to physical tape and you will be using hardware or software compression, you should set the maximum capacity to 15% less than the uncompressed capacity of the selected media. (A 15% reduction is the default value). This is because the compression algorithm can vary depending upon the dataset; certain file types (ZIP, PDF, GIF, RAR, etc.) are already compressed and cannot be compressed further.

Capacity On Demand – Enable *Capacity On Demand* to create small resources for your tapes and then automatically allocate additional space when needed. This can save considerable amounts of disk space without affecting system performance.

Initial Size/Incremental Size - The minimum value for the incremental size is (Maximum Capacity – Initial Tape Size) / 63. Leave the default settings.

Encrypt Data - Indicate if you want to use encryption for this virtual tape library. When encryption is enabled, each new tape that is created in the library is encrypted with the selected key. Each encrypted tape always retains its key, even if it is moved to another library.

Tapes moved to/from this library will preserve their encryption status. This means that unencrypted tapes moved to this library will not be encrypted and encrypted tapes will not change their key to the key used by the library.

If encryption is ever disabled for this library, tapes created afterward will not be encrypted. Therefore, each library can have both encrypted and unencrypted tapes. Also, if the library properties are changed to use a different key, existing tapes will retain their key and new tapes will be created with the newly designated key.

3. Do not enable any service options on the *Tape Services* tab unless you are using a physical tape library and want to export virtual tape data to physical tapes. Refer to *Configure Cloud Archive* section below to export tape data to cloud.
4. If you are using a physical tape library, on the *Physical Tape Duplication* tab, specify if you want to make up to five duplicate copies of a physical tape whenever virtual tape data is exported to a local physical tape.
5. Click *Create* when done.

Create virtual tapes

1. Select a virtual tape library, click the *Manage* icon  , and select *Create Tapes*.

   You can also create virtual tapes for a library from the *Tapes* tab.
2. On the *General* tab, specify general tape properties.

   Name Prefix - The prefix can be up to 25 characters excluding spaces or special characters < > & $ " / \' * ? | %

   Size - Because *Capacity on Demand* is enabled for the library, the default is the initial size.

   Storage Type - Select your storage pool or physical device(s) from which to create the tapes.

   Number of Tapes - The maximum number of tapes is the number of slots in the library. You cannot specify if Physical Tape Archive for Replica is configured for the library or if you are matching physical tapes in a Physical Tape Archive configuration.

   Starting Barcode - If desired, specify the starting barcode. The default is to generate a barcode based on the available barcode range for the library.

   WORM - Indicate if you want the virtual tape to be a write-once-read-many (WORM) tape. A WORM tape cannot be overwritten by backup software. Only virtual tapes that support ULTRIUM5 media type and above can be configured as WORM tapes. WORM tapes are not supported with Veritas NetBackup WENCR media (WORM media on which NetBackup encrypts data).
3. On the *Service* tab, specify what services you want to enable for the virtual tape and select the policy to use.

   Physical Tape Archive for Replica automatically exports the contents of a replicated virtual index tape (VIT) to a physical tape at the remote site.

   If Physical Tape Archive is enabled for the virtual library, specify if you want to create virtual tapes that match physical tapes in the physical library.
4. Click *Create* when done.

Add virtual tape drives to a library

You can add virtual tape drives to a virtual tape library. The number of virtual tape drives translates into the number of concurrent backup jobs that can run. Backup software licensing considerations may affect the number of tape drives you wish to present to each client server. This number can exceed the standard number of drives for the library as long as the backup software supports it.

1. Select a virtual tape library, click the *Manage* icon  , and select *Add Drive*.

   If desired, specify a drive name prefix up to 59 characters excluding < > & $ " / \'.

   Indicate how many tape drives to add.
2. Click *Create* when done.

Configure Cloud Archive for a virtual tape library

Habanero's Cloud Archive can be used to archive virtual tape data to the cloud. Policies can be set on virtual tape libraries to automatically migrate tape data to object storage as soon as a backup is complete and a tape is ejected to the vault.

You can only use Cloud Archive if you are not currently using Physical Tape Archive on the virtual tape library.

You need an object storage account from a provider where buckets are created to hold tape data.

If replication has been enabled for a tape that is configured for migration to object storage, data replication occurs according to the policy as long as the tape has not been moved to the vault. Once the tape is moved to the vault, it is added to the migration job queue. When migration starts, replication will be performed only when migration is complete, before the tape becomes a stub tape. The replica tape will not be automatically migrated. Migration only applies when a tape is ejected to the vault.

If a tape migrated to object storage is deleted, it will also be deleted from object storage. However, if the object storage used for migrated tapes is deleted on the provider, those tapes need to be manually deleted. There is no automatic tape deletion on Habanero.

While Cloud Archive is set at the virtual tape library level, it can be modified for an individual tape while it is in the virtual library (not while it is in the vault).

If you enable Cloud Archive on a virtual tape library with existing tapes, those tapes will inherit the library's Cloud Archive properties when they are ejected to the vault.

If you are using Cloud Archive on a Habanero system with existing tapes that need to be migrated, you can manually trigger migration to object storage from the *Vault* tab.

1. Inform FalconStor of the retention and immutability options that you need for your buckets.
2. Refer to Object Storage to add an object storage account using credentials and bucket name provided to you by FalconStor.
3. Select a virtual tape library and click the Cloud Archive icon  .
4. Enable *Cloud Archive* checkbox.
5. Select an existing object storage account.

   Each library can use a different object storage account.

   While Cloud Archive is set at the virtual tape library level, it can be modified for an individual tape while it is in the virtual library (not while it is in the vault).

   If you enable Cloud Archive on a virtual tape library with existing tapes, those tapes will inherit the library's Cloud Archive properties when they are ejected to the vault.
6. Select the mode.

   Migration can be configured in *Copy* mode or *Move* mode.
   In *Copy* mode, the source tape stays in the vault after migration and is not automatically converted to a stub tape. Virtual tapes can still be manually converted to stub tapes.

   In *Move* mode, the source virtual tape is converted to a stub tape and all related disk space used by the virtual tape is freed once migration to object storage completes successfully. Since the tape is converted to a stub tape once the migration completes, any replication

policies that include the tape become obsolete. If you select *Move* mode, specify the number of days (up to 90, where 0 means immediate) to keep virtual tapes before they get converted to stub tapes.

7. Click *Save* when you are done.

Tape Drives

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, and click the *Drives* tab. This page lists virtual tape drives on the selected Habanero server.

Select to view standalone tape drives or tape drives in a specific tape library from the *Library* drop-down list.

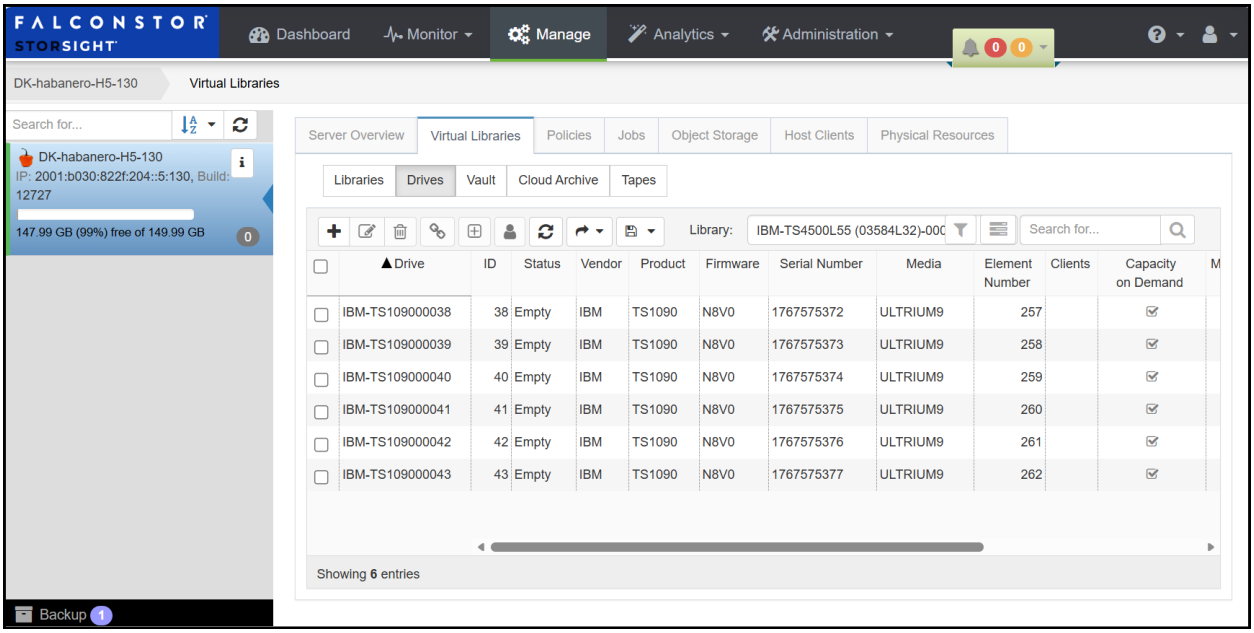

Create a standalone virtual tape drive

You can create standalone virtual tape drives that emulate your physical tape drives. Each virtual tape drive can be assigned to one or more backup application servers (clients).

1. Select *Standalone Drives* as the *Library*. Click the "+" (*Create*) icon.
2. Enter information about the standalone virtual tape drive.

   Standalone Tape Drive Type - Select the type of drive.

   Number of Drive to Create - Specify the number of virtual tape drives to create. This translates into the number of concurrent backup jobs that can run.

   Media Type - Informational.

   Drive Name Prefix - The prefix is combined with a number to form the name of the virtual drive. The drive name can be up to 59 characters excluding < > & $ " / \'

   Maximum Tape Capacity - Indicate the maximum tape size, up to 8,192 GB.
3. Click *Create* when done.

   To create a tape for this standalone virtual drive, select the drive and click *Create Tape*.

Tapes

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, and select the *Tapes* tab. This page lists virtual tapes on the selected Habanero server.

Select the tape location from the *Location* drop-down list. You can view tapes in a specific tape library, replica tapes, or all tapes on this server. If you select to view tapes in a specific library, select the library from the drop-down list.

Replica tapes store data from local and remotely replicated virtual tapes. Clients do not have access to replica tapes.

The total number of tapes is displayed at the bottom of the page. You can also select how many tapes to display per page.

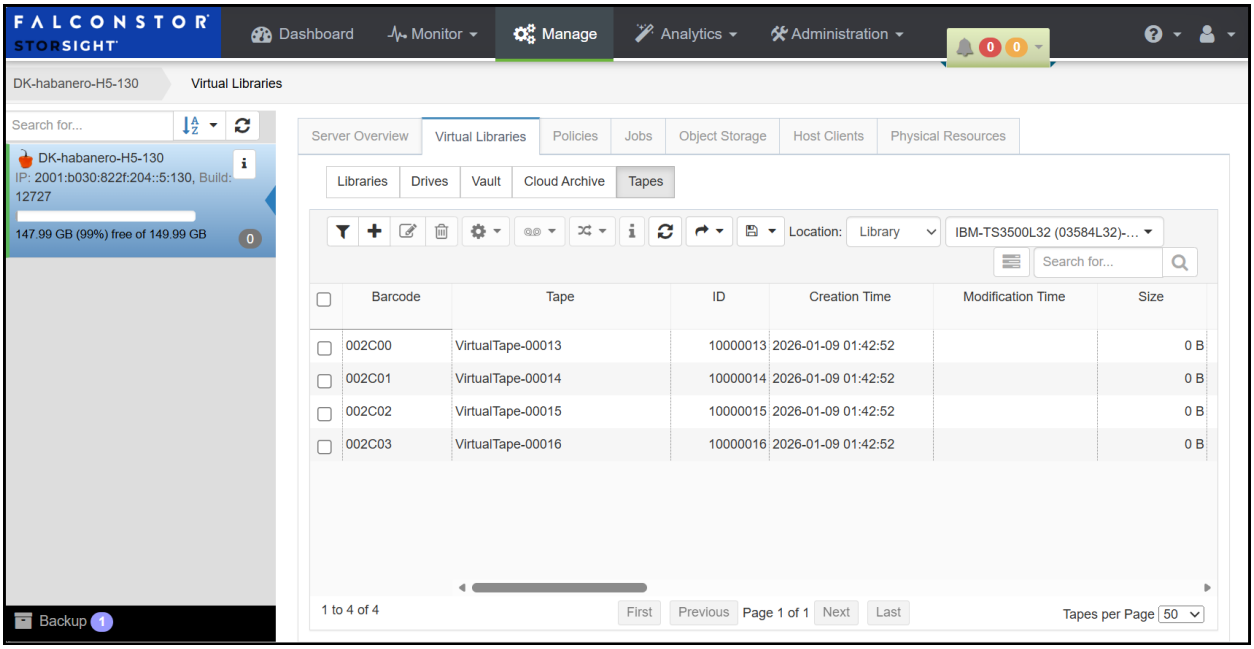

Filter tapes

Because it is possible to have a large number of virtual tapes, there are several tools you can use to locate tapes.

In addition to using the *Search for* box to locate a tape on the current page, you can filter the list of tapes by different criteria, including barcode, type of tape, features enabled, media type, etc. for a specific library, replica resources, or the whole server.

To filter tapes:

1. Select a location and click the *Filter Tapes* icon. 
2. On the *General* tab, you can indicate the type of tape(s) you are looking for
3. On the *Range* tab, you can enter a range of barcodes and/or sizes.

   For example, you can enter an individual barcode, a barcode prefix, or a range of barcodes.

If you want to specify a barcode range, select *Start With* or *End With* in the *From/To* fields. You can then type the number in the box. You can also specify to start with the first/last barcode.

If you want to specify a range for used or available tape size, select *Start With* or *End With* in the *From/To* fields. You can then type the number in the box. You can also specify to start with the lowest or highest size.

You can use multiple filters to further narrow your search. For example, you may want to locate empty tapes (select on the *General* tab) within a specific barcode range.

4. On the *Time* tab, you can enter a specific time or a range of times based on when a tape was created or modified.

   If you want to specify a particular date/time, select *Start With* or *End With* in the *From/To* fields.
5. Click *Apply*.

   The tapes that match the selected criteria will be displayed. You can click the *Reset* link when you are done.

Create virtual tapes

Select a library and click the '+' icon. Refer to Create virtual tapes for more information.

Update virtual tapes

Select a tape and click the *Update* icon  .

You can change the following properties:

- Change a virtual tape barcode
- Write protect a tape
- Configure Cloud Archive, Physical Tape Archive, and/or Physical Tape Duplication - By default, tapes inherit their configuration from the virtual tape library

Manage virtual tapes

Select a tape, click the *Manage* icon,  and select the action you want to perform.

You can perform the following functions for a virtual tape:

- Rename - The new name can be up to 32 characters excluding spaces or special characters < > & $ " / \' * ? | %
- Move:
  - Move one or more tapes to the vault from a library or a standalone drive. You cannot move replica tapes, direct access tapes, tapes being scanned, tapes being replicated, or any tapes if the tape database is not online.
  - Move a single tape to a library slot from another library slot or drive. You cannot move replica tapes or any tapes if the tape database is not online.

- Move a single tape to a library drive from a library slot. You cannot move replica tapes, tapes being scanned, tapes being replicated, or any tapes if the tape database is not online.
- Move a single tape to a library from the virtual vault. You cannot move replica tapes, direct access tapes, stub tapes, or any tapes if the tape database is not online.
- Move a single tape to a standalone drive from the virtual vault. You cannot move replica tapes, direct access tapes, stub tapes, tapes being scanned, tapes being replicated, tapes being shredded, or any tapes if the tape database is not online.

Delete virtual tapes

You can delete one or more virtual tapes. All data on the tapes will be lost. If you delete a VIT (stub tape) that has been migrated to object storage, data in object storage will also be deleted.

If replication is configured for a virtual tape, deleting the virtual tape removes the replication configuration on the source and replica. You can choose to promote the replica to become a usable virtual tape before deletion.

To delete a tape, select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a Habanero server, click the *Tapes* tab, select a tape, and click the *Delete* icon. 

Select *Promote Replica*, if applicable.

Select *Force Deletion* to force deletion of tapes without deleting the replication configuration on a replica tape when replica server or replica tape is not available or when the replication target server cannot be connected to.

Suspend/resume replication for a virtual tape

Suspend the schedule of a replication policy or resume a suspended policy.

Suspending replication prevents future replications from automatically being triggered by your replication policies (watermark, interval, time). This will not stop replication that is currently in progress. You can still manually start the replication process while the schedule is suspended.

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select your server, click the *Tapes* tab, select a tape with replication enabled, click the *Replication* icon,  and select the appropriate action.

Display virtual tape information

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select your server, click the *Tapes* tab, select a tape, and click the *Information* icon. 

General displays tape summary information, including sizing, tape capacity on demand information, media type, and date of creation and last modification.

Physical Layout displays ACSL, sector, and size information for the virtual header and tape segments.

If the tape is configured for replication, the *Outgoing Replication* tab displays the replica server, replica tape ID, and replication status.

Vault

Click the *Virtual Libraries* tab, and click the *Vault* tab. This page lists virtual tapes in the vault on the selected Habanero server.

The vault is comparable to the I/E slots in a physical tape library and is a storage area for tapes that are not inside a virtual tape library. Virtual tapes appear in the vault after they have been moved from a virtual tape library. Local virtual index tapes (LVITs) of tapes can also be in the vault on the target replication server after replication is complete.

There is no limit to the number of tapes that can be in the vault. The total number of tapes in the vault is displayed at the bottom of the page. You can also select how many tapes to display per page.

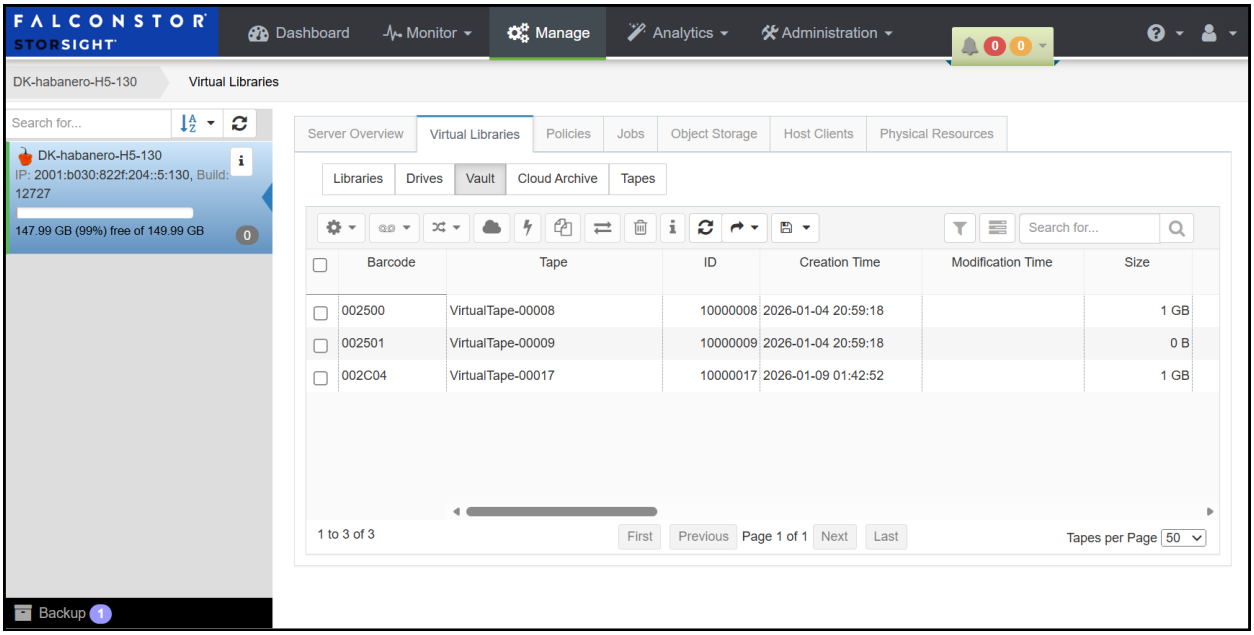

Manage tapes in the vault

Select a tape, click the *Manage* icon  , and select the action you want to perform. You can perform the following functions for tapes in the vault:

- Rename - The new name can be up to 32 characters excluding spaces or special characters < > & $ " / \' * ? | %
- Write protect
- Move multiple tapes to a virtual tape library or move a single tape to a standalone virtual tape drive or a virtual tape library.
- Shred/stop shredding

Shred a tape

Select one or more tapes, click the *Manage* icon  , and select *Shred*. If desired, select the option to delete the tape after shredding it.

Note that when a WORM tape is shredded, it will automatically be deleted.

Just as deleting a file from your hard drive does not completely destroy the file, deleting a virtual tape does not completely destroy the data on the tape. If you want to ensure that the data is unrecoverable, you must shred the tape.

Shredding a virtual tape destroys all data on the tape, making it impossible to recover the data. Tape shredding uses a military standard to destroy data on virtual tapes by overwriting it with a random pattern of bits, rendering the data unreadable.

Migrate virtual tape data to object storage

Migration is automatically triggered when a virtual tape is ejected to the vault. If you enabled Cloud Archive with existing tapes that can be migrated, you can manually trigger migration to object storage.

Select a tape and click the Cloud Archive icon  .

You can watch the status of the migration job on the *Import/Export Jobs* page.

Recover data from object storage

When migration occurs in *Copy* mode, the source tape remains in the vault after migration and is not automatically converted to a stub tape. If the virtual tape was not manually converted to a stub tape, you do not need to do anything to recover from object storage because you still have the original virtual tape.

When migration occurs in *Move* mode, the source virtual tape is converted to a stub tape after migration completes. In order to recover data, the stub tape must be recovered back to a virtual tape so that it can be accessed by backup software.

If stub tapes are lost due to a disaster, you must recover them from backup before any data restoration from object storage can be done.

You must make sure there is enough backup cache capacity available for the recovered data before you begin the recovery process. The recovered tape will be write-protected after recovery. You need to remove the write-protected flag from the tape.

To recover data from object storage:

1. Select a stub tape and click the *Recover Cloud Archive* icon. 
2. Select the virtual library to which you want the recovered tape to reside.
3. Select the virtual tape library slot and recovery mode.

   Copy mode leaves the tape contents in object storage while *Move* mode deletes the tape contents from object storage.

   If you configure tape recovery with *Move* mode, the migrated objects that are associated with the tape will be purged from object storage once the tape recovery job is complete. Further, the barcode of a virtual tape that has been recovered in *Move* mode can be changed after recovery; the barcode cannot be changed if recovery is done in *Copy* mode.

4. Confirm the information and click *OK*.

   You can watch the status of the recovery job on the *Import/Export Jobs* page. If the recovery job fails, it will be retried based on the import/export jobs retry settings. It can also be manually restarted from the *Import/Export Jobs* page.

Convert a virtual tape to a stub tape

If object migration is configured in *Copy* mode, the source tape is not automatically converted to a stub tape after migration. To manually convert a virtual tape to a stub tape, select a tape that has completed migration, and click the *Convert to Stub* icon  .

Export data to physical tape

Select one or more tapes and click the *Export* icon  to manually export data from a virtual tape to a physical tape in a physical tape library or standalone drive.

Display virtual tape information

Select a tape and click the *Information* icon. 

General displays tape summary information, including sizing, tape capacity on demand information, media type, and date of creation and last modification.

Physical Layout displays ACSL, sector, and size information for the virtual header and tape segments.

Cloud Archive

Click the *Virtual Libraries* tab, and click the *Cloud Archive* tab. This page lists virtual tapes exported to the cloud on the selected Habanero server.

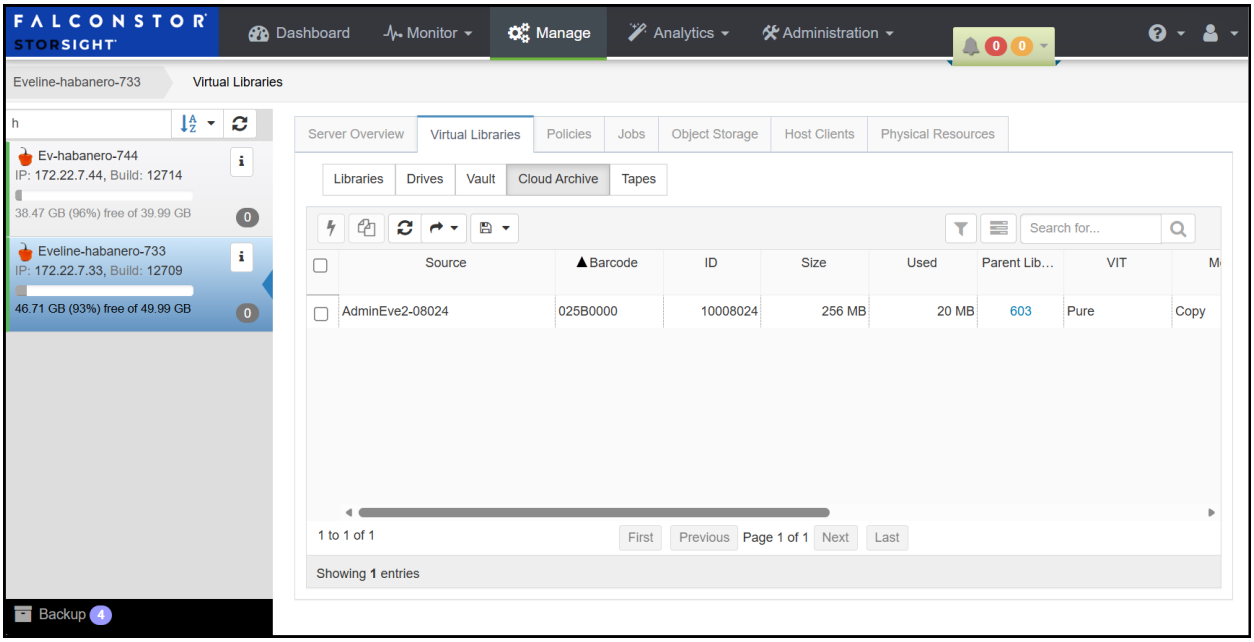

Recover data from Cloud Archive

When migration occurs in *Copy* mode, the source tape remains in the vault after migration and is not automatically converted to a stub tape. If the virtual tape was not manually converted to a stub tape, you do not need to do anything to recover from object storage because you still have the original virtual tape.

When migration occurs in *Move* mode, the source virtual tape is converted to a stub tape after migration completes. In order to recover data, the stub tape must be recovered back to a virtual tape so that it can be accessed by backup software.

If stub tapes are lost due to a disaster, you must recover them from backup before any data restoration from object storage can be done.

You must make sure there is enough backup cache capacity available for the recovered data before you begin the recovery process. The recovered tape will be write-protected after recovery.

To recover data from object storage:

1. Select a stub tape and click the *Recover from Object Storage* icon. 
2. Select the virtual library to which you want the recovered tape to reside.
3. Select the virtual tape library slot and recovery mode.

   Copy mode leaves the tape contents in object storage while *Move* mode deletes the tape contents from object storage.

   If you configure tape recovery with *Move* mode, the migrated objects that are associated with the tape will be purged from object storage once the tape recovery job is complete.

Further, the barcode of a virtual tape that has been recovered in *Move* mode can be changed after recovery; the barcode cannot be changed if recovery is done in *Copy* mode.

4. Confirm the information and click *OK*.

   You can watch the status of the recovery job on the *Import/Export Jobs* page. If the recovery job fails, it will be retried based on the import/export jobs retry settings. It can also be manually restarted from the *Import/Export Jobs* page.

Convert a virtual tape to a stub tape

If object migration is configured in *Copy* mode, the source tape is not automatically converted to a stub tape after migration. To manually convert a virtual tape to a stub tape, select *Manage* from the menu bar, select your server, click the *Virtual Libraries* tab, click the *Cloud Archive* tab, select a tape, and click the *Convert to Stub* icon. 

Host Clients

Host clients are the backup application servers that are assigned to virtual tape libraries and drives for backup. iSCSI protocol is used to communicate with host clients. iSCSI clients use iSCSI Target Mode.

Before a backup application server can communicate with a Habanero server, the two entities need to mutually recognize each other. By default, when a client machine is added as an iSCSI client, it becomes an iSCSI initiator. One iSCSI target needs to be created in StorSight for each iSCSI client initiator.

Create an iSCSI target for each client

One iSCSI target should be created for each iSCSI client initiator:

1. Click Manage from the menu bar, select your server, and click the Host Clients tab.
2. Select the iSCSI client for which you want to add an iSCSI target.
3. Click the Manage icon and select *Create iSCSI Target*. Enter a new target name for the client (maximum of 223 characters) or accept the default.
   The Microsoft iSCSI initiator can only connect to an iSCSI target if the target name contains less than 221 characters.

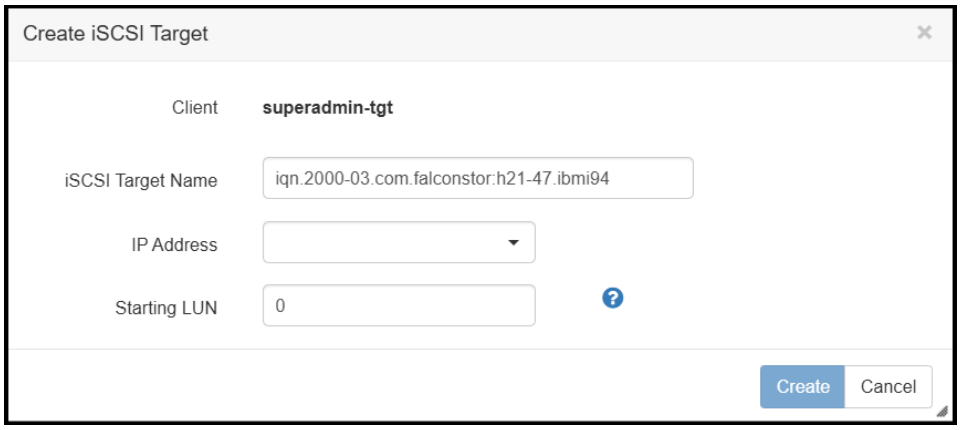

4. Select the IP address of the Habanero server that you want to use for communication with the iSCSI clients.
5. Leave the starting LUN as 0.
6. Once the iSCSI target is created for a client, LUNs can be assigned under the target using available virtual iSCSI devices.
7. Confirm all information and click *Create*.

Configure iSCSI on host clients

iSCSI target mode is supported for the following iSCSI client platforms, IBM i, AIX. Linux, and Windows.

Configure iSCSI on your host client based on the operating system.

Note: For improved performance, if your network supports it, turn on jumbo frames and set the maximum transfer unit (MTU) of each IP packet to 9000.

IBM i configuration

1. Install IBM iSCSI packages and required PTF files on your IBM I PowerVS instance.

   One iSCSI target is created in StorSight for each iSCSI client initiator. The PowerVS instance is running Ethernet under an IBM Cloud private network that belongs to a VLAN subnet; no extra VLAN is set on the Habanero server.

2. Check the *Type-Model*.

   Confirm the IBM PTF is installed on the client by checking the *Type-Model,* which should display as `298A-001`. This indicates that the iSCSI bus IOP resource is operational on the client. If it does not exist, contact IBM to install the required PTF file.

3. Create an iSCSI target/initiator.

   For IBM i version 7.3 or higher, use the IBM Navigator for i GUI to create an iSCSI target. Select *System Services –> iSCSI* Tab *–> Action –> Create iSCSI Target –>* Enter target iSCSI Qualified Name (IQN), target IP address, or hostname.

   For older IBM i versions, use the SQL service commands as described below.

   SQL commands

   - Confirm the SQL service is operational by running the run `STRSQL` command. This service is used for configuring communication between the client iSCSI initiator and StorSight iSCSI target.

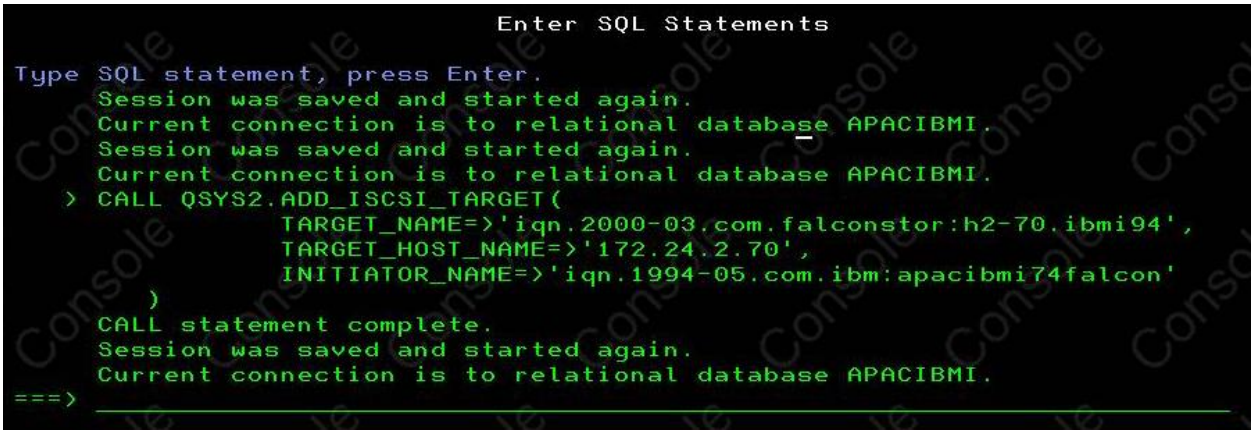

- Run the SQL command to add the iSCSI target and the initiator information on the client; you can set the target to any value matching IQN patterns, for example: `CALL QSYS2.ADD_ISCSI_TARGET( TARGET_NAME=>'iqn.2000-03.com.falconstor:h21-47.ibmi94', TARGET_HOST_NAME=>'172.22.21.47', INITIATOR_NAME=>'iqn.1994-05.com.ibm:apacibmi74falcon');`

- Run the SQL command with the IPL I/O processor option that will send a login request from the client iSCSI initiator to the Habanero server with a nonexistent target: `CALL QSYS2.CHANGE_IOP(IOP=>'ISCSI', OPTION=>'IPL');`

- Confirm the IPL I/O processor successfully completes by checking the Habanero system log, `/var/log/messages`. Make a note of the target name in the log; you will need to use it when configuring the iSCSI target in StorSight. You will see a message similar to the following:
  `May 30 14:31:36 h21-47 fsiscsid[9033]: IPSTOR||1653892296||E||0x0000c352||Login request to nonexistent target %1 from initiator %2||`**`iqn.2000-03.com.falconstor:h21-47.ibmi94`** `(ip 172.22.21.47)||`**`iqn.1994-05.com.ibm:apacibmi74falcon`**

4. Note the client iSCSI initiator.
5. From StorSight, add an iSCSI client. Refer to Add iSCSI host clients to Habanero for more information.
6. Assign a virtual tape library to your client. Refer to Assign virtual tape libraries to a client for more information.
7. Discover the assigned devices from the client.
   - Run the IPL I/O processor option to send a login request. You can run the IPL using a SQL command or the Start System Service Tools command (STRSST):

     SQL command

     Run the SQL command with the IPL I/O processor option:
     `CALL QSYS2.CHANGE_IOP(IOP=>'ISCSI', OPTION=>'IPL');`

     STRSST command

     Run `STRSST` on the client to check the system bus resource and execute a bus reset to the iSCSI bus resource:
     I/O debug to 298A-001 IOP resource (option 6):

Run the IPL I/O processor option that will send a login request from the client iSCSI initiator to the Habanero server with a nonexistent target:

If everything is successful, the Habanero server receives the iSCSI login request with the following sample messages in the system log:

```
May 30 14:34:23 h21-47 fsiscsid[9033]: IPSTOR||1653892463||I||0x0000c351||Login
to the target %1 from the initiator %2||iqn.2000-03.com.falconstor:h21-
47.ibmi94||iqn.1994-05.com.ibm:apacibmi74falcon
May 30 14:34:23 h21-47 kernel: FSISCSI client 14 initiator iqn.1994-
05.com.ibm:apacibmi74falcon type 2 login request to target iqn.2000-
03.com.falconstor:h21-47.ibmi94 from 172.24.2.94, conn 4020293, new tcp session
May 30 14:34:23 h21-47 kernel: FSISCSI conn 4020293 create new session 62603.
May 30 14:34:24 h21-47 kernel: svdp_get_cpu: vdev 536 cpu 31.
May 30 14:34:24 h21-47 kernel: svdp_get_cpu: vdev 537 cpu 32.
May 30 14:34:25 h21-47 kernel: svdp_get_cpu: vdev 538 cpu 33.
May 30 14:34:25 h21-47 kernel: [vtl_tde_537|4018142] TLE_INFO: VDrive 537
bPowerOnReset is set to 1, CDB[0]=0h vtape=-1 [n/a]
May 30 14:34:25 h21-47 kernel: IOCORE1 [kworker/31:1|3523459] release_vdev,
releasing vdev 536 without a reservation
May 30 14:34:26 h21-47 kernel: [vtl_tde_538|4018147] TLE_INFO: VDrive 538
bPowerOnReset is set to 1, CDB[0]=0h vtape=-1 [n/a]
```

- Confirm the iSCSI device is now available to the client. For example, the 3584-403 is displayed as a FalconStor vendor device Type-Model configured for the client:

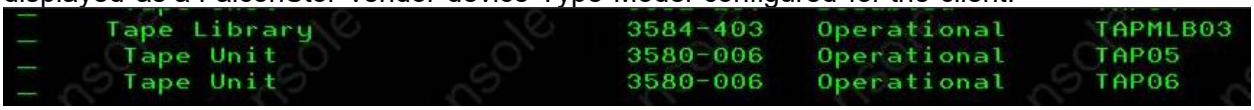

AIX configuration

AIX only supports tape drives via iSCSI, not tape libraries. You can use the `atape` driver on AIX to recognize a tape drive that is attached via iSCSI and use the standard operating system commands, such as `tar` or `dd`, to write to the `/dev/rmt0` device. You will have to manually add the tape to the drive in order to write data.

1. Install iSCSI packages on your AIX client that will access the Habanero server.
2. Collect iSCSI initiator information:

   On AIX, run the following command:

   ```
   lsattr -El iscsi0 -a initiator_name
   ```

   The resulting output will look like this:

   ```
   initiator_name iqn.localhost.hostid.7f000001 iSCSI Initiator Name
   True
   ```

   In this example, the iSCSI initiator name is `iqn.localhost.hostid.7f000001`.
3. Add an iSCSI client. Refer to Add iSCSI host clients for more information.
4. Assign a virtual tape library to your client. Refer to Assign virtual tape libraries to a client for more information.
5. On AIX, run the following command to set the correct discovery policy:

```
chdev -l iscsi0 -a disc_policy=odm
```

6. Connect the iSCSI initiator to the iSCSI target:

```
mkiscsi -g static -l iscsi0 -t 'target name from above step' -i IPaddress -n 3260
```

7. Rescan and connect to connect to the tape drive, ignoring the error about the changer device:

```
cfgmgr
```

8. Run the following command to display the tape drive to use:

```
ls -l /dev/rmt0
```

Linux configuration

1. Install iSCSI packages on your Linux client that will access the Habanero server.
2. For Red Hat Enterprise Linux on a client connected to the internet, you can install the iSCSI package by running the following as root:
   `yum install iscsi-initiator-utils`
   For Debian-based distribution on a client connected to the internet, you may be able to install the iSCSI package by running the following as root:
   `apt-get install open-iscsi`
   For other distribution of Linux, or your client is not connected to the Internet, contact your administrator for help in downloading and installing the iSCSI initiator package.
3. Edit the `/etc/iscsi.conf` file and add the following line to the end of the file:

   ```
   DiscoveryAddress=IP address of Habanero server
   ```

4. Start the initiator by typing:

   ```
   /etc/init.d/iscsi start
   ```

5. Add an iSCSI client. Refer to Add iSCSI host clients to Habanero more information.
6. Assign a virtual tape library to your client. Refer to Assign virtual tape libraries to a client for more information.
7. Log the client onto the target.

   On the client machine, type the following command to log the client onto the target:

   ```
   /etc/init.d/iscsi reload
   ```

Windows configuration

1. Install iSCSI packages on your client that will access the Habanero server.
2. Run Microsoft iSCSI Initiator on the client.
3. Click the *Discovery* tab, then click *Add* under the *Target Portals* group box.
4. Enter the Habanero server's IP address or name (if resolvable). Use the default port (3260) and then click OK to add the client.
5. Add an iSCSI client. Refer to Add iSCSI host clients to Habanero for more information.
6. Assign a virtual tape library to your client. Refer to Assign virtual tape libraries to a client for more information.
7. Log the client onto the target. To see the iSCSI targets from the client machine, run *Microsoft iSCSI Initiator* again.
8. Select the added target and click *Log On*. If it is desirable to have a persistent target, select *Automatically restore this connection* when the system boots.
9. Click *OK* to log on to the target. The status for the target will change from *Inactive* to *Connected*.

Add iSCSI host clients to Habanero

Create a host client for your iSCSI clients.

1. Click *Manage* from the menu bar, select your server, and click the *Host Clients* tab.
2. Click the "+" icon.
3. From the *General* tab, specify the client name (maximum of 64 characters) and select iSCSI as the client's protocol(s).
4. Click the *iSCSI* tab to select the initiators that this client will use.

   If the initiator does not appear, you may need to rescan by clicking the *Rescan iSCSI* Initiators tab. The client iSCSI initiator name is the same as the one configured on the client side, for example, *iqn.2000-03.com.falconstor:h21-47.ibmi94.* You can also manually add it by clicking the "+".
5. Do not enable CHAP, to allow unauthenticated access.
6. Confirm all information and click *Create* to add this client.

Assign virtual tape libraries to a client

You need to assign a virtual tape library or drive to the target of a backup application server so that the backup application server can then access the assigned virtual tape library/drive(s).

Note: To avoid disrupting backup operations, you should wait until backup application servers are finished with backup or other I/O activities before assigning them additional devices.

1. Click *Manage* from the menu bar, select your server, and click the *Host Clients* tab.
2. Select a client and click the *Assign* icon. 
3. Indicate if you are assigning virtual tape libraries or virtual tape drives.

   If you select a virtual tape library, all tape drives in the library will be assigned to the selected client.

   If you select an individual virtual tape drive, the server and backup application server will treat each individually assigned drive as if it were a standalone tape drive.

   You will need to select the iSCSI target and you must specify the starting LUN in the range of 0 (default) - 255.

   For IBM i clients, the virtual tape library type must be FalconStor FALCON TS3500L32 (03584L32) or FALCON TS3500L32 (03584L32) and the media type must be ULTRIUM3 (LT03) or newer.
4. Click *Assign*.
5. Use the backup application server's operating system to discover the Habanero server. The steps to do this vary according to the backup application server's operating system.
6. Use your backup software to discover the library. The steps to do this vary according to your backup software.

Manage Host clients

Select *Manage* from the menu bar and click the *Host Clients* tab. The *Host Clients* page lists clients on the selected server.

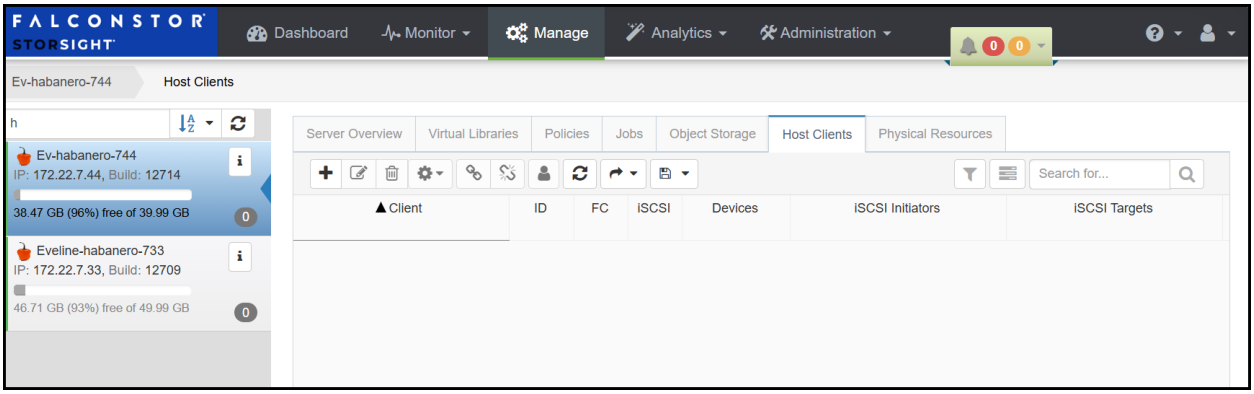

Update host client name

Select the client and click the *Update* icon , to update the host client name.

Delete a host client

Select the client and click the *Delete* icon , to delete the client.

Update host client iSCSI initiators

Select the client, click the *Manage* icon , and select *Update iSCSI Initiators*.

Assign virtual tape libraries and drives to a client

Select the client, click the *Assign* icon . Select one or more virtual tape libraries or drives to assign to the client and click *Assign*.

Unassign virtual tape libraries and drives from clients

Select a client and click the *Unassign* icon. . Select one or more virtual tape libraries or drives to assign to the client and click *Unssign*.

Assign user access to a host client

(Administrators only) The domain administrator can assign host clients to a specific domain user who will have access to a client.

1. Select a client and click the *Access Control* icon. .
2. Select a user who will manage this client and click *Assign*.

   To unassign access for a user, clear the user checkbox and click *Assign*.

Policies

Policies specify tape data replication policies. You can have a maximum of 64 policies.

Unique data replication from the source server to the target server occurs via TCP or Fibre Channel.

If cascaded or parallel (Concurrent Fan-Out or Serial Fan-Out) is configured, data can be replicated to an additional remote location.

Replicating a tape involves copying the virtual index tape and the missing data from the repository to the target server.

Note: Replication will occur only during the period of time you specified when you created the policy. Any replication jobs that have not completed by the end of this period will be stopped and will be run during the next replication period.

Replication phases

1. During the *Index* phase of replication, the Virtual Index Tape (VIT) from the source server is copied to the target server and becomes a Foreign Virtual Index Tape (FVIT).
2. During the *unique* phase of replication, the FVIT is scanned to determine whether or not the data blocks it uses exist locally. Missing data blocks are replicated from the source server to the target server. After all missing data blocks are replicated, the target server has all the data blocks used by the FVIT.
3. During the *final* phase, the tape is "*resolved*", and the target server automatically creates a Local Virtual Index Tape (LVIT) and puts it in the target server's vault or in a virtual tape library, depending upon how the policy was configured. The LVIT is now a write-protected replica of the source VIT and contains pointers to the replicated blocks of data.
    - Replication is complete when you see the LVIT on the target server in the vault or in the virtual tape library. The name of the LVIT corresponds to the name of the FVIT.
    - The FVITs are listed when you select *Replica* tapes on the *Tapes* page.
    - Note that this final step may not occur immediately after the initial replication of data and can take some time to complete, depending on the availability of tape drives on the target server and the amount of data on the FVIT.

Replication requirements

The following are the requirements for setting up a replication configuration:

1. You must have at least two Habanero servers.
2. You must finalize all IP addresses before replication is configured.
3. You must have enough space on the target server for the replica data.
4. Each virtual tape you want to replicate must be included in a policy.
5. *At the time of configuration*, each virtual tape that will be configured for replication must be in a slot, not a virtual library tape drive.

Manage policies

Select *Manage* from the menu bar and click the *Policies* tab. The *Policies* page lists policies on the selected server.

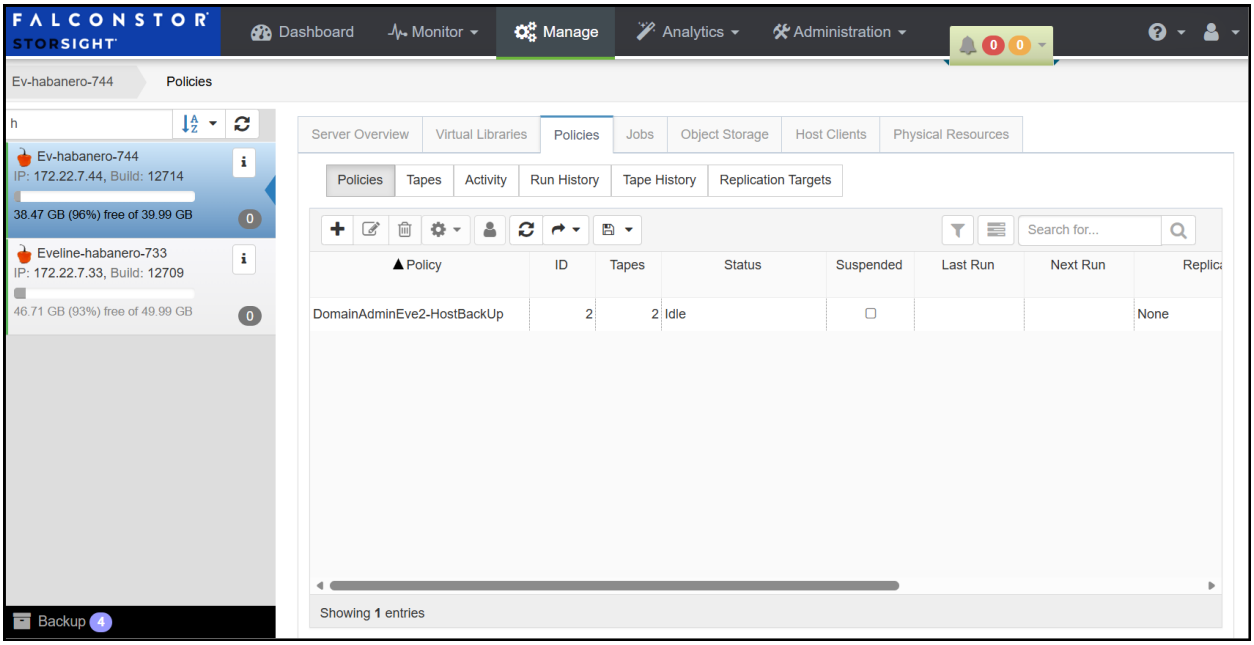

Create a policy

Policies specify tape data replication policies. You can have a maximum of 64 policies.

1. Click the “+” (*Create*) icon.

2. On the *General* tab, specify the policy name, and leave the default values for the priority level and retry parameters.

   The maximum length of the policy name is 128. The following characters are not valid: <>"&$/\'.

3. On the *Replication* tab, define the replication mode for the VITs in this policy, source IP address and target server(s).

   Single mode replicates tape data to a single target server.

   Cascaded mode replicates data to one target server, which in turn replicates to a second target server. Configuring a policy with replication in *Cascaded* mode automatically creates the same policy on the second target server, where the name of the policy is the server name followed by the policy name.

   Concurrent Fan-out - Replicates to two target servers at the same time.

   Serial Fan-out - Replicates to one target server and then to a second target server.

   Confirm/specify the IP address of the source server and select a target server, which must exist in the *Replication Targets* page. For the target server, specify the IP address used to transmit data, the storage pool from which to create the replica resource, and a remote library ID to which the LVIT will be moved after the tape is resolved. The LVIT can also remain in the vault.

For *Cascaded, Concurrent Fan-out*, and *Serial Fan-out* mode, you must also provide information for the second target server.

If you are using network address translation (NAT), set the NAT port forwarding address in the 1000-65000 network subnet range.

4. For a replication configuration, specify replication options on the *Replication Options* tab.

   Compress Data - Provides enhanced throughput during replication by compressing the data stream.

   Encrypt Data - Secures data transmission over the network during replication. Initial key distribution is accomplished using the authenticated Diffie-Hellman exchange protocol. Subsequent session keys are derived from the master shared secret, making it very secure. Compression/encryption for transmission over a network should not be set if the source tapes are already encrypted.

   Window Start Time - Replication window start time. The format is hh:mm and the default is 00:00.

   Window End Time - Replication window end time. The format is hh:mm and tThe default is 23:59.

   Remove replica when it is full - Automatically delete a tape's virtual index tape from the target server when the tape is full, thereby reducing the total tape count.

   Retry Count - The number of times to retry a failed job. The range is 0 to 12 and the default is 1. This is specific to index replication failures and will only retry index replication-specific functions.

   Retry Interval - The amount of time between retries. The range is 0-3600 minutes and the default is one second.

5. Click *Create* when done.

   Tapes can be added to the policy from the *Tapes* tab under *Policies*.

Update a policy

Select a policy and click the *Update* icon.  . Update a policy according to the following rules:

1. You can enable or disable replication but you cannot change the replication mode.
2. You can modify a policy with cascaded replication on the source server as well as on Target 1. However, if you modify the policy on the source server, updates are made in the policy on Target 1 automatically. But, if you modify the policy on Target 1, updates are made only in the policy on Target 1.
3. If you disable cascaded replication in a policy on the source server, replication is automatically disabled in the policy on Target 1.
4. If you disable cascaded replication, you have the option to skip cleanup on the replica server, in case the replica server is not available.

If you are re-configuring replication, an LVIT will be reused if all of the following criteria are met:

- Has the same name and barcode as the replica resource
- Is not an LVIT to an existing FVIT

- Is not part of another policy
- Is a pure VIT
- Is in the vault

If any of the above criteria is not met, a new tape will be created.

Manage a policy

Select a policy, click the *Manage* icon  , and select the action you want to perform. You can start, stop, suspend, and resume policies.

Starting or stopping a policy starts/stops replication for all tapes in the policy.

When you suspend a policy, future jobs are suspended; the currently running job is not affected.

When you suspend or resume Cascaded, Concurrent Fan-out, or Serial Fan-out replication, you must select a replication target.

Delete a policy

Select a policy and click the *Delete* icon. 

If the policy includes cascaded replication, you must delete it on the source server; the policy on Replication Target 1 will be deleted automatically.

Manage tapes in a policy

Select *Manage* from the menu bar, click the *Policies* tab, and click the *Tapes* tab. Select a policy from the *Policy* drop-down list to display virtual tapes in the policy on the selected Habanero server.

The total number of tapes is displayed at the bottom of the page. You can also select how many tapes to display per page.

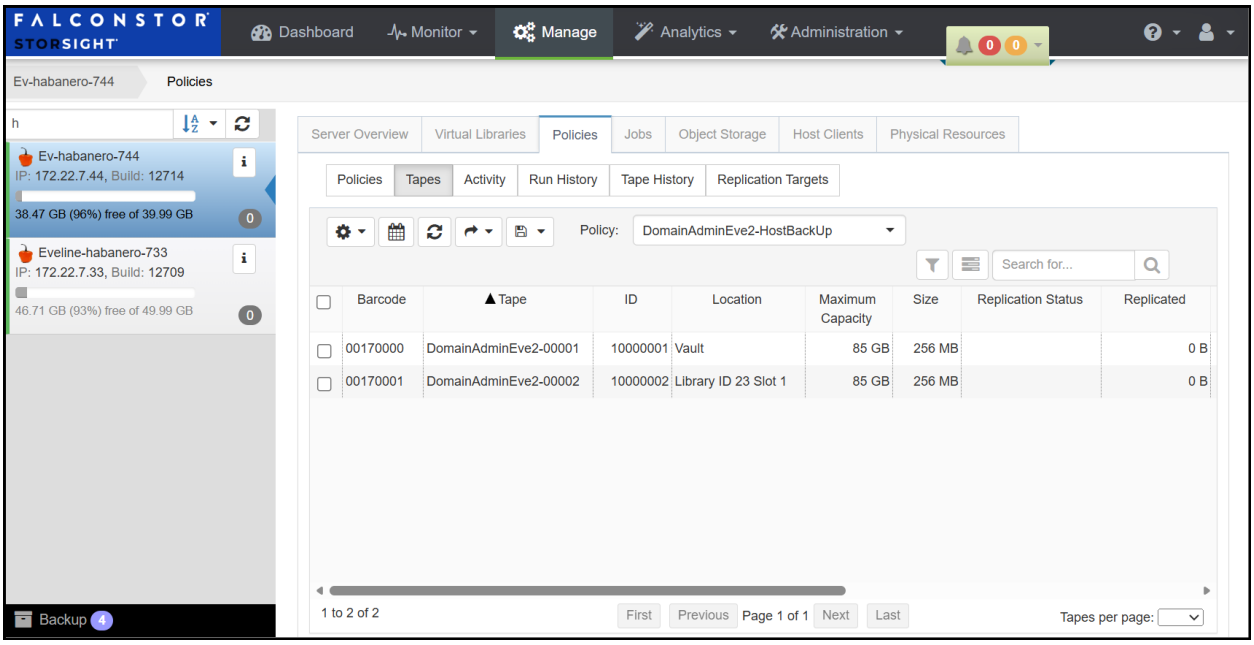

Add tapes to a policy

1. Select a policy.
2. Click the Manage icon and select *Add Tapes*.
3. Select a virtual tape library.
4. Select tapes and click *Add*.

Remove tapes from a policy

To remove tapes from a policy:

1. Select a policy.
2. Select one or more tapes.
3. Click the *Manage* icon and select *Remove Tapes*.
4. Confirm and click *Remove*.

Run replication for tapes in a policy

1. Select a policy.
2. Select one or more tapes for which you want to run replication.
3. Click the *Manage* icon  and select *Run Tapes*.
4. Confirm and click *Run*.

Schedule report of a tapes in policy

A report can be scheduled for tapes in a policy to be emailed to you on a regular basis. Refer to the Analytics for more information.

Policy Activity

Select *Manage* from the menu bar, click the *Policies* tab, and click the *Activity* tab. Select a policy from the *Policy* drop-down list to display staus of active replication jobs for tapes in the policy.

A report can be scheduled for replication activity to be emailed to you on a regular basis. Refer to Analytics for more information.

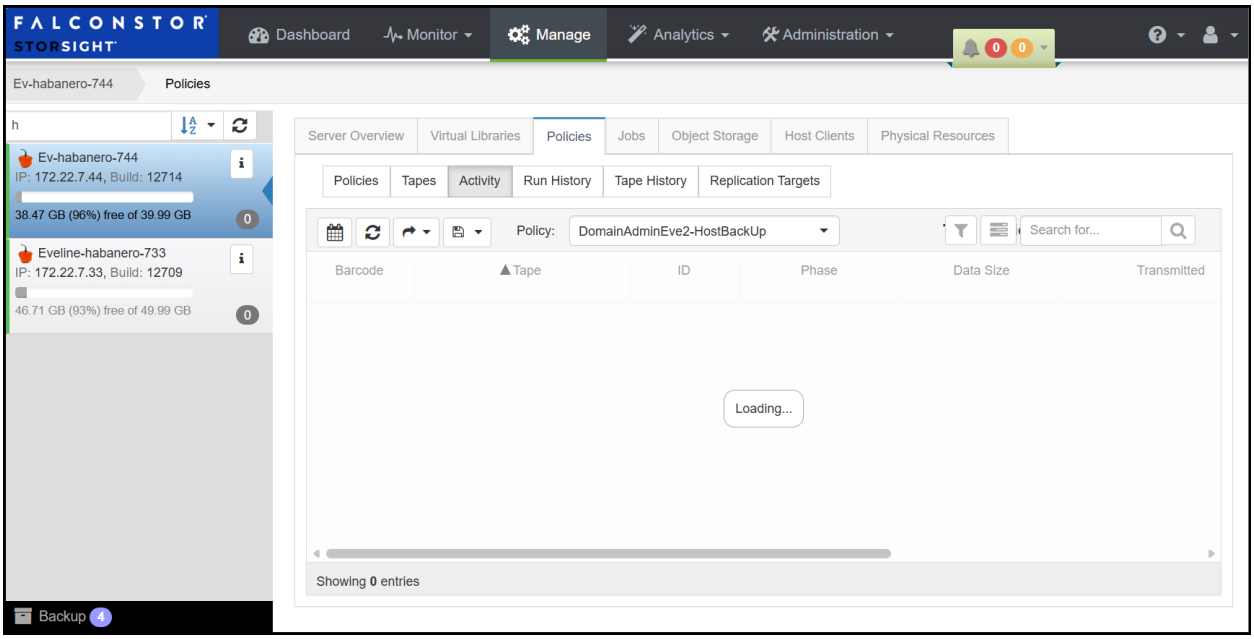

Policy Run History

Select *Manage* from the menu bar, click the *Policies* tab, and click the *Run History* tab. Select a policy from the *Policy* drop-down list. This page lists the replication history of the selected policy on the selected server.

A report can be scheduled for replication history to be emailed to you on a regular basis. Refer to Analytics for more information.

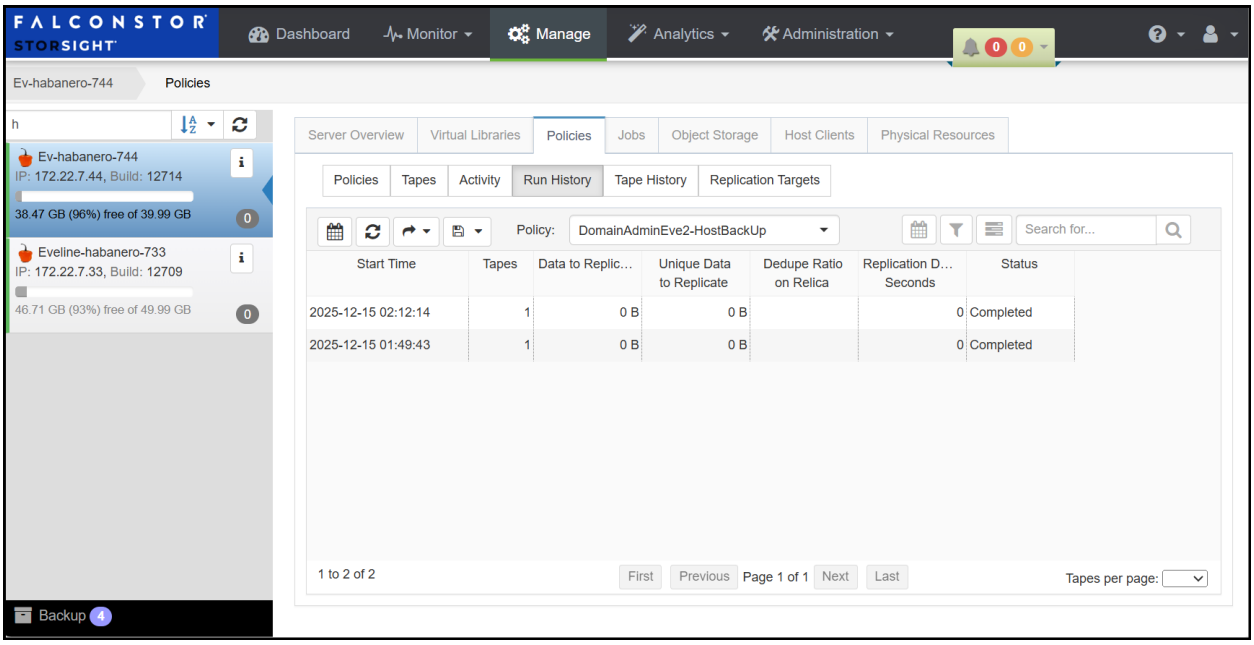

Policy Tape History

Select *Manage* from the menu bar, click the *Policies* tab, and click the *Tape History* tab. Select a policy from the *Policy* drop-down list. This page lists the history of replication for tapes in the selected policy on the selected Habanero server.

A report can be scheduled for tape replication history to be emailed to you on a regular basis. Refer to Analytics for more information.

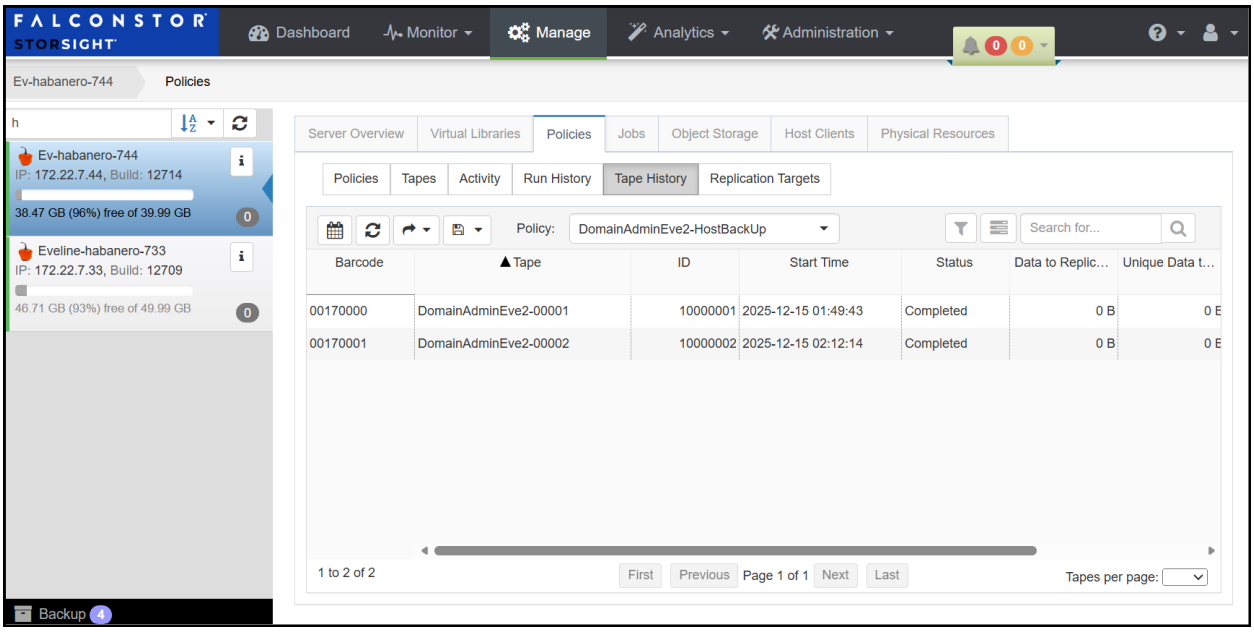

Replication Targets

Select *Manage* from the menu bar, click the *Policies* tab, and click the *Replication Targets* tab. This page lists the replication targets for the selected server.

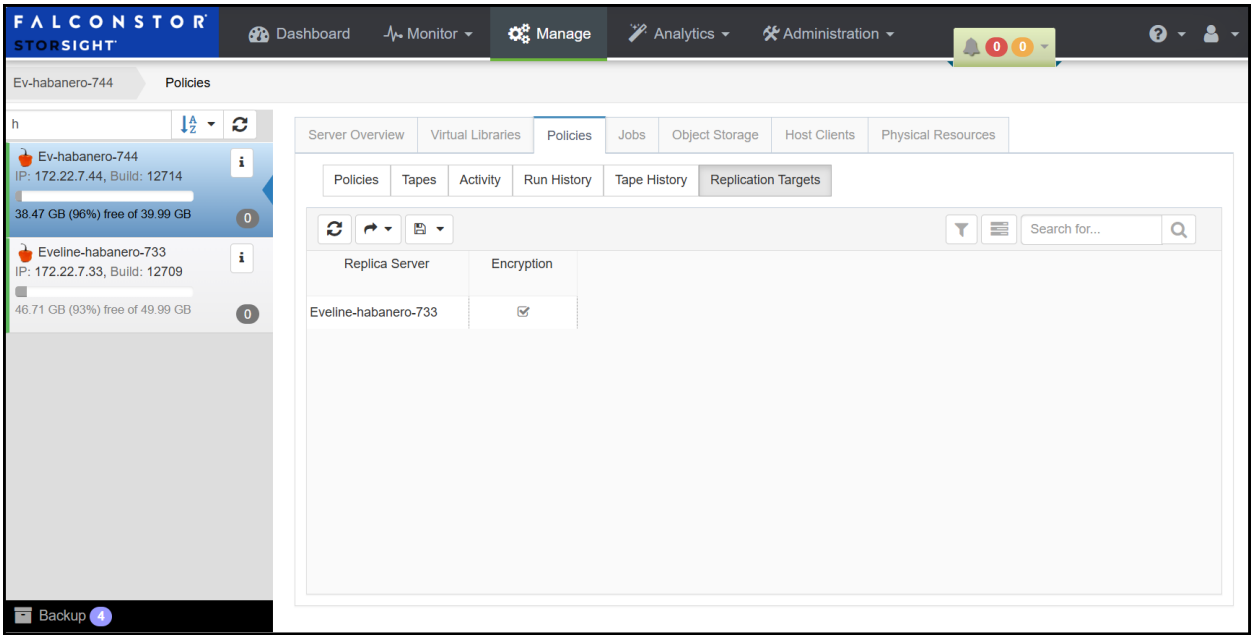

Object Storage

Organizations have been increasingly adopting object storage as a storage tier for archiving data. Because archived data is accessed infrequently and often has long retention periods, object storage is an ideal solution to extend storage capacity while minimizing costs.

Habanero's *Cloud Archive* can be used to archive virtual tape data to the cloud. At the end of each backup job, the tape is ejected to the vault, where the tape data gets migrated to the object storage of a cloud provider. You can export virtual WORM tapes and they will maintain the WORM property. If they are imported, they will become virtual WORM tapes. Refer to Configure Cloud Archive for a virtual tape library for more information.

Before you can enable Cloud Archive for a virtual tape library, you need an object storage with buckets to hold data. Inform FalconStor of the retention and immutability options that you need for your buckets. Then, FalconStor will provide you the object storage credentials and the bucket name.

Select *Manage* from the menu bar, select a Habanero server and click the *Object Storage* tab.

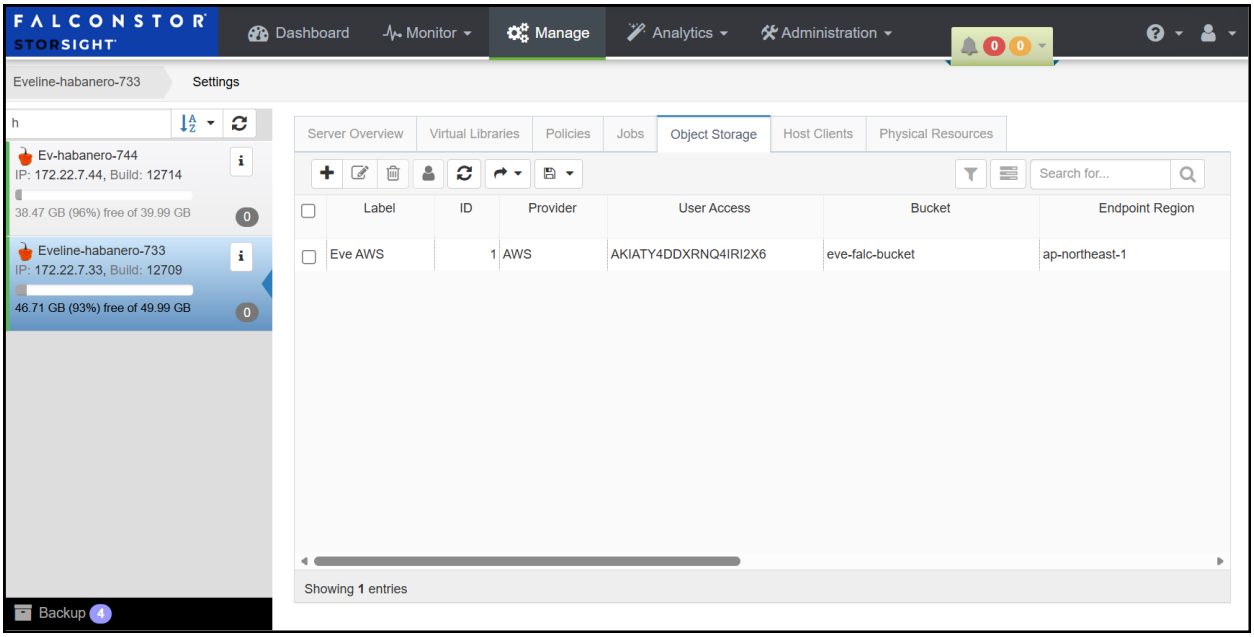

Add an object storage account

1. Click the “+” (*Add*) icon.

2. Select the IBMCOS provider and enter the configuration information provided to you by FalconStor.

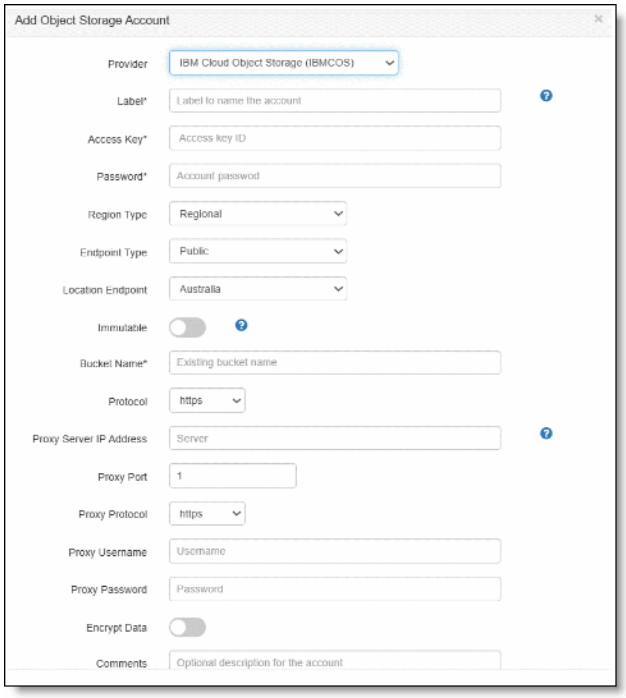

Label - Specify a name for this account.

Access Key - Specify the access key ID, up to 128 characters.

Password - Specify the password for authentication, up to 4,096 characters.

Region Type - The scope of the geographic area in which your data is distributed: regional, cross-region, single site.

Endpoint Type - Specify whether the endpoint is public, private, or direct.

Location Endpoint - Specify a location, based on *Region Type*.

Immutable - Specify if you want to make the bucket data non-rewriteable and non erasable. Tapes that get migrated to this bucket remain intact during the retention period, protecting data from malicious activity or accidental deletion. Refer to Immutable object storage for more information.

Bucket Name - Specify the name of an existing bucket. Objects created by tape migration will be kept in this bucket.

Protocol - Indicate whether object storage is accessed via *https* or *http*. The default is *https*.

Proxy Server IP Address - Optionally, specify a proxy server to access the object storage.

Proxy Port - Specify a proxy server port number, in the range of 1 to 65,535.

Proxy Protocol - Indicate whether the proxy server is accessed via *https* or *http*. The default is *https*.

Proxy Username/Password - Specify a username and password for proxy authentication.

Encrypt Data - Specify whether to use end-to-end encryption. When enabled, all data is encrypted before saving it to object storage; data is always encrypted in-flight and at-rest regardless of the protocol (http/https).

Comments - Optionally, specify a comment for this object storage account, up to 128 characters.

Proxy Username/Password - Specify a username and password for proxy authentication.

Encrypt Data - Enable or disable end-to-end encryption. When enabled, all data is encrypted before saving it to object storage; data is always encrypted in-flight and at-rest regardless of the protocol (http/https).

Comments - Optionally, specify a comment for this object storage account. The maximum length is 128 characters.

Manage object storage accounts

You can update a single object storage account and you can delete one or more object storage accounts.

Select *Manage* from the menu bar, select a server and click the *Object Storage* tab.

To update a single object storage account, select the account and click the *Update* icon. 

To delete one or more object storage accounts, select the accounts and click the *Delete* icon. 

Migration and recovery jobs

When tape data is to be migrated to object storage, Habanero automatically starts a migration job. Similarly, when tape data is to be recovered from object storage, a recovery job is started.

All object migration and recovery jobs are listed on the *Import/Export Jobs* page. While a job is running, you can see its status there.

As migration occurs, the allocated space in the cloud object storage bucket will grow based on the required space for the migrated tape.

Data upload to S3 is via multiple parts of 16 MB. Incomplete multi-part uploads will be cleaned up after seven days if there are no new request to those parts. The URI is used as prefix for objects.

A migration/recovery job will fail if the object storage runs out of its per-account space limit. It can also fail if the network connection to the object storage is down.

Should a job fail, it will be retried based on the retry settings for import/export jobs, as long as the tape remains in the vault. When a migration job is retried, it is restarted from the beginning. Migration and recovery jobs share the same job properties with import/export jobs (number of retries and retry interval for failed jobs).

All completed migration and recovery jobs are purged after 30 days. Failed and canceled migration jobs are also purged after 30 days if all retries failed. Job purging occurs only when a new export job is launched. Incomplete data objects associated with the purged jobs are deleted from object storage at that time.

When end-to-end encryption is configured for the object storage account, all jobs for that account use end-to-end encryption. This is configured when you add an object storage account. When end-to-end encryption is enabled, all data is encrypted before being saved to object storage; data is always encrypted in-flight and at-rest.

Similarly, the migration process will keep compressed data compressed before unloading it to object storage. When a tape is recovered from the object storage, it will be compressed like it was before migration.

Immutable object storage

Immutable object storage locks data to provide a safe backup and to maintain data integrity. Using Habanero with IBM Cloud Object Storage (COS), you can enable an object lock at the bucket level to get secure backups and long-term data retention. Retention policies ensure that data is stored in a non-erasable and non-rewritable manner for a specified time frame. Data cannot be changed until the retention period has expired. Once the retention period is over, data can be unlocked for further actions, according to your company policies.

Physical Resources

Select *Manage* from the menu bar and then click the *Physical Resources* tab.

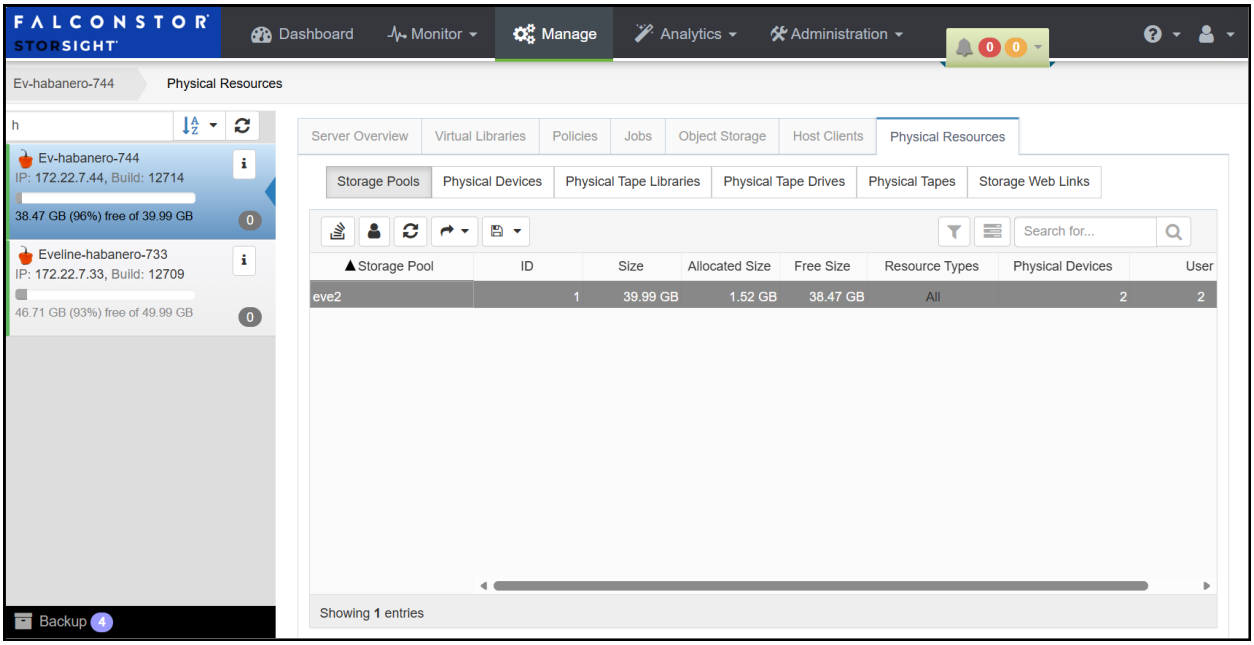

Storage pools

A storage pool is a group of one or more virtualized physical devices. The pool capacity-on-demand functionality automatically allocates storage space from a specific pool when storage is needed for a specific use.

Assign user access to a storage pool

(Administrators only) The domain administrator can assign access to a specific domain user who will have use this storage pool.

1. Select a pool and click the *Access Control* icon. 

2. Select a user that can access this storage pool and click *Assign*.

To unassign access for a user, clear the user checkbox and click *Assign*.

Physical devices

Select *Manage* from the menu bar, click the *Physical Resources* tab, and then the *Physical Devices* tab.

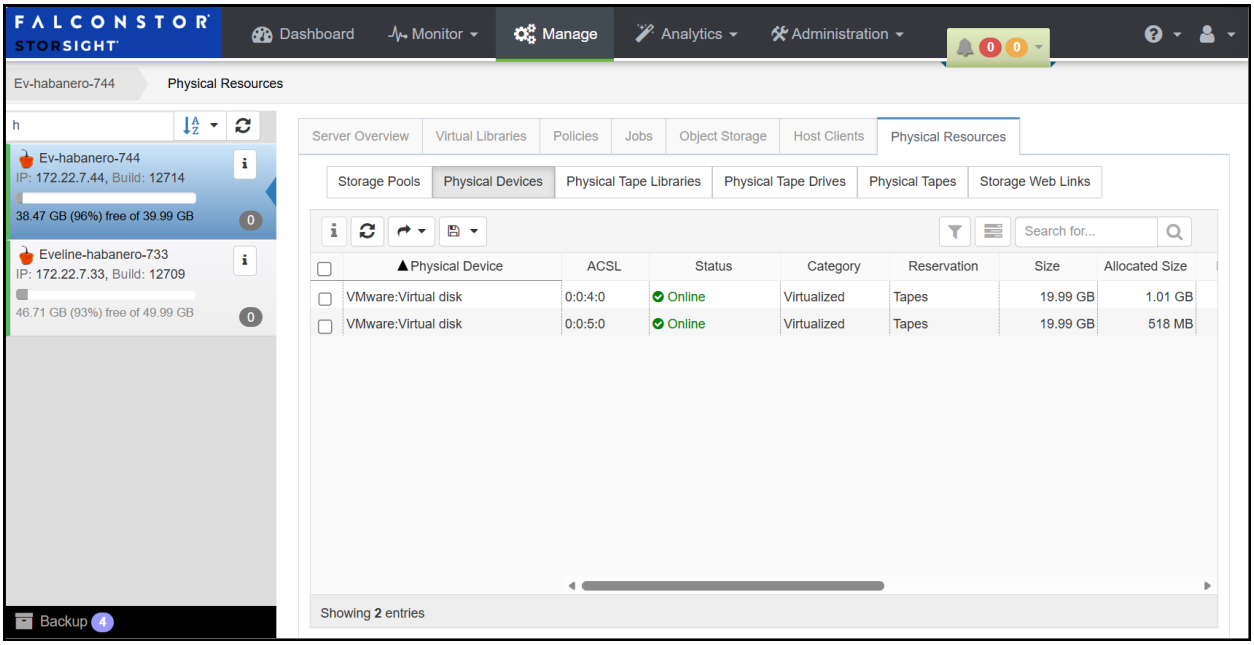

View information about a device

Highlight a device and click the Information icon to display physical device name, status, size, owner, queue depth, ACSL (adapter number, channel number, SCSI ID, LUN), category, firmware, total number of segments on the device, and physical segment/layout information.

Physical tape libraries

Select *Manage* from the menu bar, click the *Physical Resources* tab, and then the *Physical Tape Libraries* tab.

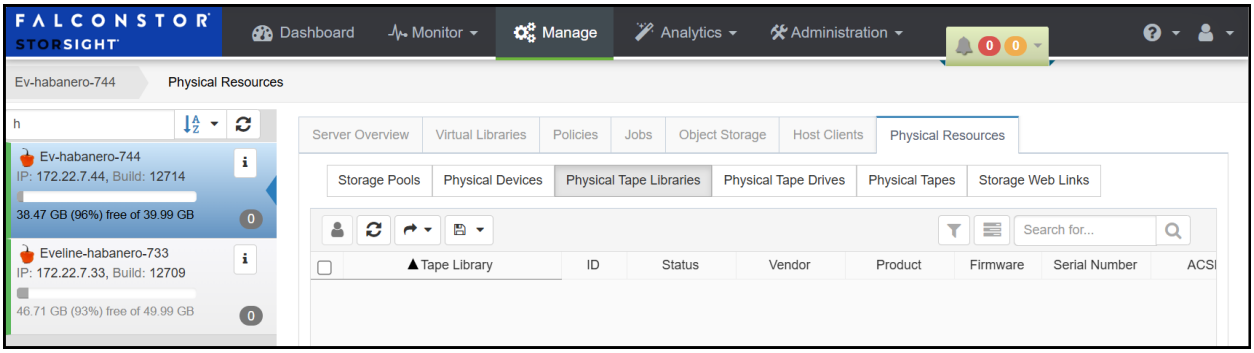

Assign user access to a physical tape library

(Administrators only) The domain administrator can assign access to a specific domain user who will have access to a physical tape library.

1. Select a physical tape library and click the *Access Control* icon. 

2. Select a user that can access this library and click *Assign*.

To unassign access for a user, clear the user checkbox and click *Assign*.

Physical tape drives

Select *Manage* from the menu bar, click the *Physical Resources* tab, and then the *Physical Tape Drives* tab. Select to view standalone tape drives or tape drives in a specific tape library from the *Tape Library* drop-down list.

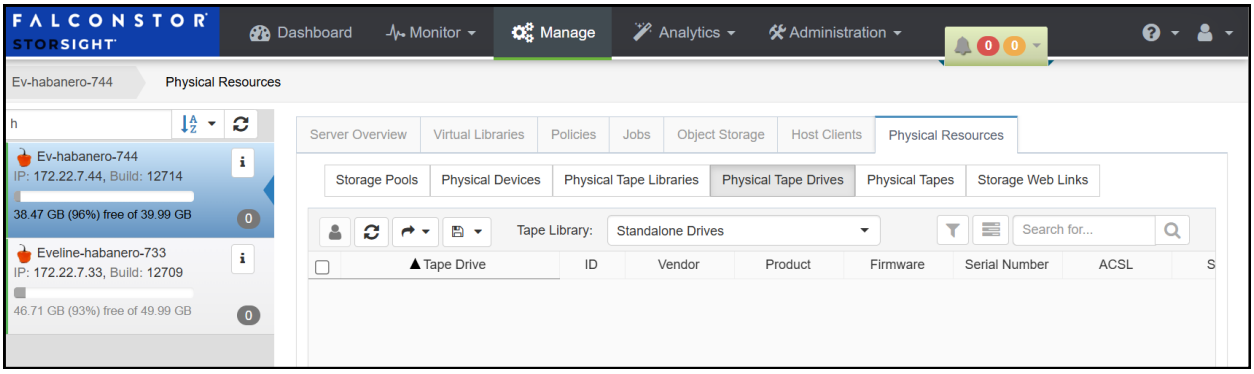

Assign user access to a physical tape drive

(Administrators only) The domain administrator can assign access to a specific domain user who will have access to a physical tape drive in a library or a standalone physical tape drive.

1. Select a physical tape drive and click the *Access Control* icon. 

2. Select a user that can access this drive and click *Assign*.

To unassign access for a user, clear the user checkbox and click *Assign*.

Physical tapes

Select *Manage* from the menu bar, click the *Physical Resources* tab, and then the *Physical Tapes* tab to see a list of physical tapes. Select the physical tape library from the *Tape Library* drop-down list.

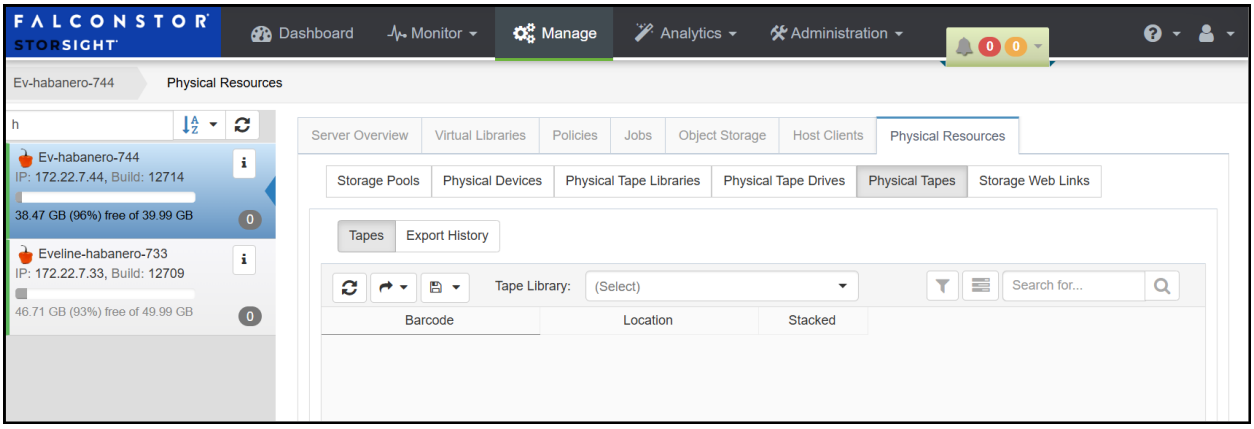

Physical Tape Export History page

Select *Export History* tab to see a list of physical tapes that were used for export jobs.

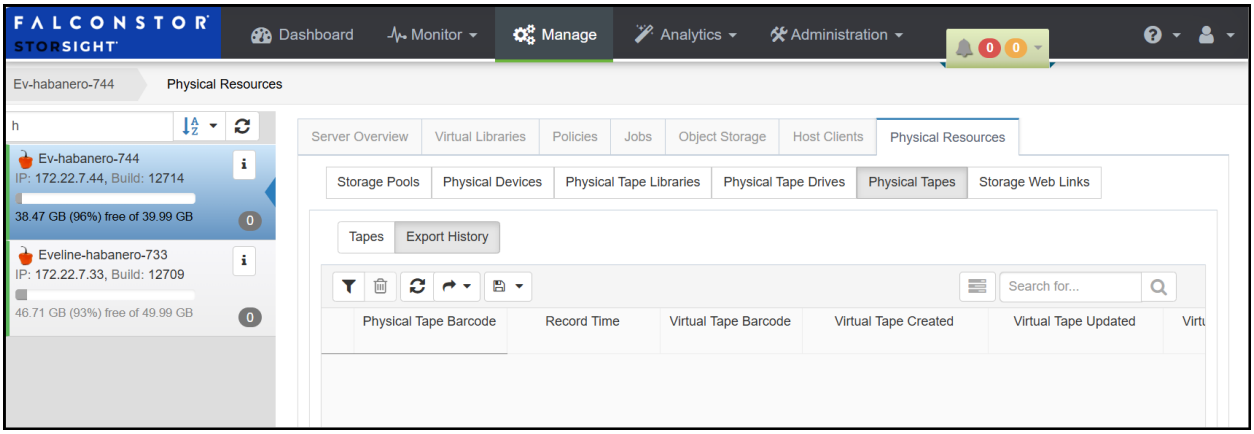

NAS

Network-attached storage provides file-level data storage to host clients using the Server Message Block (SMB) and Network File System (NFS) protocols.

NAS resources

NAS resources are virtualized disks that are used to create NAS shares and are used for standard file-based backup.

Create a NAS resource

1. Select *Manage* from the menu bar and click the *NAS* tab.
2. Click the "+" (*Create*) icon.
3. Enter information about the NAS resource.

   Name - Specify the NAS resource name up to 64 characters. excluding < > & $ " / '. The name cannot start with a dot and cannot contain consecutive dots.

   Size - Specify the NAS resource size in MB. The range is 1,000 MB to 16,776,191

   Storage Type - Select a single storage pool or physical device(s) from which to create the NAS resource.

   Format Options - These options are used when the NAS resource is formatted. The default settings include:

   - `-F` - Force the format regardless of what is on the drive.
   - `-I 512` - Increase the inode size from 128 bytes (default) to 512 bytes.
   - `-m 0` - Reserve 0% of the file system blocks for the super-user.
   - `-v` - Format in verbose.
   - `-j` - Creates the file system with an ext4 journal.
   - `-E resize=16383G` - Preserve the maximum file system metadata space (16 TB) for later file system resize.
   - `-J size=128` - Set the journal size to 128 MB for resize purposes.
   - `-b size=4096` - Specify the block size of 4 KB, which determines the minimum amount of space to use for each file.

   Mount Options - These options are used when the drive is mounted. The settings include:

   - `rw` - Allow read/write access.
   - `nosuid` - Disallow set-user-id execution.
   - `user_xattr` - Support "user" extended attributes.
   - `acl` - Enable Windows domain mode Access Control Lists (ACL), which are specific Windows permissions that can be set on files and folders when using Domain mode with Active Directory. Instead of assigning SMB clients permissions at the share level, ACLs allow the permissions to be applied to the files and directories beneath the share.

4. Click *Create* when done.

Delete a NAS resource

To delete a NAS resource, click the *Delete* icon. 

Expand a NAS resource

You cannot expand a NAS resource when the resource is being created, expanded, or formatted. We recommend that you detach any clients currently attached to the NAS resource before expansion.

To expand a NAS resource, click the *Manage* icon  and select *Expand*. You will need to specify the expansion size and the storage pool or physical device to use for expansion.

Unmount/mount a NAS resource

You can unmount the file system on a NAS resource. If unmounted, you can mount the file system.

To unmount/mount a NAS resource, click the *Manage* icon  and select the appropriate action.

Format a NAS resource

You can format a mounted file system on a NAS resource.

To format a NAS resource, click the *Manage* icon  and select *Format*. You can specify the format options and mount options to use.

Check the file system on a NAS resource

You can check the integrity of mounted file systems on a NAS resource. This operation will unmount the file system before running the integrity check. Make sure there are no jobs running on the NAS resource or they will fail.

To format a NAS resource, click the *Manage* icon  and select *Check File System*. You can specify to force the integrity check even if the file system seems clean.

Exclude/include a NAS resource from deduplication

By default, all NAS files are automatically included for deduplication except for files smaller than 8 KB, which are excluded because the file that replaces a deduplicated file and points to data in the repository (called the stub file) is at least 8 KB in size. If necessary, you can exclude specific NAS resources from deduplication.

To exclude a specific NAS resource from deduplication, click the *Exclusion* icon  , and select *Exclude Deduplication*.

To include a specific NAS resource that was previously excluded from deduplication, click the *NAS Resources* tab, select your server, select a NAS resource, click the *Exclusion* icon  , and select *Include Deduplication*.

Exclude/include a NAS resource from global replication

When you configure a single global policy for outgoing replication, all folders are included by default. If necessary, you can exclude specific NAS resources from replication.

To exclude a specific NAS resource from replication, click the *Exclusion* icon  , and select *Exclude Global Replication*.

To include a specific NAS resource that was previously excluded from replication, select *Manage* from the menu bar, click the *NAS Resources* tab, select your server, select a NAS resource, click the *Exclusion* icon  , and select *Include Global Replication*.

Copy NAS files and directories using Fast Copy

You can copy files and directories between NAS resources within the same server. If the source files have been deduplicated, the stub files will be copied so that the files do not have to be restored. File attributes, such as access permissions, user/group ownerships, and modification time, are carried over. Copied files will be automatically excluded from replication, if replication is configured. Only one fast copy job can run at a time.

Click the *Fast Copy* icon  , and select *Start*. Specify the source and target paths. You can click *Browse* and expand a folder to find a directory. Specify if you want to overwrite target data if it exists. The job will fail if the target directory exists and this option is not selected.

Display information about a NAS resource

You can display physical layout info for a NAS resource. Click the *Info* icon. 

Assign user access to a NAS resource

(Administrators only) The domain administrator can assign access to a specific domain user who will have access to a NAS resource.

1. Click the *Access Control* icon  .
2. Select a user that can access this resource and click *Assign*.

   To unassign access for a user, clear the user checkbox and click *Assign*.

Shares

NAS shares are directories on the server that can be used by NAS clients to store files.

Create a NAS share

1. Select *Manage* from the menu bar, click the *NAS* tab, and click *Shares*.
2. Click the “+” (*Create*) icon.
3. On the *General* tab, specify share properties.

   Parent Directory - Specify the parent directory path if the share directory being created is not under the NAS resource.

   Share/Folder Directory - Specify the directory to create on the NAS resource.It will become a NAS share when it gets assigned to an NFS or SMB client, otherwise it will remain a folder. The directory name cannot exceed 238 characters and cannot have any leading or trailing spaces, cannot start with a dot, cannot have blanks or < > / \ " % # : | * ? ~ special characters. It also cannot have any of the following reserved words: "com1"…"com9", "lpt1"…"lpt9", "com", "nul", "prn", "aux".
4. Select NFS clients who will access the share and specify permissions for each client.

   Access - Specify access rights to the share: read-write (default), read only

   Squash - Option to reduce the privileges of certain users by mapping user IDs to *nobody:*

   - root (default): UID=0 (*root* user) is remapped to *nfsnobody:nasgrp*
   - All: All UIDs are remapped to *nfsnobody:nasgrp*
   - none

   Insecure - Option to enable insecure mode for clients that do not use a reserved port for NFS (an Internet port less than IPPORT_RESERVED -- 1024), such as AIX clients:

   Kerberos - Option to enable Kerberos v5 protocol for the client.

   Kerberos with Integrity -Option to enable Kerberos v5 protocol with integrity checking using checksums.

   Kerberos with Privacy - Option to enable Kerberos v5 protocol with privacy service.

   System Protocol - Option to enable the system protocol for client security.
5. On the *SMB Clients* tab, select SMB clients who will access the share and specify share name and options.

   Type - Client type is set as *User* (default) or *Group*.

   Access - Specify access rights to the share: read-write (default), read only

   Share Name - A unique share name can be specified for SMB clients. The default is the same as the directory name. The name cannot exceed 64 characters and cannot have any leading or trailing spaces, cannot start with a dot, cannot contain 2 or more consecutive spaces or / \ " * : ? < > | # % , ; [ ] = + special characters. It also cannot have any of the following reserved words: "com1"…"com9", "lpt1"…"lpt9", "com", "nul", "prn", "aux", “homes”.

   Description - Option to enter a description for the SMB share.

Max Number of Connections - Option to set a limit on the number of SMB clients connecting to the share. The range is 1-1089, the default is 0, which means unlimited.

6. On the *SMB Properties* tab, set properties for the SMB share.

   By default, `nt acl support` is set to `no`. Click *Add* to specify other properties and values.

7. Click *Create* when done.

   If NFS or SMB clients are selected, a NAS share will be created; if no clients are selected, a folder without a share will be created.

Update a NAS share

Click the *Update* icon  to update client settings for a NAS share.

Delete a NAS share

Click the *Delete* icon  to delete a NAS share if it does not have any files or subfolders.

Exclude/include a NAS share from deduplication

By default, all NAS files are automatically included for deduplication except for files smaller than 8 KB, which are excluded because the file that replaces a deduplicated file and points to data in the repository (called the stub file) is at least 8 KB in size. If necessary, you can exclude specific NAS shares from deduplication.

To exclude a specific NAS share from deduplication, click the *Exclusion* icon  , and select *Exclude Deduplication.*

To include a specific NAS share that was previously excluded from deduplication, select a share, click the *Exclusion* icon  , and select *Include Deduplication*.

Exclude/include a NAS share from global replication

When you configure a single global policy for outgoing replication, all folders are included by default. If necessary, you can exclude specific NAS shares from replication.

To exclude a specific NAS share from replication, select a share, click the *Exclusion* icon  , and select *Exclude Global Replication*.

To include a specific NAS share that was previously excluded from replication, select a share, click the *Exclusion* icon  , and select *Include Global Replication*.

Set properties for a NAS share

You can set specific permissions for owners, groups, and other users at the directory level of a share or folder on a NAS resource.

1. Click the *Set Properties* icon. 

2. Set the appropriate properties.

Owner - Directory owner user account name.

Group - User group with permissions on the directory.

Sticky Bit - Setting a directory with a *sticky bit* gives it additional security by requiring that users own the file or directory, have write permissions, or be the *root* user if they want to remove or rename a file.

Apply changes to subdirectories - Applies the permissions to all subdirectories beneath the current one.

3. Click *Set* when done.

Show files for a NAS share

You can display the folders and files on a NAS share. You will see the size of each and deduplication and replication status, if enabled. Select a share and click the *Show Files* icon.

Display statistics for a NAS share

You can display deduplication and replication statistics, if enabled. This includes data size, number of files and folders, number of deduplicated files, space used by deduplicated files, number of files awaiting deduplication, space used by files awaiting deduplication, number of replicated files, number of files awaiting replication, space use by files awaiting replication, replicated data size, and size of unique replicated data. Select a share and click the *Show Files* icon. 

Map/mount shares

Windows clients

You should map a share for each Windows client to allow access to the share. Do the following on each Windows client:

1. Open *Windows Explorer* (or *My Computer*).
2. Select *Tools --> Map Network Drive*.
3. Set the path to the shared folder.

   The path is: *hostname**sharename*

   where *hostname* is the Habanero server's name or IP address and *sharename* is the name of the shared folder. For example: \\server1\Data1
4. Enter login information.

   For *User* mode, enter the user name and password that was set when the user was created.

   For *Domain* mode, enter the user's full account name (*Domain**username*) and the user's password that is defined at the Active Directory level.

NFS clients

You should mount a share for each NFS client to allow access to the share. Do the following on each NFS client:

1. Create a directory to mount the NFS share to.

   For example: `/mnt/share`

2. Locally, mount the share.

   `mount hostname:/nas/resource/fds/share/mount_point /mnt/share`

   where *hostname* is the Habanero server's name or IP address, *resource* is the name of the NAS resource, *share* is the share name, and *mount_point* is a directory the NFS share can be accessed from, in this case */mnt/share,* as set in step 1.

   Note: In the path above, /nas/ and /fds/ are not variables and must be included in the path.

   For example:

   `mount HABANERO35:/nas/NAS-00002/fds/Data1 /mnt/share`

3. For NFSv4 mounts, it is necessary to set the NFS domain in the `/etc/idmapd.conf` file. After adding the domain, run the following command: `nfsidmap -c .`

   Then, unmount and re-mount the filesystem.

   Refer to the following link for more information:
   https://access.redhat.com/solutions/2791811

NFS Client mount options

We recommend using the following mount options for NFS clients:

Note: UDP protocol is not supported as a mount option for NAS shares.

AIX 6.x/7.x

NFSv4 (with krb5i authentication):

```
mount -v nfs -o vers=4,sec=krb5i,combehind 172.22.21.69:/nas /mnt/nfs
```

Refer to the following link for additional information about mounting NFSv4 on an AIX client: https://www.ibm.com/support/pages/ibm-aix-how-setup-nfsv4-mount-clinet-and-server

NFSv3:

```
mount -v nfs -o proto=tcp,vers=3,intr,hard,llock, combehind NAS@IP:/NAS_path
/NAS_mount_point
```

Solaris 10 and 11

NFSv4:

```
mount -F nfs -o vers=4 172.22.21.69:/NAS_path / NAS_mount_point
```

NFSv3:

```
mount -F nfs -o hard,llock,intr,vers=3,proto=tcp NAS@IP:/NAS_path /NAS_mount_point
```

Note: We highly recommend using the "`llock`"" mount option with Solaris NFS clients.

Linux

NFSSv4.1:

```
# mount -o vers=4.1,sec=sys 172.22.21.69:/nas/NAS-00007/fds/nfs-1   /mnt/nfs
```

NFSv4.1 (with krb5 authentication):

```
# mount -t nfs4 -o minorversion=1,sec=krb5 172.22.21.69:/   /mnt/nfs
```

NFSSv4.1 when the NAS server is part of a failover configuration (client mount path needs to include `/nas` in the source path):

```
# mount -t nfs4 -o minorversion=1,sec=sys 172.22.21.69:/nas /mnt/nfs
```

NFSv3:

```
mount -t nfs -o hard,nolock,intr,nfsvers=3,tcp,bg NAS@IP:/NAS_path /NAS_mount_point
```

Note: We highly recommend using the "`nolock`" mount option with Linux NFS clients.

Oracle Recovery Manager (RMAN)

RMAN provides a foundation for backing up and recovering an Oracle database. Additional mount options are required by RMAN:

Operating System	Options
AIX	`cio,bg,rsize=65536,wsize=65536,noac,timeo=600`
Solaris	`bg,rsize=32768,wsize=32768,noac,forcedirectio,suid`
Linux	`rsize=524288,wsize=524288,actimeo=0,timeo=600`

Folders

Select *Manage* from the menu bar, click the *NAS* tab, and click the *Folders* tab. This page lists first-level folders and files on the selected NAS resource. Select which NAS resource to view from the *NAS Resource* drop-down list.

A TV icon indicates a shared folder. An X indicates a resource/share is excluded from deduplication. Click on a share or folder to expand and display files.

Manage folders and shares

You can perform the following functions for folders and shares:

- Create a NAS share or folder/Update client settings - Refer to Create a NAS share for more information.
- Delete a NAS share or folder if it does not have any files or subfolders.

Select a share, click the *Manage* icon  , and select the action you want to perform.

Exclude/include a NAS folder or share from deduplication

By default, all NAS files are automatically included for deduplication except for files smaller than 8 KB, which are excluded because the file that replaces a deduplicated file and points to data in the repository (called the stub file) is at least 8 KB in size. If necessary, you can exclude specific NAS shares from deduplication.

To exclude a specific NAS share from deduplication, select a share, click the *Exclusion* icon  , and select *Exclude Deduplication.*

To include a specific NAS share that was previously excluded from deduplication, select a share, click the *Exclusion* icon  , and select *Include Deduplication*.

Set properties for a NAS folder or share

You can set specific permissions for owners, groups, and other users at the directory level of a share or folder on a NAS resource.Select a share and click the *Set Properties* icon. 

Display statistics for a NAS folder or share

You can display deduplication and replication statistics, if enabled. This includes data size, number of files and folders, number of deduplicated files, space used by deduplicated files, number of files awaiting deduplication, space used by files awaiting deduplication, number of replicated files, number of files awaiting replication, space use by files awaiting replication, replicated data size, and size of unique replicated data. Select a share, and click the *Show Files* icon. 

SMB clients - Users

Select *Manage* from the menu bar, click the *NAS* tab, and click the *SMB Clients* tab. This page lists SMB clients on the selected server and the NAS shares assigned to a client. SMB clients are Windows users and groups that use the SMB protocol to access NAS shares. From this page, you can also access the SMB clients *Groups* page.

Depending upon the columns that have been selected for display, the properties for each SMB client and assigned NAS share are visible at a glance

Select a user from the *User* drop-down box to view properties for each NAS share assigned to the user.

Create an SMB user

You can only create an SMB user to access NAS shares if the authentication mode is *User*. Otherwise, client authentication is controlled by a Windows Active Directory domain controller.

1. Click the "+" (*Create*) icon.
2. Enter information about the SMB user.

   Name - Specify a Windows user name up to 32 ASCII characters starting with a letter or underscore and excluding blanks or ~`!@#$%^&*()+={}[]|\\\/?:;<>,\\'"

   Password/*Confirm Password* - Specify a Windows user password up to 256 characters.

   Description - Optionally, specify a description for the Windows user up to 256 characters without consecutive colons.
3. Click *Create* when done.

Update an SMB user

You can update the description and change the password for an SMB user if the authentication mode is *User*. Select a user and click the *Update* icon. 

Delete an SMB user

You can delete an SMB client if the authentication mode is *User*.

Select a user and click the *Delete* icon. 

Set the authentication mode

There are two security modes that you can use to authenticate users trying to access NAS shares:

- User mode - Authentication is controlled by passwords that are set for each Windows user.
- Domain mode - Authentication is controlled by a Windows Active Directory (AD) Domain Controller. The Habanero server and all clients must belong to the domain controlled by this Domain Controller.

Notes:

- If you change the NAS authentication method, you will lose all of your share assignments.
- If you change the organizational units (OU) list, you will lose all of your share assignments.
- The NAS authentication method cannot be changed if failover has been configured.
- The currently used UID/GID ranges cannot be deleted except when changing the authentication method.

1. Click the *Manage* icon  , and select *Set Authentication Mode.*

2. Enter information on the *General* tab.

   Authentication Mode - Specify the mode. *User* authentication is controlled by passwords that are set for each Windows user. *Domain* authentication is controlled by a Windows AD Domain Controller. The maximum length of the host name must be 15 characters for Domain mode due to the netbios name limitation.

   Work Group - (User mode) Specify a work group name.

   Server Description - Optionally, specify a server description to display in the comment field of Windows Explorer, up to 128 characters.

   Primary Domain Controller - (Domain mode) Specify a primary Domain Controller host name. The name must be a resolvable short name, not an IP address or fully qualified domain name. The Habanero server and Active Directory server must have synchronized clocks within five minutes of each other. Habanero will get user account information from this Domain Controller.

   Secondary Domain Controller - (Domain mode) Optionally, specify a secondary Domain Controller host name to use in case the primary controller is not available. The name must be a resolvable short name, not an IP address or fully qualified domain name.

   User Name/Password - (Domain mode) Specify the login name and a valid password of the authorized user account for the server to browse AD.

   Kerberos User Name/Password - (Domain mode) Optionally, specify a user name and password with AD administrative privileges when Kerberos authentication is required for NAS clients. This user name/password is used one time to create a computer account for the server in the domain so the server can become a trusted member in AD. By default, the

server tries to join the domain as an AD member server and if that fails, it will try to join the domain as a legacy Windows NT 4 server.

3. Reserve a range of User IDs (UIDs) and Group IDs (GIDs) on the *UID/GID range* tab.

   Click *Add* to select a list of ranges of UIDs and GIDs to reserve on the server for NAS users and groups. The ranges can only be changed when changing the authentication mode.

4. (Domain mode) Enter information on the *Organizational Units* tab.

   Bind Point - Optionally, specify a *Bind Point* to mark where the server will start reading OUs. This is useful if the Habanero domain login account does not have access to the entire OU tree. The default is "/", which means the root of the OU tree.

   Organizational Units - Click *Add* to select the OUs to which you will offer NAS shares. If you specified a bind point, the OU tree begins at that point; if you did not specify a bind point, the OU tree begins at the root of the OU tree.

5. Click *Set* when done.

Update the domain controller

(Domain mode) To update the primary or secondary domain controller and specify a new user/password, click the *Manage* icon  and select *Update Domain Controller*. Refer to Set the authentication mode for more information.

Sync SMB Clients

(Domain mode) To synchronize SMB clients with Active Directory, click the *Manage* icon  , and select *Sync SMB Clients*.

Set Windows ACLs

(Domain mode) Access Control Lists (ACLs) are specific Windows permissions that can be set on files and folders instead of assigning Windows users to NAS shares. After enabling ACL support, all file systems are remounted on the server, which may take a while.

All NAS resources with the mounted status will be unmounted and mounted again. For those that are not currently mounted, the mount option will be updated but the resource will not be mounted. You will then need to connect as a user with an administrative role to the Windows machine and configure ACL permissions of different users on different shares.

Click the *Manage* icon  , and select *Set Windows ACL*. Select users with administrative roles to support ACL permissions of users on NAS resources.

Set SMB options

1. Click the *Manage* icon  , and select *Set SMB Options.*
2. Enter information on the *General* tab.

   Server Description - Optionally, specify a server description to display in the comment field of Windows Explorer, up to 128 characters.

   Alias Names - Optionally, specify a list of alternative names for the server that SMB clients can use when connecting to the server.

   Domain Sync Interval - Frequency to sync users with the Domain Controller in minutes. The range is 1-1440 minutes and the default is 30 minutes.

   Samba Options - Double-click a field to change the default global Samba options. A few are described below. Refer to the Samba manual for information about the other settings.

   - *name resolve order* - The resolver order to look up hostnames. This can be useful if there is no DNS server configured and the server is on a different subnet than the SMB clients.
   - *root preexec* and *root postexec* - Changing the values will allow specified commands to be executed when clients connect to or disconnect from a service, such as common maintenance procedures.
   - *strict sync and sync always* – Setting the values to *yes* will enable the SMB write-through mode. By default, the write-through mode is disabled.

3. Change/reserve a range of User IDs (UIDs) and Group IDs (GIDs) on the *UID/GID range* tab.

   Refer to Set the authentication mode for more information.

4. Specify which users should get administration (*root*) access rights on the server from the *Admin Users* tab.

   If Windows ACL is enabled, there must be at least one admin user.

Display SMB connection status

To display active connections, shares, and locked files, click the *Manage* icon  , and select *SMB Connection Status*.

Update access rights for SMB client shares

To update the access rights for SMB clients, select a user in the *NAS Shares* section of the page, and click the *Update* icon.  For each path, you can change access rights between *Read* and *Read-Write*.

Assign user access to an SMB user

(Administrators only) The domain administrator can assign access to a specific domain user who will have access to an SMB user.

1. Select a client and click the *Access Control* icon. 
2. Select a user that can access this client and click *Assign*.

To unassign access for a user, clear the user checkbox and click Assign.

SMB clients - Groups

Select *Manage* from the menu bar, click the *NAS* tab, click the *SMB Clients* tab, and click the *Groups* tab. This page lists SMB groups on the selected server and the NAS shares assigned to a group. SMB groups are groups of Windows SMB users that use the SMB protocol to access NAS shares.

Create an SMB group

(User mode) To create an SMB group:

1. Click the "+" (*Create*) icon.
2. Enter information about the SMB group.

   Name - Specify a Windows group name up to 32 ASCII characters starting with a letter or underscore and excluding blanks or ~`!@#$%^&*()+={}[]|\\\/?:;<>,\\'"

   Members - Specify users to include in this group.

   Description - Optionally, specify a description for the Windows group up to 256 characters without consecutive colons.
3. Click *Create* when done.

Update an SMB group

(User mode) You can update the description and add/delete members of an SMB group.

Click the *Groups* tab, select a user, and click the *Update* icon. 

Delete an SMB group

(User mode) You can delete an SMB group. Select a user, and click the *Delete* icon. 

Update access rights for SMB group shares

(User mode) To update the access rights for SMB groups, select a group in the *NAS Shares* section of the page, and click the *Update* icon.  For each path, you can change access rights between *Read* and *Read-Write*.

Assign user access to an SMB group

(Administrators only) The domain administrator can assign access to a specific domain user who will have access to an SMB group.

1. Select a group and click the *Access Control* icon. 
2. Select the users that can access this group and click *Assign*.

To unassign access for a user, clear the user checkbox and click *Assign*.

NFS clients

Select *Manage* from the menu bar, click the *NAS* tab, and click the *NFS Clients* tab. This page lists NFS clients on the selected server and the NAS shares assigned to a client. NFS clients are Unix machines that use the NFS protocol to access NAS shares.

Create an NFS user

1. Click the "+" (*Create*) icon.
2. Enter information about the NFS client.

   Name - Specify a unique NFS client name up to 36 ASCII characters.

   Machine - Specify a unique resolvable name or IP address of the client machine. It can also be a subnet and netmask to grant NAS share access to an entire subnet using the following format: `subnet/netmask`, where netmask can be an IP or a prefix, for example, `192.168.0.0/255.255.255.0` or `172.30.123.0/24`.

   Comment - Optionally, specify a description for the NFS client up to 128 ASCII characters.
3. Click *Create* when done.

Update an NFS client

Select a client and click the *Update* icon. 

Delete an NFS client

Select a client and click the *Delete* icon.  You can force the deletion if NAS shares are assigned to the client.

Update access rights for NFS client shares

To update the access rights for NFS clients, select a client in the *NAS Shares* section of the page, and click the *Update* icon.  For each path, you can set permissions.

Assign user access to an NFS user

(Administrators only) The domain administrator can assign access to a specific domain user who will have access to an NFS user

1. Select a client and click the *Access Control* icon. 
2. Select a user that can access this client and click *Assign*.

To unassign access for a user, clear the user checkbox and click *Assign*.

Replication - Configuration

Select *Manage* from the menu bar, click the *NAS* tab, and click the *Replication* tab. This page lists the replication configurations and current replication sessions on the selected server. If this server is a target replication server, the source server(s) will be listed.

Replication protects data on NAS resources by maintaining a copy of data on another Habanero server. At prescribed intervals, new data from the source server is transmitted to the target server so that the NAS resources are synchronized.

The target Habanero server is usually located at a remote location. Only deduplicated data is sent over the WAN, providing bandwidth savings.

If the replica is needed, the administrator can share the replica folder that is on the target server, map/mount the appropriate share(s) and recover the necessary files. Data restore is quick and efficient from native format files rather than from tape backup formats.

Enable outgoing replication

The following are the requirements for setting up a replication configuration:

- You must have two Habanero servers with NAS.
- You must have enough space on the target server for the replica data.
- The target server must have been added into the portal.

1. Click the *Outgoing Replication* icon and select *Enable Outgoing Replication*.

   Target Server - Specify the remote replica server that will hold the replicated data.

   Transmission IP Address - Select the IP address to use for replication.

   Encryption - Determine if you want to use encryption of in-flight data.

   Global Replication Policy - Specify if you want to define a global replication policy that replicates all NAS directories instead of individual policies for some directories. For a global policy you will need to manually exclude directory paths that you do not want to replicate.

2. For a global replication policy, configure policy settings.

   Replicate only deduplicated files - Determine if you want to replicate only deduplicated files instead of all files. Files smaller than 8 KB are included for replication even though they are not deduplicated;

   Allow data replication back to source - Determine if you want to allow replication of replica data back to the source for disaster recovery purposes. There is no need to select this option for an initial replication configuration. For example, if server A normally replicates to server B and then server A is destroyed, you can set this option on server B when needed and all of the data will be replicated back to server A.

   Replication Method - Determine how replication should occur, manually or via a schedule.

   Start Time/End Time - For scheduled replication, specify when it should begin and, optionally, when it should end.

   Frequency/Interval - For scheduled replication, specify how often it should run, hourly, daily, weekly, or monthly.
   - Hourly - Run every *n* hours and *m* minutes.
   - Daily - Run every *n* days.

- Weekly - Run every n weeks. Select which days to run.
- Monthly - Run every day or week. Select which day and/or week in the month to run.

3. For a global replication policy, select when replication should not occur.

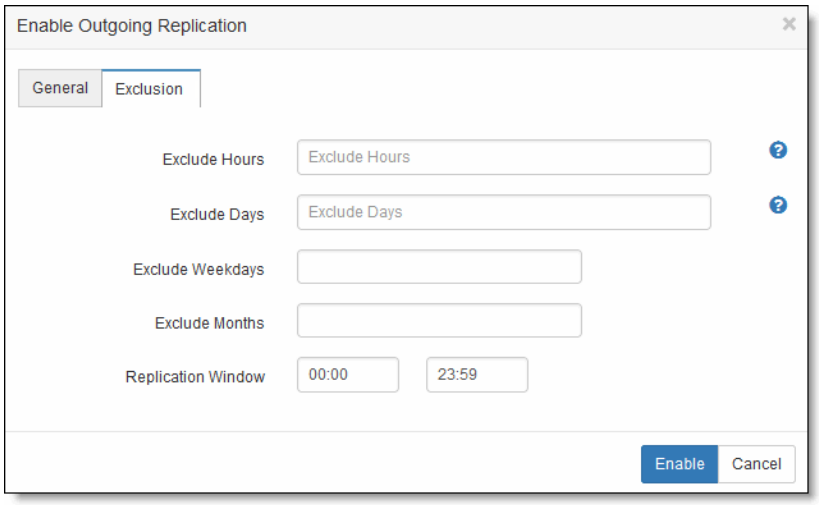

In the *Exclude Hours* field, you can:

- Specify one or more hours (0-23), separated with commas, such as 1,3,5
- A time range, such as 2-8
- Mix of hours and range, such as 0,2-8,18

In the *Exclude Days* field, you can specify one or more days (1-31), separated with commas. For example: 1,3,21

You can also exclude specific days of the week or specific months.

You can specify a window to limit when replication can run.

4. Click *Enable* when done.

Set outgoing replication throttle

Throttling allows you to limit the amount of available IP network bandwidth that is used for outgoing replication on the source server side. If throttling is not used, replication will use the maximum bandwidth that is available.

To set a throttle for outgoing replication, click the *Outgoing Replication* icon  , select *Set Outgoing Replication Throttle*, and select *Throttle Replication*.

Click *Add* to specify the time range during which throttling should occur and the maximum number of Kbs per second that should be used for replication. Transmission will not exceed the set value. This is a global server parameter and affects all resources. Valid input for bandwidth is 10-100,000,000 Kb/s.

Disable outgoing replication configuration

To disable outgoing replication, click the *Outgoing Replication* icon  , and select *Disable Outgoing Replication*.

Configure incoming replica

Configure a server to store replicated files from a source server. You should have created at least one NAS resource before you can configure incoming replication.

1. Click the *Incoming Replication* icon  and select *Configure Incoming Replica*.

   Use Any Resource - Specify if you want the system to automatically select any available NAS resource for incoming replication the first time replication is started.

   NAS Resource - Select a specific NAS resource if you did not select the *Use Any Resource* option.

2. Click *Configure* when done.

Purge replica files

You can purge old replica files from the target server. Only one purge job can run at a time.

1. Click the *Incoming Replication* icon  and select *Purge Replica Files*.

   Retention Days - Specify the number of days that files have not been accessed to identify old files to purge, in the range of 1 and 2,000. The default is 365 days.

   Source - Specify whether files should be purged from all source servers or just a specific source server.

   Path - For *All Servers*, specify if you want to purge all files or if you want to start from a specific directory path. Blank means purge all replica files.

   Save purge results - Specify if you want to save the list of purge files in a CSV file.

   Do not delete old files - Select if you *only* want to save purge files in a CSV file and do not want to actually delete them at this time.

2. Click *Start* when done.

View purge results

If you selected to save purge results when running *Purge Replica Files*, click the *Incoming Replication* icon,  and select *Get Purge Results* to download a CSV file containing the purge files.

Refresh replication session status

To refresh the status of active replication session, click the *Refresh Replication Sessions* icon. 

Replication - Policies

Select *Manage* from the menu bar, click the *NAS* tab, click the *Replication* tab, and click the *Policies* tab. This page lists the replication policies on the selected server and the replication status of each policy. If this server has a global replication policy for NAS, only the global policy will be listed.

Create a NAS replication policy

You can only create a replication policy if there is no global replication policy on the *Configuration* tab. The maximum number of replication policies is 32.

1. Click the "+" (*Create*) icon.
2. On the *General* tab, specify policy properties.

   Policy name - Specify a name up to 128 characters, excluding '"&<>`/\,

   Replicate only deduplicated files - Determine if you want to replicate only deduplicated files instead of all files. Files smaller than 8 KB are included for replication even though they are not deduplicated;

   Allow data replication back to source - Determine if you want to allow replication of replica data back to the source for disaster recovery purposes. There is no need to select this option for an initial replication configuration. For example, if server A normally replicates to server B and then server A is destroyed, you can set this option on server B when needed and all of the data will be replicated back to server A.

   Paths to Replicate - Select NAS resources or share paths to include in replication. You can select up to 64 paths.

   Replication Method - Determine how replication should occur, manually or via a schedule.

   Start Time/End Time - For scheduled replication, specify when it should begin and, optionally, when it should end.

   Frequency/Interval - For scheduled replication, specify how often it should run, hourly, daily, weekly, or monthly.
   - Hourly - Run every *n* hours and *m* minutes.
   - Daily - Run every *n* days.
   - Weekly - Run every n weeks. Select which days to run.
   - Monthly - Run every day or week. Select which day and/or week in the month to run.
3. For scheduled replication, on the *Exclusion* tab, select when replication should not occur.

   In the *Exclude Hours* field, you can:
   - Specify one or more hours (0-23), separated with commas, such as 1,3,5
   - A time range, such as 2-8
   - Mix of hours and range, such as 0,2-8,18

   In the *Exclude Days* field, you can specify one or more days (1-31), separated with commas. For example: 1,3,21

   You can also exclude specific days of the week or specific months.

   You can specify a window to limit when replication can run.
4. Click *Create* when done.

Update a NAS replication policy

You can update a NAS replication policy as long as replication of the policy is not in progress.

You cannot update a global replication policy. Select a policy, and click the *Update* icon. 

Delete a NAS replication policy

You can delete a NAS replication policy as long as replication of the policy is not in progress. You cannot delete the global replication policy.Click the *Policies* tab, select a policy, and click the *Delete* icon. 

Start replication of a NAS policy

You can start replication for a manual or scheduled policy. Only one outgoing replication or file synchronization job can be active at a time. For a global replication policy, you can select a NAS resource, share, or folder to be replicated instead of replicating all resources.

Click the *Policies* tab, select a policy in the bottom section of the screen, click the *Manage* icon  , and select *Start Replication*.

Synchronize files

You can start synchronization of files/directories between the source and replica server. If a replicated file has been deleted from the source, the file will be deleted from the replica.

Only one outgoing replication or file synchronization job can be active at a time. For a global replication policy, you can exclude directory paths that you do not want to synchronize.

Select a policy in the bottom section of the screen, click the *Manage* icon  , and select *Synchronize Files*.

Suspend/resume replication schedule

You can suspend the schedule of a replication policy. If suspended, you can resume the schedule.

Select a policy in the bottom section of the screen, click the *Manage* icon  , and select *Suspend Replication Schedule* or *Resume Replication Schedule*.

Integrity Check

Select *Manage* from the menu bar, click the *NAS* tab, and click the *Integrity Check* tab. This page lists the configuration for integrity checking as well as the status for the selected server. Integrity checking confirms that the stub files and the unique data stored in the repository are valid for deduplicated NAS files.

Configure integrity checking

During deduplication, the system analyzes blocks of data and determines whether the data is unique or has already been copied to the repository. This process passes single instances of unique data to the repository and replaces each deduplicated file with a small *stub* file, whose function is to point to the repository and is used to retrieve stored data.

You can run an integrity check to confirm that the stub files and the unique data stored in the repository are valid and that the stub file can generate source data correctly.

The integrity check can be run on an as-needed basis or can be scheduled to occur automatically.

Notes:

- The integrity check will only be able to validate stub files and source data deduplicated **after** integrity checking has been enabled. Therefore, integrity verification will only apply to those files deduplicated after integrity checking has been enabled.
- Performing an integrity check is a resource-intensive operation that may affect the performance of other operations on the server.

1. Click the *Configure* icon  and select *Enable Integrity Check*.

   Validate Source Data - Specify if you want to validate source data as well. If selected, source data will be generated from the stub file to validate the checksum of the source data too.

   Check Specific Paths - Determine if you want to check all paths or just specific NAS resources, shares, or folders. If enabled, select one or more paths.

   Method - Determine how integrity checking should occur, manually or via a schedule.

   Start Time/End Time - For scheduled integrity checking, specify when it should begin and, optionally, when it should end.

   Frequency/Interval - For scheduled integrity checking, specify how often it should run, hourly, daily, weekly, or monthly.

   - Hourly - Run every *n* hours and *m* minutes.
   - Daily - Run every *n* days.
   - Weekly - Run every n weeks. Select which days to run.
   - Monthly - Run every day or week. Select which day and/or week in the month to run.

2. For scheduled integrity checking, on the *Exclusion* tab, select when integrity checking should not occur.

   In the *Exclude Hours* field, you can:

- Specify one or more hours (0-23), separated with commas, such as 1,3,5
- A time range, such as 2-8
- Mix of hours and range, such as 0,2-8,18

In the *Exclude Days* field, you can specify one or more days (1-31), separated with commas. For example: 1,3,21

You can also exclude specific days of the week or specific months.

3. Click *Enable* when done.

Update the integrity check configuration

To update the integrity check configuration, click the *Configure* icon  and select *Update Integrity Check*.

Disable the integrity check configuration

To disable the integrity check configuration, click the *Configure* icon  and select *Disable Integrity Check*.

Start/suspend integrity check

You can manually run an integrity check. If running, you can suspend it.Click the *Manage* icon  and select *Start* or *Suspend*.

Refresh status of integrity checking

To refresh the status of integrity checking, click the *Refresh Integrity Status* icon. 

Jobs

Select *Manage* from the menu bar, select your server, and click the *Jobs* tab, to display running jobs.

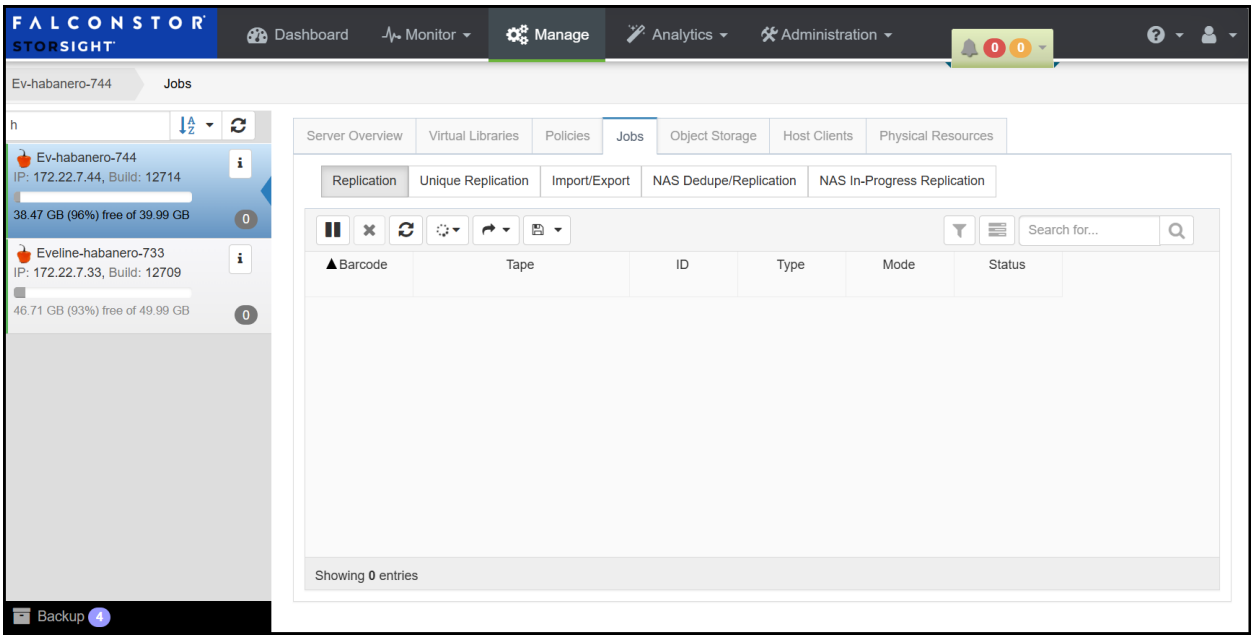

Replication jobs

The *Replication* tab lists the virtual tapes that are currently being replicated on the selected server. The following information is displayed for each tape:

- Barcode
- Tape
- Tape ID
- Replication mode
- Status - Running, Waiting for retry, Waiting for a slot.

Note: The page does not automatically refresh by default. Click the *Refresh* icon above the tape list to manually refresh the tape information. Select the *Status Auto Refresh* icon to enable automatic refreshing.

Cancel replication job of a tape

You can cancel a replication job that is running. Select a tape and click the *Cancel Job*.

Incoming Replication jobs

Select *Manage* from the menu bar, click the *Jobs* tab and then the *Incoming Replication* tab. This page displays information for incoming replication jobs for deduplicated tapes. It includes tapes currently replicating (after the index has been replicated) and those awaiting replication.

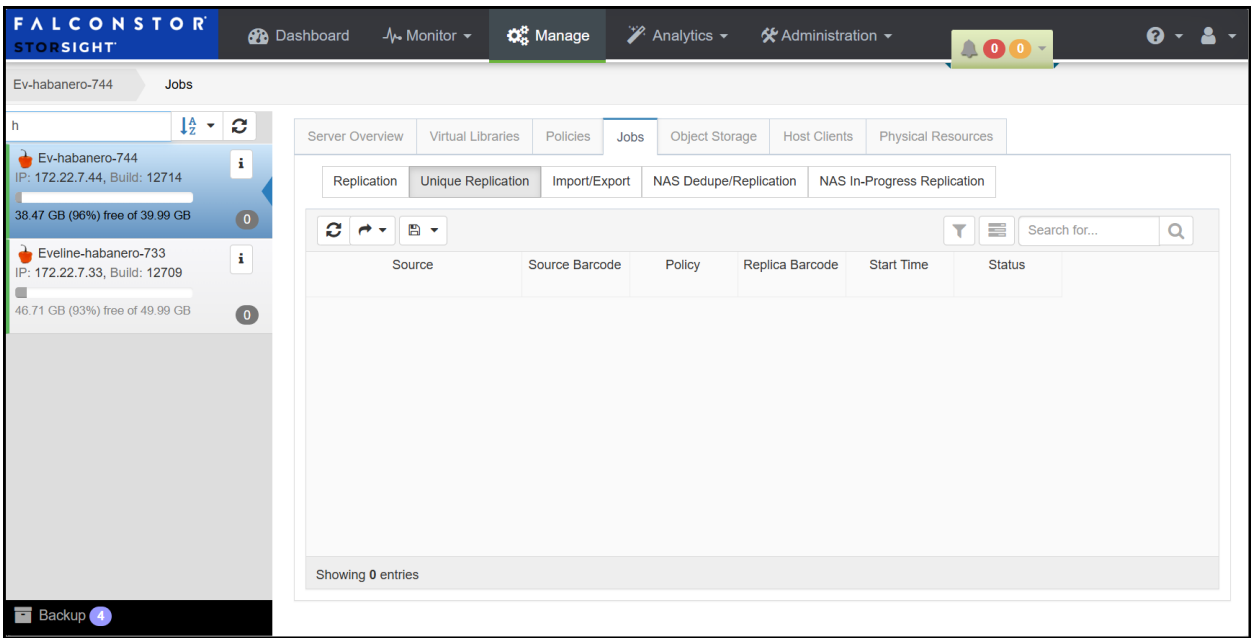

The following information is displayed for each job:

- Source server
- Source tape barcode
- Policy
- Replica barcode
- Start time
- Status - Running, waiting for retry.

Note: The page does not automatically refresh by default. Click the *Refresh* icon above the tape list to manually refresh the tape information. Select the *Status Auto Refresh* icon to enable automatic refreshing.

Import/Export jobs

Select *Manage* from the menu bar, click the *Jobs* tab and then the *Import/Export* tab. This page lists all import, export, migration to object storage, and recovery from object storage jobs on the selected server. Up to 10 migration and 5 recovery jobs can be running at the same time. Completed jobs are purged after 30 days. Migration to object storage jobs that do not complete (failed or canceled jobs) are also purged after 30 days and the data at the object storage provider is deleted at that time.

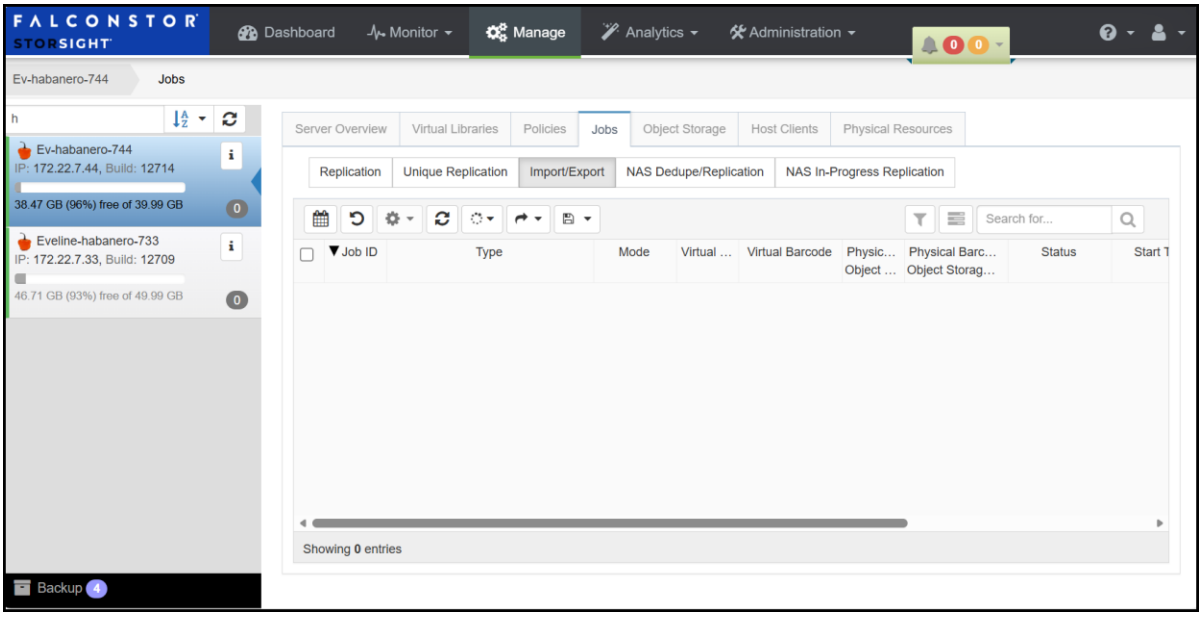

The following information is displayed for each job:

- Job ID
- Job type - Export to standalone drive/library, import from standalone drive/library, move to import/export slot, export to standalone drive/library with stacking, import from library with stacking, move to import/export slot with stacking, scan physical tape in standalone drive/library, export tape to object storage, import from object storage
- Cloud Archive mode - Copy or Move
- Virtual tape ID
- Virtual tape barcode
- Physical tape barcode or object storage
- Status - Waiting for tape/drive, Preparing for Object Storage Import/Export, Running, Completed, Cancelled, Failed, Canceling, Waiting for I/E slot, Waiting to be deleted, Hold, Resumed.
- Start time
- Elapsed time in seconds for running jobs
- Amount of data transferred
- Progress percentage

Note: The page does not automatically refresh by default. Click the *Refresh* icon above the tape list to manually refresh the tape information. Select the *Status Auto Refresh* icon to enable automatic refreshing.

Manage import/export jobs

You can restart one or more canceled or failed jobs, cancel a running job, or delete one or more completed, canceled, or failed jobs.

Select *Manage* from the menu bar, select your server, click the *Jobs* tab, click the *Import/Export Jobs* tab, select a job, and then click the *Manage* icon. 

Set retry parameters

Should a job fail, it will be retried based on the import/export retry settings, as long as the tape remains in the vault. To specify rules for retrying failed jobs:

1. Select *Manage* from the menu bar, select your server, click the *Jobs* tab, and then the *Import/Export Jobs* tab.
2. Click *Set Retry Parameters*.

   Select *Retry failed jobs* and specify the number of retries (1-10) and the interval between retries (1-1,440 minutes).
3. Click the *Set* button.

Schedule Import/Export Jobs report

The Import Export Jobs report can be scheduled to be emailed to you on a regular basis. Refer to Schedule a report for more information.

NAS Dedupe/Replication page

Select *Manage* from the menu bar, click the *Jobs* tab and then the *NAS Dedupe/Replication tab.* This page displays deduplication and outgoing replication statistics.

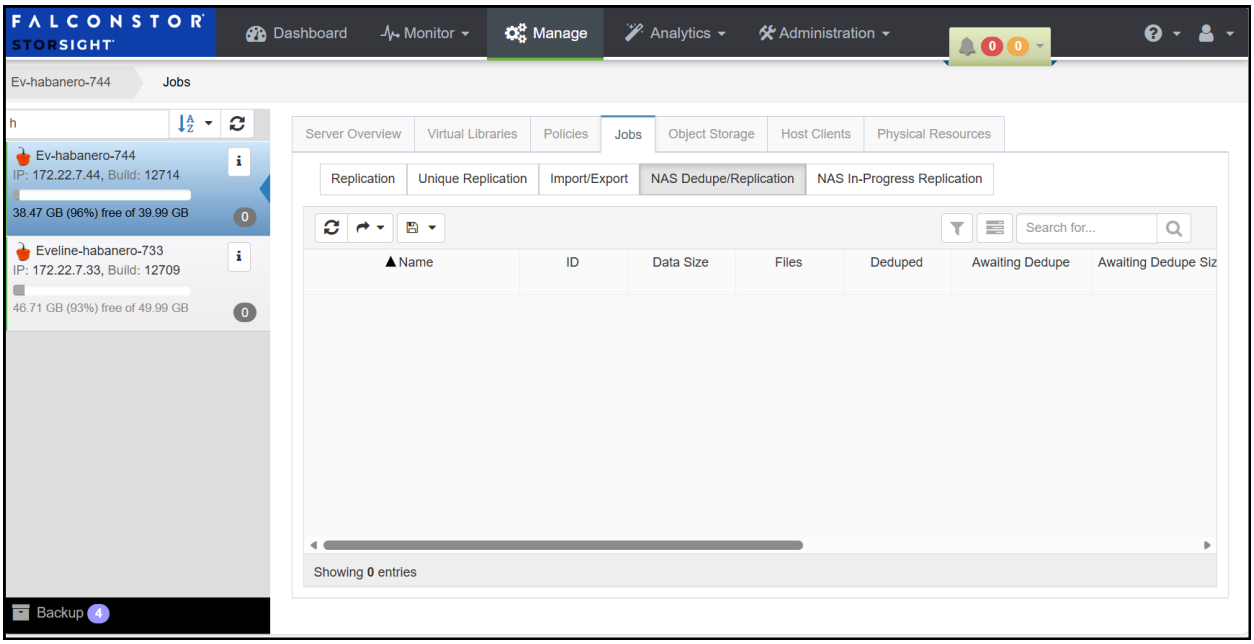

The following is displayed for each NAS resource:

- Name
- ID
- Data Size - Total size of files copied to all file systems (including those that have not yet been deduplicated)
- Number of files - Includes files that have been deduplicated, files that were excluded, and files awaiting deduplication
- Number of deduplicated files - Does not include excluded files (smaller than 8 KB or encrypted/compressed files with known file extensions)
- Number of files awaiting deduplication
- Size of files awaiting deduplication
- Number of files replicated
- Number of files awaiting replication
- Size of files awaiting replication
- Size of replicated data
- Size of unique replicated data

Note: The page does not automatically refresh by default. Click the *Refresh* icon above the tape list to manually refresh the tape information. Select the *Status Auto Refresh* icon to enable automatic refreshing.

NAS In-Progress Replication page

Select *Manage* from the menu bar, click the *Jobs* tab and then the *NAS In-Progress Replication* tab. This page lists current outgoing replication sessions at the server level for all NAS resources.

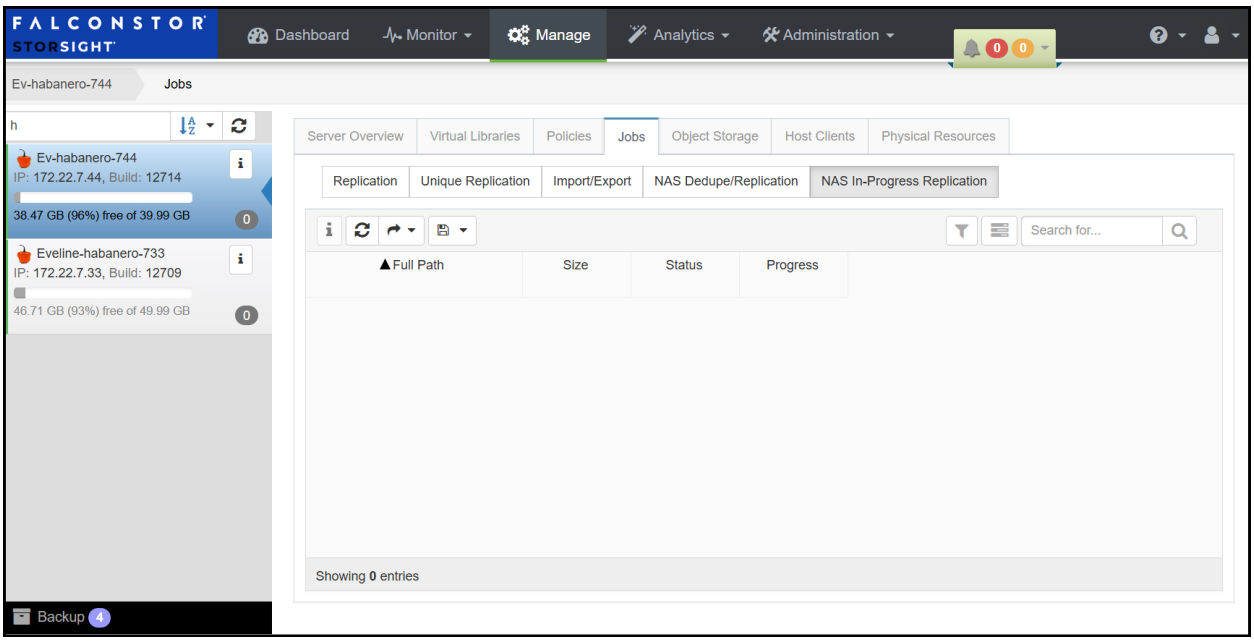

The following is displayed for each file being processed:

- Full path of the file
- File size
- Status
- Progress percentage - Click Refresh to update the data

Note: The page does not automatically refresh by default. Click the *Refresh* icon above the tape list to manually refresh the tape information. Select the *Status Auto Refresh* icon to enable automatic refreshing.

Analytics

Analytics offers centralized monitoring, reporting, and analytics across multiple Habanero servers.

Select *Analytics* from the menu bar and click one of the options below:

Reports - Display a wide variety of predefined and custom reports that help you manage your servers.

Detailed reports – Generate a detailed virtual tape information.

Sc - Manage and view scheduled reports.

Smart rules – Allow to create rules to check performance of virtual tape libraries and FC clients.

Alerts - Display alerts generated based on configured rules

Au - Tracks and logs actions performed by users.

Reports

You can also search for a report, sort the information in the columns, use the *Filter*  or the *Arrange*  icons to select which parameters you want displayed in each report and in which order.

Inventory reports

Inventory reports provide the list of resources on all servers. The following inventory reports are available:

Clients

List of host clients on all servers. For each client, this report displays:

- Name
- Client ID
- Server
- Server IP address
- Server type
- Number of assigned devices (click to see a list)
- Protocol(s) - FC and/or iSCSI
- Volume Set Addressing (VSA) for FC clients

Virtual Tape Libraries

List of all virtual tape libraries on all servers. For each virtual tape library, this report displays:

- Name
- ID
- Server
- Server IP address
- Vendor
- Product
- Firmware
- Serial Number
- Media type
- Number of slots
- Number of drives
- Number of tapes
- Number of loaded tapes
- Number of clients (click to see a list)
- Starting/ending barcodes
- Storage pool
- Capacity on Demand (COD) option
- COD initial allocation size
- COD increment size
- COD maximum capacity
- Physical tape archive option
- Cloud archive option
- Cloud archive export mode
- Physical Tape duplication option
- Encryption option

Virtual Tape Drives

List of all virtual tape drives in virtual tape libraries and all standalone virtual tape drives on all servers. For each virtual tape drive, this report displays:

- Name
- ID
- Virtual tape library
- Server
- Server IP address
- Status - Empty, Loaded, Ejected
- Loaded tape – tape name if the status is loaded
- Barcode of loaded tape
- Vendor
- Product
- Firmware

- Serial Number
- Media type
- Element number
- Number of clients (click to see a list)
- COD option
- Max tape capacity

Physical Tape Libraries

List of all physical tape libraries on all servers. For each physical tape library, this report displays:

- Name
- ID
- Server
- Vendor
- Product
- Firmware
- Serial number
- ACSL (Adapter:Channel:SCSI ID:LUN)
- Number of drives (click to see a list)

Physical Tape Drives

List of all physical tape drives in physical tape libraries on all servers. For each physical tape drive, this report displays:

- Name
- ID
- Server
- Physical tape library
- Vendor
- Product
- Firmware
- Serial number
- ACSL (Adapter:Channel:SCSI ID:LUN)

NAS Resources

List of all NAS resources on all servers. For each NAS resource, this report displays:

- Name
- ID
- Status - Online/offline
- Server
- Server IP address
- Mount point - The directory from which the share can be accessed
- Mount status - Mounted, Unmounted
- Capacity
- Allocated and free size
- Deduplication - Included, excluded, partial (some sub-folders are excluded)
- Replication - Included, excluded, partial (some sub-folders are excluded), empty (not configured or configured with multiple policies)
- Number of shares
- Storage pool
- Physical device (click to see the ACSL)

Capacity Management reports

Capacity management reports provide capacity information for resources on all servers. The following capacity management reports are available:

Storage Pools

Storage pool usage on all servers. For each storage pool, this report displays:

- Name
- Server
- Server IP address
- Server type
- Resource types (if a number is displayed, click to see a list)
- Size - Total, allocated, and free
- Number of physical devices (click to see a list with ACSL and GUID)
- Amount used for NAS
- Amount used for tapes

Daily Source Data

Daily amount of source data for a server during a selected month. This report displays:

- Date
- Total amount of source data

Monthly Source Data Summary

Total amount of source data for a selected month for each server. This report displays:

- Server
- IP address
- Amount of source data (TB)

NAS SMB Shares

This report shows usage of SMB shares on all servers. For each SMB share, this report displays:

- Timestamp for the record; data is recorded every 30 minutes
- Server
- Server IP address
- SMB share name
- Used size

Performance Monitoring reports

Performance monitoring reports provide I/O performance information for resources on all servers.

You can click on the blue links to display a graph that shows information for the last 15, 30, 45, and 60 minutes. You can mouse over any point on the graph for more information. Use the slider at the bottom to adjust the granularity of the graph.

The following performance monitoring reports are available:

Virtual Tape Library Performance

This report shows performance of virtual tape libraries on all servers. Library performance is the sum of performance values of each tape drive in the library. If compression is enabled, the read/write throughput and read/write compressed values will be different. For each virtual tape library, this report displays:

- Name
- ID
- Server
- Server IP address
- Read throughput
- Write throughput
- Compressed read
- Compressed write

Virtual Tape Drive Performance

This report shows performance of virtual tape drives on all servers. Tape drive performance values are cumulative from the time a tape is loaded on the drive. Unloading or reloading a tape will reset the value to 0; if the current values are less than the previous values, a tape has been reloaded or unloaded in the drive. If compression is enabled, the read/write throughput and read/write compressed values will be different.

For each virtual tape drive, this report displays:

- Name
- ID
- Virtual tape library
- Server
- Server IP address
- Read throughput
- Write throughput
- Compressed read
- Compressed write

NAS Resource Performance

This report shows performance of NAS resources on all servers. For each NAS resource, this report displays:

- Name
- ID
- Server
- Server IP address
- Read, write, and total throughput

Client Performance

This report shows performance values for Fibre Channel clients on all servers. Performance values are not available for iSCSI clients and NAS clients and display as 0. For each client, this report displays:

- Name
- ID
- Server
- Server IP address
- Server type
- Number of assigned devices (click to see a list)
- Read, write, and total throughput for FC clients
- Read, write, and total IOPS for FC clients
- Read, write, and mean latency for FC clients

Service Status reports

Service status reports provide information for deduplication and replication services for NAS resources. The following service status reports are available:

NAS Deduplication and Replication

This report shows deduplication and replication statistics for NAS resources for a specific server. For each NAS resource, this report displays:

- Name
- ID
- Total size of files, including those that not yet been deduplicated
- Total number of files, including files that have been deduplicated, files that were excluded, and files awaiting replication
- Number of deduplicated files - Does not include excluded files (smaller than 8 KB or encrypted/compressed files with known file extensions)
- Number of files awaiting deduplication
- Size of files awaiting deduplication
- Number of files replicated
- Number of files awaiting replication
- Size of files awaiting replication
- Size of replicated data
- Size of unique replicated data

NAS Replication Sessions

This report shows NAS incoming and outgoing replication sessions for the specified period of time on all servers. For each NAS replication session, this report displays:

- Server
- Record timestamp
- Replication session ID
- Replication type - Incoming or outgoing
- Replication start and end times
- Duration of replication session
- Replication source server
- Replication trigger - Schedule or manual
- Replication status - Completed, In progress, Cancelled, or Failed
- Option to replicate deduplicated files only
- Number of files analyzed for replication
- Number of skipped files
- Number of aborted files
- Number of replicated files
- Size of replication data
- Size of transmitted data
- Replication throughput for incoming replication per source server - If multiple servers are replicating to this server, this represents the sum of throughput from each.

Trend reports

Trend reports show performance and capacity trends based on real-time and historical data.

The following trend reports are available:

- Virtual Tape Library Performance Trend
- Virtual Tape Drive Performance
- NAS Resource Performance Trend
- FC Client Performance Trend

Trend reports can be presented in line or area format.

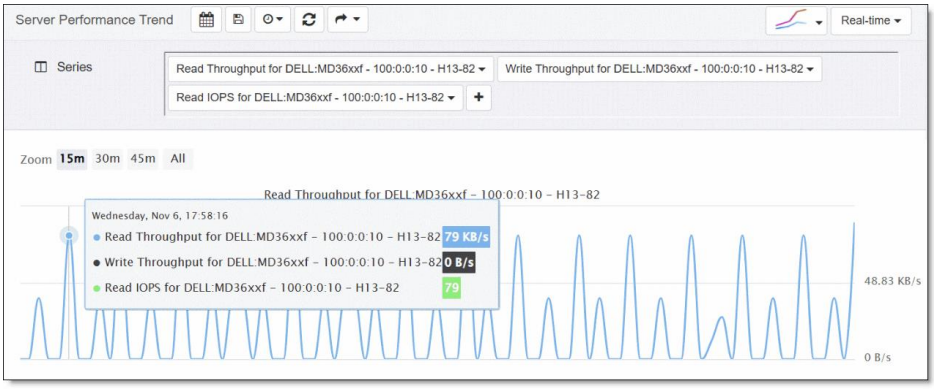

The number of data points on a report is determined by the time period selected. When you highlight a data point on the graph, the average value for the time between data points is displayed. The resolution for data points is as listed below:

Time period	Resolution
Past hour	10 seconds
Two hours to seven days	5 minutes
More than 7 days but less than 30 days	1 hour
More than 30 days	1 day

By default, data for the last 15 minutes is displayed. You can change the resolution by selecting a different *Zoom* time period. When a larger time period is selected, a higher resolution is used to retrieve the average value for that time period.

The report does not automatically refresh by default. Click the *Refresh* icon above the report to manually refresh the information. Select the *Status Auto Refresh* icon to enable automatic refreshing every 30 seconds, one minute, or five minutes.

You can change the displayed statistics by clicking on the *Real-time* drop-down menu. For historical usage, you can select a time frame or date range.

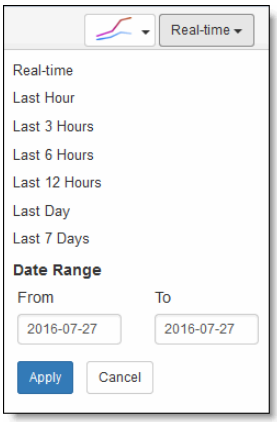

Use the slider at the bottom to further adjust the granularity of the data displayed.

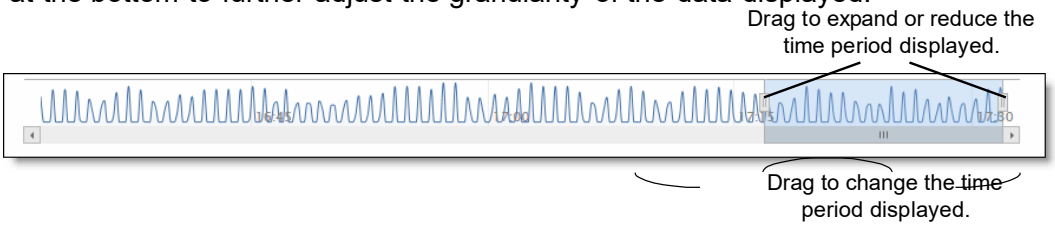

The number of data points on a trend report is determined by the time period selected. Low-resolution data points get purged regularly from the database. For example, 10-second data points get purged every two hours and the data is rolled up to a higher resolution. If there are not enough data points to plot, the trend graph may become empty and you will need to select a larger time period for display.

Each trend report consists of one or more graphs, each called a *Series*. The first time you view a trend report, each series is displayed for the first standalone server. You can customize trend reports to suit your needs.

1. Select a trend report.
2. Click the arrow by a series to modify it or click the “+” icon to add a series.
3. Make your selections.
4. Specify a *Series Name* or accept the default name.
5. Click *OK*.
6. Repeat for the other series in this report.
7. You can mix and match categories, components, and objects in a single trend report.

8. Click the *Save* icon  to save your report.

   Saving a report makes a copy of the report. Saved reports are available under *Custom* in the left pane.

   Note that if you do not save the report, your last series selections will still be visible the next time you return to the report.

Top N Components reports

The Top Virtual Tape Libraries by Throughput report displays the average value of top 10 libraries on all servers.

The report can be generated in pie chart format and bar chart format and can be generated with values from the past 5 minutes, 1 hour, or 1 day. Report values can be displayed for *Read*, *Write*, or *Total*.

Note that all report values are rounded. Therefore, if the average value is very low and results in a fraction for the selected time frame, the report can display zero even though there was some activity during the time frame.

In pie chart format, you can remove pieces of the pie for what-if scenarios. To remove a pie piece, click the corresponding name under the pie chart to remove that section. Click the name again to return that pie section back to the chart.

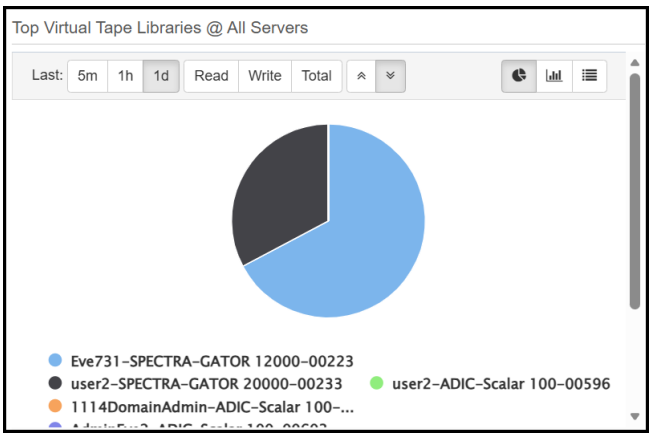

Detailed reports

You can generate a detailed CSV report for virtual tape information on a specific server. The report generation job runs in the background and can take a while. The status of the report job can be queried within 24 hours. You can get status regularly and once the report is completed, download it. The report remains available on the server for 7 days, after which it is automatically deleted to free up space.

To create a report, select a server, click the *Generate Detailed Report* icon  , and select the Virtual Tape Information report.

You can filter report data by specifying barcode information: tapes within a specified bar code range, with a specified prefix, or containing a specified text string.

On the *Filtering Components* tab, you can also select the virtual tape libraries and policies that you want to include. All libraries and polices are included by default.

Because of the amount of information available for virtual tapes, you can select one or more types of information to include:

- *Overall summary view* - For each tape, the report displays information about usage, location, policy, enabled options and services.
- *Vault view* - This view is like the *Overall Summary View*, for tapes that are in the vault.
- *Replica resources view* - This view presents information about usage, source server, and FVIT or LVIT type for replica tapes.
- *Detailed tape view* - This view is like the *Overall Summary View*, for tapes that are in a specific library.
- *Migration view* - This view presents information about usage, export time, and object storage for exported tapes.

Scheduled reports

Reports can be scheduled to be emailed to you on a regular basis. Select *Scheduled Reports* from the *Analytics* menu bar to see a list of existing scheduled reports. For each report, you will see:

- Report name
- Type of report
- Format – PDF or CSV
- Indication of whether the schedule is suspended
- Time period for the report
- Frequency of the report
- Start time
- Email recipient
- Date last sent

From this page, you can:

- Update the report schedule and email settings
- View the report as PDF or CSV
- Suspend the report schedule
- Delete the report and its schedule

The following types of reports can be scheduled to be emailed to you on a regular basis:

- Capacity Management - Storage Pool
- Capacity Management – Daily Source Data
- Capacity Management – Monthly Source Data Summary
- Service Status - NAS Deduplication and Replication
- Trend
- Deduplication Tapes
- Deduplication Activity
- Deduplication Run History
- Deduplication Tape History
- Import/Export Jobs

Schedule a report

Some reports are scheduled from the *Reports* page, some from the *Policies* tab for *Tapes, Activity, Run History, Tape History,* and from the *Jobs* tab for *Import/Export*. To schedule a report:

1. Select a report.
2. Click the *Schedule Report* icon  .
3. Specify a name for the report.
4. Select email frequency - Daily, Weekly, Monthly.
5. Start time to send the email for each frequency
6. Enter the mail address to send the report to. You can specify several addresses separated by comma (no space).
7. Specify the email subject.
8. Optionally, specify email message contents.
9. Select the report format: PDF or CSV.
10. Optionally, enter a name for the report file; if not specified, the report name will be used.
11. Click *Create*.

Smart rules

Smart rules allow you to monitor performance of virtual tape libraries and FC clients.

When a rule is created, monitoring begins immediately (assuming new data is available for the components in the rule) and the rule condition is compared against the average values for the last five minutes of metrics. When a rule condition is reached for first time, an alert is created. The *Alerts* page will show the *First Occurrence* as the current time. The *Last Occurrence* will not show a time until the condition no longer exists. Monitoring occurs continuously and there is no update to the alert until the condition is resolved. At that time, the *Last Occurrence* will be updated with current time. If the rule condition is reached again, a new alert is created.

Select *Smart Rules* from the *Analytics* menu bar to see a list of existing rules, along with the alert severity, and if they are active or not. From this page, you can create a rule, modify or delete an existing rule.

Note that you will see an exclamation icon ⚠ next to any rule that has an obsolete/missing object. You should update the rule with a new object or delete the rule if it is no longer needed.

Create a smart rule

1. Click the *"+"* icon.

2. Specify a name and select the *Rule Type:*
   Virtual Tape Library Performance - Monitors virtual tape libraries. You can select multiple server/virtual tape library pairs.
   Client Performance - Monitors specific FC clients. You can select multiple server/client pairs.

3. Click *Create*.

   The rule template displays.

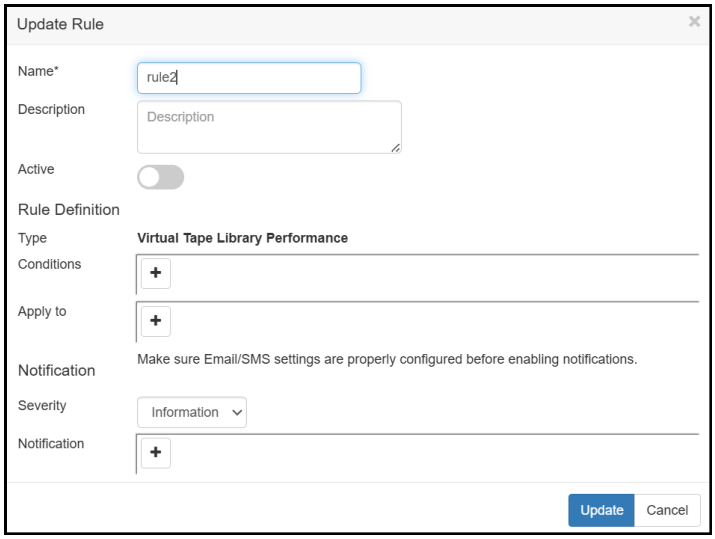

4. Select *Active* to enable the rule.

5. Click '+' in *Conditions* to define rule conditions as appropriate for the rule type:
   Field - The performance metric

   Operator - =,!=, <, >, <=, >=, between, not between

Value – Performance value and applicable unit: KBps, MBps, GBps, μs and ms

When you add multiple conditions, an "and/or" box displays, in which "and" is the default. This means an alert will be triggered when all of the rule conditions apply. The last condition is the one that actually triggers the alert and is the one that displays in the list of alerts. If you specify "or", whichever condition triggers the alert is the one that displays.

6. Click '+' in *Apply to* specify on which servers and libraries or clients the rule applies to.
7. Set the severity level - *Information*, *Warning*, or *Critical*.

Alerts

Alerts display as a result of the rules configured. When the condition occurs, an alert is triggered. Monitoring occurs continuously and there is no update to the alert until the condition is resolved. After alert resolution, if the rule condition is reached again, a new alert is created.

Select *Alerts* from the *Analytics* menu bar to see all alert notifications.

By default, all unacknowledged alerts are listed. You can filter the display by criteria listed on the top. To clear the search criteria, click the *Clear Search* icon  .

The total number of alerts is displayed at the bottom of the page. You can also select how many alerts to display per page.

Each alert includes:

- Server name
- Alert severity
- Rule triggering alert
- Resource name on which the rule condition has occurred
- Timestamp of the first occurrence
- Timestamp of the last occurrence of the alert
- Alert details

Alert general rules

The following table shows examples of the types of alerts you may see:

Rule Name - Alert	Description	Interval	Level
High Storage Pool Usage	Available disk space usage in a storage pool is over a specified threshold. Parameter: Threshold (default 90%)	Hourly	Warning
NAS Low Free Space	Free disk space on a NAS resource is below a specified threshold (default 200 GB) for any resource larger than a specified size (default 1024 GB)	Hourly	Warning

Service Failure	An event is recorded in the event log indicating a service failure since the last check.	Every minute	Error
Unresponsive Server	Main server modules (including the alert module) are not running resulting in the server becoming unresponsive.	Every minute	Critical

Manage alerts

You can acknowledge, unacknowledge, or delete one or several alerts. Click the *Manage* icon  and select an operation.

When you acknowledge an alert, you are “marking it as read” and it will not be displayed by default. However, you can change the *Acknowledgment* type to show all alerts or only acknowledged alerts.

Acknowledged alerts are retained for 7 days; non-acknowledged alerts are retained for 14 days.

Audit log

The audit log tracks and logs actions performed by administrators and users on all servers. Information in the audit log is maintained for 365 days.

Select *Analytics* from the menu bar and select *Audit Log* to see all activities.

By default, all actions are listed. You can filter the display by criteria listed on the top. To clear the search criteria, click the *Clear Search* icon  .

The total number of items is displayed at the bottom of the page. You can also select how many to display per page.

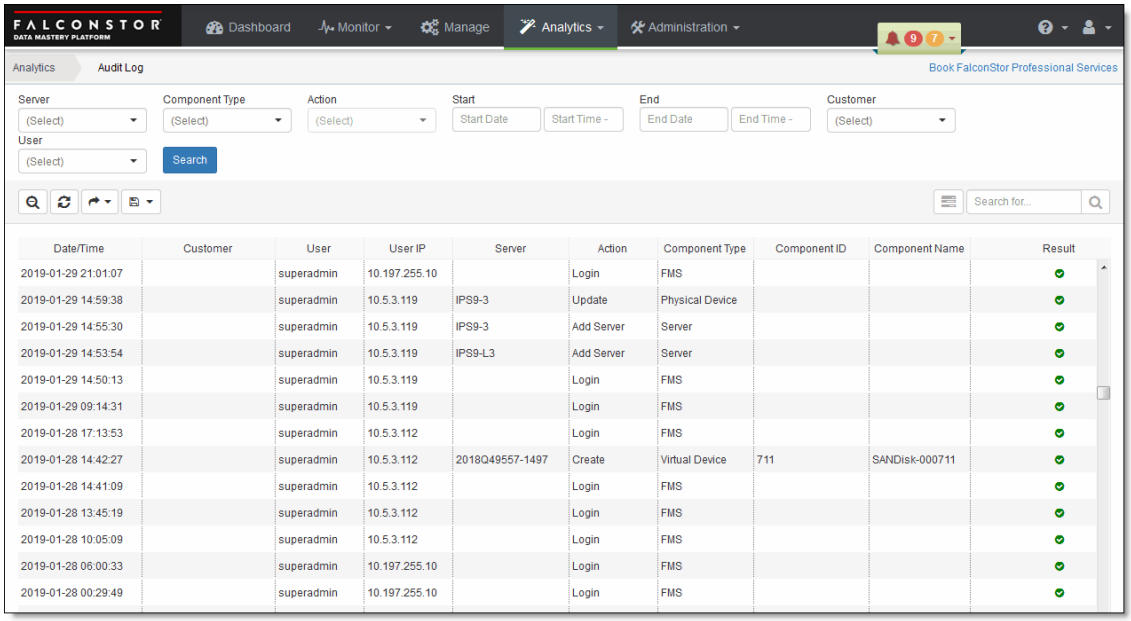

You will see the following for each entry:

- Date/time of occurrence
- User who performed the action
- IP address from user
- Server name, if a configuration change was performed.
- Action performed
- Component type - NAS, Physical Tape Drive, Physical Tape Library, Policy, Replica Tape, Replication, Replication Job, Tape Import/Export Job, Virtual Tape, Virtual Tape Drive, Virtual Tape Library.
- Component name
- Component ID
- Action result – Successful or Failed; click to see details.

Monitor

Monitor provides a read-only view of information about your resources. This is meant for monitoring the setup without making any configuration changes. The user Viewer role has access to this page but not to the *Manage* page.

Overview page

The *Overview* page allows you to view storage usage, performance, and the number of alerts for your servers.

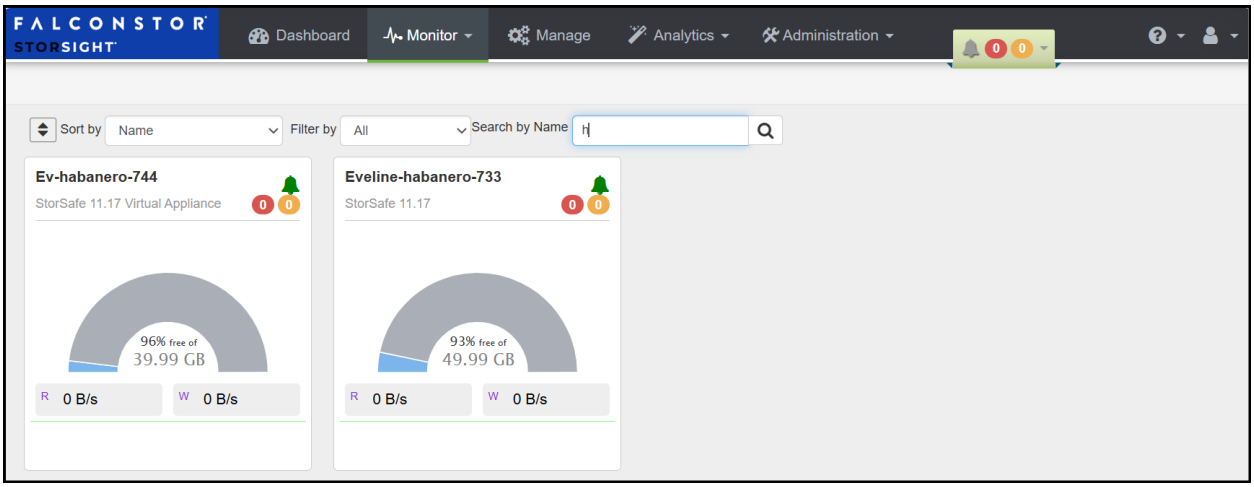

Server View page

This page displays same tabs that you can see in the *Manage* page, without any configuration change options.

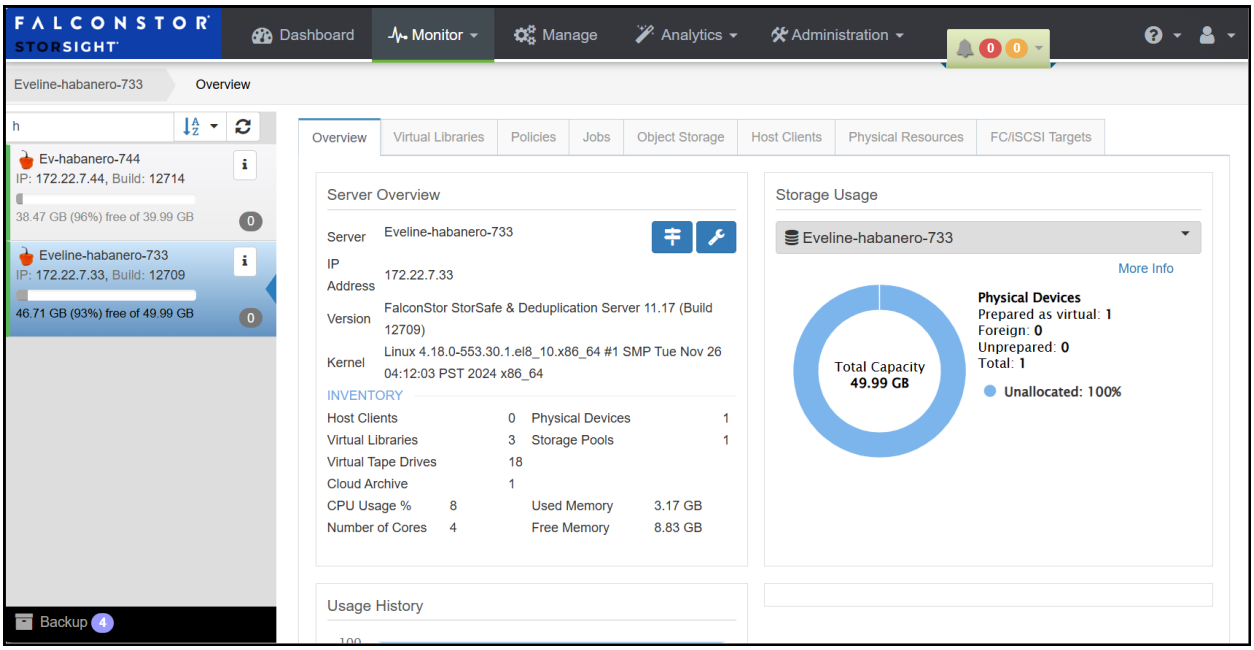

Administration

The *Administration* tab displays only for domain administrators to see information about servers, domain users, and AD/LDAP settings.

Servers

The *Servers* page displays information about all of your servers, including server name, IP address, installed version and build number, status, and capacity.

Users

The *Users* page lists all system users who have access to this StorSight server.

Local user accounts can be added manually. If Active Directory or LDAP authentication methods are used, a user account is created automatically the first time the user logs into the StorSight portal.

Add a user

1. Click the “+” icon.
2. Enter information for the new user:

   Active – Select to create an active account.

   Username – Enter the user account name.

   Password/Retype Password - Enter the user account password.

   Role - Select the user role:
   - *Admin* – An administrator who can create *User* or *Viewer* accounts and assign resources to other users.
   - *User* - A user with read-write access to resources that have been created by that user or assigned to that user by the administrator.
   - *Viewer* - A user with read-only access who can view configuration and reports; this account cannot make any configuration changes.

   First Name/Last Name - Specify the user first and last name.

   Email/Phone Number - Specify an email address and telephone number for the user.
3. Click *Save*.

Update a user

Accounts with a higher or equal role can change the properties of other accounts. For example, a domain administrator can change the password of administrator, user, and viewer accounts; users and viewers can only change their own password.

If a user originated from Active Directory or LDAP, you cannot update the password from StorSight; the password needs to be updated on the Active Directory or LDAP server.

Synchronize users

Domain administrators can use this option to synchronize users between the StorSight and Habanero servers. This is useful if a Habanero server was down when a new local, Active Directory, or LDAP user was added to the Habanero server.

To synchronize users, click the *Sync Users with Managed Servers* icon. 

Delete a user

To delete a user, highlight the user and click the *Delete* icon. 

The default domain administrator account cannot be deleted.

If the user being deleted had defined any smart rules, those rules will be deleted when the user is deleted.

AD/LDAP settings

There are three different methods that can be used for authenticating and authorizing StorSight domain users and administrators:

- Active Directory - Uses an Active Directory (AD) authentication server
- LDAP - Uses an LDAP authentication server
- Local users - Users added on the *Users* page

You can use only one authentication method. For example, if Active Directory is configured, all users need to log in with Active Directory user names and passwords.

When AD or LDAP are used, a user account is automatically created the first time each AD/LDAP user logs into the StorSight portal. If a user's password is changed on the AD/LDAP server, the user's password must be manually updated in the portal.

Both Active Directory/LDAP methods offer the option to use the local authentication method as backup if the Active Directory/LDAP server fails with an I/O or communication error. If this occurs, only users added locally will be able to log in. Typically, the domain administrator will log in. If needed, this administrator can create additional users to manage the system until the Active Directory/LDAP server is accessible again.

Active Directory and LDAP can only be used if strong passwords are not enabled.

Active Directory

Once Active Directory is configured, all users need to log in with Active Directory user names and passwords.

1. Click the *Active Directory* button.

   Once you enable Active Directory, you will see the following fields:

   Use the local authentication method as backup - Use local users if Active Directory authentication fails with an I/O or communication error. For all other errors returned from Active Directory, local authentication will not be provided.

   Active Directory Settings

   Primary Server - IP address or hostname of the primary Active Directory domain controller. If your Active Directory/LDAP server uses a self-signed certificate, make sure the primary server matches the one that has been specified in the certificate. For example, if host name is specified in the certificate, you must put the exact name.

   Secondary Server - (Optional) IP address or hostname of the secondary Active Directory domain controller.

   Port - Port number for Active Directory services (default 389) used by the primary and secondary servers. If you are using TLS, set the port to 636.

   Use TLS - Use Transport Layer Security, aka *SSL*, for additional security (required if you will be using a self-signed certificate).

   Domain - Domain name for the Active Directory server.

Active Directory Authorization

Select Active Directory group(s) that contain users for a given role (Administrator, User, or Viewer). These users will be granted the associated privileges. Note that StorSight only supports administrator/user groups at the parent level in the Active Directory tree.

- *Group Name for Admin Role* - Group used for Administrators; all Active Directory users who belong to this group will be granted Administrator privileges in StorSight for the customer domain.
- *Group Name for User Role* - Group used for Users; all Active Directory users who belong to this group will be granted User privileges in StorSight for the customer domain.
- *Group Name for Viewer Role* - Group used for Viewers; all Active Directory users who belong to this group will be granted Viewer privileges in StorSight for the customer domain.

Separate multiple groups with commas. If the same group is added in multiple fields, the role of the highest field will be used.

Active Directory users who do not belong to any of the specified groups will not be granted access to StorSight.

2. Click *Save*.

   Before the configuration can be saved, you must provide a username and password to test Active Directory authentication. This user must belong to a valid authorization group with Administrator.

LDAP

Once LDAP is configured, all users need to log in with LDAP user names and passwords.

1. Click the *LDAP* button.

   Once you enable LDAP, you will see the following fields:

   Use the local authentication method as backup - Use local users if LDAP authentication fails with an I/O or communication error. For all other errors returned from LDAP, local authentication will not be provided.

   ### LDAP Settings

   Primary Server - IP address or hostname of the primary LDAP server.

   Secondary Server - (Optional) IP address or hostname of the secondary LDAP server.

   Port - Port number for LDAP services (default 389) used by the primary and secondary servers. If you are using SSL, set the port to 636.

   Bind DN - LDAP bind authentication user distinguished name (not required if the LDAP server provides unauthenticated access).

   Bind Password - LDAP bind authentication password (not required if the LDAP server provides unauthenticated access).

Use SSL - Use SSL for additional security (required if you will be using a self-signed certificate).

LDAP Authentication

User Search Base DN - Starting point for the user search (e.g., ou=Users,dc=example,dc=com)

User Search Attribute - Attribute name that contains the username (e.g., uid).

LDAP Authorization

Group Search Base DN - Starting point for the group search (e.g., ou=Groups,dc=example,dc=com)

Group Search Attribute - Attribute name of the group in which the member 000user is specified (e.g., member).

Group Name Attribute: Attribute name of the group from which we can get the group name (e.g., cn).

For each, select LDAP group(s) that contain users for a given role (Administrator, User, or Viewer). These users will be granted the associated privileges. Note that StorSight only supports administrator/user groups at the parent level in the LDAP tree.

- *Group Name for Admin Role* - Group used for Administrators; all Active Directory users who belong to this group will be granted Administrator privileges in StorSight.
- *Group Name for User Role* - Group used for Users; all Active Directory users who belong to this group will be granted User privileges in StorSight.
- *Group Name for Viewer Role* - Group used for Viewers; all Active Directory users who belong to this group will be granted Viewer privileges in StorSight.

Separate multiple groups with commas. If the same group is added in multiple fields, the role of the highest field will be used.

LDAP users who do not belong to any of the specified groups will not be granted access to StorSight.

2. Click *Save*.

   Before the configuration can be saved, you must provide a username and password to test LDAP authentication. This user must belong to a valid authorization group with Administrator rights.

Local

Local users are defined on the Users page. If you enable *Local*, only locally defined users will be able to log in; Active Directory and LDAP users will be disabled.

1. Select *Administration* from the menu bar and then click *AD/LDAP Settings*.
2. Click the *Local* button.
3. Click *Save*.