DSI® Optimized Backup & Deduplication for VTL and NAS User Guide

Version 8.20

User Guide content may change between major product versions in order to reflect product updates released via patches. In the guide and its table of contents, the heading for content changed within the past six months will be followed by “(updated *Month Year)*”.

The document code at the bottom of the page includes the guide publication date and the associated software build number, in the format *date.build*.

Dynamic Solutions International, LLC

1 Inverness Drive East

Englewood, CO 80112
Phone: 303-754-2000

Technical Support: 1-800-332-9200

Fax: 303-754-2009

Web site: www.dynamicsolutions.com

FalconStor Software reserves the right to make changes in the information contained in this publication without prior notice. The reader should in all cases consult FalconStor Software to determine whether any such changes have been made.

This product is protected by United States Patents Nos. 7,093,127 B2; 6,715,098; 7,058,788 B2; 7,330,960 B2; 7,055,008 B2; 7,469,337 and additional patents pending.

082015.9309

Contents

Introduction

VTLVA Configuration

Getting Started

Deduplication Configuration

VTL Tape Configuration

NAS Configuration

Console

Multi-Node Groups

Tape Libraries, Tape Drives, and Tapes

Deduplication

Encryption

VTL Server Failover

Redundant Node Failover

Data Replication

Automated Tape Caching

Fibre Channel Configuration

iSCSI Configuration

IBM 3494 Configuration

IBM System i Configuration

Hosted Backup

NDMP Backup Support

ACSLS and Library Station Configuration

Server Maintenance

Reports

Email Alerts

Command Line

SNMP Integration

Troubleshooting

Appendix

Index 725

Introduction

The DSI Optimized Backup & Deduplication solution combines industry-leading Virtual Tape Library (VTL) technology and NAS functionality in a unified, simple-to-use product, providing disk-based backup and deduplication at the fastest available speeds.

The solution offers high-speed backup/restore, global data deduplication, enterprise-wide replication, and tape integration, without requiring changes to the existing environment. NAS functionality provides block-level deduplication for data on network-based file shares using the standard NAS protocols Common Internet File System (CIFS) and Network File System (NFS).

Data can be backed up from third-party tape backup software, third-party disk backup software, database backup utilities, archiving applications, and any other mechanism for delivering data to a network share, such as DSI FileSafe™.

With its integrated deduplication, the solution removes redundant copies of data, thereby reducing capacity requirements and minimizing replication time.

Since VTL technology uses disk to back up data, it eliminates the media and mechanical errors that can occur with physical tapes and drives. And, because VTL can emulate more tape drives than your physical tape library really has, more backup streams can run simultaneously, enabling organizations to easily complete their backups within the allotted backup window.

Because you may already have physical tapes that you would like to protect, you can import data from physical tapes into your virtual tape system. If you ever need to recover files from a physical tape, you can access those tapes for immediate recovery.

For additional data protection, the data on virtual tapes can be exported to physical tapes for long-term data archiving. Data can also be copied to physical tapes using your backup application’s copy function.

Optimized Backup & Deduplication components

There are several components:

- Optimized Backup & Deduplication server(s) - Appliances providing VTL, deduplication, and/or NAS functionality.
- VTL and deduplication storage - Physical devices used for tapes, database, and deduplicated data.
- DSI Management Console - The graphical administration tool, installed on a separate workstation, lets you configure and manage VTL, NAS, SIR, and VTL-S servers.

- Clients - The backup application servers that use VTL via Fibre Channel, iSCSI, CIFS, or NFS. If the Hosted Backup feature is available, you can install a certified backup application directly on the VTL appliance, eliminating the need for a dedicated backup application server.
- (Optional) Deduplication appliance - The appliance that manages deduplication.

Additional resources

VTL is designed to work with FalconStor OpenStorage Option (FSOST). For more information, refer to the *FalconStor OpenStorage Option User Guide*.

You can download software builds, patches, and other documentation related to your DSI product from the DSI Customer Support Portal at *support.dynamicsolutions.com* (account required). Click the *View Builds, Patches, & Documentation* link in the *GA Releases* area to complete a simple search form and display available downloads.

Note that the product release notes and patch descriptions can include information that may not appear in the user guide. Be sure to review all available documents.

Optimized Backup & Deduplication supports a broad range of server hardware, tape devices, and backup software. The Certification Matrix at *www.dynamicsolutions.com* identifies all certified hardware and softwares.

If you need technical support, create a support ticket on the DSI Customer Support portal.

VTLVA Configuration

DSI VTL virtual appliance software is available from the DSI Customer Support Portal. This chapter explains how to install DSI VTL on a VMware virtual appliance.

Minimum ESXi server hardware requirements

A VTL virtual appliance with a 1 TB repository (base system) requires the following:

Component	Configuration
CPU	2 core
Memory	10 GB memory (2 GB for ESXi server, 8 GB for VTLVA)
Disk	At least 2.5 TB free space
Network	Two physical network adapters

The pre-defined OVF file for VTLVA provides the following virtual machine configuration:

Component	Virtual Machine Configuration
CPU	Two virtual processors
Memory	8 GB
Disk	Configuration repository and DB: 1 x 20 GB Backup cache: 1 x 300 GB Deduplication repository: 1 x 1 TB Deduplication index: 1 x 50 GB Boot LUN: 1 x 70 GB
Network	One 1GbE port

The VTL virtual appliance base system includes 1 TB of repository capacity and is expandable to 10 TB with additional capacity license key(s). Refer to ’Expand repository capacity’.

Before you install

Verify requirements

Run the VMware CPU identification utility to verify the processor capabilities, Virtual Technology (VT) support, and BIOS settings, and ensure that the following requirements have been met:

- System memory on the ESXi server is at least 10 GB.
- The CPU and BIOS of the ESXi server supports 64-bit operating systems.
- *Virtual Technology* (VT) is enabled in the system BIOS. To set it, go to the motherboard's BIOS configuration, go to the CPU's advanced settings, and select *Enable VT*. After enabling VT, you must power off (not reset) the ESXi server completely, power it on, and reboot to ESXi.
- Install the appropriate version of VMware vSphere Client.
- VMware ESXi 5.x requires the corresponding VMware vSphere Client version.
- Use the VMware vSphere Client to verify the VMware tools are up to date; update if needed.

Verify ESXi server time

Use the VMware vSphere Client to verify that the ESXi server time is correct.

1. On the VMware vSphere Client, select the ESXi server.
2. Select the *Configuration* tab, select *Software* --> *Time Configuration* menu.
3. Check that the date and time of the ESXi server is correct or click *Properties* to change it.

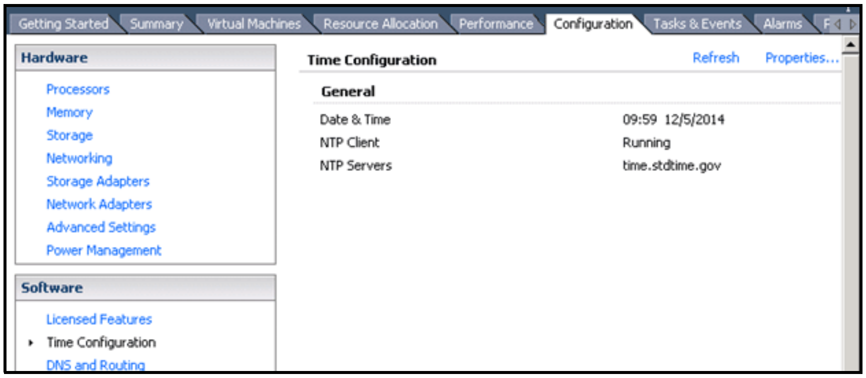

Time Configuration

General

Date & Time	09:59 12/5/2014
NTP Client	Running
NTP Servers	time.stdtime.gov

Installation and configuration

Import VTLVA from the VMware Infrastructure Client

1. Launch the VMware vSphere Client and connect to the ESXi server with root privileges.
2. Navigate to the directory containing the zip file `FalconStor-VTLVA-v8.n-bxxxx-1TB.zip`

   Note that the file name includes the VTLVA version (represented with an `n`) and the build number (represented as `xxxx`).
3. Unzip the file:

   `unzip FalconStor-VTLVA-v8.n-bxxxx-1TB.zip`

   The files will be automatically extracted to the following directory: `FalconStor-VTLVA-v8.n-bxxxx-1TB`*.*
4. Select *File --> Deploy OVF Template* to run the deployment wizard.
5. For *Import Location*, click *Browse* and browse to the folder into which the zip file was extracted*.*

   Expand the folder and select *FalconStor-VTLVA.ovf*.

   Click *Next* to continue.
6. If desired, change the default appliance name, *DSI-VTLVA*, which is displayed in the *Name* and *Location* fields.

   This is the virtual machine name displayed on the vSphere client.

   This change will not be applied to the actual appliance name (the host name of VTLVA that is displayed by the operating system when using the Linux `hostname` command).

   > **Note:** Virtual appliance folders are not visible when connected directly to the host, so you will not be able to select the appliance location.

7. Choose a *Datastore* to indicate where to store the files for the virtual appliance being imported.

   The Datastore must contain at least 2,073,060 MB of space for the VTLVA system import.
8. For each network specified on the *Network Mapping* screen, right-click and set up network mapping for the VTLVA virtual Ethernet adapter.

   The *Destination Networks* column lists all destination networks in the vSphere inventory that can be used for virtual machines in the OVF template.
9. On the *Ready to Complete* screen, review all settings and click *Finish* to start the virtual appliance import process.

   The import status window displays the completion percentage; import will complete in several minutes.

Check VTLVA resource reservations

It is important to make sure the VTLVA has enough available resources, especially in a shared architecture with other virtual machines. Resource allocation is set using the VMware vSphere Client. Do the following to verify the resource reservations:

1. Launch the VMware vSphere Client and connect to the ESXi server with an account that has root privileges.

2. Right-click the installed DSI-VTLVA and select *Edit Settings*.

   On the *VTLVA Virtual Machine Properties* screen, select the *Resources* tab and then select *CPU* in the *Settings* list.

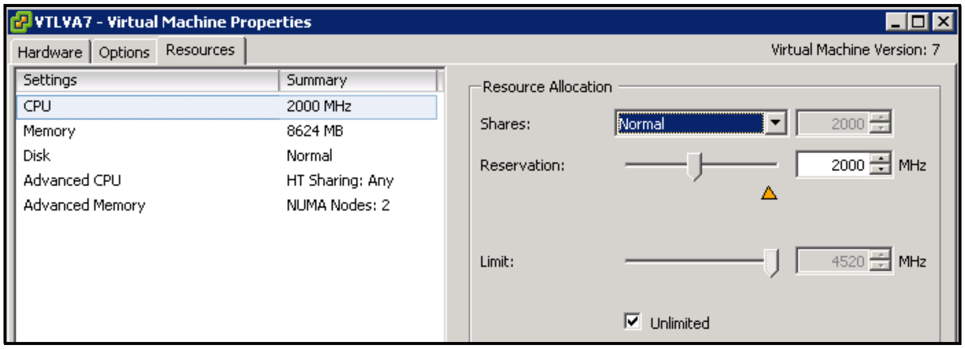

VTLVA7 - Virtual Machine Properties

Settings	Summary
CPU	2000 MHz
Memory	8624 MB
Disk	Normal
Advanced CPU	HT Sharing: Any
Advanced Memory	NUMA Nodes: 2

Resource Allocation

Shares: Normal 2000

Reservation: 2000 MHz

Limit: 4520 MHz

☑ Unlimited

   In the *Resource Allocation* screen, the required *Reservation* setting (2000 MHz) is set by default.

3. Select *Memory* in the *Settings* list.

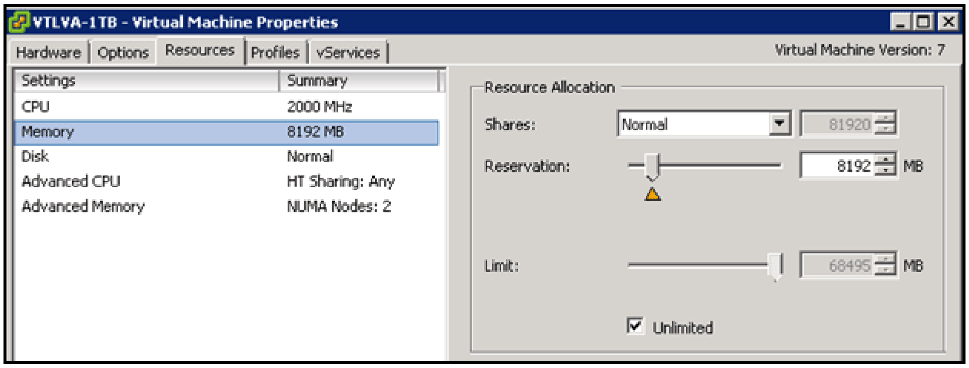

VTLVA-1TB - Virtual Machine Properties

Settings	Summary
CPU	2000 MHz
Memory	8192 MB
Disk	Normal
Advanced CPU	HT Sharing: Any
Advanced Memory	NUMA Nodes: 2

Resource Allocation

Shares: Normal 81920

Reservation: 8192 MB

Limit: 68495 MB

☑ Unlimited

   In the *Resource Allocation* screen, the required *Reservation* setting (8192 MB) is set by default.

Check the virtual network adapter setting

The DSI VTL virtual appliance is pre-configured with one virtual network adapter. Do the following to verify the network connection setting:

1. In the VMware vSphere Client, right-click the installed DSI-VTLVA and then click *Edit Settings*.

2. On the *VTLVA Virtual Machine Properties* screen, select the *Hardware* tab and then select the Network Adapter in the *Hardware* list.

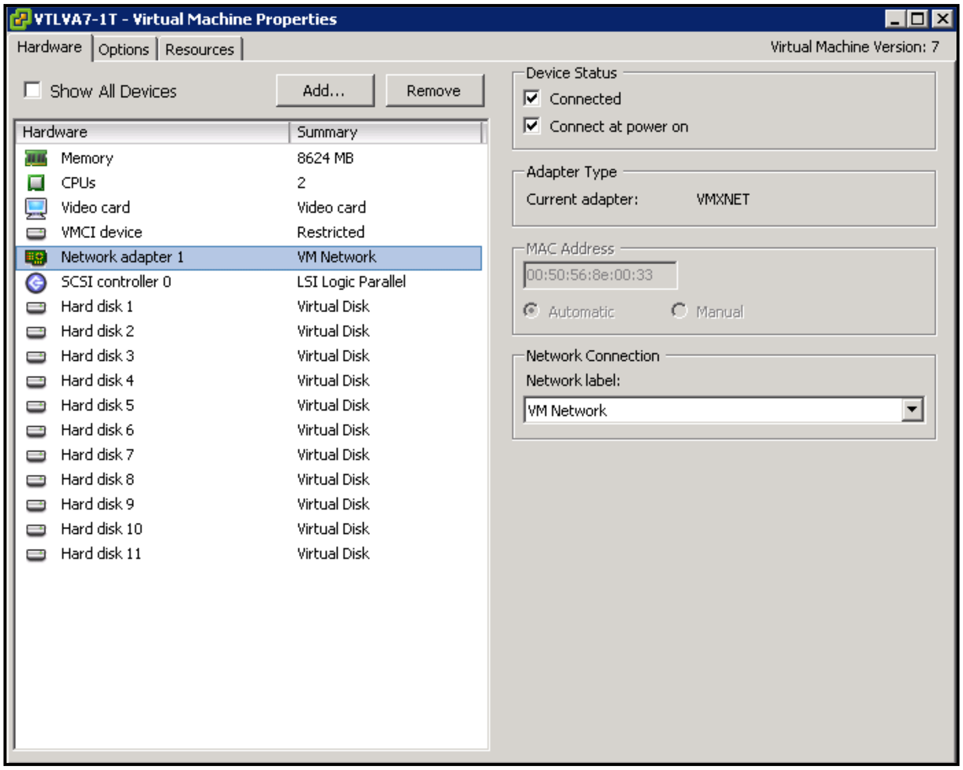

Virtual Machine Version: 7

☐ Show All Devices

Hardware	Summary
Memory	8624 MB
CPUs	2
Video card	Video card
VMCI device	Restricted
Network adapter 1	VM Network
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
Hard disk 2	Virtual Disk
Hard disk 3	Virtual Disk
Hard disk 4	Virtual Disk
Hard disk 5	Virtual Disk
Hard disk 6	Virtual Disk
Hard disk 7	Virtual Disk
Hard disk 8	Virtual Disk
Hard disk 9	Virtual Disk
Hard disk 10	Virtual Disk
Hard disk 11	Virtual Disk

Device Status
☑ Connected
☑ Connect at power on

Adapter Type
Current adapter: VMXNET

MAC Address
00:50:56:8e:00:33
● Automatic ○ Manual

Network Connection
Network label:
VM Network

The default *Network Label* for the *Network Connection* setting is *VM Network*. If this is not correct for your environment, select another network from the drop-down list, which lists all networks configured for virtual machine use on the host.

Configure the ESXi server to use Fibre Channel with a VTL virtual appliance

Do the following on the ESXi server to configure Fibre Channel connectivity to VTL clients:

1. Configure passthrough devices on an ESXi host:
   - Select an ESXi host from the Inventory of VMware vSphere Client.
   - On the *Configuration* tab, click *Advanced Settings*. The passthrough Configuration page lists all available pass-through devices.
   - Click *Edit*.
   - Select the Fibre Channel devices and click *OK*.
   - When the devices are selected, they are marked with an orange icon. Reboot for the change to take effect. After rebooting, the devices marked with a green icon are enabled.

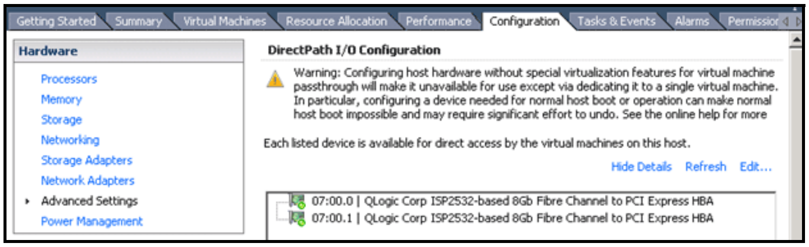

DirectPath I/O Configuration

Warning: Configuring host hardware without special virtualization features for virtual machine passthrough will make it unavailable for use except via dedicating it to a single virtual machine. In particular, configuring a device needed for normal host boot or operation can make normal host boot impossible and may require significant effort to undo. See the online help for more

Each listed device is available for direct access by the virtual machines on this host.

07:00.0 | QLogic Corp ISP2532-based 8Gb Fibre Channel to PCI Express HBA
07:00.1 | QLogic Corp ISP2532-based 8Gb Fibre Channel to PCI Express HBA

   For more information, refer to Configuring VMDirectPath I/O pass-through devices on a VMware ESX or VMware ESXi host (1010789).

2. Configure a PCI device on a virtual machine:
   - From the Inventory in vSphere Client, right-click the virtual machine and click *Edit Settings*.
   - Click the *Hardware* tab.
   - Click *Add*.
   - Select *PCI Device* and click *Next*.
   - Select the Fibre Channel devices and click *Next.*
   - Follow the wizard to finish the configuration.

Start the virtual appliance

Do the following to start the virtual appliance for the first time:

1. Launch the *VMware Infrastructure Client* and connect to the ESXi server with an account that has root privileges.

2. Power on the virtual machine: select *VM --> Power On*. (Alternatively, you can right-click the virtual machine and select *Power On* or click the *Power On* button).

3. Click the *Console* tab and log into the VTL server with user name *root* and password (case-sensitive) *IPStor101*.

After you log in, the *DSI Virtual Appliance Setup* utility will run automatically.

Configure the virtual appliance using the setup utility

The *DSI Virtual Appliance Setup* utility lets you configure a variety of settings provided that you are using the pre-defined configuration created by the OVF file for VTLVA. If you have modified the configuration (i.e., added disks), you will not see the step to select a deduplication product in the setup utility. In this case, you must configure the appropriate deduplication product through the DSI Management Console after completing the setup utility.

System keyboard

Select the appropriate keyboard for the system.

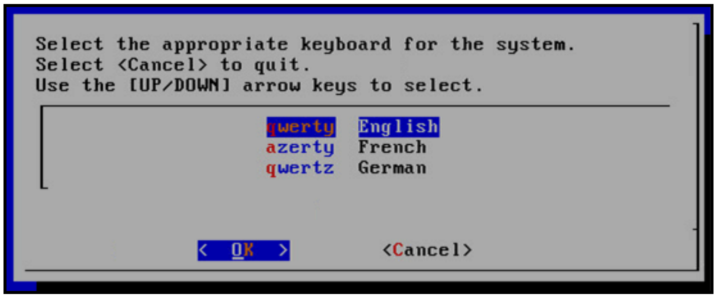

```
Select the appropriate keyboard for the system.
Select <Cancel> to quit.
Use the [UP/DOWN] arrow keys to select.

          qwerty  English
          azerty  French
          qwertz  German
```

Time zone

The default system clock is set to *UTC.*

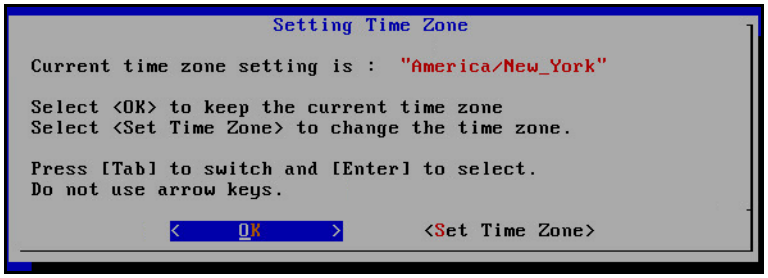

```
                    Setting Time Zone

Current time zone setting is :  "America/New_York"

Select <OK> to keep the current time zone
Select <Set Time Zone> to change the time zone.

Press [Tab] to switch and [Enter] to select.
Do not use arrow keys.
```

Network settings

You can change the IP address, subnet mask, and gateway for eth0. The default settings for FS-VTLVA-1T are:

- eth0 - 10.0.0.2
- subnet mask - 255.255.255.0
- gateway - 10.0.0.254

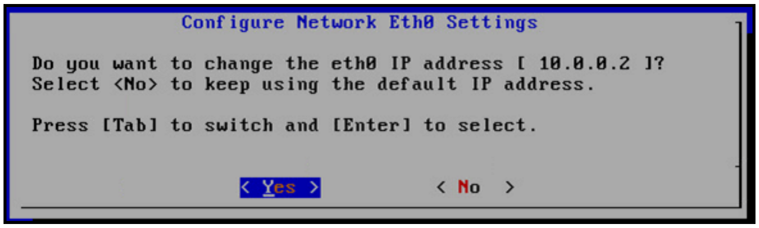

```
              Configure Network Eth0 Settings

Do you want to change the eth0 IP address [ 10.0.0.2 ]?
Select <No> to keep using the default IP address.

Press [Tab] to switch and [Enter] to select.
```

Domain name

Set the domain name.

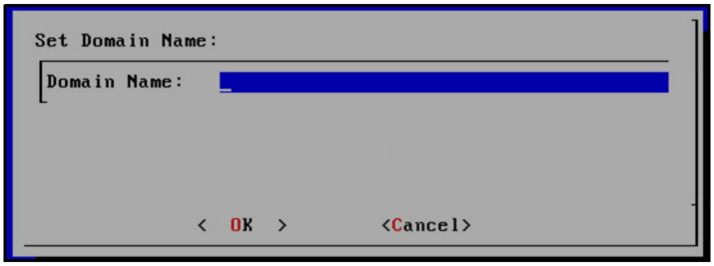

```
Set Domain Name:
Domain Name: _
```

DNS servers

You can add up to two DNS servers.

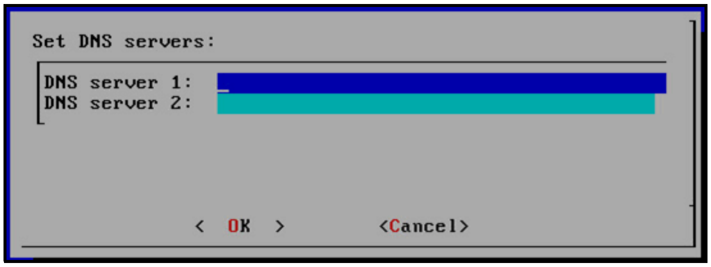

```
Set DNS servers:
DNS server 1: _
DNS server 2:
```

Network time protocol servers

You can add up to four NTP servers.

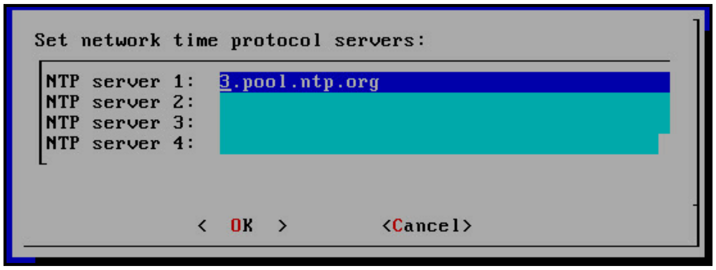

```
Set network time protocol servers:
NTP server 1: 3.pool.ntp.org
NTP server 2:
NTP server 3:
NTP server 4:
```

Hostname

You can change the default hostname.

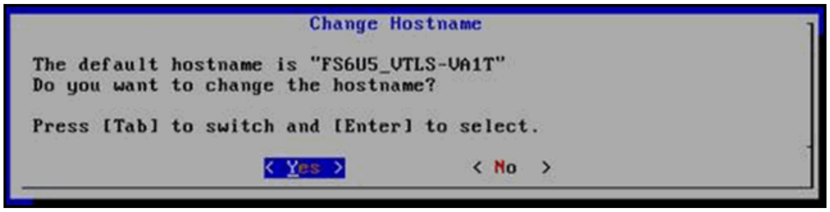

```
                    Change Hostname

The default hostname is "FS6U5_VTLS-VA1T"
Do you want to change the hostname?

Press [Tab] to switch and [Enter] to select.
```

Deduplication product

Select the deduplication product you are configuring.

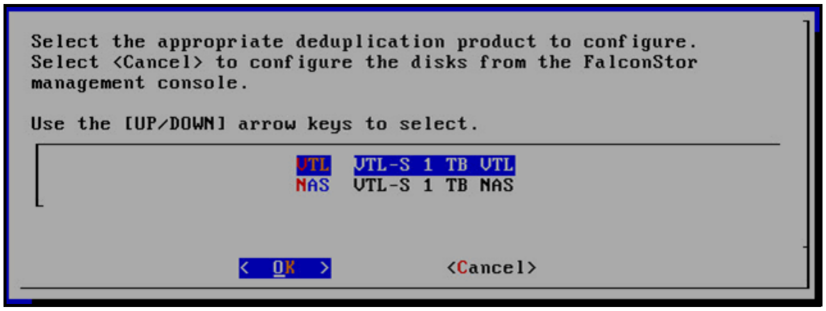

```
Select the appropriate deduplication product to configure.
Select <Cancel> to configure the disks from the FalconStor
management console.

Use the [UP/DOWN] arrow keys to select.

                    VTL  VTL-S 1 TB VTL
                    NAS  VTL-S 1 TB NAS
```

Configure VTLVA to use Fibre Channel

Do the following to configure Fibre Channel connectivity to VTL clients:

1. If your HBA was added after the first time the virtual appliance was started, type the following at the command line and select the type of HBA you are using:

   ```
   # vtl configtgt
   ```

   You do not need this step if the HBA was installed before the virtual appliance was started.

2. From the DSI Management Console, enable FC Target Mode, create an FC client, and assign a tape library to the client.

   These steps are performed in the configuration wizard and are described in the "VTL Tape Configuration" chapter.

Expand repository capacity

This section describes how to expand the repository for the VTL virtual appliance. When you are done, additional resources will be configured and are available for use.

The VTL virtual appliance includes 1 TB of repository capacity and is expandable to 10 TB by purchasing a license for additional VTL index cache/deduplication repository capacity in 1 TB increments.

Each 1 TB increase in capacity requires 50 GB index and folder disk.

Expanding the VTLVA system requires the following steps:

- Add RAM to the ESXi server.
- Add license keycodes for the expansion capacity.
- If necessary, increase memory allocated for the virtual machine in the VMware vSphere Client.
- Add disks to the ESXi server in the VMware vSphere Client.
- Add data disks to VTL via the DSI Management Console. Refer to ’Add data disks’ for more information.

Add license keycode(s) for the expansion capacity

Add the keycode(s) for the additional capacity via the DSI Management Console. Whether your expansion capacity license includes one or several keycodes, you can add all keycodes at this time.

Increase memory allocated to the VTL virtual appliance

Depending upon your expansion, you may need to increase the memory allocated to the VTL virtual appliance. The following memory is required:

Total Capacity	Total RAM
2 TB	16 GB
3 TB	16 GB
4 TB	16 GB
5 TB	16 GB
6 TB	32 GB
7 TB	32 GB
8 TB	32 GB
9 TB	32 GB
10 TB	32 GB

To increase allocated memory:

1. Power off the VTL virtual appliance.
2. Launch the VMware vSphere Client and connect to the ESXi server as a user with root privileges.
3. Right-click the VTL virtual appliance and select *Edit Settings*.
4. Select the appropriate amount of memory.

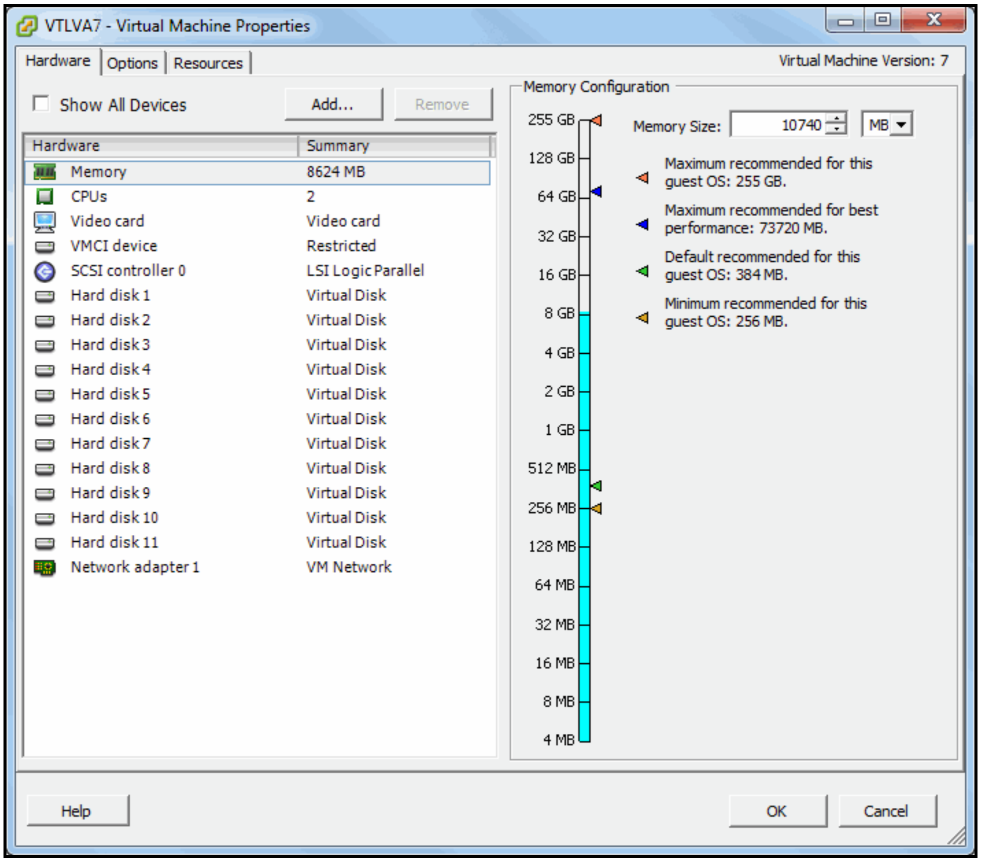

VTLVA7 - Virtual Machine Properties

Virtual Machine Version: 7

□ Show All Devices

Hardware	Summary
Memory	8624 MB
CPUs	2
Video card	Video card
VMCI device	Restricted
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
Hard disk 2	Virtual Disk
Hard disk 3	Virtual Disk
Hard disk 4	Virtual Disk
Hard disk 5	Virtual Disk
Hard disk 6	Virtual Disk
Hard disk 7	Virtual Disk
Hard disk 8	Virtual Disk
Hard disk 9	Virtual Disk
Hard disk 10	Virtual Disk
Hard disk 11	Virtual Disk
Network adapter 1	VM Network

Memory Configuration

Memory Size: 10740 MB

Maximum recommended for this guest OS: 255 GB.

Maximum recommended for best performance: 73720 MB.

Default recommended for this guest OS: 384 MB.

Minimum recommended for this guest OS: 256 MB.

Make sure that the VTLVA has enough memory resource reservations. Refer to 'Check VTLVA resource reservations' for more information.

Next, add disks to the ESXi server.

Add disks to the ESXi server

Each 1 TB increase in capacity requires:

- Deduplication repository - 1 x 1 TB virtual disk for deduplicated data
- Deduplication index - 1 x 50 GB virtual disk
- Optionally, backup cache - 1 x 300 GB virtual disk

Add the disks needed for 1 TB of licensed expansion. Remember that the amount of disk resources you create must be consistent with the capacity key codes you have added.

1. Select the *Add* button to launch the *Add Hardware* wizard.

2. For *Device Type*, select *Hard Disk*.

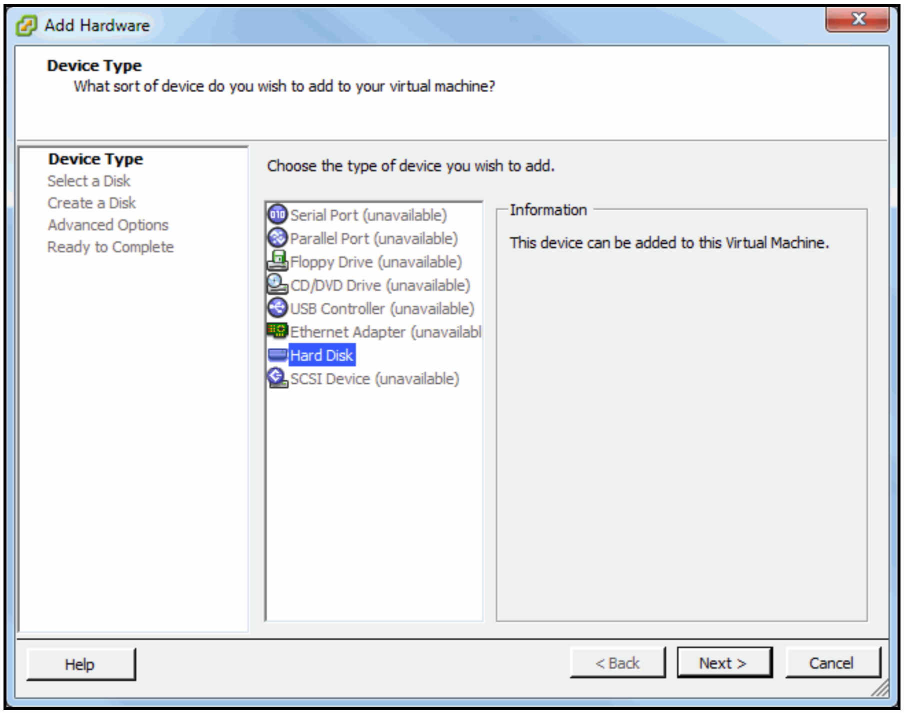

Device Type
What sort of device do you wish to add to your virtual machine?

Choose the type of device you wish to add.

Information
This device can be added to this Virtual Machine.

Add Hardware
Device Type
Select a Disk
Create a Disk
Advanced Options
Ready to Complete
Serial Port (unavailable)
Parallel Port (unavailable)
Floppy Drive (unavailable)
CD/DVD Drive (unavailable)
USB Controller (unavailable)
Ethernet Adapter (unavailabl
Hard Disk
SCSI Device (unavailable)
Help
< Back
Next >
Cancel

3. On the *Select a Disk* screen, select *Create a new virtual disk*.

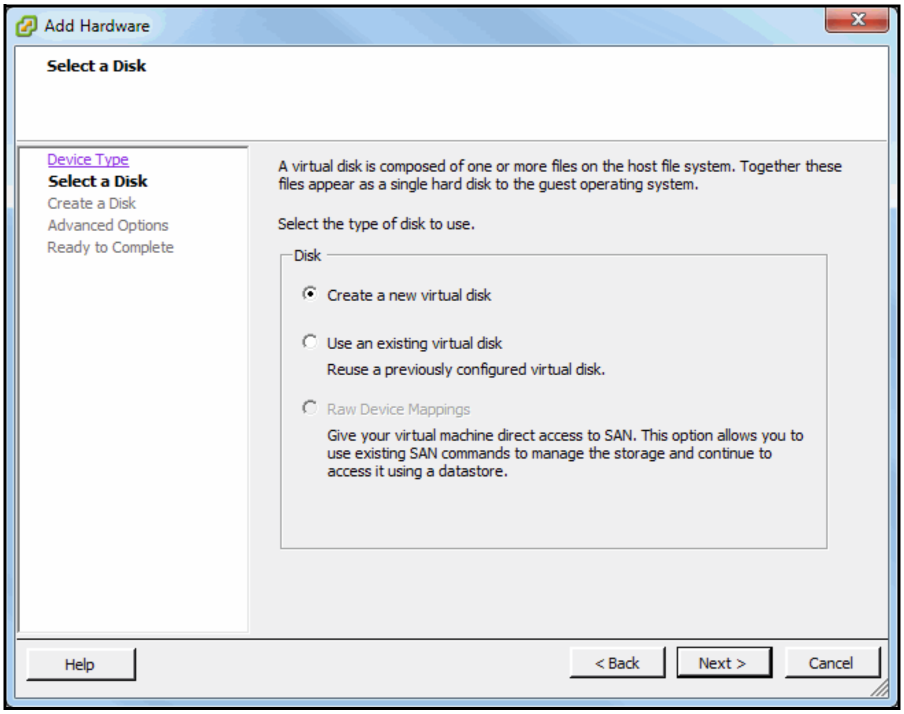

Select a Disk

A virtual disk is composed of one or more files on the host file system. Together these files appear as a single hard disk to the guest operating system.

Select the type of disk to use.

Disk

● Create a new virtual disk

○ Use an existing virtual disk
Reuse a previously configured virtual disk.

○ Raw Device Mappings
Give your virtual machine direct access to SAN. This option allows you to use existing SAN commands to manage the storage and continue to access it using a datastore.

4. On the *Create a Disk* screen, enter the capacity as defined in the expansion table for *Disk Size*.

   For *Location*, select *Store with the virtual machine*; click *Next*.

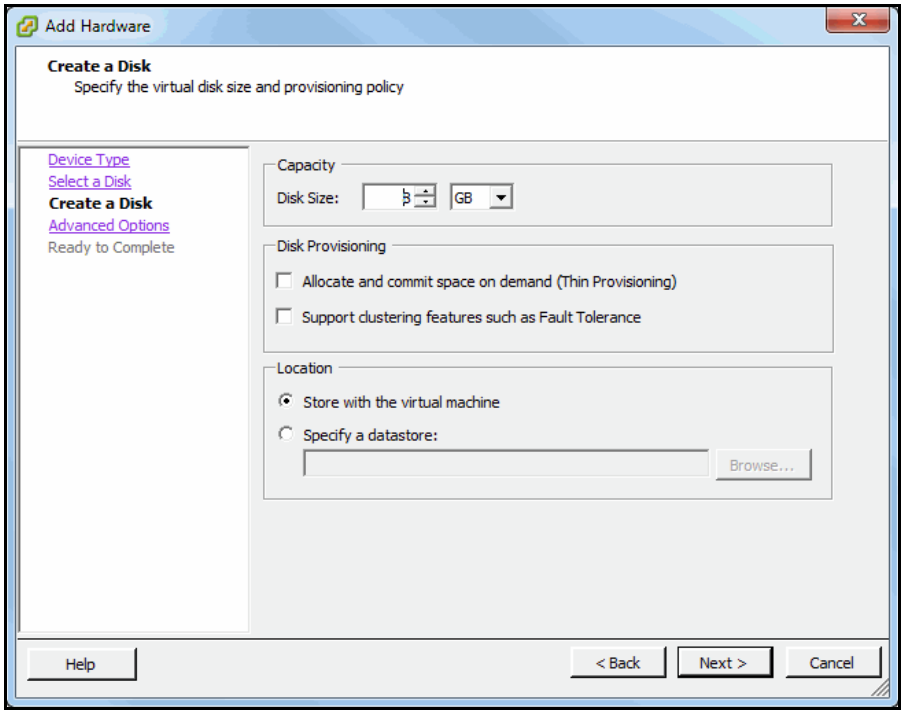

Create a Disk
Specify the virtual disk size and provisioning policy

Capacity
Disk Size: 3 GB

Disk Provisioning
□ Allocate and commit space on demand (Thin Provisioning)
□ Support clustering features such as Fault Tolerance

Location
● Store with the virtual machine
○ Specify a datastore:

5. On the last screen, review the options you selected and click *Finish.*

A new disk with the specified capacity will be displayed in the *Hardware* tab.

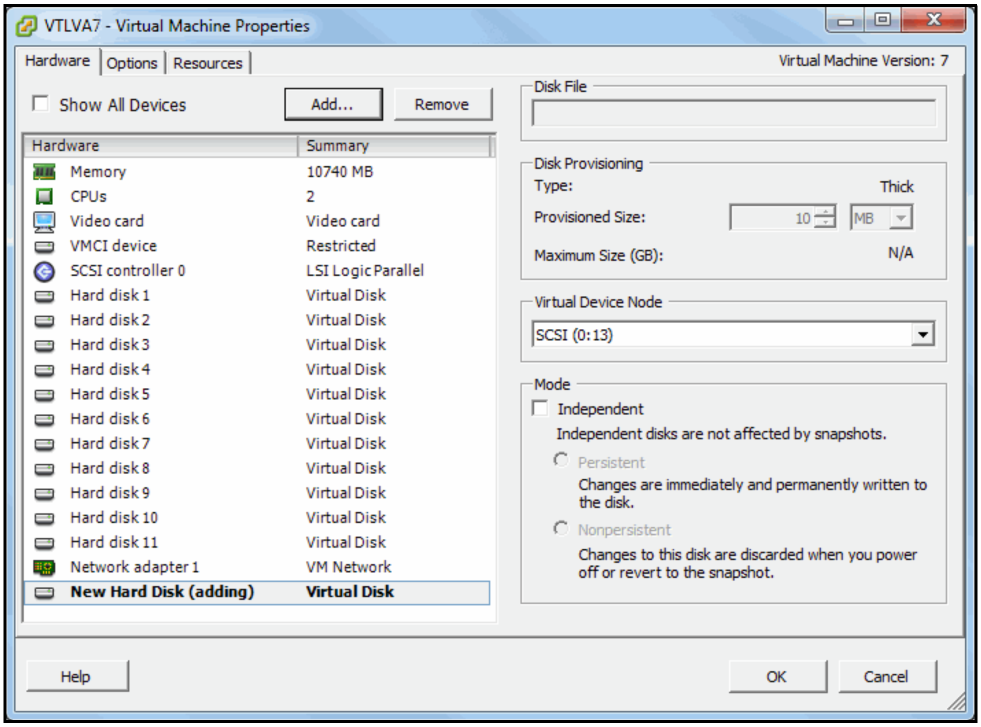

VTLVA7 - Virtual Machine Properties

Hardware | Options | Resources

Virtual Machine Version: 7

□ Show All Devices

Add... | Remove

Hardware	Summary
Memory	10740 MB
CPUs	2
Video card	Video card
VMCI device	Restricted
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
Hard disk 2	Virtual Disk
Hard disk 3	Virtual Disk
Hard disk 4	Virtual Disk
Hard disk 5	Virtual Disk
Hard disk 6	Virtual Disk
Hard disk 7	Virtual Disk
Hard disk 8	Virtual Disk
Hard disk 9	Virtual Disk
Hard disk 10	Virtual Disk
Hard disk 11	Virtual Disk
Network adapter 1	VM Network
New Hard Disk (adding)	**Virtual Disk**

Disk File

Disk Provisioning

Type: Thick

Provisioned Size: 10 MB

Maximum Size (GB): N/A

Virtual Device Node

SCSI (0:13)

Mode

□ Independent

Independent disks are not affected by snapshots.

○ Persistent

Changes are immediately and permanently written to the disk.

○ Nonpersistent

Changes to this disk are discarded when you power off or revert to the snapshot.

Help | OK | Cancel

6. Click *OK* when you are done.

   All new hard disks will be listed.

Getting Started

Each VTL server can have one of the following roles:

- VTL and/or NAS - Backup server with VTL (tape) interface and/or NAS interface
- SIR - Deduplication repository
- VTL-S - Backup server with VTL (tape) interface and/or NAS interface, combined with a deduplication repository

Once you configure your virtual appliance or connect your physical appliance(s) to your storage network (as described in the *Hardware QuickStart Guide*), you should launch the console and configure each of them using the configuration wizard.

Run the DSI Management Console

The DSI Management Console is the graphical administration tool that enables you to manage your VTL, NAS, SIR, and VTL-S servers. The computer that runs the console needs connectivity to the network segment where VTL is running, because it communicates directly with the server and backup application servers.

Install and launch the console on an administrative computer

To install the DSI Management Console software, go to the DSI Customer Support Portal at *support.dynamicsolutions.com* (account required) and download it from the *GA Releases* area.

You can install the DSI Management Console onto any machine, as long as that machine has a Graphical User Interface. Note that if you are installing the console on a Windows machine, you must be a Power User or Administrator.

To launch the console after installation, select *Start* --> *Programs* --> *DSI*. Select the version of the VTL console that corresponds to your installation and then click *VTL Console*.

Launch the Java Web Start console

To launch the Java Web Start version of the console, open a browser from any machine and enter the IP address of the server **(**for example: http://10.0.0.2:81). Web Start allows you to download and run the DSI Management Console from the web. You should have the right version of Java Runtime Environment (JRE) on the machine. If the required version is not present, you will be prompted to install it before downloading it from the server.

Note: In order to launch the console using the IP address of the VTL server, you must first download `US_export_policy.jar` and `local_policy.jar` from http://www.oracle.com/technetwork/java/javase/downloads/jce-7-download-432124.html and replace the versions in your java runtime directory:

- Unix - <java-runtime-home>/lib/security
- Win32 - <java-runtime-home>\lib\security

If you are having problems connecting to a server from the console, you may need to go to the Java Control Panel (from the Windows Control Panel) and de-select the option "Keep temporary files on my computer".

In some environments, in order to allow the console to be launched, you will need to add it to the security exception site list in the Java Control Panel (*Control Panel* --> *Java* --> *Security* tab) with port 81 (i.e., http://10.1.1.1:81).

Connect to your server

If your server already appears in the tree, right-click it and select *Connect*. For a multi-node group, right-click the group and select *Connect* to connect to all of the servers in the group.

If your server does not appear in the tree, do the following to add it:

1. Right-click the *Servers* object and select *Add.*

   If you are running on a Windows machine, you can right-click the *Servers* object and select *Discover* to detect VTL servers in a range of IP addresses. You should then specify the subnet range of your VTL server and wait for the server hostname to appear in the navigation tree. For DSI appliances, the default hostname has the format DSI-*xxxxx*, where *xxxxx* is a unique number for your appliance, which is displayed on a label on your appliance. When the hostname appears in the navigation tree, right-click it and select *Connect*.

2. Type the server name or address and enter a valid user name and password (both are case sensitive) to log in.

   If you purchased an appliance from DSI, you can log in with *root* as the User Name and *IPStor101* as the password*.* Note that the user name and password are case sensitive.

   Once you are connected to a server, the server icon will change to show that you are connected.

   After connecting to the server, the configuration wizard launches.

Configure servers

Servers should be configured in the following order:

- SIR deduplication server - Refer to "Deduplication Configuration" for more information.
- VTL/VTL-S server - Refer to "VTL Tape Configuration" for more information.
- NAS - Refer to "NAS Configuration" for more information.

After you have configured your server using the configuration wizard, you can do the following:

- Install and configure *FalconStor OpenStorage Option* (FSOST) - A software interface between VTL and your Symantec™ NetBackup™ or Backup Exec™ servers that allows a NetBackup server to perform high-speed backup and recovery to intelligent disk devices. You can download FSOST and its user guide from the FalconStor Customer Support Portal at *support.falconstor.com* (account required).
- Configure *NIC Port Bonding* - A load balancing/path redundancy feature that enables you to load balance network traffic across two or more network connections. Refer to 'NIC Port Bonding' for more information.

Deduplication Configuration

This section describes how to configure a deduplication server.

Prepare a deduplication server via the configuration wizard

If your deduplication server has not been configured yet, the configuration wizard will be launched when you connect to it.

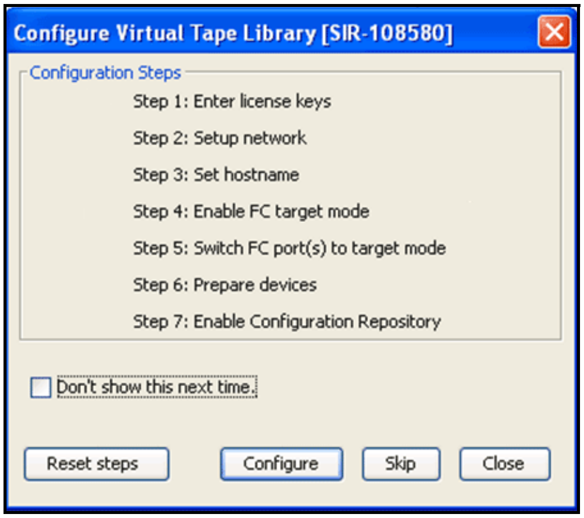

Step 1: Enter license keys

Step 2: Setup network

Step 3: Set hostname

Step 4: Enable FC target mode

Step 5: Switch FC port(s) to target mode

Step 6: Prepare devices

Step 7: Enable Configuration Repository

□ Don't show this next time.

Select *Configure* to begin the steps in the wizard.

> **Note:** The actual wizard you see will depend upon the options you have licensed and activated.

Enter license keys

Be sure to enter keycodes for any options you have purchased. Each option requires that a keycode be entered before the option can be configured and used. Refer to 'Add and register licenses' for more information.

> **Configuration note:** After completing the configuration wizard, if you need to add license keys, you can right-click your server and select *License*.

Set up network

(Optional for all servers) This step allows you to set/change your network configuration.

1. Enter information about your network configuration.

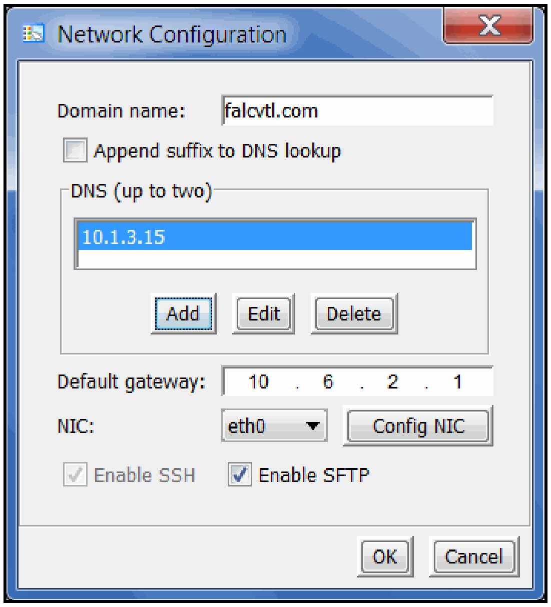

Network Configuration

Domain name: falcvtl.com

□ Append suffix to DNS lookup

DNS (up to two): 10.1.3.15

Default gateway: 10 . 6 . 2 . 1

NIC: eth0

☑ Enable SSH ☑ Enable SFTP

Domain name - Internal domain name.

Append suffix to DNS lookup - If a domain name is entered, it will be appended to the machine name for name resolution.

DNS - IP address of your Domain Name Server.

Default gateway - IP address of your default gateway.

NIC - List of Ethernet cards in the server.

Enable SSH - Enable/disable the ability to use the SSH protocol. The VTL server must have “openssh” installed in order to use SSH.

Enable SFTP - Enable/disable the ability to securely FTP into the server. SFTP can be used to save the server configuration.

2. Click *Config NIC* to configure each network interface card (NIC).

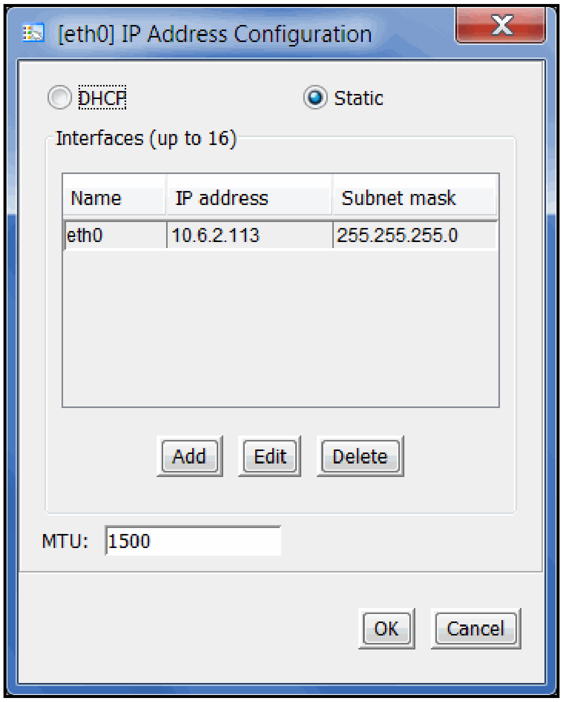

[eth0] IP Address Configuration

○ DHCP ● Static

Interfaces (up to 16)

Name	IP address	Subnet mask
eth0	10.6.2.113	255.255.255.0

MTU: 1500

If you select *Static*, you must click the *Add* button to add IP addresses and subnet masks.

MTU - Set the maximum transfer unit of each IP packet. If your card supports it, set this value to 9000 for jumbo frames.

Configuration note: After completing the configuration wizard, if you need to change these settings, you can right-click your server and select *System Maintenance --> Network Configuration*

Set hostname

(Optional for all servers) Enter a valid name for your deduplication appliance.

Valid characters are letters, numbers, underscore, or dash. The server will automatically reboot when the hostname is changed.

Configuration note: After completing the configuration wizard, if you need to change the name again, you can right-click your server and select *System Maintenance --> Set Hostname*.

Enable FC Target mode

To use Fibre Channel, click *Configure*. This step takes just a few seconds and there are no additional screens to go through.

Configuration note: Before you enable FC Target Mode, verify that your Fibre Channel configuration is set properly. Refer to the “Fibre Channel Configuration” section for information.

Switch FC port(s) to target mode

This is necessary to allow the server to access the repository in order to restore data. In a multi-node deduplication cluster, the target port is also used by each deduplication server to access the data repository of the other deduplication servers.

You will get a *Loop Up* message on your server if a QLogic port has successfully been placed in target mode.

In order to identify your ports, you need to know the WWPN of each. One way to find the WWPN is through the SNS table at your Fibre Channel switch.

Alternatively, for QLogic HBAs, you can find the WWPN in the BIOS (press Ctrl+Q during boot up).

(Single-ID HBAs) Select which ports should be in target mode.

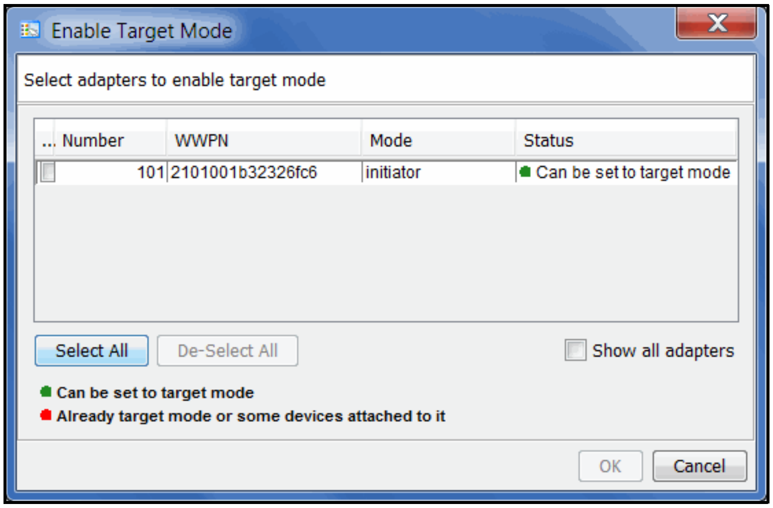

Select adapters to enable target mode

...	Number	WWPN	Mode	Status
□	101	2101001b32326fc6	initiator	🟩 Can be set to target mode

□ Show all adapters

🟩 Can be set to target mode

🟥 Already target mode or some devices attached to it

(Multi-ID HBAs) Click *Ok*.

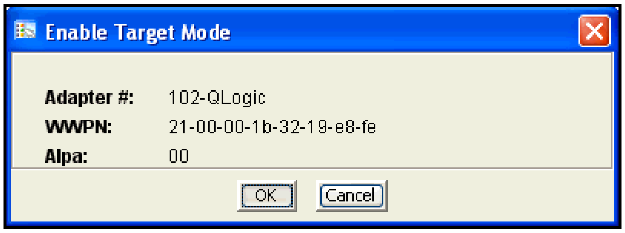

Adapter #: 102-QLogic
WWPN: 21-00-00-1b-32-19-e8-fe
Alpa: 00

Note that the adapter mode changes from **initiator** to **target** or **dual** (for multi-ID HBAs).

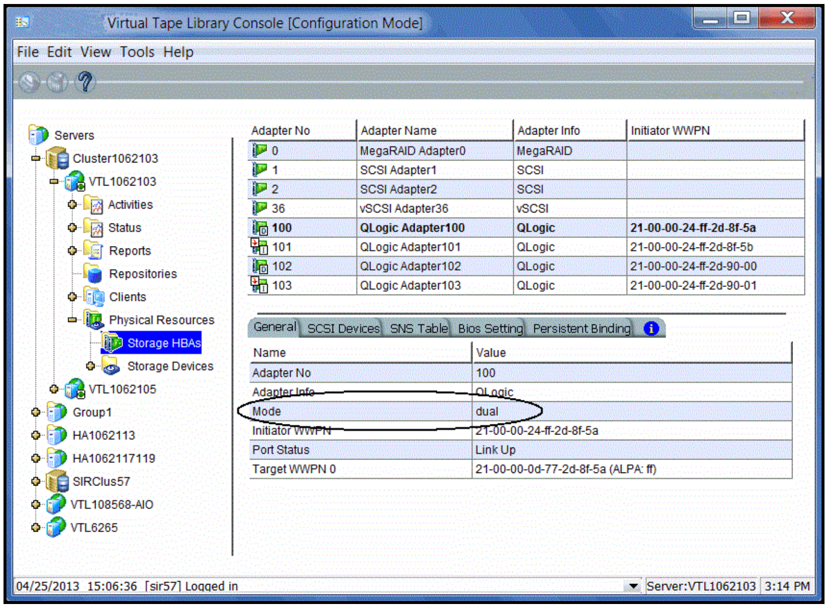

Adapter No	Adapter Name	Adapter Info	Initiator WWPN
0	MegaRAID Adapter0	MegaRAID	
1	SCSI Adapter1	SCSI	
2	SCSI Adapter2	SCSI	
36	vSCSI Adapter36	vSCSI	
100	**QLogic Adapter100**	**QLogic**	**21-00-00-24-ff-2d-8f-5a**
101	QLogic Adapter101	QLogic	21-00-00-24-ff-2d-8f-5b
102	QLogic Adapter102	QLogic	21-00-00-24-ff-2d-90-00
103	QLogic Adapter103	QLogic	21-00-00-24-ff-2d-90-01

Name	Value
Adapter No	100
Adapter Info	QLogic
Mode	dual
Initiator WWPN	21-00-00-24-ff-2d-8f-5a
Port Status	Link Up
Target WWPN 0	21-00-00-0d-77-2d-8f-5a (ALPA: ff)

04/25/2013 15:06:36 [sir57] Logged in | Server:VTL1062103 | 3:14 PM

Configuration note: After completing the configuration wizard, if you need to switch a port’s mode, you can right-click the adapter and select *Enable/Disable Target Mode*.

Prepare storage devices for each deduplication server

You must prepare your storage devices to define how each will be used.

Disk is used for the configuration repository (configuration information for each server) and deduplication repository (index, folder, and data disks). Memory is used for the deduplication index cache.

The following are some considerations you should be aware of before preparing your storage devices from the console:

- Each LUN used for the deduplication repository cannot be larger than 16 TB.
- Each LUN used for index or folder resources cannot be larger than 2 TB.
- The minimum size disk recommended for index or folder resources is at least 4.3% of the total size of the deduplication data disks per deduplication server.
- If you will be configuring the deduplication data disks in *Classic* mode (which requires organizing your disks into columns/rows), all data LUNS in the same row should be the same size. Refer to ’Classic configuration’ for more information about this mode.

To prepare storage devices from the console:

1. Expand *Physical Resources* --> *Storage Devices* --> *Fibre Channel Devices* so that you can see the available physical disks.

2. Right-click the *Storage Devices* object (or *Fibre Channel Devices*) and select *Prepare Devices*.

   The Physical Devices Preparation Wizard launches.

3. In the *Select Preparation Operation* dialog, select *Virtual Device*.

4. In the *Select Physical Devices* dialog, select and reserve the unassigned disks to be virtualized.

   The reservation type determines what kind of resources can be created on this device. On a deduplication server, devices can be reserved for:

   - Configuration repository - The configuration repository contains configuration information for each server. A maximum of two devices can be reserved for a deduplication server (including mirror devices).
   - Deduplication repository - includes deduplication index, folder, and data disks, as well as the associated mirror devices.

- None - the device will not be allocated. If there are existing resources on this device, they can still be accessed; however, new resources will not be created on this device.

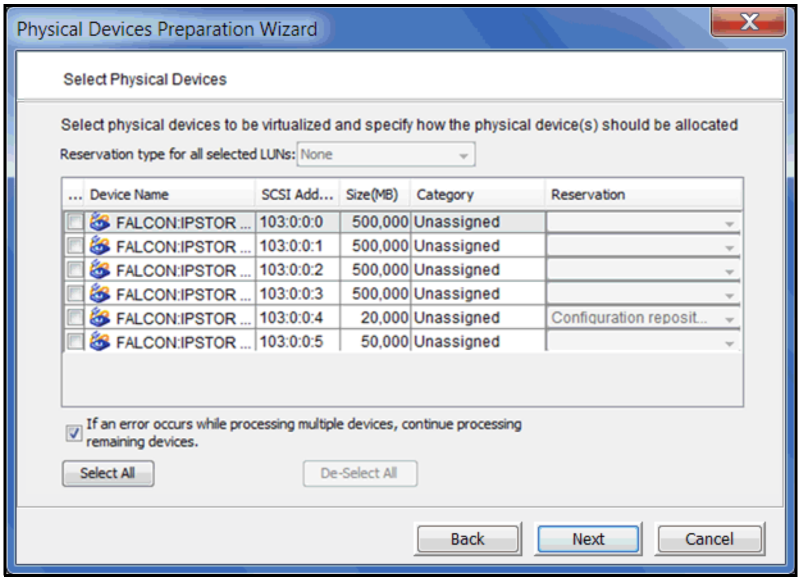

Physical Devices Preparation Wizard

Select Physical Devices

Select physical devices to be virtualized and specify how the physical device(s) should be allocated

Reservation type for all selected LUNs: None

...	Device Name	SCSI Add...	Size(MB)	Category	Reservation
☐	FALCON:IPSTOR ...	103:0:0:0	500,000	Unassigned	
☐	FALCON:IPSTOR ...	103:0:0:1	500,000	Unassigned	
☐	FALCON:IPSTOR ...	103:0:0:2	500,000	Unassigned	
☐	FALCON:IPSTOR ...	103:0:0:3	500,000	Unassigned	
☐	FALCON:IPSTOR ...	103:0:0:4	20,000	Unassigned	Configuration reposit...
☐	FALCON:IPSTOR ...	103:0:0:5	50,000	Unassigned	

☑ If an error occurs while processing multiple devices, continue processing remaining devices.

5. When prompted, type *Yes* to acknowledge the warning and then click *OK*.

Enable the configuration repository

The configuration repository contains configuration information for each server and it must be enabled on each deduplication server.

1. Right-click on the deduplication server and select *Options --> Enable Configuration Repository*.

2. Select the device to be used for the configuration repository.

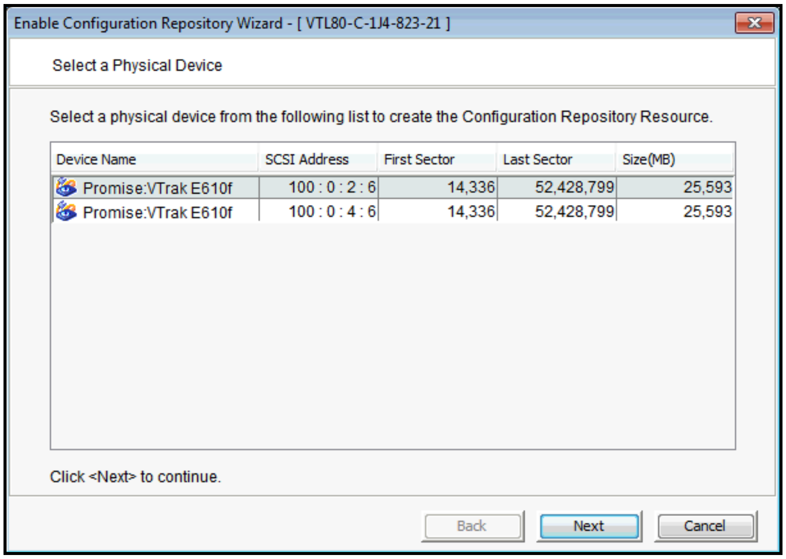

Enable Configuration Repository Wizard - [VTL80-C-1J4-823-21]

Select a Physical Device

Select a physical device from the following list to create the Configuration Repository Resource.

Device Name	SCSI Address	First Sector	Last Sector	Size(MB)
Promise:VTrak E610f	100 : 0 : 2 : 6	14,336	52,428,799	25,593
Promise:VTrak E610f	100 : 0 : 4 : 6	14,336	52,428,799	25,593

Click <Next> to continue.

Devices must have been reserved for this purpose. One device will be used for the primary configuration repository. If a second device was reserved, it can be used as the mirror. Refer to 'Mirror repository disks to protect your configuration' for more information.

The wizard is now complete. Continue with the steps on the subsequent pages to create a deduplication cluster. If you need to restart the wizard, right-click the server and select *Configuration Wizard*.

Create a deduplication cluster

Your deduplication server can be configured using unassigned physical devices. Creating a deduplication cluster is required, regardless of the number of nodes you have.

1. Right-click the deduplication server and select *Enable Deduplication Cluster*.

2. Enter cluster information.

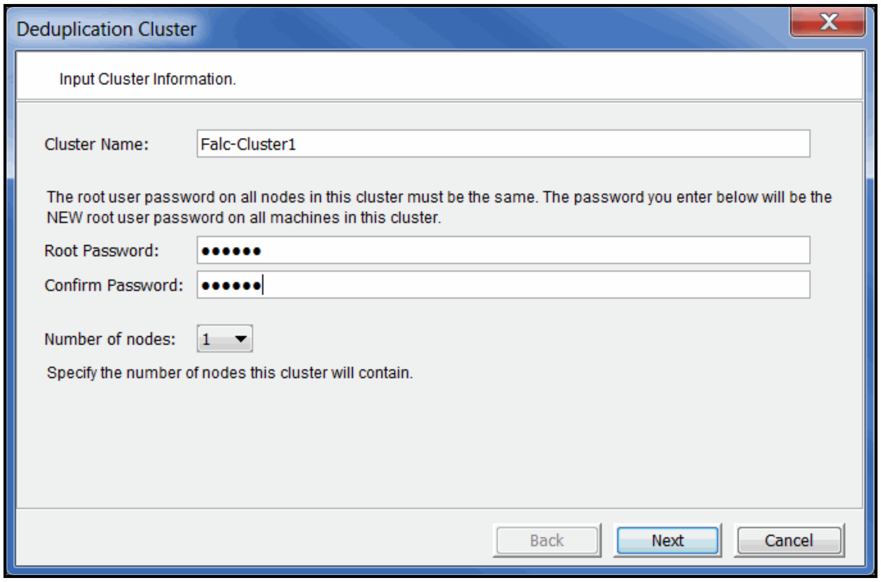

Input Cluster Information.

Cluster Name: Falc-Cluster1

The root user password on all nodes in this cluster must be the same. The password you enter below will be the NEW root user password on all machines in this cluster.

Root Password: ●●●●●●

Confirm Password: ●●●●●●

Number of nodes: 1

Specify the number of nodes this cluster will contain.

You need to provide a unique name for this cluster; the cluster cannot have the same name as another deduplication cluster or server. The cluster name is not case-sensitive.

Type a root password and then confirm the password by typing it again. If you are using the *Strong Password* option, the password must adhere to the rules for strong passwords.

The root user password for all nodes in the cluster must be the same. The root password you enter will replace any pre-existing password for all machines in the cluster.

Specify the number of servers that this cluster will contain. You can select 1, 2, or 4 nodes. Do not include any standby nodes.

If you will be configuring the deduplication data disks in your cluster in *Classic* mode (which requires organizing your disks into columns/rows), you must have at least as many data devices as the maximum you specify (255 max per cluster). Therefore, if you specify four nodes, you must have at least four data devices. This is not necessary if you will be configuring your cluster in *Flexible* method (which allows any number of data disks to be used).

3. Select the nodes for this cluster and the IP address that should be used on that node.

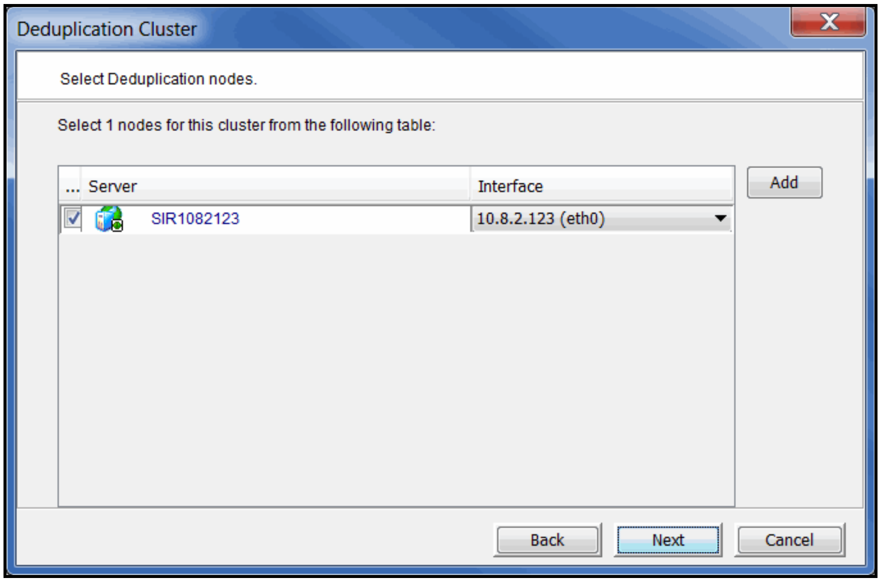

Deduplication Cluster

Select Deduplication nodes.

Select 1 nodes for this cluster from the following table:

...	Server	Interface
☑	SIR1082123	10.8.2.123 (eth0)

4. Select if you want to use deduplication repository encryption.

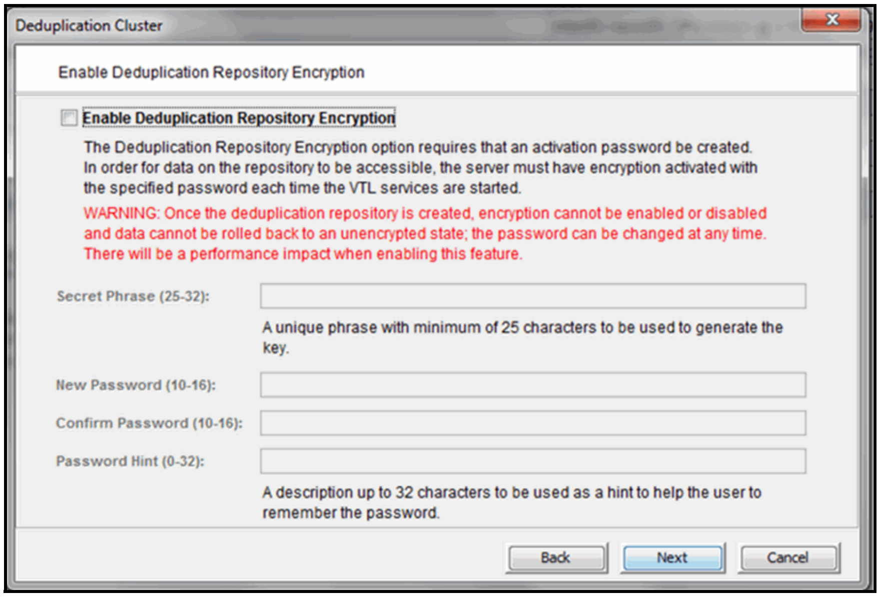

Deduplication Cluster

Enable Deduplication Repository Encryption

☐ Enable Deduplication Repository Encryption

The Deduplication Repository Encryption option requires that an activation password be created. In order for data on the repository to be accessible, the server must have encryption activated with the specified password each time the VTL services are started.

WARNING: Once the deduplication repository is created, encryption cannot be enabled or disabled and data cannot be rolled back to an unencrypted state; the password can be changed at any time. There will be a performance impact when enabling this feature.

Secret Phrase (25-32):

A unique phrase with minimum of 25 characters to be used to generate the key.

New Password (10-16):

Confirm Password (10-16):

Password Hint (0-32):

A description up to 32 characters to be used as a hint to help the user to remember the password.

For new VTL 8.20 and higher installations, you will only see this dialog if encryption was unlocked from the command line of your deduplication server.

Encryption can be used to ensure that the data stored in the deduplication repository is confidential and secure. In order for data on the repository to be accessible, the cluster servers must have encryption activated with the specified password each time the VTL services are started. All servers in the cluster use the same encryption activation password.

In the *Secret Phrase* text box, type a phrase (25–32 characters, including numbers and spaces) that will be used to create an internal key to encrypt the data.

In the *New Password* and *Confirm Password* text boxes, type a password that will be used as the encryption activation password (10–16 characters). You will need to provide this password in order to activate encryption or change the password.

In the *Password Hint* text box, type a hint (0–32 characters) that will help you remember the password. This hint appears when you type an incorrect password and request a hint.

Notes:

- This is your only opportunity to enable encryption. Once the deduplication repository is created, encryption cannot be enabled or disabled and data cannot be rolled back to an unencrypted state.
- NAS cannot be enabled for servers associated with a deduplication cluster repository encryption.
- When encryption is enabled, there will be an overall performance decrease for read and write operations. The actual impact will depend upon a number of factors, including the number of CPU cores and speed, number of concurrent IO operations, data compression ratio, and data deduplication ratio. Faster server processors with more cores can be used to minimize the impact.

5. Determine which method to use to configure deduplication data disks, *Flexible* or *Classic*.

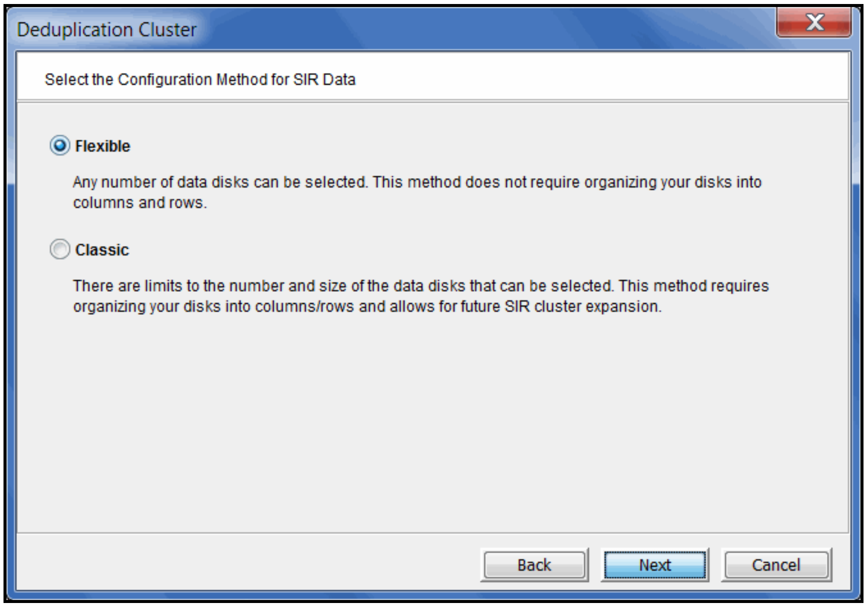

Deduplication Cluster

Select the Configuration Method for SIR Data

● Flexible

Any number of data disks can be selected. This method does not require organizing your disks into columns and rows.

○ Classic

There are limits to the number and size of the data disks that can be selected. This method requires organizing your disks into columns/rows and allows for future SIR cluster expansion.

Flexible - This method allows you to create a cluster with fewer, yet a more flexible number of data drives (255 max per cluster). This method does not require organizing your disks into columns and rows and allows for the use of all deduplication data drives in a cluster, rather than a row-by-row utilization of drives. With *Flexible* storage, each node in a cluster has a single hash range and all data drives for the node store data for the node's hash range. Note that if you select this method, you will not be able to add deduplication nodes in the future to increase your cluster size.

Classic - There are limits to the number and size of the data disks that can be selected. This method requires organizing your disks into columns/rows and allows for future deduplication cluster expansion.

> **Note:** If you select the *Classic* method, you can switch to Flexible at a later date. If you select the *Flexible* method, it is not reversible.

Flexible configuration

1. Select the device(s) that will be used for data storage by each server.

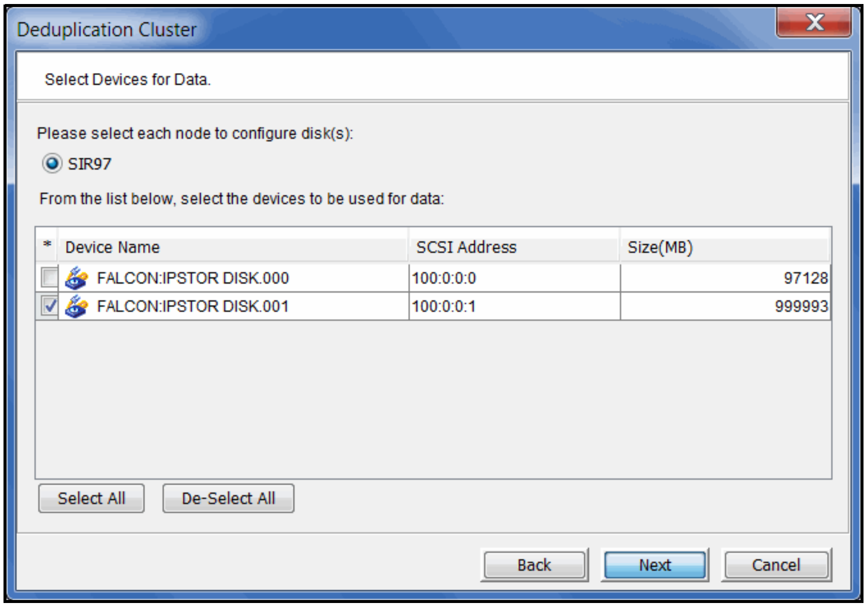

Deduplication Cluster

Select Devices for Data.

Please select each node to configure disk(s):

● SIR97

From the list below, select the devices to be used for data:

*	Device Name	SCSI Address	Size(MB)
□	FALCON:IPSTOR DISK.000	100:0:0:0	97128
☑	FALCON:IPSTOR DISK.001	100:0:0:1	999993

2. Select the device(s) that will be used for the deduplication index and folder resources by each server.

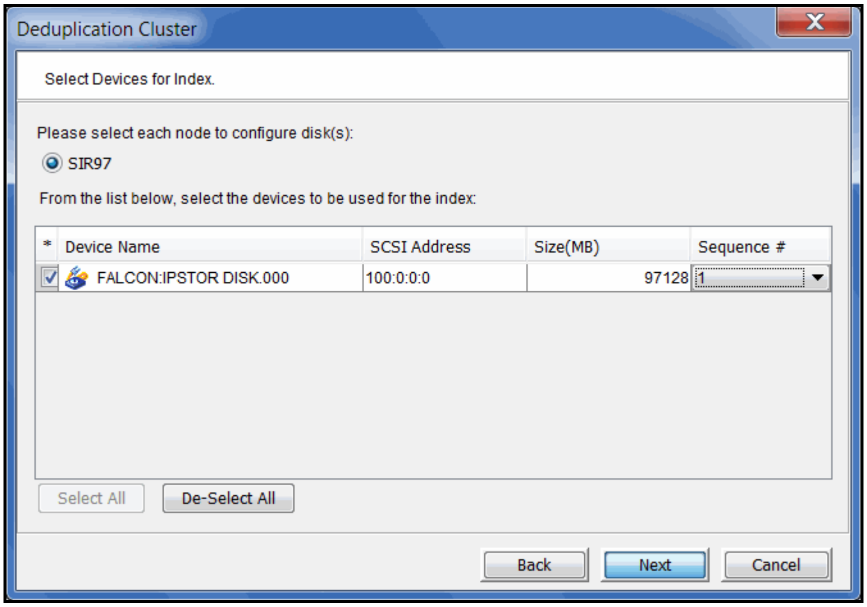

Select Devices for Index.

Please select each node to configure disk(s):

● SIR97

From the list below, select the devices to be used for the index:

*	Device Name	SCSI Address	Size(MB)	Sequence #
☑	FALCON:IPSTOR DISK.000	100:0:0:0	97128	1

The folder resource is used to store information related to the contents of each deduplication session.

The selected device(s) must be of sufficient size to store these resources. The minimum size disk recommended for these resources is at least 4.3% of the total size of the deduplication data disks.

3. In the confirmation dialog, select *Finish* to complete the wizard.

Classic configuration

1. For performance purposes, select how your data devices will be organized.

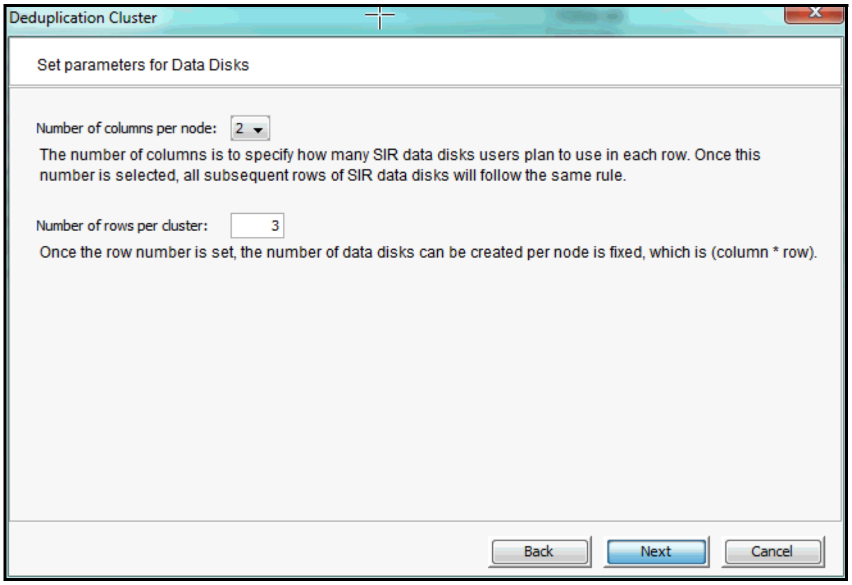

Set parameters for Data Disks

Number of columns per node: 2

The number of columns is to specify how many SIR data disks users plan to use in each row. Once this number is selected, all subsequent rows of SIR data disks will follow the same rule.

Number of rows per cluster: 3

Once the row number is set, the number of data disks can be created per node is fixed, which is (column * row).

You need to divide the number of storage devices you have per server into columns and rows. The purpose of doing this is to fully utilize the bandwidth of your storage to maximize performance.

A *column* represents the *y* axis in a graph while a *row* represents the *x* axis.

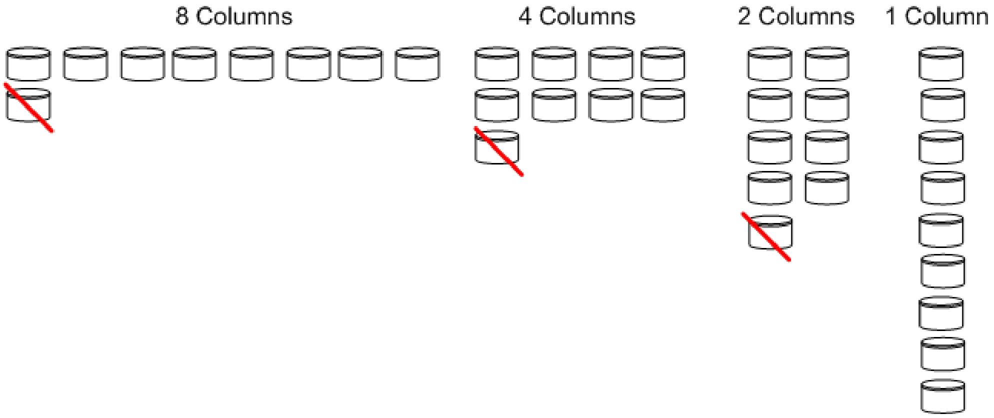

In the above example, there are nine disks and these are the possible groupings for a single node. The ninth disk is not used in the 2, 4, and 8 column configurations.

If you have two cluster nodes, the single column configuration is no longer viable because each node needs at least one column.

How do the number of columns affect performance?

Typically, if you have one RAID group, a single column is best. For multiple RAID groups, performance will be better with multiple columns if each column is on its own RAID controller. With multiple columns, the writes will be distributed between the storage devices, limiting the number of writes to each RAID group.

Write caching also affects performance. Performance will be better with multiple nodes if write caching is enabled on all disks.

2. If you prepared physical disks for deduplication storage, select *Physical Device*. If you are using existing virtualized resources for deduplication storage, select *Virtual Device*.

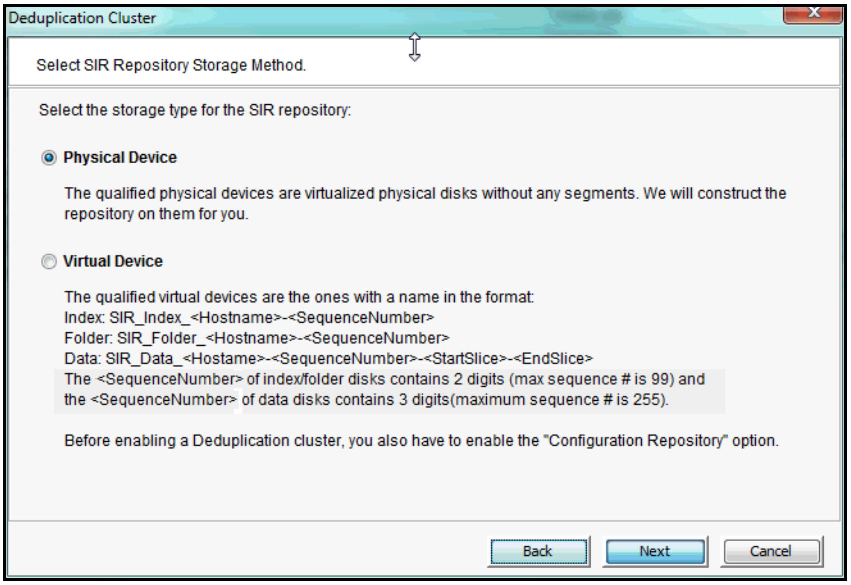

Select SIR Repository Storage Method.

Select the storage type for the SIR repository:

● **Physical Device**

The qualified physical devices are virtualized physical disks without any segments. We will construct the repository on them for you.

○ **Virtual Device**

The qualified virtual devices are the ones with a name in the format:
Index: SIR_Index_<Hostname>-<SequenceNumber>
Folder: SIR_Folder_<Hostname>-<SequenceNumber>
Data: SIR_Data_<Hostame>-<SequenceNumber>-<StartSlice>-<EndSlice>
The <SequenceNumber> of index/folder disks contains 2 digits (max sequence # is 99) and the <SequenceNumber> of data disks contains 3 digits(maximum sequence # is 255).

Before enabling a Deduplication cluster, you also have to enable the "Configuration Repository" option.

3. Select the device(s) that will be used for data storage by each server and assign the rows and columns.

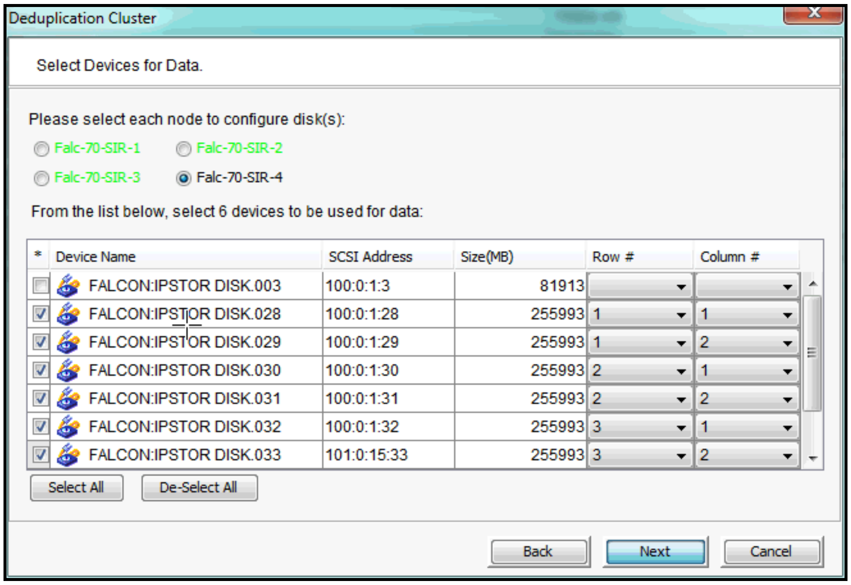

Select Devices for Data.

Please select each node to configure disk(s):

○ Falc-70-SIR-1 ○ Falc-70-SIR-2
○ Falc-70-SIR-3 ● Falc-70-SIR-4

From the list below, select 6 devices to be used for data:

*	Device Name	SCSI Address	Size(MB)	Row #	Column #
□	FALCON:IPSTOR DISK.003	100:0:1:3	81913		
☑	FALCON:IPSTOR DISK.028	100:0:1:28	255993	1	1
☑	FALCON:IPSTOR DISK.029	100:0:1:29	255993	1	2
☑	FALCON:IPSTOR DISK.030	100:0:1:30	255993	2	1
☑	FALCON:IPSTOR DISK.031	100:0:1:31	255993	2	2
☑	FALCON:IPSTOR DISK.032	100:0:1:32	255993	3	1
☑	FALCON:IPSTOR DISK.033	101:0:15:33	255993	3	2

Select a server, and then, based on the number of rows and columns you set earlier in the wizard, select the disks to be used. The row and column assignment will be automatically defined by LUN but you may want to override these settings.

For example, as you can see in the diagram below, you have three LUNs per storage tray. The bus to each tray supports a maximum of 200 MB/s. By default, the system will group LUN1-LUN4 together. However, that limits storage throughput to 400 MB/s (Tray 1 gives 200 MB/s and Tray 2 gives 400 MB/s).

However, if you grouped all of the LUNs in Row 1 (LUN1, LUN4, LUN7, and LUN10) together, you would be able to get 800 MB/s (Tray 1 - 200 MB/s, Tray 2 - 400 MB/s, Tray 3 - 200 MB/s, and Tray 4 - 200 MB/s). To do this, you would set your rows and columns as shown.

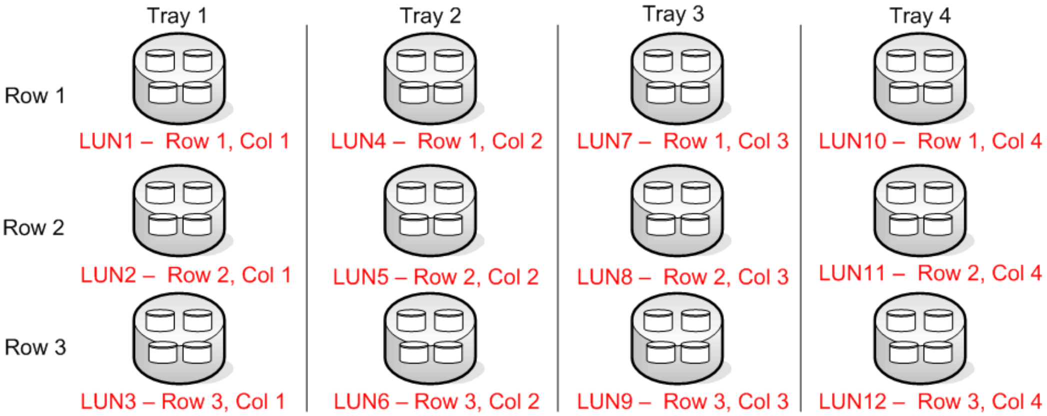

	Tray 1	Tray 2	Tray 3	Tray 4
Row 1	LUN1 – Row 1, Col 1	LUN4 – Row 1, Col 2	LUN7 – Row 1, Col 3	LUN10 – Row 1, Col 4
Row 2	LUN2 – Row 2, Col 1	LUN5 – Row 2, Col 2	LUN8 – Row 2, Col 3	LUN11 – Row 2, Col 4
Row 3	LUN3 – Row 3, Col 1	LUN6 – Row 3, Col 2	LUN9 – Row 3, Col 3	LUN12 – Row 3, Col 4

When you finish assigning rows and columns for one server and then select the next one, the server name will be displayed in green to show that the assignment was successful. If there are any issues, they will be presented to you when you try to select another server.

4. Select the device(s) that will be used for the deduplication index and folder resources on each server.

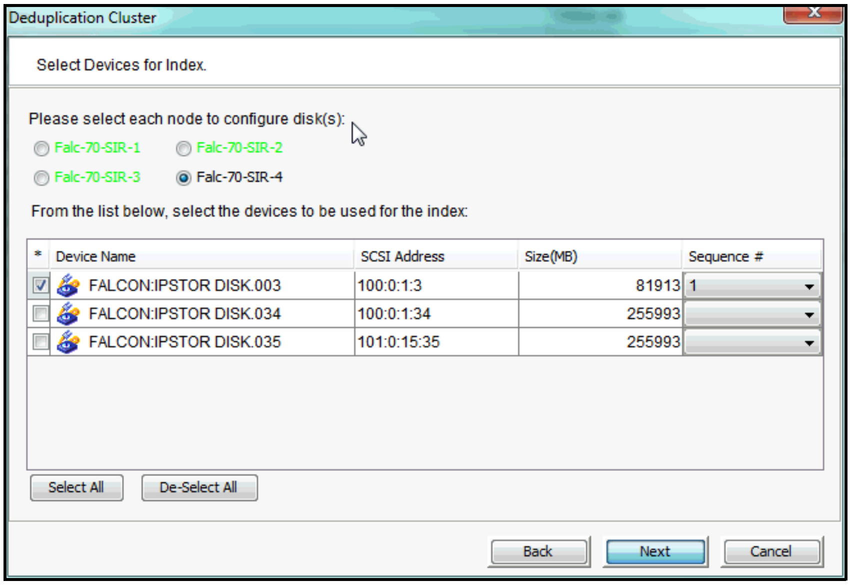

Deduplication Cluster

Select Devices for Index.

Please select each node to configure disk(s):

○ Falc-70-SIR-1 ○ Falc-70-SIR-2

○ Falc-70-SIR-3 ● Falc-70-SIR-4

From the list below, select the devices to be used for the index:

*	Device Name	SCSI Address	Size(MB)	Sequence #
☑	FALCON:IPSTOR DISK.003	100:0:1:3	81913	1
□	FALCON:IPSTOR DISK.034	100:0:1:34	255993	
□	FALCON:IPSTOR DISK.035	101:0:15:35	255993	

The folder resource is used to store information related to the contents of each deduplication session.

The selected device(s) must be of sufficient size to store these resources. The minimum size disk recommended for these resources is at least 4.3% of the total size of the deduplication data disks.

5. In the confirmation dialog, select *Finish* to complete the wizard.

VTL Tape Configuration

This section describes how to configure a VTL or VTL-S server for tape backups and/or NAS.

Prepare a VTL/VTL-S server via the configuration wizard

Notes:

- If you are using VTL in a Fibre Channel environment, refer to the "Fibre Channel Configuration" section first before beginning the wizard.
- If you are setting up a failover configuration, refer to the "VTL Server Failover" section first before beginning the wizard.

If your VTL server has not been configured yet, the configuration wizard will be launched when you connect to it.

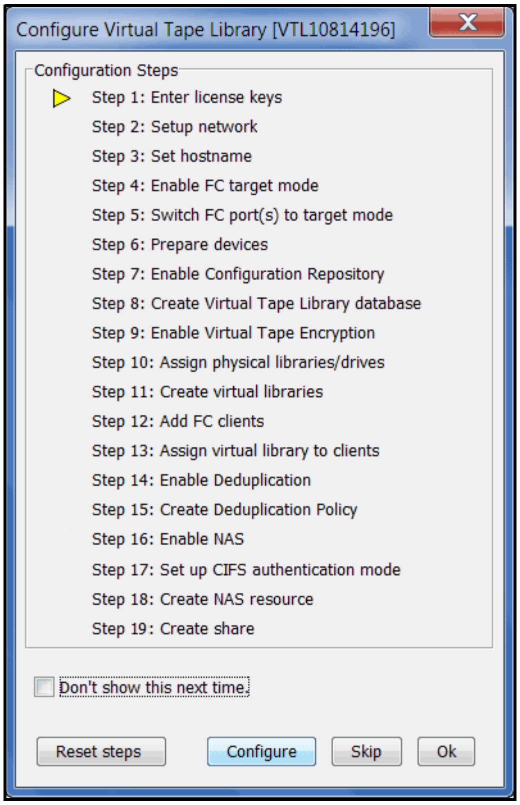

Configure Virtual Tape Library [VTL10814196]

Configuration Steps

- Step 1: Enter license keys
- Step 2: Setup network
- Step 3: Set hostname
- Step 4: Enable FC target mode
- Step 5: Switch FC port(s) to target mode
- Step 6: Prepare devices
- Step 7: Enable Configuration Repository
- Step 8: Create Virtual Tape Library database
- Step 9: Enable Virtual Tape Encryption
- Step 10: Assign physical libraries/drives
- Step 11: Create virtual libraries
- Step 12: Add FC clients
- Step 13: Assign virtual library to clients
- Step 14: Enable Deduplication
- Step 15: Create Deduplication Policy
- Step 16: Enable NAS
- Step 17: Set up CIFS authentication mode
- Step 18: Create NAS resource
- Step 19: Create share

☐ Don't show this next time.

Click *Configure* to begin the steps in the wizard.

Note: The actual wizard you see will depend upon the options you have licensed and activated.

Enter license keys

(Required for all servers) Click the *Add* button and enter your keycodes, one at a time.

Be sure to enter keycodes for any options you have purchased. Each VTL option requires that a keycode be entered before the option can be configured and used. Refer to 'Add and register licenses' for more information.

Note: After completing the configuration wizard, if you need to add license keys, you can right-click your server and select *License*.

Set up network

(Optional for all servers) This step allows you to set/change your network configuration.

Note: Changing the IP address is allowed only before you enable replication or configure replication. To change the IP address after you configure replication, you must remove the replication target, change the IP address, and then configure replication again.

1. Enter information about your network configuration.

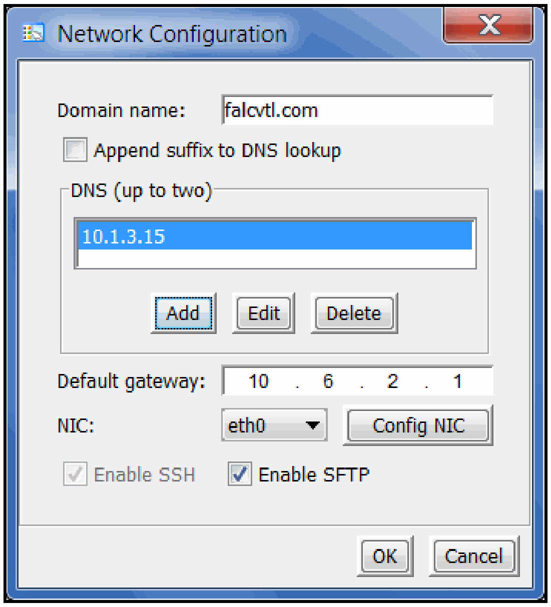

Network Configuration

Domain name: falcvtl.com

☐ Append suffix to DNS lookup

DNS (up to two)

10.1.3.15

Default gateway: 10 . 6 . 2 . 1

NIC: eth0

☑ Enable SSH ☑ Enable SFTP

Domain name - Internal domain name.

Append suffix to DNS lookup - If a domain name is entered, it will be appended to the machine name for name resolution.

DNS - IP address of your Domain Name Server.

Default gateway - IP address of your default gateway.

NIC - List of Ethernet cards in the server.

Enable SSH - Enable/disable the ability to use the SSH protocol. The VTL server must have “openssh” installed in order to use SSH.

Enable SFTP - Enable/disable the ability to securely FTP into the server. SFTP can be used to save the server configuration.

2. Click *Config NIC* to configure each network interface card (NIC).

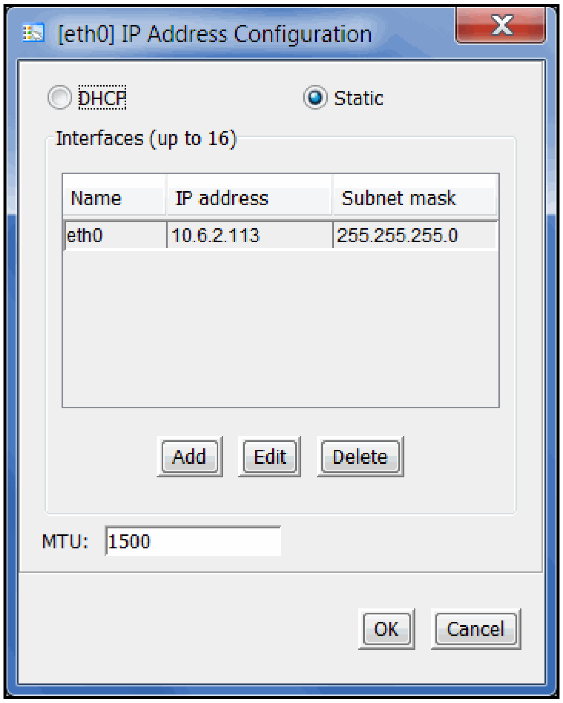

[eth0] IP Address Configuration

○ DHCP ● Static

Interfaces (up to 16)

Name	IP address	Subnet mask
eth0	10.6.2.113	255.255.255.0

MTU: 1500

If you select *Static*, you must click the *Add* button to add IP addresses and subnet masks.

MTU - Set the maximum transfer unit of each IP packet. If your card supports it, set this value to 9000 for jumbo frames.

Configuration note: After completing the configuration wizard, if you need to change these settings, you can right-click your server and select *System Maintenance --> Network Configuration*

Set hostname

(Optional for all servers) Enter a valid name for your VTL appliance.

Valid characters are letters, numbers, underscore, or dash. The server will automatically reboot when the hostname is changed.

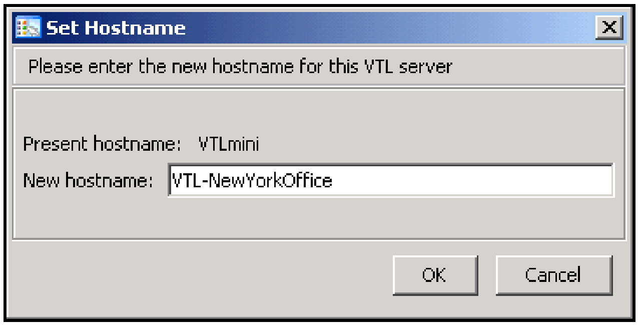

Set Hostname

Please enter the new hostname for this VTL server

Present hostname: VTLmini

New hostname: VTL-NewYorkOffice

Configuration notes:

- Do not use an underscore ('_') in the hostname if this VTL server will be associated with a deduplication cluster.
- After completing the configuration wizard, if you need to change the name again, you can right-click your server and select *System Maintenance* --> *Set Hostname*.

Enable FC Target Mode

(Required for tape backups) To use Fibre Channel, click *Configure*. This step takes just a few seconds and there are no additional screens to go through.

Note: Before you enable FC Target Mode, verify that your Fibre Channel configuration is set properly. Refer to the “Fibre Channel Configuration” section for information.

Switch FC port(s) to target mode

(Required for tape backups) Target mode allows a port to receive requests from your backup application server(s).

If you haven't already done so, you will need to switch any initiator zoned with a backup application server to target mode so that the backup application server can see the VTL server. You will then need to select the equivalent adapter on the secondary server and switch it to target mode.

You will get a *Loop Up* message on your VTL server if a QLogic port has successfully been placed in target mode.

In order to identify your ports, you need to know the WWPN of each. One way to find the WWPN is through the SNS table at your Fibre Channel switch.

Alternatively, for QLogic HBAs, you can find the WWPN in the BIOS (press Ctrl+Q during boot up).

For VTL Failover, if you are not using multi-ID HBAs, you minimally need two ports for client connectivity (one for normal operation, one for standby) and one initiator port to connect to storage. If you are using multi-ID HBAs, you need a minimum of one port for client connectivity and one for storage.

(Single-ID HBAs) Select which ports should be in target mode.

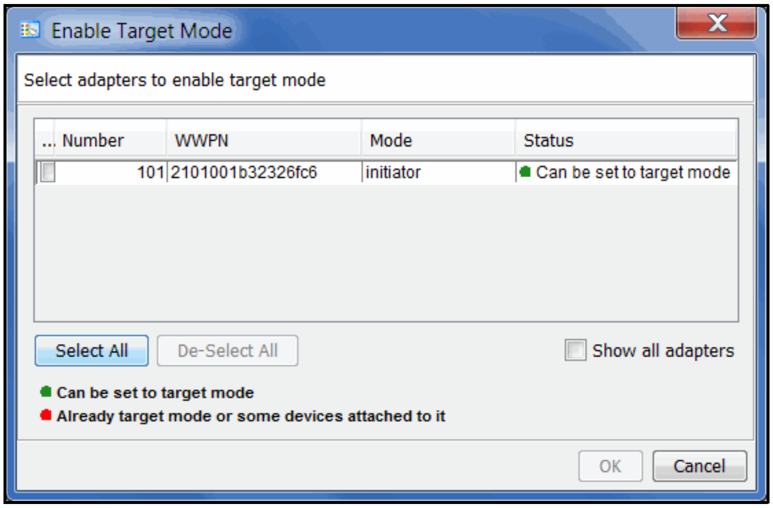

Select adapters to enable target mode

...	Number	WWPN	Mode	Status
□	101	2101001b32326fc6	initiator	🟩 Can be set to target mode

□ Show all adapters

🟩 Can be set to target mode
🟥 Already target mode or some devices attached to it

(Multi-ID HBAs) Click *Ok*.

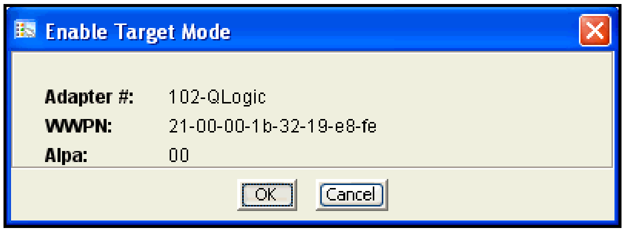

Adapter #: 102-QLogic
WWPN: 21-00-00-1b-32-19-e8-fe
Alpa: 00

Configuration note: After completing the configuration wizard, if you need to switch a port's mode, you can right-click the adapter and select *Enable/Disable Target Mode*.

Prepare devices

(Required for all servers) Select each physical device that you want to virtualize and use with VTL and reserve how each device can be used. The reservation type determines what kind of resources (tapes, NAS, deduplication data disk, etc.) can be created on this device. Devices can be reserved for:

- None - The device will not be allocated. If there are existing resources on this device, they can still be accessed; however, new resources will not be created on this device.
- Configuration repository - The configuration repository contains configuration information for each server. A maximum of four devices can be reserved for the configuration repository and VTL database (including mirror devices). After the configuration repository is created, you cannot change the reservation of devices used for the configuration repository.
- Deduplication repository - (VTL-S only) Includes deduplication index, folder, and data disks, as well as the associated mirror devices.
- Tapes - Used only for tape storage (VTL servers and VTL-S servers).
- NAS resources - Used only for NAS resources (VTL servers and VTL-S servers).

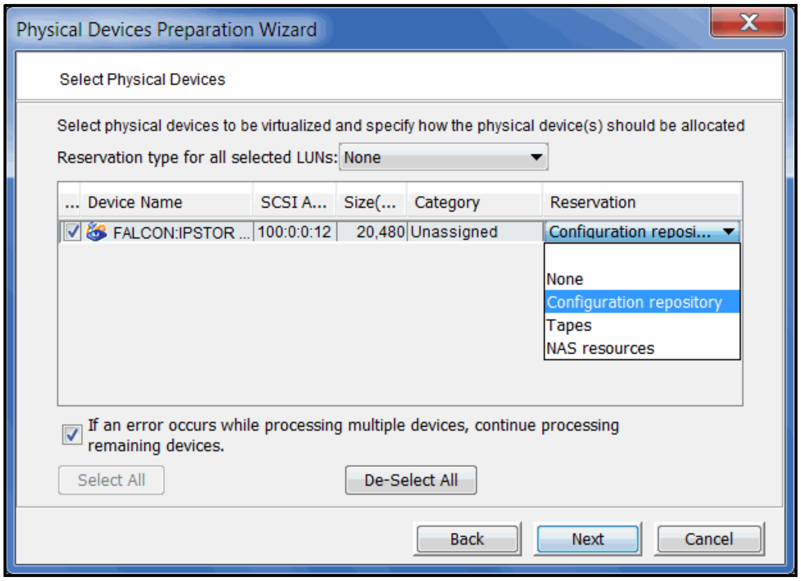

Physical Devices Preparation Wizard

Select Physical Devices

Select physical devices to be virtualized and specify how the physical device(s) should be allocated

Reservation type for all selected LUNs: None

...	Device Name	SCSI A...	Size(...	Category	Reservation
☑	FALCON:IPSTOR ...	100:0:0:12	20,480	Unassigned	Configuration reposi...

☑ If an error occurs while processing multiple devices, continue processing remaining devices.

Configuration note: After completing the configuration wizard, if you add new hardware that you need to prepare, you can right-click the *Physical Resources* node (or one of the nodes below it, including *Storage Devices, Fibre Channel Devices* and *SCSI Devices*) and select *Prepare Devices.* Note that the console display mode must be set to *Configuration* in order to see the *Physical Resources* node.

Enable Configuration Repository

(Required for all servers) The configuration repository contains VTL configuration information. To enable it, click *Configure*. You will need to select the physical device(s) reserved for this purpose. If mirror devices were reserved, refer to ’Mirror repository disks to protect your configuration’ for more information.

Create Virtual Tape Library database

(Required for tape backups) The database contains information about libraries, clients, and replication setup. If your VTL appliance has been preconfigured, this step takes just a few seconds and there are no additional screens to go through. If your VTL appliance has not been preconfigured, you must have already prepared storage resources for use with VTL.

1. Select how you want to create the Virtual Tape Library database.

   Custom lets you select which physical device(s) to use and lets you designate how much space to allocate from each.

   Express automatically creates the resource for you using an available device(s), but is not recommended unless your appliance has been preconfigured.

2. If desired, enable disk compression for VTL.

   This can save disk space because it compresses the data that is being written to your virtual tapes. It does not affect data that is exported to physical tapes because that is controlled by the tape drive.

3. (Tape Caching only) Set the global *disk capacity* threshold.

   When this threshold value is reached, migration from virtual tape to physical tape will be triggered for all virtual libraries using the disk capacity threshold.

4. Click *Finish* to create the database.

   If you know that you have a disk available, you can create a mirror for the VTL database in order to protect your VTL configuration. Even if you lose your VTL server, the data on your tapes will be maintained. Mirroring the database is highly recommended. Refer to ’Mirror repository disks to protect your configuration’ for more information.

Configuration note: After completing the configuration wizard, if you want to enable disk compression, right-click the *Virtual Tape Library System* object and select *Properties*. (If the server is a member of a group, right-click the group and select *VTL Properties*.) To mirror the database at a later time, you must be in console configuration mode (set under *Console Options*). Then, right-click the *Database* object (under the *Repositories* object) and select *Mirror* --> *Add*.

Enable Virtual Tape Encryption

For new VTL 8.20 and higher installations, you will only see this step if encryption was unlocked from the command line of your virtual tape server.

Encryption can be used to ensure that data backed up on virtual tapes is confidential and secure.

Encryption requires that an activation password be created. In order for data on encrypted virtual tapes to be accessible, the server must have encryption activated with the specified password each time the VTL services are started. Encrypted virtual tapes will not be accessible for backup or restore without the activation password.

After enabling virtual tape encryption, you will need to create one or more encryption keys and enable encryption for your virtual tape libraries. Refer to ’Virtual tape encryption’ for more information.

Configuration note: After completing the configuration wizard, if you want to enable virtual tape encryption, right-click your VTL tape backup server and select *Options --> Enable Virtual Tape Encryption*.

Assign physical libraries/drives

(Optional for tape backups) If you will be importing data from physical tapes into your virtual tape library or exporting virtual tapes to physical tapes, you must assign your physical tape libraries/drives to VTL.

This step also inventories the physical tapes in your library/drive so that you can create virtual tapes that match your physical tapes.

Note: VTL does not support physical libraries when tape drive numbering does not start with 0 or is not sequential.

Create virtual libraries

(Required for tape backups) Select the tape library that you are emulating.

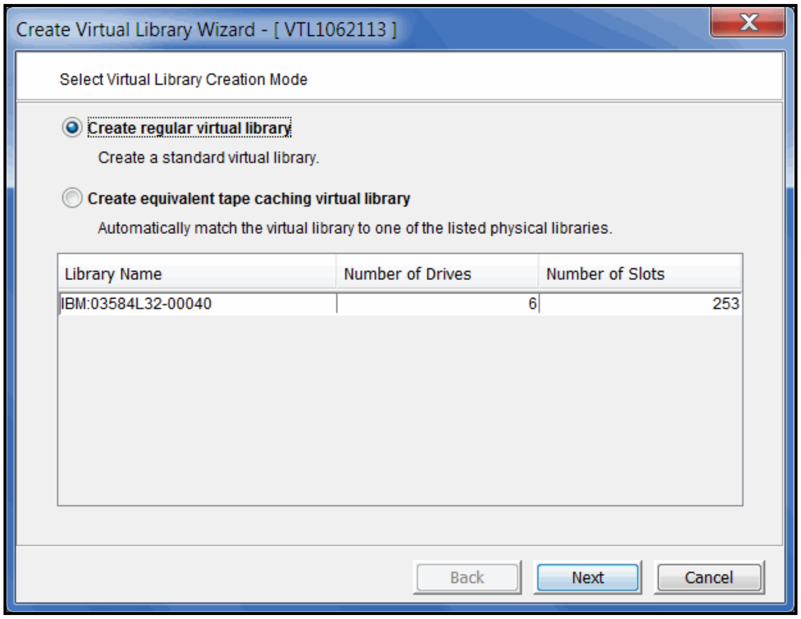

Select Virtual Library Creation Mode

● Create regular virtual library

Create a standard virtual library.

○ Create equivalent tape caching virtual library

Automatically match the virtual library to one of the listed physical libraries.

Library Name	Number of Drives	Number of Slots
IBM:03584L32-00040	6	253

If you have a physical tape library, you need to create a virtual tape library that resembles it. This way the virtual tapes will use the same format as those of the physical tapes. This is important for importing and exporting functions and guarantees that your backup application will accept the tapes.

If you assigned a physical tape library to your VTL, you will only see your physical tape libraries listed. Select the check box and the system will automatically match your physical library.

You will have to enter information about the tape drives in your library, including:

- Barcode information
- Tape properties such as Tape Capacity On Demand and maximum tape capacity.
- If you are using Automated Tape Caching, you will have to select the type of data migration triggers that you want to set and specify when the data that has been migrated to physical tape can be deleted to free up cache disk space.
- If you are not using Automated Tape Caching, you need to determine if you want to use auto archive *or* auto replication for this virtual library.

Refer to 'Create virtual tape libraries' for detailed information about creating virtual tape libraries and 'Automated Tape Caching' for detailed information about configuring Automated Tape Caching.

After you create a virtual tape library you will be prompted to create virtual tapes. Refer to ’Create virtual tapes’ for detailed information about creating them. After you create virtual tapes, you will be prompted to create more virtual libraries or to continue with the next step.

Configuration note: After completing the configuration wizard, if you need to create virtual tape libraries, you can right-click the *Virtual Tape Libraries* object and select *New*. If you need to add drives to an existing virtual tape library, you can right-click the library and select *New Drive(s)*.

Add FC clients

(Required for tape backups) This step allows you to select the clients (backup servers) to which you will be assigning a tape library.

Refer to ’Add tape backup clients’ for detailed information.

Configuration note: After completing the configuration wizard, if you need to add new clients, you can right-click *Virtual Tape Library System* and select *Configuration wizard* or you can right-click the *FC Clients* object and select *Add*.

Assign virtual library to clients

(Required for tape backups) If you added backup servers, do the following:

1. Select a backup server to assign.
2. Click *Finish* when you are done.

Refer to ’Assign virtual tape libraries and drives to backup servers’ for detailed information.

Configuration note: After completing the configuration wizard, if you need to assign new virtual libraries, you can right-click *Virtual Tape Library System* and select *Configuration wizard* or you can click a virtual tape library or a client and select *Assign*.

Enable Deduplication

(Required for all servers that will use deduplication) To use deduplication, click *Configure*.

You must associate each VTL server with one deduplication cluster.

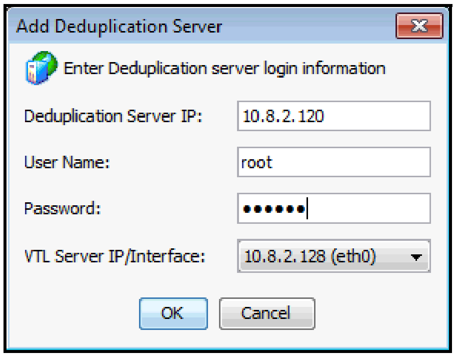

Add Deduplication Server

Enter Deduplication server login information

Deduplication Server IP: 10.8.2.120

User Name: root

Password: ••••••

VTL Server IP/Interface: 10.8.2.128 (eth0)

Note: You should not associate a VTL server while any I/O is running on the VTL system. As part of the configuration, the deduplication server will rescan all of the Fibre Channel HBA ports on the system and this will cause the I/O jobs running on the clients to fail.

Enter the IP address, user name, and password of any server in the deduplication cluster as well as the IP address on the VTL server that should be used for VTL/ deduplication communication.

If the deduplication server has its repository encrypted, virtual tape encryption must be enabled before deduplication can be configured. This protects data on virtual tapes before deduplication occurs, ensuring that data in the system is always secure.

In order to enable deduplication when the deduplication server has its repository encrypted, you need to enter the encryption activation password set on the deduplication repository.

For VTL-S servers, if virtual tape encryption is enabled when you enable deduplication, deduplication repository encryption will be enabled by default. You will need to enter the secret phrase for the encryption key.

You do not need to associate your VTL server with each deduplication server in a cluster. Multiple VTL servers can be associated with the same deduplication server.

If this VTL server will be a replication target, be sure to associate it with the *target* deduplication cluster.

Configuration notes:

- The VTL server that is being associated with the deduplication cluster cannot have an underscore ('_') in its name.
- If you see a connectivity error message, check the storage connectivity between your VTL and deduplication servers and then rescan.
- After completing the configuration wizard, if you need to enable Deduplication, you can right-click the server and select *Options* --> *Deduplication* --> *Enable Deduplication*.

Create Deduplication Policy

(Required for tape backups with deduplication) If you enabled deduplication and would like to create an additional deduplication policy at this time, complete this step, which launches the deduplication policy wizard (refer to ’Create tape deduplication policies’).

Configuration note: After completing the deduplication wizard, if you wish to create additional deduplication policies, right-click the *Deduplication Policies* object in the navigation tree and select *New*.

Enable NAS

(Required for NAS) To use NAS, click *Configure*. This step takes just a few seconds and there are no additional screens to go through.

Configuration note: After completing the configuration wizard, if you need to enable NAS, you can right-click the server and select *Options* --> *Enable NAS*.

Set up CIFS authentication mode

(Required for NAS) This step allows you to set the security mode to use to authenticate users/groups trying to access NAS shares.

Refer to ’Set the security mode for Windows clients’ for detailed information.

Configuration note: After completing the configuration wizard, if you need to change the authentication mode, you can right-click the *CIFS Client* object and select *Set Security Mode*.

Create NAS resource

(Required for NAS) This step allows you to create a NAS resource.

Refer to ’Create NAS resources’ for detailed information.

Configuration note: After completing the configuration wizard, if you need to create NAS resources, you can right-click the *NAS Resources* object and select *New*.

Create share

(Required for NAS) This step allows you to create a NAS shared folder.

Refer to ’Create shared folders’ for detailed information.

Configuration note: After completing the configuration wizard, if you need to create a share, you can right-click a NAS Resource and select *New Share*. You can also select *New Folder*. Any time after creating a folder, you can assign clients to it by right-clicking and selecting *Sharing*

The wizard is now complete. Continue with the steps on the subsequent pages to continue your configuration. If you need to restart the wizard, right-click the server and select *Configuration Wizard*.

Add tape backup clients

You need to add a client for each tape backup application server.

Note: Refer to "NAS Configuration" for information about CIFS and NFS clients.

1. Right-click the *FC Clients* or *iSCSI Clients* object and select *Add.*

2. Enter a unique client name or IP address (maximum length is 32 characters).

   The client name cannot be the same as any current or initial client name already in the system. The initial name will be preserved for this client even if you rename the client in the future.

3. If you started the wizard from the configuration wizard, select the protocol(s) being used by the client.

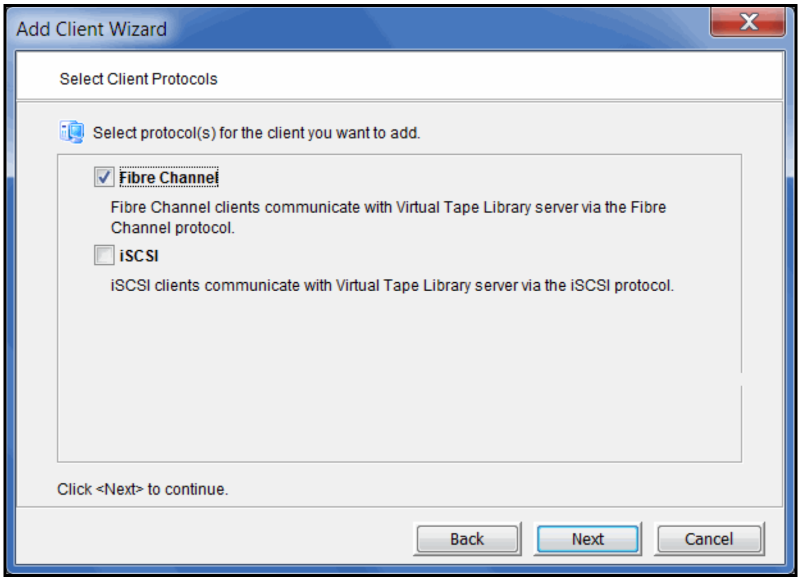

Add Client Wizard

Select Client Protocols

Select protocol(s) for the client you want to add.

☑ Fibre Channel

Fibre Channel clients communicate with Virtual Tape Library server via the Fibre Channel protocol.

☐ iSCSI

iSCSI clients communicate with Virtual Tape Library server via the iSCSI protocol.

Click <Next> to continue.

For Fibre Channel clients, click *Next* and select the *initiator* WWPN for the client. Initiator ports with a green dot are available; yellow dots indicate that the port is already assigned to a client; red dots indicate that the port is not currently available. If FC initiator ports on the backup application server are already zoned with VTL's target port and are properly connected/powered up, they are listed automatically and you can select specific initiators from that zone. In addition, if there is only one initiator WWPN in the client, VTL will automatically select it for you and the dialog will not be displayed. If no WWPNs are listed, the backup application server is not currently zoned with the VTL appliance.

Click *Next* and set Fibre Channel options.

Enable Volume Set Addressing may be required for Fibre Channel clients, that require VSA to access storage devices.

Select *Enable Celerra Support* if you have a licensed EMC Celerra client.

Select *Enable IBM System i Support* if you have an IBM System i server.

For iSCSI clients, click *Next* and select the initiator that the client uses. If the initiator does not appear, you can manually add it. For additional details on adding and managing iSCSI clients, refer to the “iSCSI Configuration” chapter.

Click *Next* and add/select users who can authenticate for this client. When you add users, you will have to enter a name and password for each.

If you select *Allow Unauthenticated Access,* the VTL server will recognize the client as long as it has an authorized initiator name. With authenticated access, an additional check is added that requires the user to type in a user name and password. More than one user name/password pair can be assigned to the client, but they will only be useful when coming from the machine with an authorized initiator name.

4. Click *Finish* when you are done.

Rename a client

To create an alias name for a client, right-click the client and select *Rename*. The initial client name is preserved for compatibility.

Configure backup server access to the VTL server

VTL uses a “Secured Access” scheme, whereby access is dictated by creating specific clients to represent specific backup application servers. A backup application server can access *only* its own designated virtual tape library or drives via a dedicated port.

FC backup servers

In order for Fibre Channel backup servers to access VTL resources, you must do the following:

1. Set QLogic HBA ports to target mode.
2. Add a FC client for each backup application server.
3. Create and assign a virtual tape library to clients.
4. Discover the virtual tape library from your backup application server.

   Refer to ’Discover the virtual tape library from your backup server’ for more information.

   Additional information about steps 1-3 can be found in the ’Fibre Channel Configuration’ chapter.

iSCSI backup servers

In order for iSCSI backup application servers to access VTL resources, you must do the following:

1. Add an iSCSI client for each backup application server.
2. Create targets for the iSCSI client to log into.
3. Create and assign a virtual tape library to the iSCSI target.
4. Register client initiators with your VTL server.
5. Log the client onto the target.
6. Discover the virtual tape library from your backup application server.

   Refer to ’Discover the virtual tape library from your backup server’ for more information.

   Additional information about steps 1-5 can be found in the ’iSCSI Configuration’ chapter.

Discover the virtual tape library from your backup server

To enable your backup server to recognize the default virtual tape library and drives, perform a device scan on your backup application server at the operating system level and then use your backup software to scan for new devices as well.

Use your operating system to scan for hardware changes

The steps to do this vary according to the backup application server's operating system.

For Fibre Channel environments, if your zoning has been correctly configured, and devices have been properly assigned to clients, a simple bus rescan performed on the client should show the new backup devices. Of course, this procedure varies depending on the OS.

Windows

To discover a tape library on a backup application server running a Windows operating system:

1. Select *Control Panel --> Administrative Tools --> Computer Management.*
2. In the left pane, under *System Tools*, select *Device Manager.*
3. In the right pane, right-click the backup application server and select *Scan for hardware changes.*

   New devices representing the specific VTL resources will appear (the library under *Medium Changers* and tape drives under *Tape Drives*) and if the appropriate tape drive and tape library device drivers are installed on the backup application server, the correct device name and type are associated and the devices will become ready for use by the backup software.

   If a new device is unknown, right-click it to display its *Properties.* Acquire and update the driver according to your Windows documentation. Your backup software may include a procedure that updates drivers.

Linux

To discover a tape library on a backup application server running a Linux operating system:

1. Rescan your host adapter.

   Rescanning in Linux is host adapter-specific. For QLogic:

   ```
   echo "scsi-qlascan" > /proc/scsi/qla<model no>/<adapter-instance>
   ```

   For Emulex:

   ```
   sh force_lpfc_scan.sh "lpfc<adapter-instance>"
   ```

2. Identify the detected devices.

   ```
   # cat /proc/scsi/scsi
   ```

3. For each identified device do the following:

```
# echo "scsi add-single-device <host> <channel> <id> <lun>" >/proc/scsi/scsi
```

where <*host*> is the host adapter number, <*channel*> is channel number <*id*> is the target id and <*lun*> is the LUN number.

AIX

To discover a tape library on a backup application server running AIX:

1. Rescan devices.

   ```
   # cfgmgr -vl fcsX
   ```

   where *X* is the number of the FC adapter.

2. Verify the new devices.

   ```
   # lsdev -Cc <disk|tape>
   ```

Solaris

1. Determine the FC channels.

   ```
   # cfgadm -al
   ```

2. Force a rescan.

   ```
   cfgadm -o force_update -c configure cX
   ```

   where *X* is the FC channel number.

3. Install device files.

   ```
   # devfsadm
   ```

Use backup software to detect new devices

The steps to do this vary according to your backup software.

After you complete the procedure, you are ready to create and run backup jobs.

> **Note:** For all other platforms, such as Unix and Linux, consult the appropriate reference material that came with your backup software for details on how to load drivers and how to perform discovery for hardware changes.

Create and run backup jobs

Once your backup application server software can discover and access the virtual tape library/drives defined in the VTL server, you can start to use the VTL as if it were a real physical tape library.

The preparation required to start a backup job successfully is identical whether you are using a real tape library or a virtual one. You simply configure the backup software to use the VTL just like you would a physical tape library.

Generally, in order to perform a backup to a newly acquired/configured tape library, you need to:

1. Add new tape media.
   - Real library: Buy new tapes and insert into the mail slot followed by a sequence of keys pressed on the keypad of the tape library.
   - VTL: Virtual tapes are typically created when you create a virtual tape library. Additional virtual tapes can be created as needed.
2. Start a "tape inventory" process in your backup software.
3. Format the tapes and assign them into various "tape pools".
4. Define backup jobs and associate tapes with each job.

   When one or more backup jobs start to kick-off, tapes are allocated by the backup software and are loaded into the tape drives. Backup data is then sent to the tapes until the backup job is done. The backup software then sends commands to unload the tapes and return them to their assigned slot within the library. All of the above actions are emulated by VTL.

   When it is time to remove a tape from a physical library and to store it onto a nearby tape shelf, the administrator must physically walk over to the library, use a key pad/console to select the tape to be removed, and then catch the tape as it is physically being ejected from the "mail slot". The above can sometimes be done via commands from within the backup software.

   For a VTL server, obviously there is no keypad or physical mail slot for this purpose. However, the DSI VTL server has a *Virtual Tape Vault* to hold all the virtually "ejected" tapes from any virtual tape library. In the case where an "eject" is performed by the backup software, the ejected virtual tape will be automatically placed in the Virtual Tape Vault. This can be confirmed using the VTL Console (select the *Virtual Tape Vault* object and verify the virtual tape is indeed there). If tape removal is not done using the backup software, the equivalent of a "keypad" is to use the VTL console and right-click the virtual tape and select *Move to Vault*.

   Typically, after the backup is complete, the backup software will automatically remove the tape from the drive and store it back in its assigned library slot. When the deduplication policy executes at the designated time, or when you click *Run* manually from the console for the selected policy, you can use the console to confirm that the deduplication policy is running. To do this, highlight the *Deduplication Job Queue* tab pane to see a list of tapes currently being processed.

Special note for NetBackup users: To prevent a backup from going to the same tape more than once, when you are configuring backup jobs for Microsoft Exchange, DO NOT span your policies across tapes.

Confirm successful backups

While a backup job is running, you can use the VTL console to verify that data is being written to virtual tapes.

1. In the VTL console, expand the *Virtual Tape Library System* object.
2. Expand *Virtual Tape Libraries*, the specific library, and then *Tapes*.
3. Under the *Tapes* object, select each tape that is included in a backup job.

   In the right-hand pane, you should see a value for *Data Written*, which updates dynamically during a backup job.

After the backup job completes, use your backup software to verify that the data was written completely and can be restored.

NAS Configuration

Now that your appliance has been configured, you can configure NAS to prepare for deduplication. Afterward, you should refer to ’Deduplication’ to manage and monitor NAS deduplication.

> **Note:** If you did not enable NAS in the configuration wizard, you can right-click the server and select *Options* --> *Enable NAS*.

Set the security mode for Windows clients

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

- User mode - Authentication is controlled by passwords that are set for each Windows user. Refer to ’CIFS clients’ for information about adding clients for user mode.
- Domain mode - Authentication is controlled by a Windows Active Directory Domain Controller. The VTL server and all clients must belong to the domain controlled by this Domain Controller.

By default, NAS uses user mode. If you want to use domain mode, you will need to follow the instructions in ’Domain mode configuration’ below.

> **Note:** It is important that you do not change your authentication mode once you begin using your VTL system. If you do change it, you will lose all of your share assignments.

User mode configuration

To configure user mode:

1. Right-click *CIFS Clients* and select *Set Security Mode*.
2. Select *User Mode.*

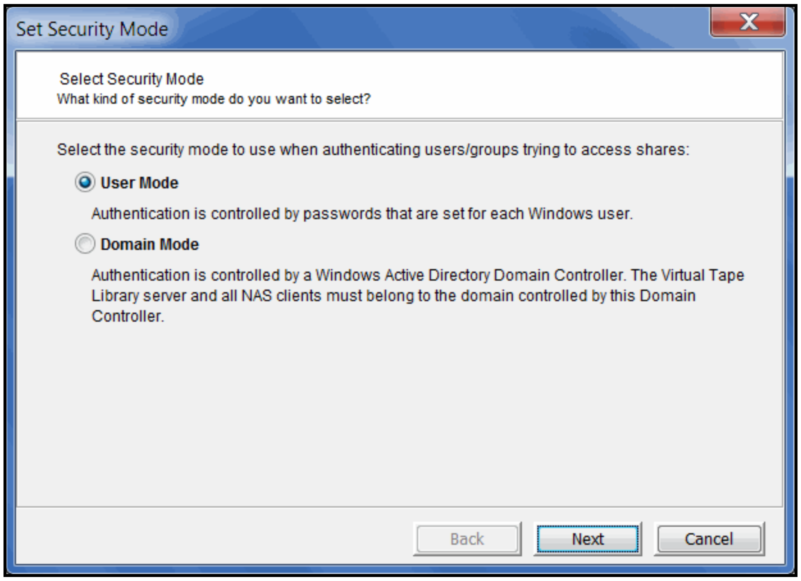

Select Security Mode
What kind of security mode do you want to select?

Select the security mode to use when authenticating users/groups trying to access 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 Domain Controller. The Virtual Tape Library server and all NAS clients must belong to the domain controlled by this Domain Controller.

3. Specify the workgroup that this server must join.

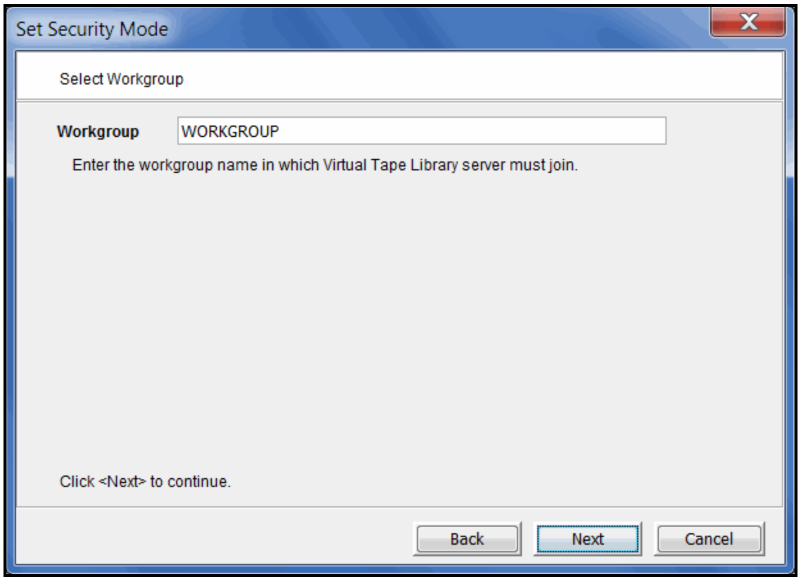

Select Workgroup

Workgroup WORKGROUP

Enter the workgroup name in which Virtual Tape Library server must join.

Click <Next> to continue.

4. If desired, specify a comment for the server.

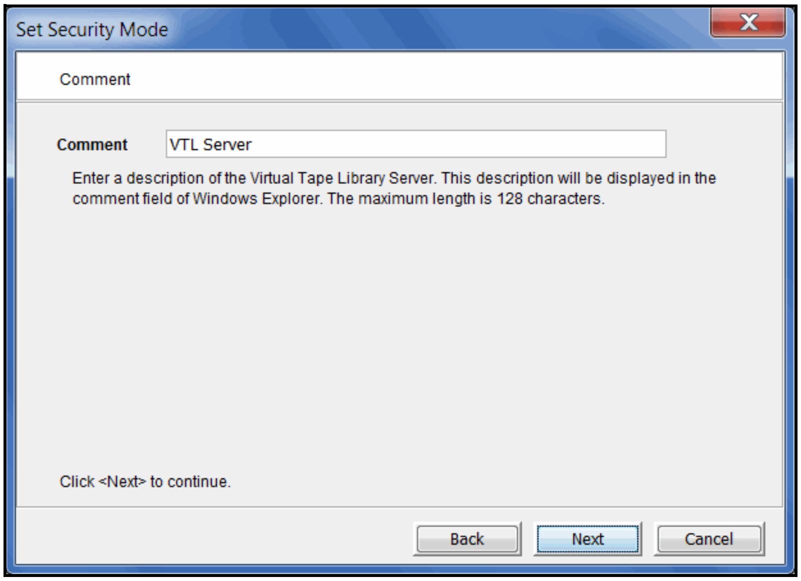

Comment

Comment VTL Server

Enter a description of the Virtual Tape Library Server. This description will be displayed in the comment field of Windows Explorer. The maximum length is 128 characters.

Click <Next> to continue.

This description of the server will be displayed in the *Comment* field of Windows Explorer, such as when you see a list of computers under *My Network Places*.

5. Reserve a range of User IDs (UIDs) for NAS users.

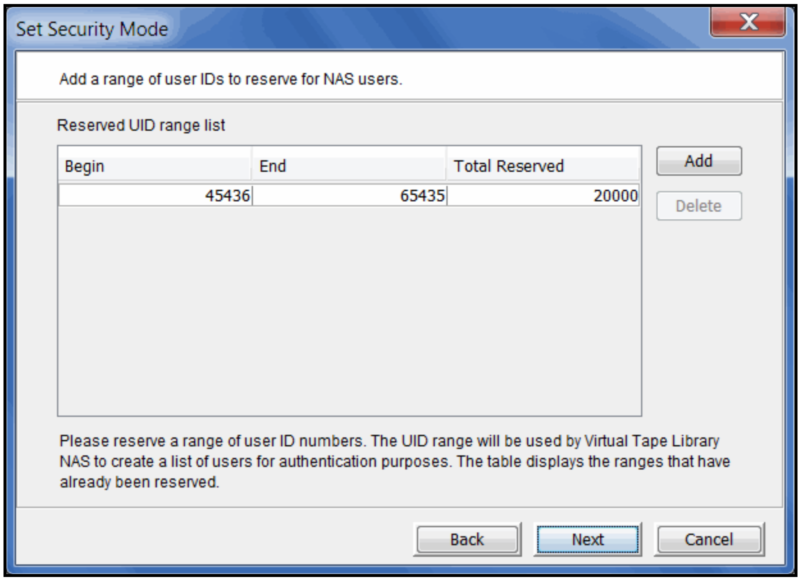

Add a range of user IDs to reserve for NAS users.

Reserved UID range list

Begin	End	Total Reserved
45436	65435	20000

Please reserve a range of user ID numbers. The UID range will be used by Virtual Tape Library NAS to create a list of users for authentication purposes. The table displays the ranges that have already been reserved.

UIDs are associated with users on your system.

Be sure to select a big enough range to handle the number of users you have.

6. Reserve a range of Group IDs (GIDs) for NAS groups.

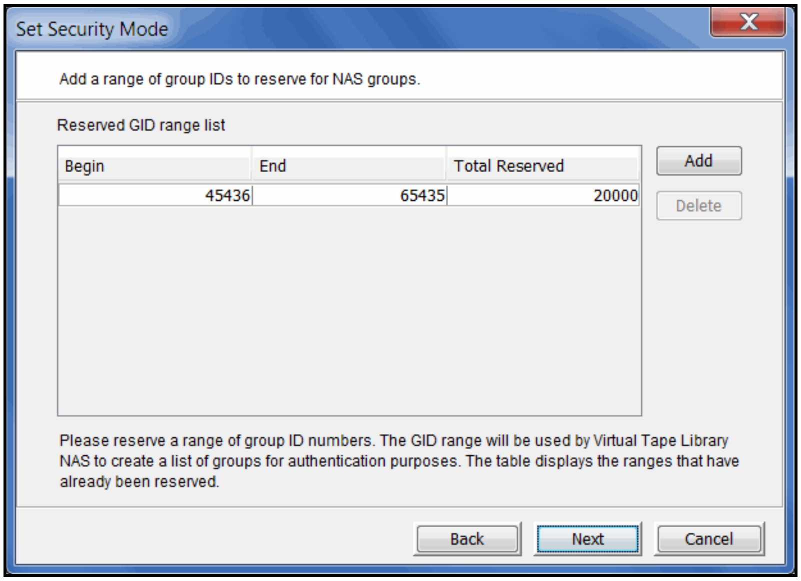

Set Security Mode

Add a range of group IDs to reserve for NAS groups.

Reserved GID range list

Begin	End	Total Reserved
45436	65435	20000

Please reserve a range of group ID numbers. The GID range will be used by Virtual Tape Library NAS to create a list of groups for authentication purposes. The table displays the ranges that have already been reserved.

GIDs are associated with groups on your system.

Be sure to select a big enough range to handle the number of groups you have.

7. Confirm all information and click *Finish*.

Domain mode configuration

You must do the following if you will be using domain mode for authentication:

Synchronize clocks

Your VTL server and your Windows domain controller must have their clocks synchronized to within five minutes of each other. If they are not synchronized, you can use the *date* command on your VTL server to adjust the date and time. However, the system clock on a PC can "drift" over time. Therefore, we recommend that you use an automated synchronization service to adjust the system's clock. Refer to the *ntpd* service on Linux and the *Windows Time* service on Windows for more information. You can add a time server through the console (right-click the server and select *System Maintenance* --> *Configure Network*). If no time server is available, you can use the IP address of the domain controller (DC). You need to `edit /etc/ntp.conf` on the VTL server and add `'server [DC IP address]'`. Also, edit `/etc/ntp/step-tickers` and add just the domain controller IP address.

Server is resolvable by DNS

Make sure each VTL server has a valid DNS entry created in the Microsoft DNS (part of the Active Directory you plan to join) with a valid PTR resource record (DNS reverse lookup zone must contain this record).

If your server's hostname is resolvable by a DNS server, you should have configured the DNS server name and IP during VTL installation. You should check the files /etc/hosts and /etc/resolv.conf to make sure the DNS server is configured correctly. Refer to 'Update /etc/hosts' below if you need to update it manually.

If the DNS server was not configured, you can add it manually by right-clicking on your server in the console and selecting *System Maintenance* --> *Configure Network.* Otherwise, refer to 'Server is not resolvable' below.

Server is not resolvable

If your server's hostname is not resolvable by a DNS machine, you need to manually add it to /etc/hosts. Refer to 'Update /etc/hosts' below for details.

Update /etc/hosts

If your server's hostname is not resolvable by a DNS machine, you need to manually add it to /etc/hosts. Even if your server is resolvable, we recommend checking /etc/hosts file and updating, if necessary. For example, your environment has the following configuration:

- VTL server name: *vtl-server*
- VTL server IP address: *192.168.15.151*
- Primary Authentication Server: *windows-domain*
- Primary Authentication Server IP address: *192.168.15.145*
- Domain being joined: *abc.net*

1. Use the command `vi /etc/hosts` to edit */etc/hosts* on your VTL server.
2. Add the following lines to */etc/hosts*:

```
192.168.15.151 vtl-server.abc.net vtl-server
192.168.15.145 windows-domain.abc.net windows-domain
```

If `vtl-server.abc.net` or `vtl-server` are already part of another entry, such as:

```
127.0.0.1            vtl-server     localhost.localdomain localhost
```

Remove the vtl-server entry from that line and add the other lines:

```
127.0.0.1 localhost.localdomain localhost
192.168.15.151 vtl-server.abc.net vtl-server
192.168.15.145 windows-domain.abc.net windows-domain
```

Active Directory

Your VTL server can be configured to utilize Microsoft’s Active Directory to obtain users and groups. If you will be using Active Directory, you will need the following:

- Account for VTL - This account should have minimal security, similar to that of the *guest* account. The account will be used by VTL to access the active directory that VTL will browse to identify the users/groups that will have access to NAS shares. For a more secure account, you can limit this account to have *read-only access* to the Organizational Units (OUs) that will be browsed by VTL.
- Your VTL server and your Active Directory Server must have their clocks synchronized to within five minutes of each other. We recommend that you use an automated synchronization service (NTP server) to adjust the system’s clock. Refer to the *ntpd* service on Linux and the *Windows Time* service on Windows for more information.

Select domain mode as the security mode

To select domain mode:

1. Right-click *CIFS Clients* and select *Set Security Mode*.
2. Select *Domain Mode* and indicate if you are using Active Directory.

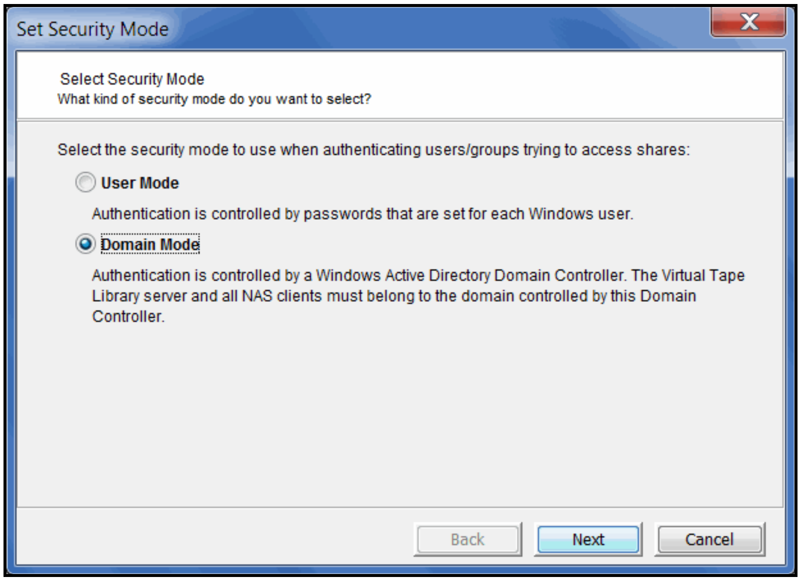

Set Security Mode

Select Security Mode
What kind of security mode do you want to select?

Select the security mode to use when authenticating users/groups trying to access 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 Domain Controller. The Virtual Tape Library server and all NAS clients must belong to the domain controlled by this Domain Controller.

3. Enter your domain controller name.

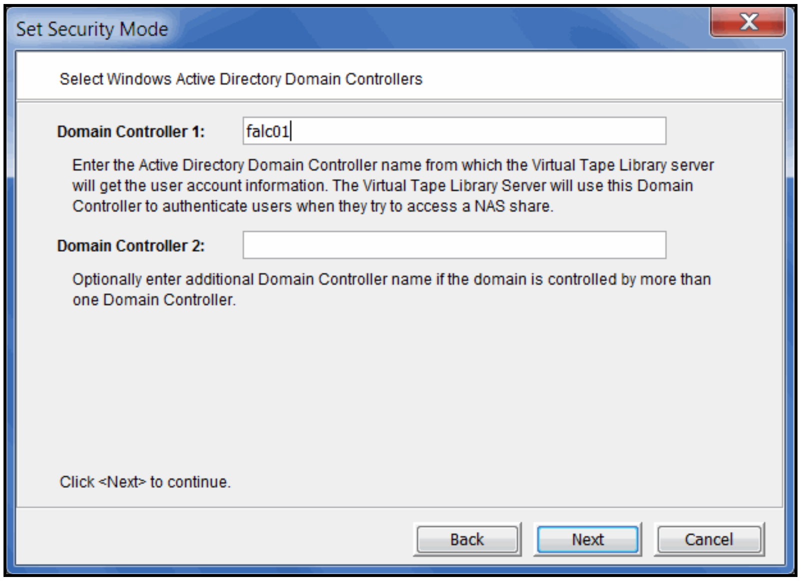

Set Security Mode

Select Windows Active Directory Domain Controllers

Domain Controller 1: falc01

Enter the Active Directory Domain Controller name from which the Virtual Tape Library server will get the user account information. The Virtual Tape Library Server will use this Domain Controller to authenticate users when they try to access a NAS share.

Domain Controller 2:

Optionally enter additional Domain Controller name if the domain is controlled by more than one Domain Controller.

Click <Next> to continue.

Domain Controller 1 - Enter the name of the server (short name, not an IP address or fully qualified domain name) from which the VTL server will get the user account information. The VTL server will use this server to authenticate users when they try to share a NAS file system. The server's name must be resolvable.

Domain Controller 2 - You can optionally enter a server name (short name, not an IP address or fully qualified domain name) to use for authentication if the primary authentication server is not available. The server's name must be resolvable.

If you need to change the domain controller in the future, right-click *CIFS Clients* and select *Update Domain Controllers.*

4. Enter information about the account VTL will use to log into Active Directory.

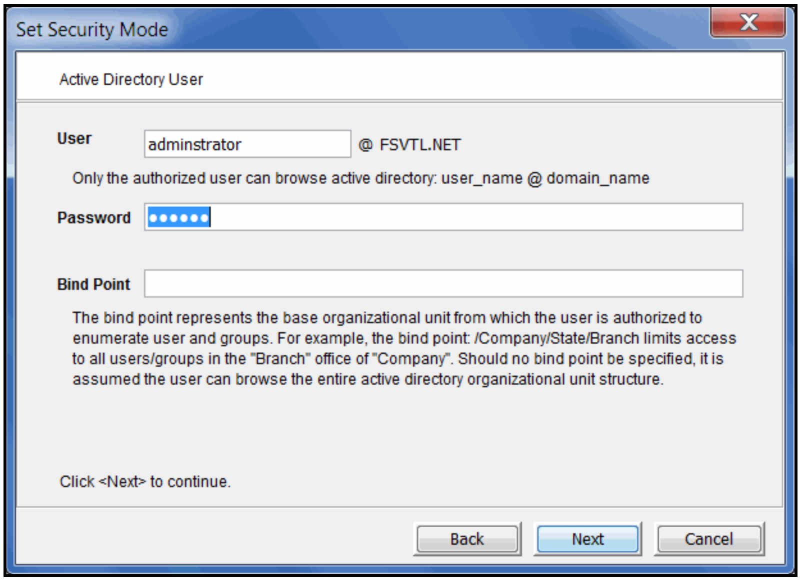

User adminstrator @ FSVTL.NET

Only the authorized user can browse active directory: user_name @ domain_name

Password ●●●●●●

Bind Point

The bind point represents the base organizational unit from which the user is authorized to enumerate user and groups. For example, the bind point: /Company/State/Branch limits access to all users/groups in the "Branch" office of "Company". Should no bind point be specified, it is assumed the user can browse the entire active directory organizational unit structure.

Click <Next> to continue.

User - Enter the account VTL will use to log into Active Directory.

Password - Enter a valid password for this account.

Bind Point - You can use the *Bind Point* to mark from where in the OU tree VTL will start reading OUs. This is useful if VTL's domain login account does not have access to the entire OU tree. Without this access, VTL cannot see anything in the tree. In this case, enter a *Bind Point* to direct VTL to a starting point or a single tree such as the /Engineering or /Accounting tree. If you leave this field blank or enter "/", VTL will start at the root of the OU tree.

Notes:

- If you see the message "*Unable to connect to active directory due to excessive clock skew. Please synchronize server and active directory clocks.*" when you click *Next*, the clocks on your VTL server, Windows Domain Controller, and your Active Directory Server are not synchronized to within five minutes of each other. Use the *date* command or an automated synchronization service to adjust the system's clock.
- If you see the message "*Failed to locate the authentication server*" when you click *Next*, the Windows primary authentication server name and IP address are not in the /etc/hosts file or are not resolvable.

5. Select the organizational units to which you will offer NAS shares.

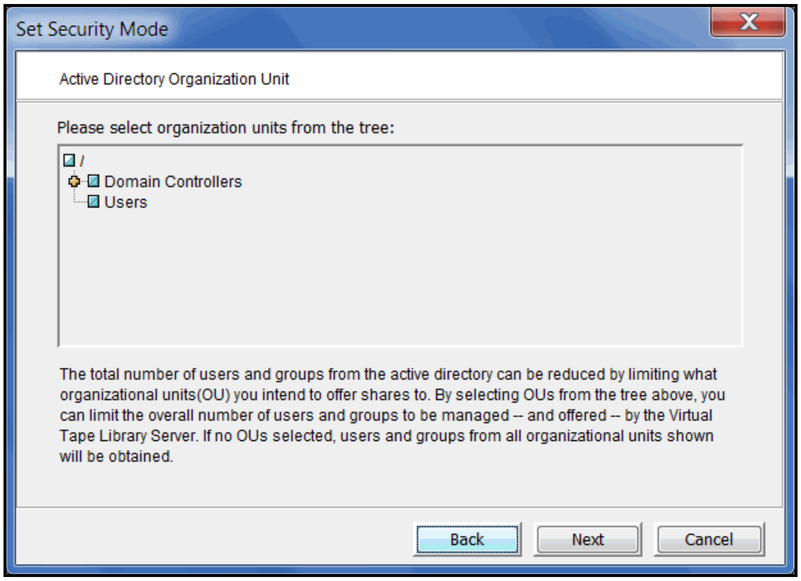

Set Security Mode
Active Directory Organization Unit
Please select organization units from the tree:
/
Domain Controllers
Users
Back
Next
Cancel

The total number of users and groups from the active directory can be reduced by limiting what organizational units(OU) you intend to offer shares to. By selecting OUs from the tree above, you can limit the overall number of users and groups to be managed -- and offered -- by the Virtual Tape Library Server. If no OUs selected, users and groups from all organizational units shown will be obtained.

Click in the checkbox next to the OUs to which you want to offer NAS shares.

If you specified a bind point in the previous dialog, 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.

By default, everything is selected. If you click the checkbox next to the OU, it will de-select that OU. In order to select only certain OUs, you should click the root first, then select the OUs to which you want to offer NAS shares.

6. Enter a descriptive comment.

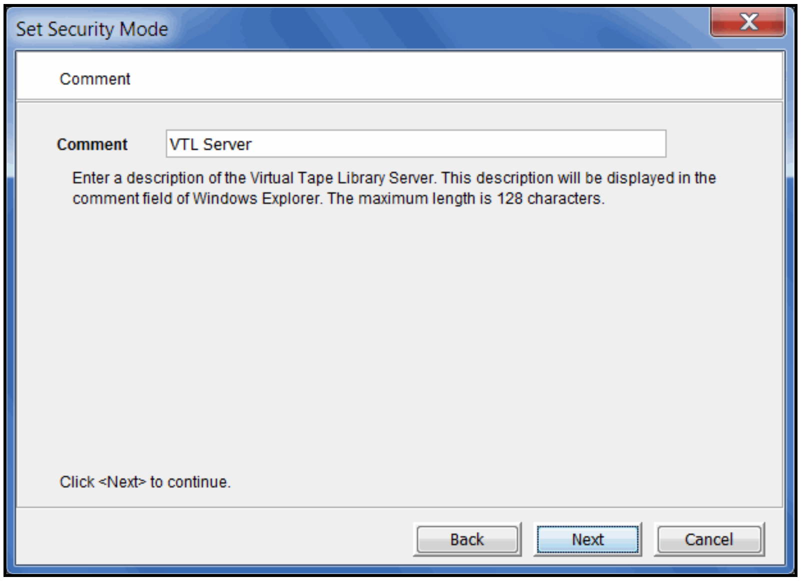

Comment VTL Server

Enter a description of the Virtual Tape Library Server. This description will be displayed in the comment field of Windows Explorer. The maximum length is 128 characters.

Click <Next> to continue.

This description of the VTL server will be displayed in the *Comment* field of Windows Explorer, such as when you see a list of computers under *My Network Places*.

7. Reserve a range of User IDs (UIDs) for NAS users.

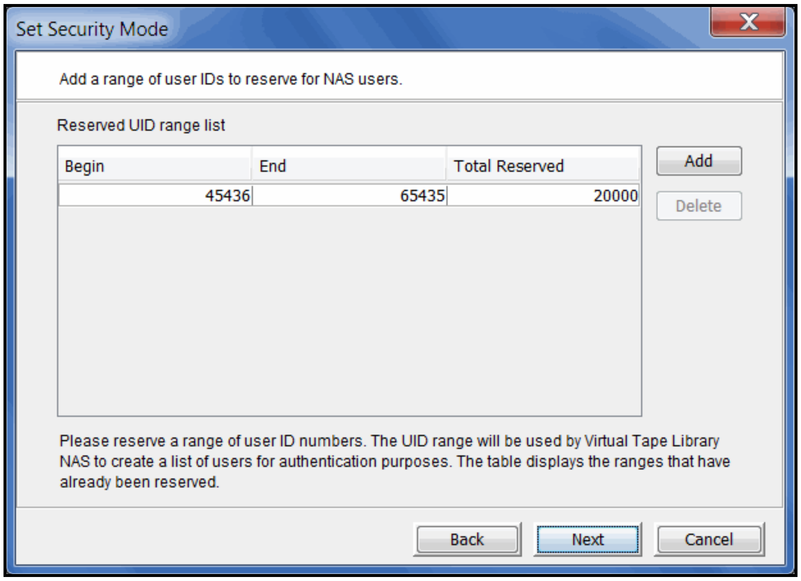

Begin	End	Total Reserved
45436	65435	20000

Please reserve a range of user ID numbers. The UID range will be used by Virtual Tape Library NAS to create a list of users for authentication purposes. The table displays the ranges that have already been reserved.

UIDs are associated with users on your system.

Be sure to select a big enough range to handle the number of users you have.

8. Reserve a range of Group IDs (GIDs) for NAS groups.

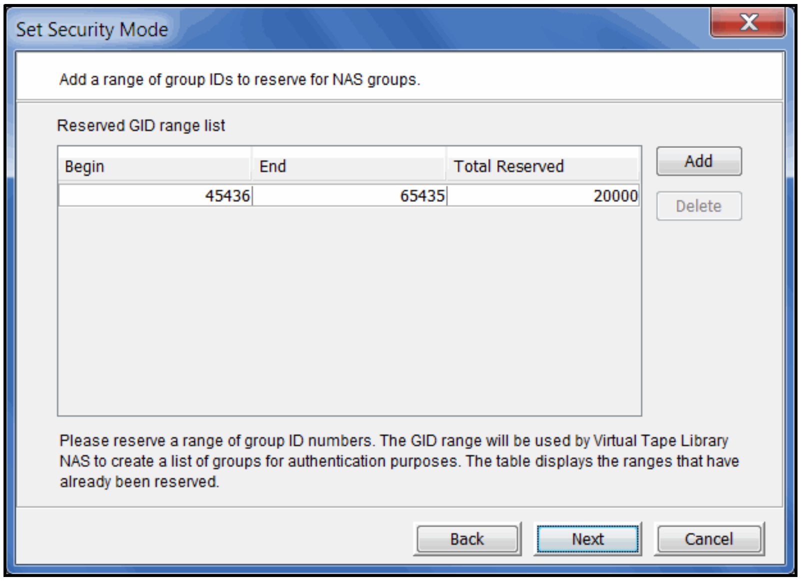

Begin	End	Total Reserved
45436	65435	20000

Please reserve a range of group ID numbers. The GID range will be used by Virtual Tape Library NAS to create a list of groups for authentication purposes. The table displays the ranges that have already been reserved.

GIDs are associated with groups on your system.

Be sure to select a big enough range to handle the number of groups you have.

9. Enter an administrative user name and password.

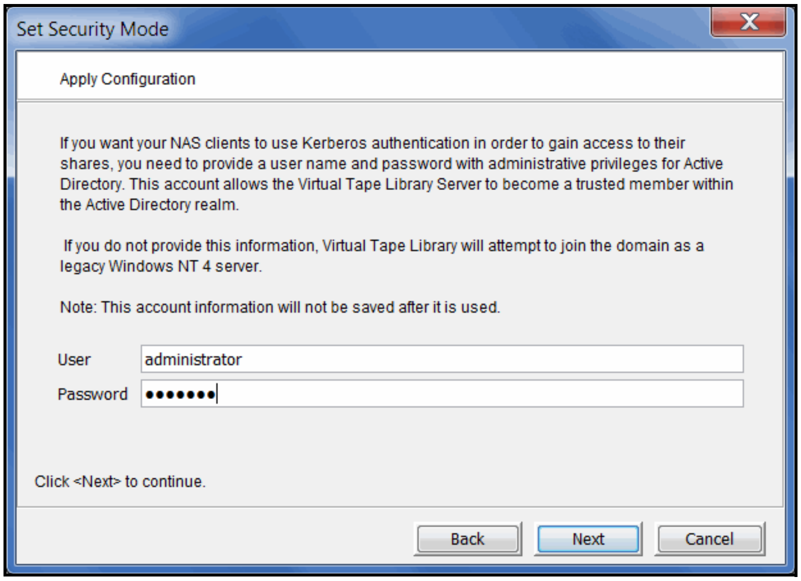

Apply Configuration

If you want your NAS clients to use Kerberos authentication in order to gain access to their shares, you need to provide a user name and password with administrative privileges for Active Directory. This account allows the Virtual Tape Library Server to become a trusted member within the Active Directory realm.

If you do not provide this information, Virtual Tape Library will attempt to join the domain as a legacy Windows NT 4 server.

Note: This account information will not be saved after it is used.

User administrator

Password •••••••

Click <Next> to continue.

This name/password will be used to create a computer account for the VTL server in the domain.

If you are using Active Directory, the Console will try to join the domain as an Active Directory Member Server even if no username/password is supplied. If that fails, VTL will try to join the domain as a legacy Windows NT 4 server.

10. Confirm all information and click *Finish*.

Access Control Lists (ACLs)

ACLs are specific Windows permissions that can be set on files and folders when using Domain mode with Active Directory. Instead of assigning CIFS clients permissions at the share level, ACLs allow the permissions to be applied to the files and directories beneath the share. VTL currently supports POSIX ACLs.

To enable ACL support and specify administrators who can set ACLs:

1. Right-click the *CIFS Clients* object and select *Enable ACL Support*.

2. Select the users who will administrate ACL permissions.

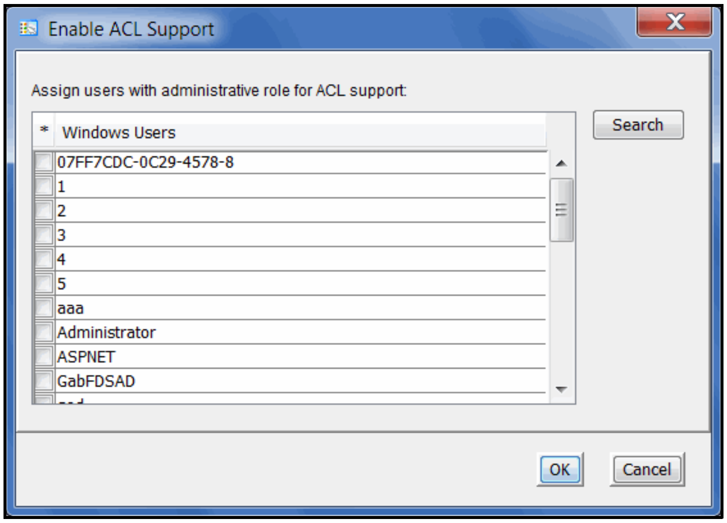

Assign users with administrative role for ACL support:

*	Windows Users
□	07FF7CDC-0C29-4578-8
□	1
□	2
□	3
□	4
□	5
□	aaa
□	Administrator
□	ASPNET
□	GabFDSAD

3. Type “yes” to remount all NAS resources.

   Remounting may take a while, depending upon the size of the NAS resource(s) being created. Larger NAS resources take significantly longer to remount than smaller ones.

4. Configure ACL permissions from the Windows machine of each user selected in step 2.

Configuring ACL attributes

In the following example, you have one share named “Data” and two users, UserA and UserB. You want both users to have full access to a common sub-directory called “Everybody” and you want each user to have full access to his/her own directory. These are the steps you would take to accomplish this:

1. Create a share named "Data" and assign the admin user (set above), UserA, and UserB to the share.

2. As the admin user, go to your Windows Explorer and map the share.

   For more information about mapping a share, refer to ’Map/mount shares’.

3. Modify the security of the base share by selecting *Properties* --> *Advanced* --> *Change Permissions* from Windows Explorer. Select 'nasgrp', click *Edit*, and remove all privileges from ‘nasgrp’. Then, check *Allow* for the *Read* permissions box only.

   ‘nasgrp’ is a group on the server side created by VTL. It contains all of the Windows users that the VTL server can see.

4. Create three directories at the root of “Data”: "UserADirectory", "UserBDirectory", and "Everybody".

5. Right-click on the "UserADirectory" directory and select *Properties* --> *Security*.

6. Add users UserA and UserB.

By default, the newly added users will only have read access.

7. To give write access, select UserA and check *Allow* for the *Full Control* box.
8. Apply these same steps for the "UserBDirectory" and "Everybody" directories.

   For the "UserBDirectory" directory, give UserB *Full Control*.

   For the "Everybody" directory, give both UserA and UserB *Full Control*.

   As a result, when UserA or UserB maps to the "Data" share, each user will have both read and write access to his/her own directory and the "Everybody" directory, but only read access to each other’s directory.

Add clients

Clients are application servers that use a NAS share to store data. There are two types of NAS clients:

- NFS clients - These are usually Unix machines using the NFS protocol.
- CIFS clients - These are Windows users and groups that use the CIFS protocol.

> **Note:** All data should be backed up/archived to VTL via the NAS clients. You should not manually connect to the VTL server locally (via SSH, telnet, etc.) to change a file or file attributes directly. If you do, the system may be unable to deduplicate or replicate the affected files.

NFS clients

To add NFS clients:

1. Expand *Clients*, right-click the *NFS Clients* object, and select *Add*.

   You can also access this dialog and add NFS clients while creating a NAS share.

2. Enter information as applicable.

   Name - This is the name displayed in the console for this group of one or more NFS clients. For example, you may want to enter *NetBackup* to identify these clients.

   Machine - The machine that will become an NFS client. You can enter an abbreviated name that can be resolved, a fully qualified domain name, or an IP address for a machine. You can also enter a subnet and netmask if you want to grant NAS share access to an entire subnet. The format to do this is: `subnet/netmask` (for example: 192.168.0.0/255.255.255.0). Wildcard characters are not supported in machine names or domain names.

   Comment - You can optionally enter a description or explanation in this field. This information will be displayed in the right pane of the console for this client.

   > **Note:** You should verify that the hostnames listed in the /etc/hosts files on both the VTL server and NFS clients (Linux, Solaris, IBM AIX) are lowercase only. This is because, when a share is accessed via NFS, the kernel converts hostnames to lower case, while NFS uses hostnames as typed (i.e. lowercase, uppercase, combo, or both). This can result in errors on the NFS clients.

CIFS clients

If you are using *User* authentication mode, you need to add Windows users who can access shares. Once you add your clients, they can be assigned to shares.

For the sake of convenience, you can also create groups of users who can then be assigned to shares.

You cannot add CIFS clients if you are using *Domain* authentication mode. In *Domain* authentication mode, the list of users comes from the authentication server.

To add Windows clients:

1. Expand *Clients*, right-click the *CIFS Clients* object, and select *User Accounts*.
2. Click the *Add* button.

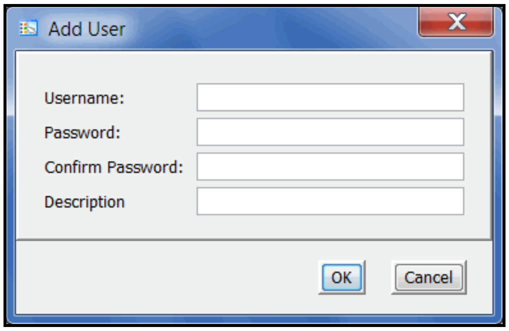

3. Enter a username and password and click *OK* to save.

   You cannot use the same user name that already exists for a VTL administrator.

Create a group

To create a group:

1. Select the *Groups* tab and click the *Add* button.

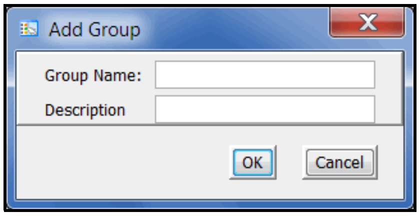

2. Enter a group name and click *OK* to save.
3. Click the *Membership* button and assign users to this group.

   Each user can belong to more than one group.

CIFS client properties

To set CIFS client properties or update the CIFS sharing settings, right-click *CIFS Clients* and select *Properties*.

- *General* – Specify a comment for the server.
- *Reserved UID/Reserved GID* – Reserve available user ID and group ID range(s) for use by CIFS clients.
- *NetBIOS Alias* – Set a NetBIOS alias for a Samba server, which allows the server to have more than one NetBIOS name.
- *Admin Users* – (Domain Authentication mode only) Give a user admin rights by making the user root on the server.
- *Advanced* – Change default global Samba options. A few are described below. Refer to the Samba manual for information about the other settings.
    - *name resolve order* - Set 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 CIFS clients.
    - *root preexec* and *root postexec* - Changing the values will allow specified commands (such as common maintenance procedures) to be executed when clients connect to or disconnect from a service.

Monitor CIFS connections

By default, information is updated on the *Connection(s)*, *Active Share(s)*, and *Locked File(s)* tabs each time you highlight the *CIFS Clients* object.

If you want the information refreshed more frequently, you can set an interval to determine how frequently the console should refresh this information. To do this, right-click *CIFS Clients* and select *Monitor CIFS Connections.* You can set the interval to be between 0 (continuously refresh) and 99 seconds. The default is 10 seconds.

Create NAS resources

NAS resource file systems are virtualized disks that are used to create NAS shares. Note that in order to create NAS resources, at least one disk must be reserved for NAS resources. Refer to 'Prepare physical storage devices' for more information.

1. Click *Next* to accept the defaults or click *Advanced* to change the settings.

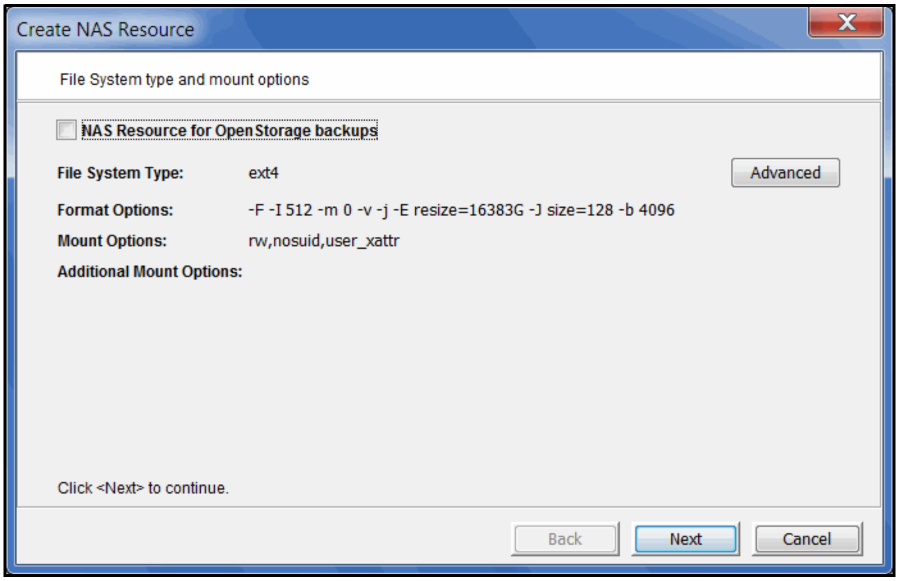

□ NAS Resource for OpenStorage backups

File System Type: ext4

Format Options: -F -I 512 -m 0 -v -j -E resize=16383G -J size=128 -b 4096

Mount Options: rw,nosuid,user_xattr

Additional Mount Options:

Click <Next> to continue.

NAS Resource for OpenStorage backups - In order to create a NAS resource for OpenStorage backups, you must enable OpenStorage in the console. To do this, right-click your VTL server and select *Options --> Enable OST*.

File System Type - VTL automatically detects the file system, *ext4*.

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

- `-F` - Forces the format regardless of what is on the drive.
- `-I 512` - Increases 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` - Preserves the maximum file system metadata space (16 TB) for later file system resize.
- `-J size=128` - Sets the journal size to 128 MB for resize purposes.
- `-b size=4096` - Block Size determines the minimum amount of space to use for each file. For example, if you keep the default of 4096, each file will minimally be 4k in size.

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

- `rw` - Allows read/write access.
- `nosuid` - Disallows set-user-id execution.

- `user_xattr` - Supports “user” extended attributes.

If you need to change the default *Format Options* or *Mount Options*, or specify additional mount options, click the *Advanced* button.

2. Select the physical device(s) from which to create this NAS resource and enter a name for the new NAS resource.

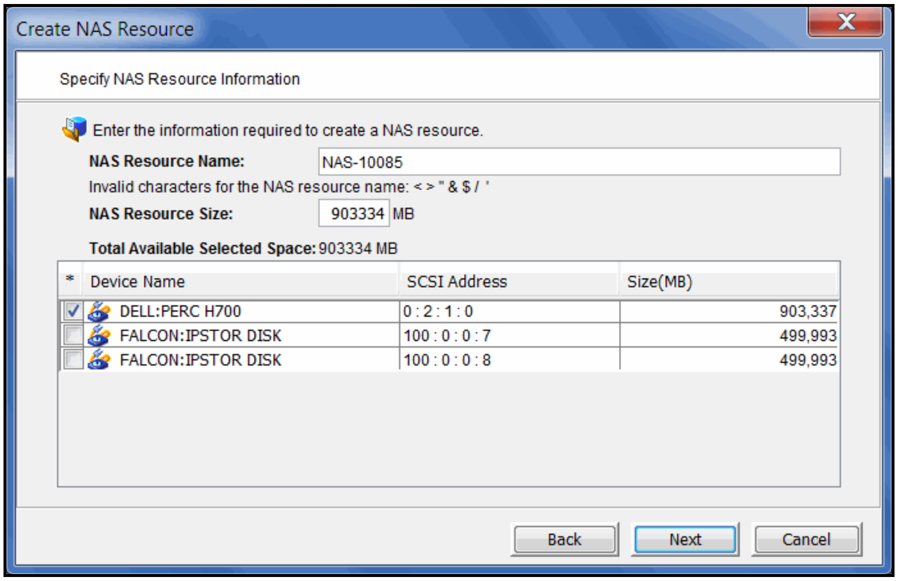

Create NAS Resource

Specify NAS Resource Information

Enter the information required to create a NAS resource.

NAS Resource Name: NAS-10085

Invalid characters for the NAS resource name: < > " & $ / '

NAS Resource Size: 903334 MB

Total Available Selected Space: 903334 MB

*	Device Name	SCSI Address	Size(MB)
☑	DELL:PERC H700	0 : 2 : 1 : 0	903,337
□	FALCON:IPSTOR DISK	100 : 0 : 0 : 7	499,993
□	FALCON:IPSTOR DISK	100 : 0 : 0 : 8	499,993

NAS resource names cannot use blanks or contain the following characters: < > / \ " % # : ; | * ? & $ ' ( ) `

3. Confirm that all information is correct and then click *Finish* to create the NAS resource.

This may take a while, depending upon the size of the NAS resource(s) being created. Larger NAS resources take significantly longer to format than smaller ones.

You should wait until the NAS resource is attached and mounted before creating shares.

Create shared folders

Create shared folders into which your backup software or other application can put files.

1. Right-click a NAS resource and select *New Share*.

2. Enter a folder name.

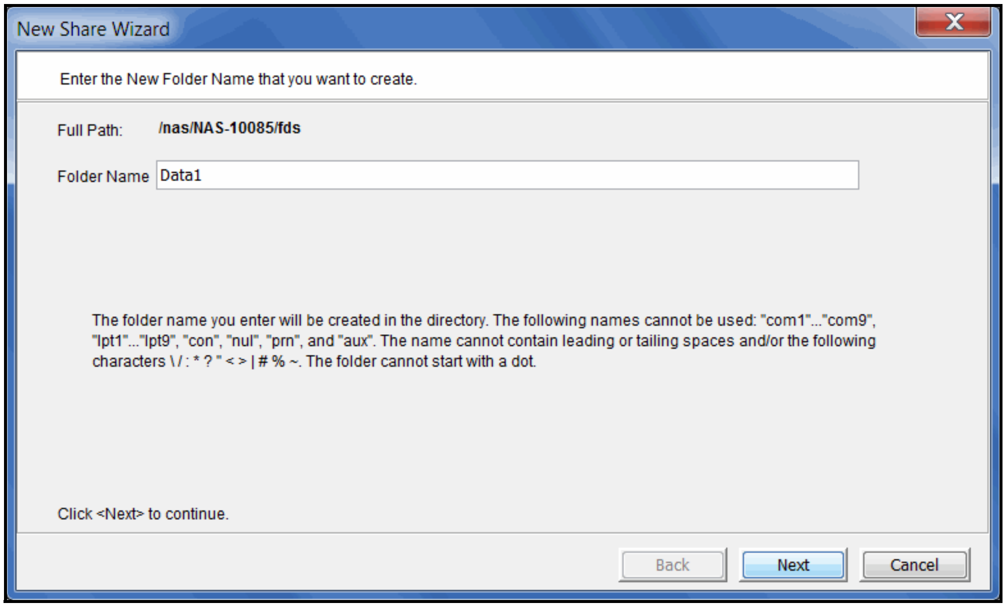

New Share Wizard

Enter the New Folder Name that you want to create.

Full Path: /nas/NAS-10085/fds

Folder Name: Data1

The folder name you enter will be created in the directory. The following names cannot be used: "com1"..."com9", "lpt1"..."lpt9", "con", "nul", "prn", and "aux". The name cannot contain leading or tailing spaces and/or the following characters \ / : * ? " < > | # % ~. The folder cannot start with a dot.

Click <Next> to continue.

The folder name cannot exceed 238 characters. If you will be assigning NFS clients to this share, the folder name cannot contain spaces.

3. If you want CIFS clients to have access to this share, enter a share name.

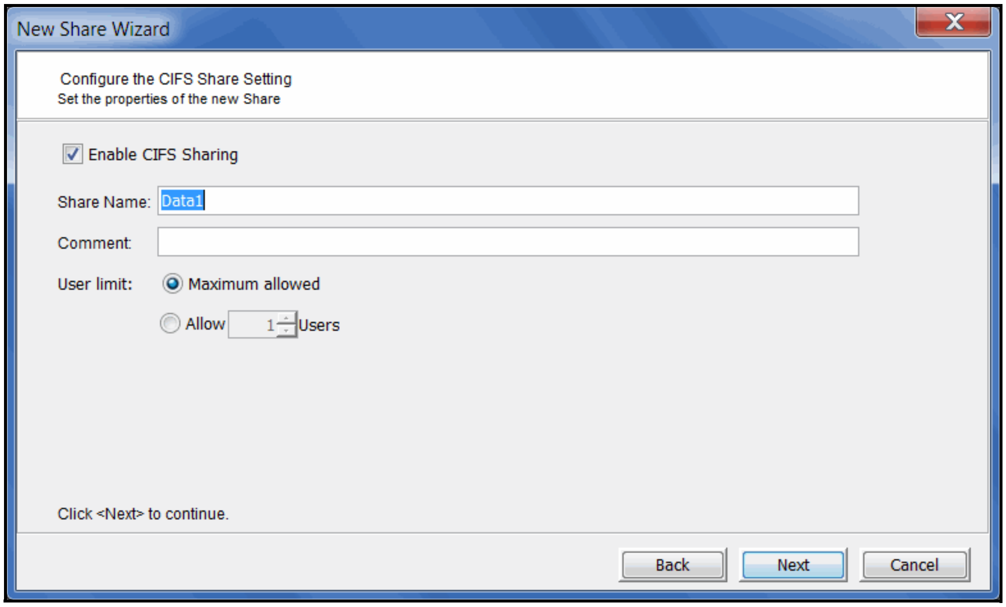

New Share Wizard

Configure the CIFS Share Setting
Set the properties of the new Share

☑ Enable CIFS Sharing

Share Name: Data1

Comment:

User limit: ● Maximum allowed

○ Allow 1 Users

Click <Next> to continue.

Share names cannot start with a dot or contain the following characters: < > / \ " % # : ; | * ? [ ] = + ,

4. (CIFS clients) Select which Windows users and groups can access the share and set access rights for each.

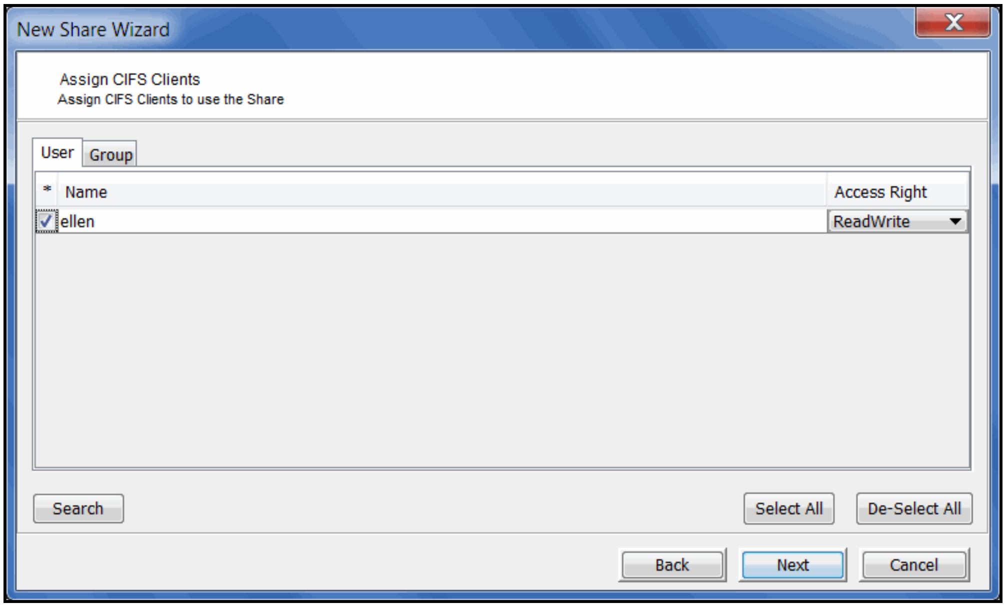

Assign CIFS Clients

Assign CIFS Clients to use the Share

*	Name	Access Right
☑	ellen	ReadWrite

5. (NFS clients) Enter permissions for the NFS clients who will access the share.

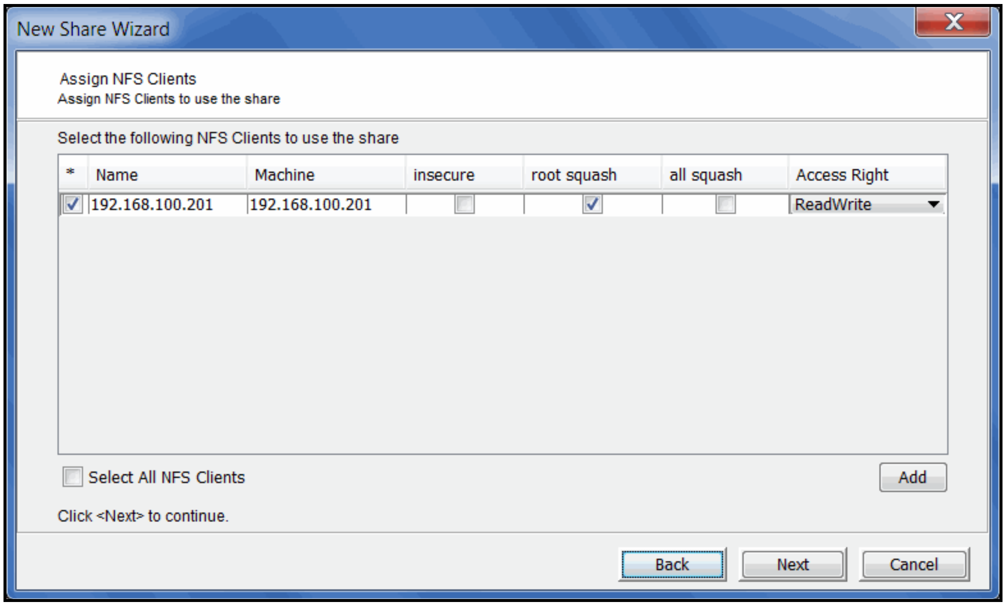

Assign NFS Clients

Assign NFS Clients to use the share

Select the following NFS Clients to use the share

*	Name	Machine	insecure	root squash	all squash	Access Right
☑	192.168.100.201	192.168.100.201	□	☑	□	ReadWrite

□ Select All NFS Clients

Click <Next> to continue.

You can click the *Add* button to add an NFS client. Refer to ’Map/mount shares’ for more information.

Select *insecure* if your client’s operating system does not use a reserved port for NFS (an Internet port less than IPPORT_RESERVED -- 1024). AIX is an example of an operating system that needs to select *insecure*.

For security purposes, *Squash* can be used to reduce the privileges of certain users by mapping user IDs to *nobody*.

root_squash	all_squash	Action
-	-	No UIDs mapped.
X	-	UID=0 (root user) is remapped to *nfsnobody:nasgrp* (default).
X	X	All UIDs are mapped to *nfsnobody:nasgrp*.

6. Click *Finish* to confirm.

Create folders

You can create folders that will be shared at a later time. Any time after creating a folder, you can assign clients to it by right-clicking and selecting *Sharing.*

1. Right-click a NAS resource and select *New Folder*.
2. Enter a folder name.

   The folder name cannot exceed 238 characters. If you will be assigning NFS clients to this share, the folder name cannot contain spaces.

Directory properties

In addition to assigning access rights to NAS clients, you can set access rights at the directory level.

1. Right-click the directory and select *Directory Properties*.

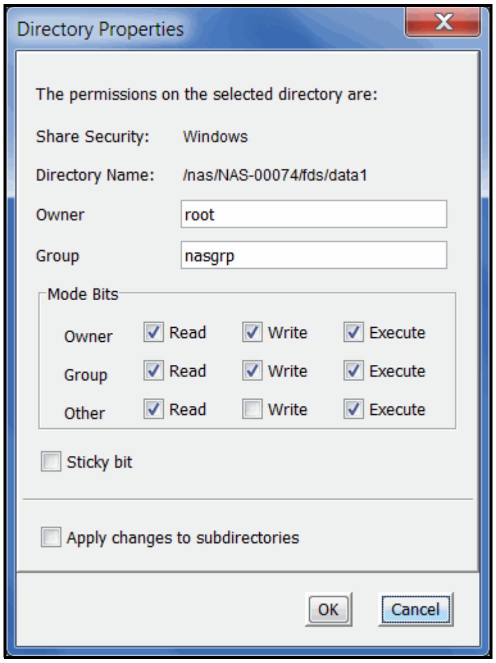

Directory Properties

The permissions on the selected directory are:

Share Security: Windows

Directory Name: /nas/NAS-00074/fds/data1

Owner: root

Group: nasgrp

Mode Bits

Owner	☑ Read	☑ Write	☑ Execute
Group	☑ Read	☑ Write	☑ Execute
Other	☑ Read	□ Write	☑ Execute

□ Sticky bit

□ Apply changes to subdirectories

2. Set the appropriate properties.

   You can set specific permissions for owners, groups, and others.

   Group must be a valid group on your VTL server.

   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.

Map/mount shares

Windows clients

Map a share

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 VTL 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

Mount a share

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 VTL 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 VTL35:/nas/NAS-00002/fds/Data1 /mnt/share
   ```

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 5.x

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

Solaris 9 and 10

```
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

```
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.

RMAN

Additional options are required by RMAN:

Operating System	Options
AIX 5.x	cio,bg,rsize=65536,wsize=65536,noac,timeo=600
Solaris 9 and 10	bg,rsize=32768,wsize=32768,noac,forcedirectio,suid
Linux	rsize=524288,wsize=524288,actimeo=0,timeo=600

Console

The DSI Management Console allows you to manage your VTL, NAS, SIR, and VTL-S servers.

Console user interface

The console displays the configuration for your servers. The information is organized in a familiar Explorer-like tree view.

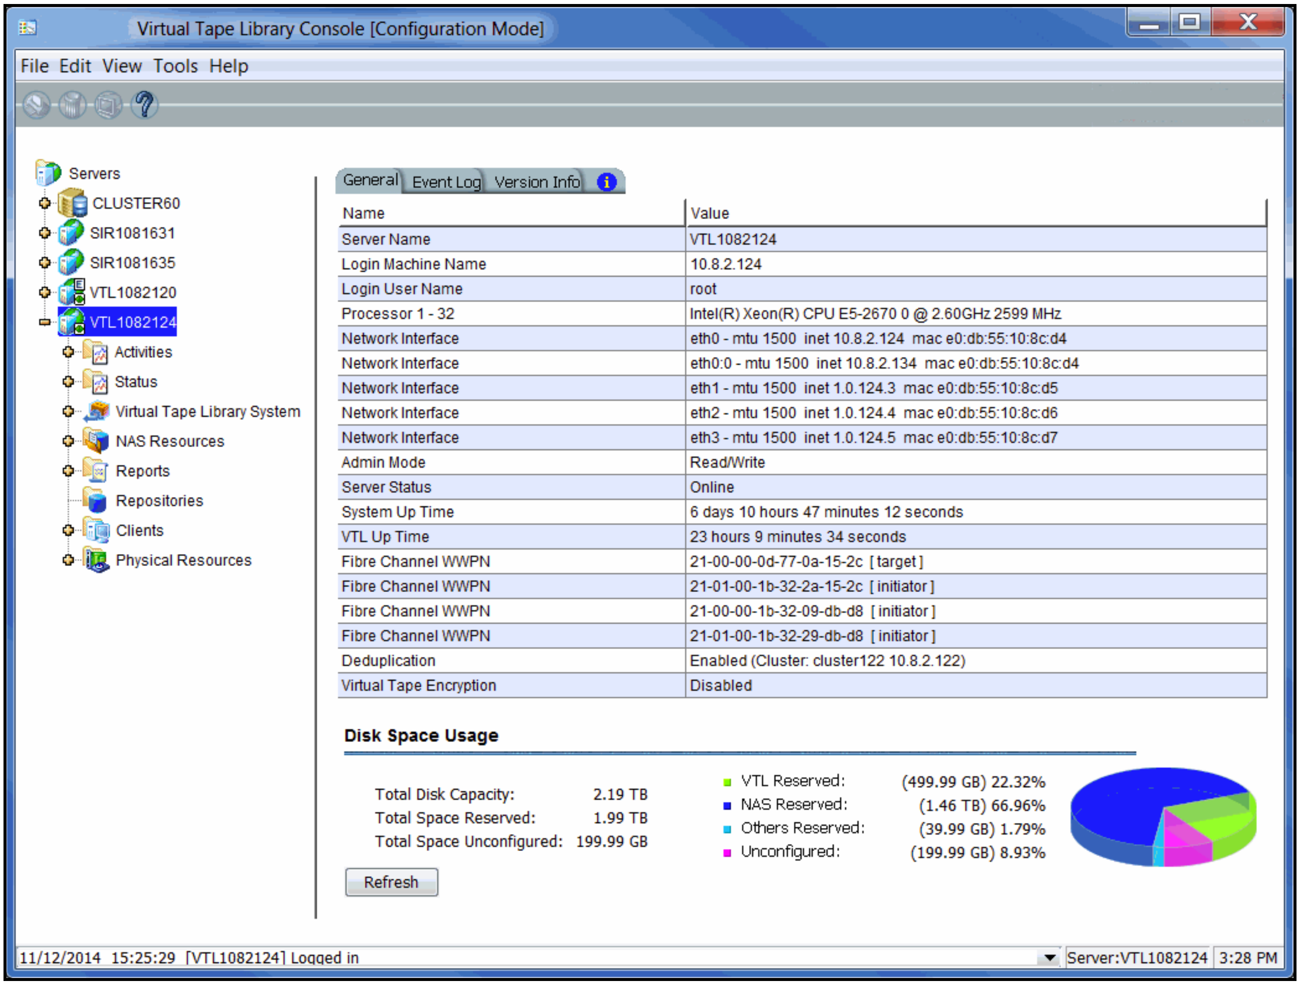

Name	Value
Server Name	VTL1082124
Login Machine Name	10.8.2.124
Login User Name	root
Processor 1 - 32	Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz 2599 MHz
Network Interface	eth0 - mtu 1500 inet 10.8.2.124 mac e0:db:55:10:8c:d4
Network Interface	eth0:0 - mtu 1500 inet 10.8.2.134 mac e0:db:55:10:8c:d4
Network Interface	eth1 - mtu 1500 inet 1.0.124.3 mac e0:db:55:10:8c:d5
Network Interface	eth2 - mtu 1500 inet 1.0.124.4 mac e0:db:55:10:8c:d6
Network Interface	eth3 - mtu 1500 inet 1.0.124.5 mac e0:db:55:10:8c:d7
Admin Mode	Read/Write
Server Status	Online
System Up Time	6 days 10 hours 47 minutes 12 seconds
VTL Up Time	23 hours 9 minutes 34 seconds
Fibre Channel WWPN	21-00-00-0d-77-0a-15-2c [target]
Fibre Channel WWPN	21-01-00-1b-32-2a-15-2c [initiator]
Fibre Channel WWPN	21-00-00-1b-32-09-db-d8 [initiator]
Fibre Channel WWPN	21-01-00-1b-32-29-db-d8 [initiator]
Deduplication	Enabled (Cluster: cluster122 10.8.2.122)
Virtual Tape Encryption	Disabled

Disk Space Usage

Total Disk Capacity:	2.19 TB
Total Space Reserved:	1.99 TB
Total Space Unconfigured:	199.99 GB

🟩 VTL Reserved:	(499.99 GB) 22.32%
🟦 NAS Reserved:	(1.46 TB) 66.96%
🟦 Others Reserved:	(39.99 GB) 1.79%
🟪 Unconfigured:	(199.99 GB) 8.93%

11/12/2014 15:25:29 [VTL1082124] Logged in | Server:VTL1082124 | 3:28 PM

The tree allows you to navigate the various servers and their configuration objects. You can expand or collapse the display to show only the information that you wish to view. To expand a collapsed item, click the ⊞ symbol next to the item. To collapse an item, click the ⊟ symbol next to the item. Double-clicking on the item will also toggle the expanded/collapsed view of the item.

You need to connect to a server before you can expand its object.

When you highlight any object in the tree, the right-hand pane contains detailed information about the object. You can select one of the tabs for more information.

The console log located at the bottom of the window displays information about the activities performed in this console. The log features a drop-down box that allows you to see activity from this console session. The local server name and time are displayed in the bottom right corner of the console.

Console modes

To simplify management tasks, two operational modes are available for the console:

- *Configuration mode* is typically used during initial system configuration and adds the *Repositories, Clients*, and *Physical Resources* objects to the tree, providing the ability to mirror repository disks, as well as configure and manage physical resources and clients.
- *Standard mode* includes sufficient objects for daily operations, including virtual tape library and NAS resource configuration and management, job and system status, and reports.

The console display option is located in the *Console Options* dialog (refer to ‘Set console options’).

Understanding the objects in the tree

The objects displayed in the navigation tree are described below. While some objects are displayed for any connected server, some are available only for specific server objects or resources.

Backup server object

A backup server can have a VTL interface, NAS interface, or both. If the server is a VTL-S server, it can also have deduplication. From a backup server object, you can manage administrator accounts for that server, add/remove licenses, change the system password, configure server-level options such as failover and email alerts, manage software licenses, perform system maintenance, enable NAS integrity checking, set tape encryption keys, generate an X-ray file, join a group, and set server properties.

When you are connected to a server, you can see the following objects: *Activities, Status, Virtual Tape Library System, NAS Resources, Reports, Repositories, Clients*, and *Physical Resources*.

You can also see the following tabs:

- *General* - Displays server configuration and status. Configuration information includes the server name and machine name, type and number of processors, network adapter information, and whether or not deduplication is enabled. Status information includes server and system status and the amount of time each has been running, storage capacity usage, and system drive usage.
- *Event Log* - Displays system events and errors.
- *Version Info* - Displays the version of the server and console software, server type, enabled options, a license summary, and installed patches.
- *Location* - Displays information about the location of this server and who is responsible for maintaining it. This tab only appears if the location information was set (via *Server Properties*).
- *Attention Required* - Appears when the system has information to report, such as physical library failures, replication errors, import/export job status.
- *Failover Information* - Displays the current status of your failover configuration, including all settings. This tab only appears if failover is configured.
- *SIR Replication* - This tab only appears for a VTL-S server, and only if replication has been configured. The tab provides the names of *primary* (source) and *replica* (target) pairs.

Deduplication server object

A deduplication server functions as the repository for a backup server and allows you to view information about deduplication activities and performance. From a deduplication server object, you can manage administrator accounts for that server, change the system password, configure server-level options such as email alerts, manage software licenses, perform system maintenance, add deduplication resources, generate an X-ray file, and set server properties.

When you are connected to a server, you can see the following objects: *Activities*, *Status*, *Reports*, *Repositories, Clients*, and *Physical Resources*.

You can also see the following tabs:

- *General* - Displays the configuration and status of the server. Configuration information includes the server name and machine name, type and number of processors, and network adapter information. Status information includes server and system status and the amount of time each has been running, storage capacity usage, and system drive usage.
- *Event Log* - Displays system events and errors.
- *Version Info* - Displays the version of the server and console software, server type, enabled options, a license summary, and installed patches.

Deduplication cluster object

A deduplication cluster includes one or more deduplication servers. From a cluster object, object, you can manage administrator accounts for the cluster, change the system password, configure cluster options such as redundant node failover and replication, configure and run reclamation, and generate an X-ray file. Initiating a procedure from this object executes the procedure for all servers in the cluster.

When you select a deduplication cluster object, you can see the following tabs:

- *Dashboard Summary* - Shows statistics about the current state of the deduplication repository and deduplication.
- *Associated Servers* - Lists the backup servers that are associated with the cluster, including their IP addresses and protocols (VTL and/or NAS).
- *Replication* - Lists the replication jobs configured to and from this cluster. For each job, this tab shows the source and target servers as well as the communication protocol.
- *Redundant Node* - Displays information about redundant node failover, if it has been configured.

Multi-Node Group object

If you have configured multi-node groups, the group object contains the tape backup servers that have been grouped together.

All of the servers in a group can be managed together. From the group level, you can manage user accounts for all servers in the group and you can set common configuration parameters, such as SNMP settings, storage monitoring trigger threshold, tape caching thresholds, compression settings, and X-rays. You can also log in to all of the servers in the group at the same time.

Activities object

This object allows you to display information for job queues and active jobs on the server. Available activities depend on server type. Click an object to display related information in the right-hand pane.

Backup server activities

For a backup server, the *Activities* object lists all possible job queues on the server. Click an object to display information about active jobs in the right-hand pane.

Deduplication Job Queue - Lists all of the tapes that are being processed. From this object, you can change the priority of tapes in the queue or cancel processing for a tape.

Unique Replication Queue - Displays information for replication jobs for deduplicated tapes. These jobs are carried out after the index has been replicated; tapes in the list are currently replicating or awaiting replication.

NAS Replication Activities - Displays information about outgoing replication of data on NAS resources.

Replication Queue - Displays information for replication jobs for virtual tapes. From here, you can suspend and resume replication.

Tape Import/Export Queue - Lists current import, export, and Automated Tape Caching jobs. From here, you can display the *Import/Export Job Properties* dialog and specify rules for retrying failed jobs. You can also delete one or more jobs from the list in the right-hand pane.

Deduplication server activities

For a deduplication server, deduplication activities are available.

Scanner Processes - Displays a list of currently running policies.

Scanner History - Displays details for deduplicated virtual tapes.

Status object

This object allows you to display status information for various server activities as well as capacity information. Available status categories depend on server type. Click an object to display information in the right-hand pane.

Backup server status

This object can have up to three objects:

- *Dashboard Summary* - This object can have up to three tabs.
  - *VTL/NAS Space Usage* - Displays capacity information for VTL storage and/or NAS file systems.
  - *VTL/NAS Performance* - Displays performance statistics for VTL and/or NAS file systems.
  - *Deduplication Repository* - On a VTL-S server, displays information about repository capacity and the usage of index cache capacity, deduplication data disks, and metadata disks.
- *NAS Integrity Check Statistics* - Displays the status (or current progress) of an integrity check job, which verifies that stub files point correctly to deduplicated data in the repository and can be used to retrieve data.
- *NAS Performance Monitoring* - Displays information about NAS performance.

Deduplication server status

Dashboard Summary - Displays the *Dashboard Summary* tab, which collects information about repository capacity and the usage of index cache capacity, deduplication data disks, and metadata disks.

Virtual Tape Library System object

The *Virtual Tape Library System* object contains all of the information about your appliance.

Virtual Tape Libraries

This object lists the virtual tape libraries that are currently available. Each virtual tape library consists of one or more virtual tape drives and one or more virtual tapes. Each virtual tape library and drive can be assigned to one or more backup application servers (clients). Each library's virtual tapes are sorted in barcode order.

For each library, you can:

- Create/delete virtual tapes
- Create/delete virtual tape drives
- Enable replication for tapes in the library
- Set Automated Tape Caching policies (if you are using this option)
- Set tape properties for the library (enable/modify tape capacity on demand, change maximum tape capacity)
- View performance statistics

For each virtual tape, you can:

- Move the virtual tape to a slot, drive, or to the virtual vault

- Enable replication for that tape or make a single remote copy
- Change tape properties (change barcode, enable/modify tape capacity on demand, enable write protection, and configure Auto Archive/Replication)
- View performance statistics

When you select a virtual tape in the list, information about that tape is displayed in the lower portion of the information pane.

Virtual Tape Drives

This object lists the standalone virtual tape drives that are currently available. Each virtual tape drive can be assigned to one or more backup application servers (clients). For each virtual tape drive, you can create/delete virtual tapes and view performance statistics.

> **Note:** If you are using deduplication, SIR tape drives will be listed when you highlight the *Virtual Tape Drives* object. These tape drives are for deduplication use only and cannot be assigned to backup application servers.

Virtual Vault

This object lists the virtual tapes that are currently in the virtual vault. The virtual vault is a tape storage area for tapes that are not inside a virtual tape library. Virtual tapes appear in the virtual vault after they have been moved from a virtual tape library. Local virtual index tapes (LVITs) of deduplicated tapes can also be in the virtual vault on the target replication server (depending upon how the deduplication policy was configured) after replication is complete. Virtual tapes in the vault can be replicated, exported to a physical tape, or moved to a virtual library or standalone drive. There is no limit to the number of tapes that can be in the virtual vault. Virtual tapes in the vault can be sorted by name, barcode, and source server. They can also be filtered to display only specific tapes.

Physical Tape Libraries

This object lists available physical tape libraries. For each physical tape library, you can assign physical tape drives, inventory the library, scan tapes, import or move a tape, reset the library, mark a library or drive disabled for maintenance purposes, or view performance statistics. For each physical tape, you can export the physical tape, copy the physical tape to a virtual tape, or link the physical tape to a virtual tape for direct access.

Physical Tape Drives

This object lists available standalone physical tape drives. For each physical tape drive, you can check for a physical tape, import a tape, mark a drive disabled for maintenance purposes, or view performance statistics. For the physical tape, you can eject a physical tape, copy the physical tape to a virtual tape, or link the physical tape to a virtual tape for direct access.

Replica Resources

This object lists the replica resources on this server. Replica resources store data from local and remotely replicated virtual tapes. Clients do not have access to Replica resources. You can sort the tapes by tape name, barcode, last replication start time, and source server. Replica resources for deduplicated tapes can also be filtered to display only tapes from a specific source server.

Physical Tape Database

The physical tape database maintains a history of all physical tapes that were used for export jobs (including stacked tapes). Physical tape entries can be removed from the database by manually purging them.

Deduplication Policies

This object lists the deduplication policies that have been set for virtual tapes. You can create or modify deduplication policies from this object, perform deduplication, and view deduplication statistics and status.

NAS Resources object

This object displays a list of NAS resources, including type, size, and number of used sectors. If reclamation is running, status information is also displayed. NAS resources are used for standard file-based backup and for OpenStorage backups when the FalconStor OpenStorage (FSOST) option is configured.

- A NAS resource from which folders can be configured and then presented to clients.
- OpenStorage-enabled resource - Resources for OpenStorage backups represent Logical Storage Units (LSUs) - virtualized disks that are used with FSOST (a software interface between a Symantec™ NetBackup™ Media or Master server or Symantec™ Backup Exec™ server and your Virtual Tape Library server). For more information about using the FSOST option, refer to your *FalconStor OpenStorage Option User Guide*.
- Folders and shares - The folder icon represents a folder that has been created on a NAS resource. If the folder icon includes a hand, this is a shared folder that can be mapped/mounted by NAS clients.
- NAS replica - If this server receives replicated data from another Virtual Tape Library server, you will see this object in the tree beneath the NAS resource being used to hold incoming replication.

Reports object

Virtual Tape Library provides reports that offer a wide variety of information:

- VTL server performance
- Physical resources - allocation and configuration
- Disk space usage
- LUN allocation
- Fibre Channel adapters configuration
- Status for deduplication and deduplication replication
- Virtual tape/library information
- Virtual library and drive assignment
- Job status
- Physical tape usage
- Deduplication policy status, tape activity and tape usage

- Deduplication repository usage, performance, reclamation
- Device usage and configuration
- NAS - resource and share usage; deduplication/replication summary information

Repositories object

This object displays information about the types of repository resources present on the selected server.

- Database - On a backup server with a VTL interface or a VTL-S server, the database stores information about libraries, clients, and replication setup.
- Configuration Repository - On a backup server, deduplication server, or VTL-S server, the configuration repository stores information about the deduplication configuration.
- Index/folder disks - On a deduplication server or VTL-S server, these resources store pointers to unique deduplicated data and information related to deduplication sessions.
- Deduplication data disks - On a deduplication server or VTL-S server, these resources store unique deduplicated data.

Repository resources can be mirrored in order to permit recovery in case of hardware failure; doing so is highly recommended. Refer to ‘Mirror repository disks to protect your configuration’ for details.

The *Repositories* object is only visible when the console is in configuration mode.

Clients object

Backup application servers that use VTL are referred to as *Clients*.

- CIFS clients (users and groups) - These clients use the Common Internet File System (CIFS) protocol to access NAS resources. From the *CIFS Clients* object, you can modify security settings, review properties, refresh the connection, add users and groups, enable ACL support, and update user accounts and authentication servers. Users and groups can also be assigned shares from the *Users/Groups* objects. If you are using Domain authentication mode, you do not need to add Windows clients as the list of users comes from the authentication server.
- NFS clients - These are usually Unix clients using the Network File System (NFS) protocol to access NAS resources. For each NFS client listed in the right pane, you can assign a share or set properties.
- Fibre Channel clients - For each FC client, you can assign/unassign virtual tape libraries/drives, set properties, and view performance statistics. If you rename an FC client, the initial name will be preserved and is listed with the client details in the right pane.

- iSCSI clients - For each iSCSI client, you can create/assign targets, set properties, and view performance statistics. If you rename an iSCSI client, the initial name will be preserved and is listed with the client details in the right pane.
- Hosted Backup Client - This client is used with the Hosted Backup and NDMP options. From the *Hosted Backup Client* object, you can assign/ unassign virtual tape libraries/drives.

The *Clients* object is only visible when the console is in configuration mode.

Physical Resources object

Physical resources are all of your SCSI adapters/FC HBAs and storage devices. Storage devices include hard disks, tape drives, and tape libraries. Hard disks are used for creating virtual tape libraries/drives, virtual tapes, and NAS file systems.

Backup application servers do not have access to physical resources, only to logical resources. Logical resources must be configured from physical resources and then assigned to clients.

From *Physical Resources*, you can ‘Rescan physical devices’ in order to identify all newly connected devices or devices connected to a single adapter and ‘Prepare physical storage devices’ in order to create logical resources for use as storage.

The *Physical Resources* object is only visible when the console is in configuration mode.

Group Reports object

If you have a multi-node group configured, you will see a *Group Reports* object. This object provides reports that can be generated for all servers in a group. This includes standard reports that are generated on each server in the group and contain data specific to that server. You can also run a consolidated *Group Disk Space Allocation for Virtual Tapes in Libraries Report* that includes every server in the group in a single report.

Understanding the icons in the tree

Virtual tape library, virtual tape drive, and virtual tape icons

The following table describes the icons that are used to describe virtual tape libraries, virtual tape drives, and virtual tapes in the console:

Icon	Description
E	The E icon on a virtual tape library indicates that the library has encryption enabled. The icon appears in red if encryption is not activated.
C	The C icon on a virtual tape drive indicates that the drive has compression enabled.
E	The E icon on a virtual tape indicates that the tape is encrypted. The icon appears in red if encryption is not activated.
T	The T icon on a virtual tape indicates that the tape was promoted from a replica in test mode. If the tape has encryption enabled, the T icon will replace the E icon. You can check the *General* tab for current information about the tape.
O A	Used with Automated Tape Caching, the A icon on a virtual tape indicates that this is a cache for a physical tape. The O icon indicates that this cached tape has unmigrated data.
S	Used with Automated Tape Caching, the S icon on a virtual tape indicates that this is a direct link tape (a link to the physical tape).
D R	The D and R icons on a virtual tape indicate the status of the last operation performed ("D" for deduplication and "R" for deduplication with replication) • Green D - The last deduplication process was successful (pure VIT). • Yellow D - Deduplication is pending or in-progress (not a pure VIT). • Red D - The last deduplication process failed (not a pure VIT). • Green R - The last replication process was successful and the tape has been successfully resolved. • Yellow R - The replication process is pending or in-progress on the source server. • Red R - The last replication process was unsuccessful. The last attempt at resolving the tape failed or the tape is currently being resolved or has not been resolved.

Icon	Description
D R	An R icon that is divided diagonally will appear for a tape in a deduplication policy with any type of Advanced Replication enabled. The upper left section represents replication to Replication Target 1; the lower-right section represents replication to Replication Target 2. Green, yellow, and red indicators apply as described above.

Physical resource icons

The following table describes the icons that are used to describe physical resources in the console:

Icon	Description
T	The T icon on an HBA indicates that this is a target port.
I	The I icon on an HBA indicates that this is an initiator port.
D	The D icon on an HBA indicates that this is a dual port.
↓ T	The red arrow on an HBA indicates that this HBA is down and cannot access its storage.
V	The V icon on a disk indicates that this disk has been virtualized.
F	The F icon on a disk indicates that this is shared storage and is being used by another server. The *Owner* field lists the other server.

Set console options

To set options for the console:

1. Select *Tools --> Console Options*.

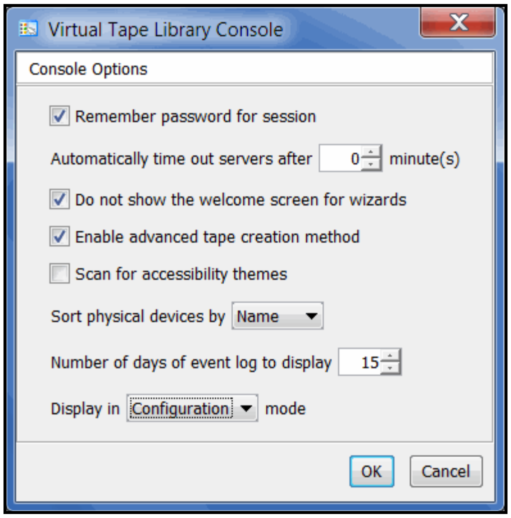

Virtual Tape Library Console

Console Options

☑ Remember password for session

Automatically time out servers after 0 minute(s)

☑ Do not show the welcome screen for wizards

☑ Enable advanced tape creation method

☐ Scan for accessibility themes

Sort physical devices by Name

Number of days of event log to display 15

Display in Configuration mode

2. Select the options you want to use.

   Remember password for session - If the console is already connected to a server, when you attempt to open a subsequent server, the console will use the credentials from the last successful connection. If this option is unchecked, you will be prompted for a password for every server you try to open. You should not remember passwords when the console is being shared by different users.

   Automatically time out servers after nn minute(s) - The console will collapse a server that has been idle for the number of minutes you specify. If you need to access the server again, you will have to reconnect to it. The default is 10 minutes. Enter 0 minutes to disable the timeout.

   Do not show the welcome screen for wizards - Each wizard starts with a welcome screen that describes the function of the wizard. Determine whether or not you want the welcome screen to be displayed.

   Enable advanced tape creation method - With *Advanced Tape Creation* enabled, you are offered advanced options when creating tapes, such as capacity-on-demand settings for virtual libraries, tape capacity of tapes, and device, name, and barcode selection for each tape that is created.

   Scan for accessibility themes - Select if your computer uses *Windows Accessibility Options.*

   Sort physical devices by - A global setting to sort physical devices by name or SCSI address. While viewing the information in the console, you can click on a column heading to re-sort the information.

   Number of days of event log to display - Specify how many days of information will be displayed in the Event Log on this console.

Display in x Mode - To simplify manageability, there are two display modes, standard and configuration. Standard mode includes sufficient objects for daily operations, including virtual tape library and NAS resource configuration and management, job and system status, and reports. Configuration mode adds the *Repositories, Clients*, and *Physical Resources* objects to the tree, providing the ability to mirror repository disks, as well as configure and manage physical resources and clients. This mode is typically used during initial system configuration.

Perform system maintenance

The console gives you a convenient way to perform system maintenance for your server.

Notes:

- The system maintenance options are hardware-dependent. Refer to your hardware documentation for specific information.
- Only the root user can access the system maintenance options.

NIC Port Bonding

NIC Port Bonding is a load balancing/path redundancy feature that enables you to load balance network traffic across two or more network connections creating redundant data paths throughout the network.

NIC Port Bonding offers a new level of data accessibility and improved performance for storage systems by eliminating the point of failure represented by a single input/output (I/O) path between servers and storage systems and permits I/O to be distributed across multiple paths.

NIC Port Bonding allows you to group network interfaces into a single group. You can think of this group as a single virtual adapter that is actually made up of multiple physical NIC adapters. To the system and the network, it appears as a single interface with one IP address. However, throughput is increased by a factor equal to the number of adapters in the group. Also, NIC Port Bonding detects faults anywhere from the NIC out into the network path and provides dynamic failover in the event of a failure.

1. To enable NIC Port Bonding, select *System Maintenance --> NIC Port Bonding.*

   Notes:

   - If you have previously set NIC Port Bonding, the system will have to remove the bonding and restart network services before continuing.
   - Mixing different types of NIC ports, such as 10 GBe with 1 GB, in a NIC Port Bonding configuration is not recommended.

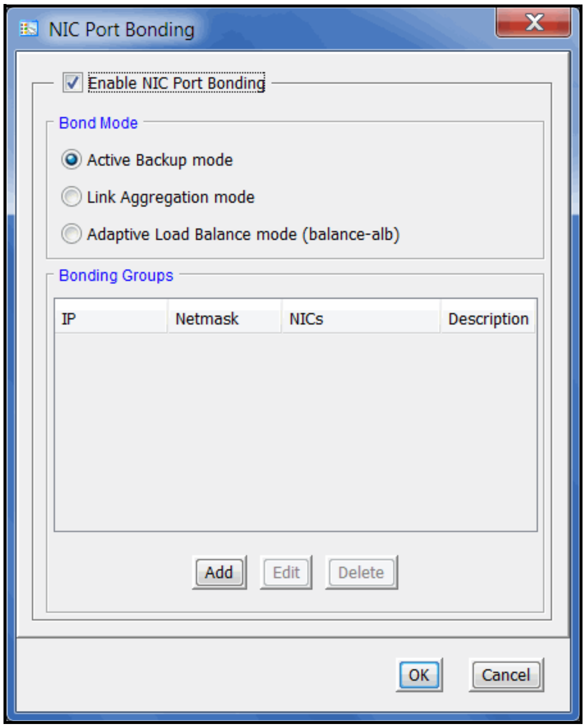

☑ Enable NIC Port Bonding

Bond Mode

● Active Backup mode

○ Link Aggregation mode

○ Adaptive Load Balance mode (balance-alb)

Bonding Groups

IP	Netmask	NICs	Description

2. Select the *Enable NIC Port Bonding* checkbox.

3. Select the bond mode.

   Active Backup mode - The default mode. Only one port is active at a time. A different port will become active only if the active port fails. The MAC address for the bond group is externally visible on only one port (network adapter).

   Link Aggregation mode - Creates groups that share the same speed and duplex settings. This is a more dedicated, tuned mode that uses IEEE 802.1AX-capable switches to optimize traffic.

   Adaptive Load Balancing mode - Provides load balancing for both incoming and outgoing traffic without requiring any special switch support. Outgoing traffic is distributed according to the current load on each slave. Incoming traffic is received by the current slave. If the receiving slave fails, another slave takes over the MAC address of the failed receiving slave.

   Regardless which mode is chosen, you may have to perform some configuration on your switch to support this. The specific configuration will depend upon the model of your switch.

4. Click *Add* and configure which ports to bond.

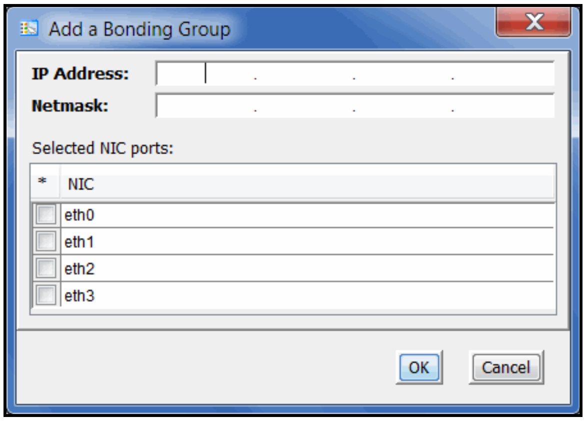

Enter an IP address and netmask to represent the bond group and select which ports to include. Depending upon how many NIC ports you have, you can repeat this step to create multiple bond groups.

> **Note:** To modify a bonded IP address of a server, you will need to remove the bonding configuration (*System Maintenance* --> *NIC Port Bonding* --> *Yes* to remove) and then rebond using the new IP addresses.

VLAN tagging

When a Virtual LAN (VLAN) spans multiple switches, VLAN Tagging helps to identify the VLAN to which data belongs and helps determine which port(s) to use for communication.

1. To configure VLAN Tagging, select *System Maintenance --> VLAN Tagging.*

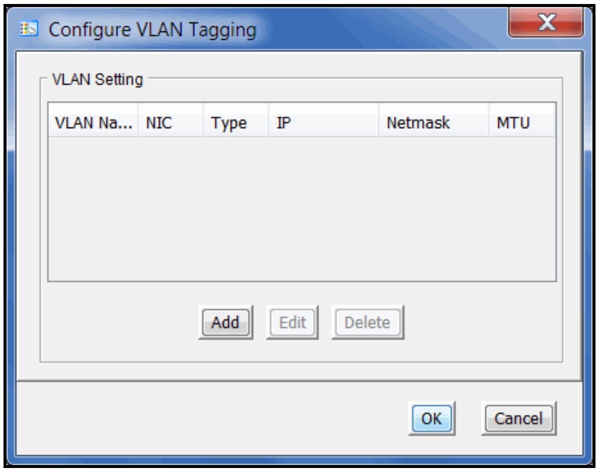

VLAN Na...	NIC	Type	IP	Netmask	MTU

2. Click *Add* and add a VLAN.

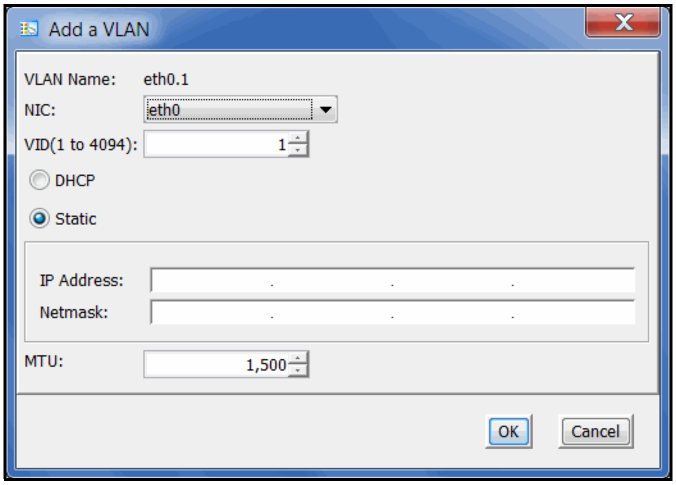

VLAN Name: eth0.1
NIC: eth0
VID(1 to 4094): 1
○ DHCP
● Static
IP Address: . . .
Netmask: . . .
MTU: 1,500

Select a NIC and specify the VLAN ID that was set in the network switch.

Specify if you are using dynamic (DHCP) or static addresses. If you select *Static*, you must add the IP address and subnet mask.

MTU - Set the maximum transfer unit of each IP packet. If your card supports it, set this value to 9000 for jumbo frames.

Network configuration

If you need to change server IP addresses, you must make these changes using *Network Configuration*. Using any other third-party utilities will not update the information correctly. Refer to ‘Set up network’ for more information.

Notes:

- If you need to change the IP address of an appliance that has replication configured for virtual tapes or deduplicated tapes, you must remove the replication configuration before changing the IP address, and then you can configure replication again.
- You cannot change the network configuration of a server that is in a multi-node group or has failover configured.
- If your backup server is associated with a deduplication server, you should not change any IP addresses.
- If you are using FSOST, you will need to update your FSOST configuration if you change the IP address. You can either manually update the *fsost.conf* file on each client being protected by your NetBackup Media/Master servers or you can use the `fsost creatests` command to re-add the VTL server with the updated IP address on each client.

Set hostname

Right-click a server and select *System Maintenance* --> *Set Hostname* to change your hostname. The server will automatically reboot when the hostname is changed.

You cannot change the hostname of a server if any of the following conditions exist:

- The server is in a multi-node group
- The server has failover configured
- The server has replication configured for deduplicated tapes
- (VTL/VTL-S) The server has deduplication enabled or is associated with a deduplication cluster
- (SIR) The server has a deduplication cluster configured

Notes:

- Make sure your storage is connected and accessible before you change the hostname. If it is not and the operation fails, you can change the hostname back to the original, fix your storage, and then try again.
- Do not change the hostname if you are using block devices. If you do, all block devices claimed by VTL will be marked offline and seen as foreign devices.
- Do not use an underscore ('_') in the hostname if a backup server will be associated with a deduplication cluster.

Set date and time

You can set the date, time, and time zone for your system, as well add NTP (Network Time Protocol) servers. NTP allows you to keep the date and time of your server in sync with up to five Internet NTP servers.

You can also access these setting by double-clicking on the time that appears at the bottom right of the console.

Notes:

- We recommend restarting VTL services if you change the date and time.
- You should only change the system time when there is no IO activity.
- Changing the date/time during operations can interfere with the scheduling functions of other processes, such as reclamation.

Restart VTL

Right-click a server and select *System Maintenance* --> *Restart VTL* to restart the server processes.

Restart network

Right-click a server and select *System Maintenance* --> *Restart Network* to restart your local network configuration.

Reboot

Right-click a server and select *System Maintenance* --> *Reboot* to reboot your server.

Halt

Right-click a server and select *System Maintenance* --> *Halt* to turn off the server without restarting it.

Add and register licenses

To license Virtual Tape Library and its options, make sure you have obtained your keycode(s) from DSI or its representatives. Once you have the license keycodes, follow the steps below:

1. In the console, right-click the server object and select *License*.

   The *License Summary* window is informational only and displays a list of the options supported for this server. You can enter keycodes for your purchased options in the *Keycodes Detail* dialog.

2. Click the *Add* button in the *Keycodes Detail* dialog to enter each keycode.

   If multiple administrators are logged into a server at the same time, license changes made from one console will take effect in other console only when the administrator disconnects and then reconnects to the server.

3. If your licenses have not been registered yet, click the *Register* button in the *Keycodes Detail* dialog.

   Select *Online* to register online if you have an Internet connection. Otherwise, select the *Offline* option.

Offline registration

Offline registration is useful when you do not have an Internet connection. When you select the *Offline* registration option, you will see the *Offline Registration* dialog:

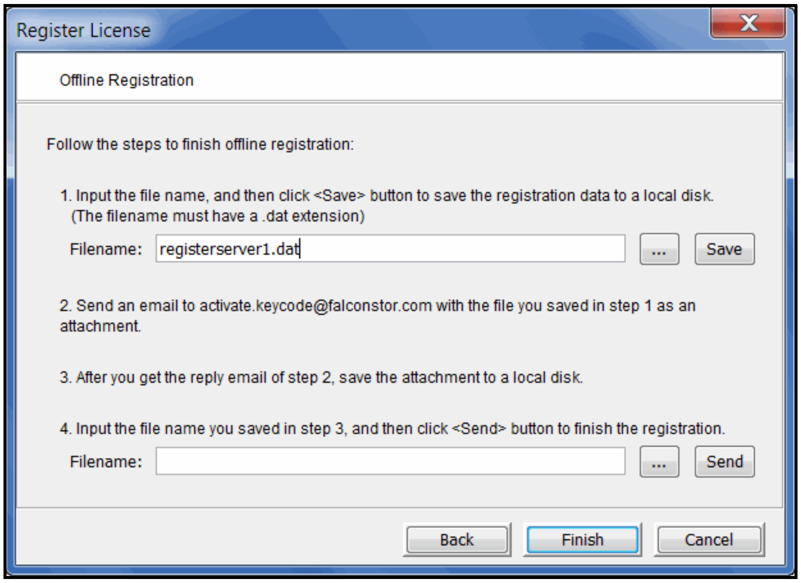

Register License

Offline Registration

Follow the steps to finish offline registration:

1. Input the file name, and then click <Save> button to save the registration data to a local disk. (The filename must have a .dat extension)

Filename: registerserver1.dat

2. Send an email to activate.keycode@falconstor.com with the file you saved in step 1 as an attachment.

3. After you get the reply email of step 2, save the attachment to a local disk.

4. Input the file name you saved in step 3, and then click <Send> button to finish the registration.

Filename:

To register offline:

1. Specify a path and file name in which to save the registration information.

   Registration information file names can only use English alphanumeric characters and must have a .dat extension. You cannot use a single digit as the name. For example, company1.dat is valid (1.dat is not valid).

2. Click the *Save* button.

   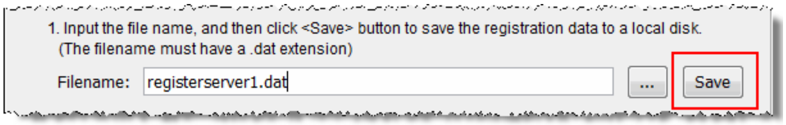

   1. Input the file name, and then click <Save> button to save the registration data to a local disk.
   (The filename must have a .dat extension)

   Filename: registerserver1.dat

   It is a good idea to keep the dialog open while you complete the remaining steps.

3. Save the generated file to portable storage media or to a shared folder on your network and email it from a computer with Internet access to FalconStor's registration server: (activate.keycode@falconstor.com).

   It is not necessary to write anything in the subject or body of the email.

   If your email is working correctly, you should receive a reply within a few minutes.

   > **Note:** In order to prevent the possibility of unsuccessful email delivery to the FalconStor registration server, disable Delivery Status Notification (DSN) before you send the activation request email.

4. When you receive an email response from FalconStor's registration server, save the attachment (with the .sig extension) to portable storage media or a shared folder and return it to the local drive of the computer where the console is running. Do not change the file name.

   > **Notes:**
   >
   > - If you do not receive a reply to your offline registration email within one hour after sending it, check your email encoding. Change it to UNICODE (UTF-8) if it is set otherwise and send the email again.
   > - If the reply email indicates that the license is successfully registered but the signature file is not attached, you may have set the name of the license information file improperly; you cannot use a single digit before the suffix in the file name. Change the registration file name to a valid alphanumeric string and then try to register again. If the issue persists, contact Technical Support.

5. Back in the *Offline Registration* dialog, specify the path and name of the .sig file.

6. Click the *Send* button to send the file (do not change the file name) to the registration server to complete your registration.

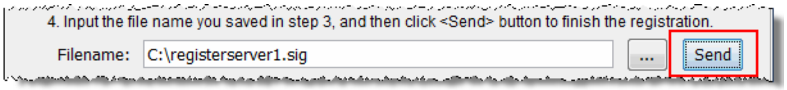

4. Input the file name you saved in step 3, and then click <Send> button to finish the registration.

Filename: C:\registerserver1.sig

7. Click the *Finish* button.

Manage physical devices

View physical devices

To see existing physical devices, expand the *Physical Resources* and *Storage Devices* objects and then select *Fibre Channel Devices* or *SCSI Devices*. The icon shown for each device describes its purpose or status (refer to 'Physical resource icons').

Filter the list

In either device list, all devices are displayed by default. To filter the list, click *Filter* to display the *Filter Options* dialog that corresponds to the type of server.

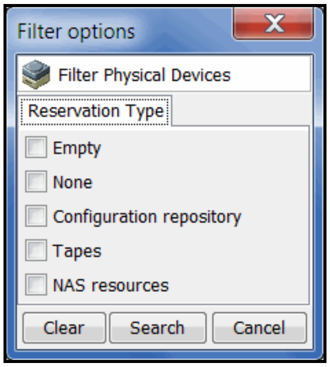

VTL server

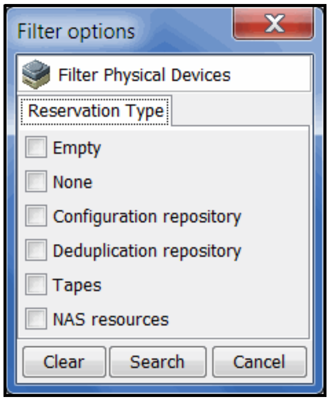

VTL-S server

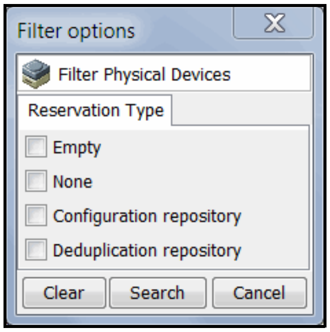

SIR server

1. Select the checkbox(es) for the types of devices you want to see in the list.
2. Click *Search* to update the list.
3. To change filter selections, de-select one or more selections or click *Clear* to de-select current choices, then click *Search*.

To clear filters and display the default list, click *Show All Devices* above the list.

View individual devices

When you select a device, details are displayed in the lower area of the right pane. Information in tab panels will vary depending upon the selected device:

- General tab - Displays attributes reported by the device (such as make, model, and firmware revision), as well as the device's SCSI address and aliases. For disks, disk size, total sectors/sector size, and status are displayed, as well as its device category: unassigned, reserved for virtual device, used by virtual device.
- Performance Statistics tab - Displays read and write throughput for the past 60 minutes.
- Layout tab - Appears only for disks and identifies the first and last sectors on the device, total size, device type (Direct Device, etc), and whether the device is used by a virtual device.
- Throughput tab - Displayed after you run the Disk Throughput Test from the device's right-click menu (refer to 'Test physical device throughput').

Rescan physical devices

If you are preparing storage manually or if you are adding storage to your pre-configured system, you can add physical drives and virtualize them for use by virtual tape libraries and deduplication (refer to ‘Prepare physical storage devices’). VTL automatically scans for devices whenever you reboot, but you can always perform a manual rescan to identify new/existing devices.

> **Note:** To ensure that the server detects attached storage, power on all storage devices before starting the appliance.

If you have not rebooted your system since attaching new storage, perform a rescan to allow the server to identify new devices. You can rescan all devices or devices on a selected adapter.

> **Note:** Rescan can take a significant amount of time to complete and can disrupt I/O activity.

1. To rescan all devices, right-click *Physical Resources* and select *Rescan*.

   If you only want to scan on a specific adapter, right-click that adapter and select *Rescan*.

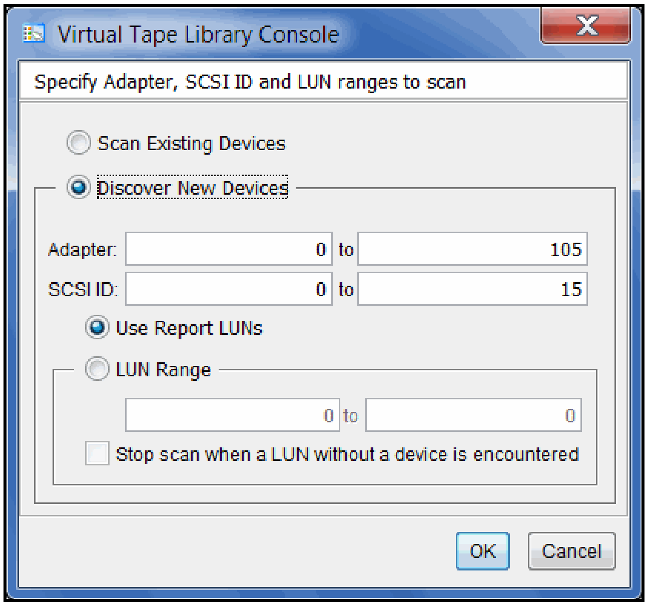

Virtual Tape Library Console

Specify Adapter, SCSI ID and LUN ranges to scan

○ Scan Existing Devices

● Discover New Devices

Adapter: 0 to 105

SCSI ID: 0 to 15

● Use Report LUNs

○ LUN Range

0 to 0

□ Stop scan when a LUN without a device is encountered

2. Determine what you want to rescan.

   If you are discovering new devices, set the range of adapters, SCSI IDs, and LUNs that you want to scan.

   Use Report LUNs - The system sends a SCSI request to LUN 0 and asks for a list of LUNs. Note that this SCSI command is not supported by all devices.

   Stop scan when a LUN without a device is encountered - This option (used with a LUN range) will scan LUNs sequentially and stop when a LUN without a device is detected. Use this option only if all of your LUNs are sequential.

Prepare physical storage devices

This procedure virtualizes physical disks so that you can create logical resources for storage.

1. Right-click *Physical Resources* and select *Prepare Devices*.

   To virtualize a single physical disk, right-click the device and select *Enlist*.

   Prepare Devices and the Device Category *Reserved for Virtual Device* are selected by default.

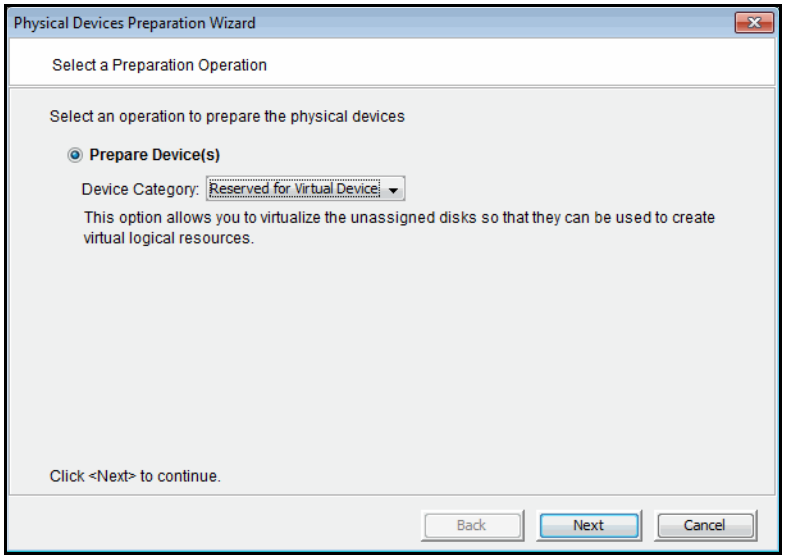

Physical Devices Preparation Wizard

Select a Preparation Operation

Select an operation to prepare the physical devices

● Prepare Device(s)

Device Category: Reserved for Virtual Device

This option allows you to virtualize the unassigned disks so that they can be used to create virtual logical resources.

Click <Next> to continue.

2. In the *Select Physical Devices* dialog, select unassigned disks to be virtualized.

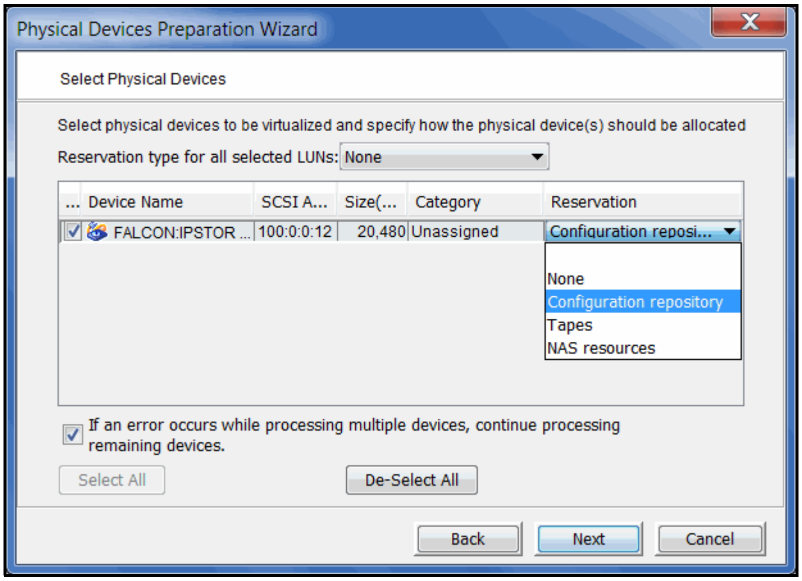

Physical Devices Preparation Wizard

Select Physical Devices

Select physical devices to be virtualized and specify how the physical device(s) should be allocated

Reservation type for all selected LUNs: None

...	Device Name	SCSI A...	Size(...	Category	Reservation
☑	FALCON:IPSTOR ...	100:0:0:12	20,480	Unassigned	Configuration reposi...

☑ If an error occurs while processing multiple devices, continue processing remaining devices.

You can also click *Select All* to select all devices. Selecting multiple LUNs or all LUNs will enable the drop-down list for *Reservation type for all selected LUNs*.

3. Select a reservation type for each device or for all selected LUNs in order to specify how the device can be allocated:
    - None - The device will not be allocated. If there are existing resources on this device, they can still be accessed; however, new resources will not be created on this device.
    - Configuration repository - The configuration repository contains configuration information for each server. A maximum of four devices can be reserved for the configuration repository and VTL database (including mirror devices). After the configuration repository is created, you cannot change the reservation of devices used for the configuration repository.
    - Deduplication repository - (VTL-S and deduplication servers only) includes deduplication index, folder, and data disks, as well as the associated mirror devices.
    - Tapes - Used only for tape storage (backup servers and VTL-S servers).
    - NAS resources - Used only for NAS resources (backup servers and VTL-S servers).
4. Confirm that all selections are correct and click *Finish*.
5. Type *Yes* at the warning message and then click *OK*.

Unassign physical storage devices

You can unassign physical devices that have been virtualized for your server, thereby removing “ownership” of these devices.

To determine how a device is being used, expand the *Physical Resources* objects until you can see the device and then select it, then review the *Layout* tab. The *Type* and *Used by Virtual Device* columns (even for a device with a category of *Reserved for Virtual Device*) let you identify the entity that “owns” the device.

In order to protect your data, this procedure will fail if a device is being used for either server database, or if tapes are currently in drives on the device, or if disk resources (such as deduplication index, folders, or data) are allocated to it.

Contact DSI Technical Support for assistance in preparing devices for unassignment.

1. In the console tree, right-click the *Physical Resources* object and select *Prepare Devices*.

    You can also right-click the *Storage Devices* object, *Fibre Channel Devices*, or *SCSI Devices*.

    To unassign a single physical disk, right-click the device and select *Discharge*.
2. Select the Device Category *Unassigned*.
3. Select the device(s) you wish to unassign from the *Select Physical Devices* dialog.
4. Confirm that all selections are correct and click *Finish*.
5. Type *Yes* at the warning message and then click *OK*.

Change a LUN reservation

You can change the LUN reservation for a virtualized device. This can be useful if you want to retire a device that is using old disk storage and you want to prevent the server from writing data to that device.

1. Right-click *Storage Devices* (under *Physical Resources*) and select *LUN Reservation* to display a list of all virtualized devices.

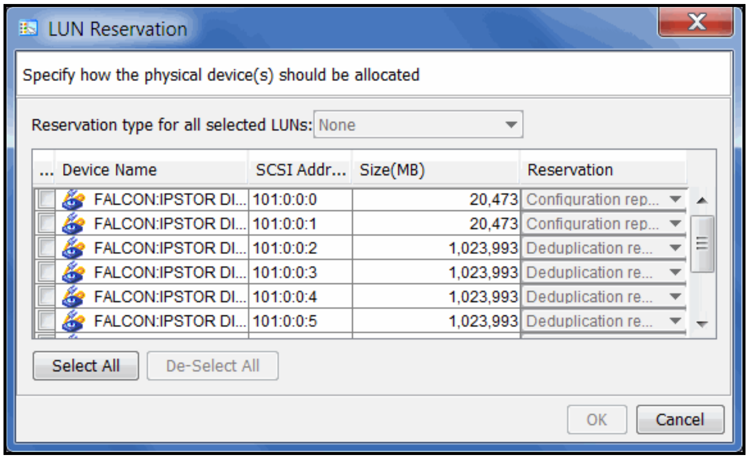

Specify how the physical device(s) should be allocated

Reservation type for all selected LUNs: None

...	Device Name	SCSI Addr...	Size(MB)	Reservation
☐	FALCON:IPSTOR DI...	101:0:0:0	20,473	Configuration rep...
☐	FALCON:IPSTOR DI...	101:0:0:1	20,473	Configuration rep...
☐	FALCON:IPSTOR DI...	101:0:0:2	1,023,993	Deduplication re...
☐	FALCON:IPSTOR DI...	101:0:0:3	1,023,993	Deduplication re...
☐	FALCON:IPSTOR DI...	101:0:0:4	1,023,993	Deduplication re...
☐	FALCON:IPSTOR DI...	101:0:0:5	1,023,993	Deduplication re...

   You can also display this dialog from the *Fibre Channel Devices* object or the *SCSI Devices* object (under *Storage Devices*) to display only those devices.

2. Select the device whose reservation you want to change.

   You can also click *Select All* to select all devices. Selecting multiple LUNs or all LUNs will enable the drop-down list for *Reservation type for all selected LUNs*.

3. Select a new reservation type for each device or for all selected LUNs. The following rules apply:
   - For a new device, the LUN can only be reserved for a single purpose. If resources already exist on the device (as might be the case with an upgraded system), the new reservation type must be compatible with the existing resource types on the LUN or can be set to *None*.
   - If the device is empty, you can choose any reservation type.
   - If the Configuration Repository has been created, you cannot change *Configuration repository* to a different reservation type.

4. Click *OK* when you are done.

Test physical device throughput

You can test the following for your physical devices:

- Sequential throughput
- Random throughput
- Sequential I/O rate
- Random I/O rate
- Latency

To check the throughput for a device:

1. Right-click the device (under *Physical Resources*).

   To test multiple devices, right-click the *Storage Devices* object.

2. Select *Test* from the menu.

   The system will test each device and then display the throughput results on a new *Throughput* tab. If you tested multiple devices from the *Storage Devices* object, aggregate results will be shown in a dialog and the *Throughput* tab will display results for each device.

Set autopathing

Autopathing gives you the ability to balance I/O to multiple LUNs by setting the first path to use to access each LUN.

To set autopathing:

1. Right-click *Storage Devices* (under *Physical Resources*) and select *Autopath*.

2. Select the autopath method you want to use.

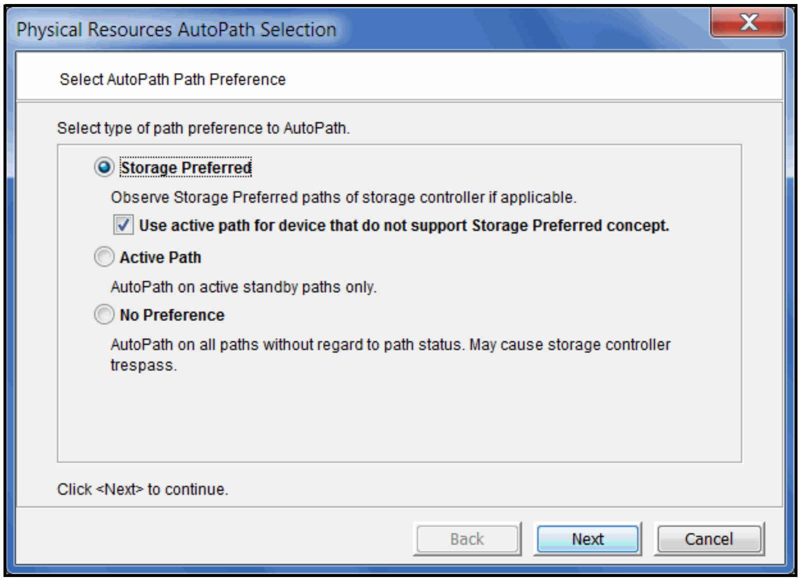

Physical Resources AutoPath Selection

Select AutoPath Path Preference

Select type of path preference to AutoPath.

● **Storage Preferred**

Observe Storage Preferred paths of storage controller if applicable.

☑ **Use active path for device that do not support Storage Preferred concept.**

○ **Active Path**

AutoPath on active standby paths only.

○ **No Preference**

AutoPath on all paths without regard to path status. May cause storage controller trespass.

Click <Next> to continue.

Storage Preferred - (Default) Detects and follows the preferred paths for each LUN as they are set by the storage controller on supported systems and uses them as preferred paths by VTL. If there is no preferred storage path, select the sub-option and VTL will use the *Active Path*.

Use active path for devices that do not support the Storage Preferred concept - Select this sub-option in case storage controllers do not have storage preferred paths. If you select this option, VTL will not trigger paths to trespass (switch).

Active Path - VTL determines the currently active paths and uses them.

No Preference - Do not use the other methods. VTL uses all available independent paths. This can cause paths to trespass.

3. Select any adapters that should be excluded.

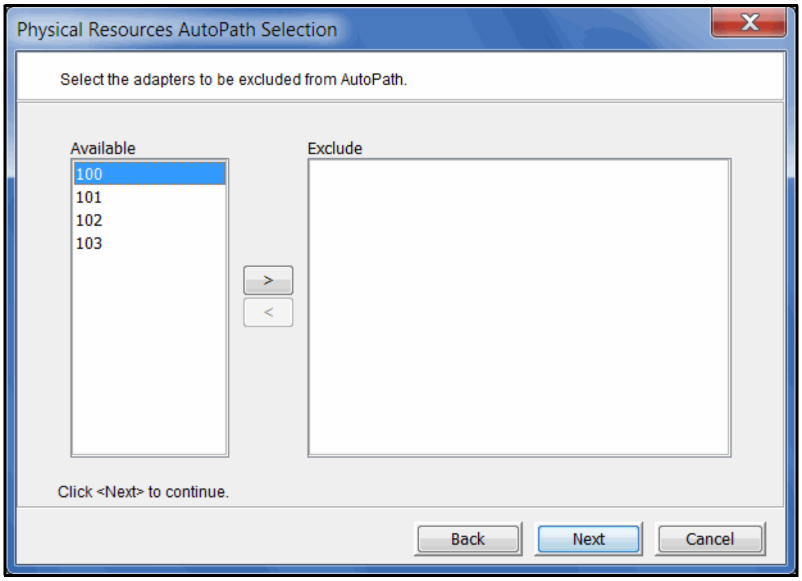

Physical Resources AutoPath Selection

Select the adapters to be excluded from AutoPath.

Available
100
101
102
103

Exclude

Click <Next> to continue.

You may need to do this if you have a specific path that you want to maintain. In that case, you would exclude the adapter, SCSI ID, and LUN for that path.

4. Select any SCSI IDs that should be excluded.

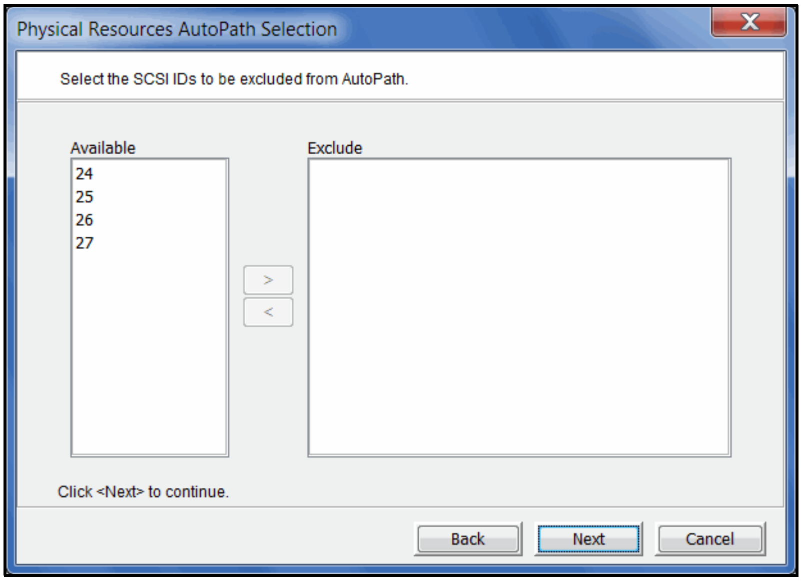

Select the SCSI IDs to be excluded from AutoPath.

Click <Next> to continue.

5. Select any LUNs that should be excluded.

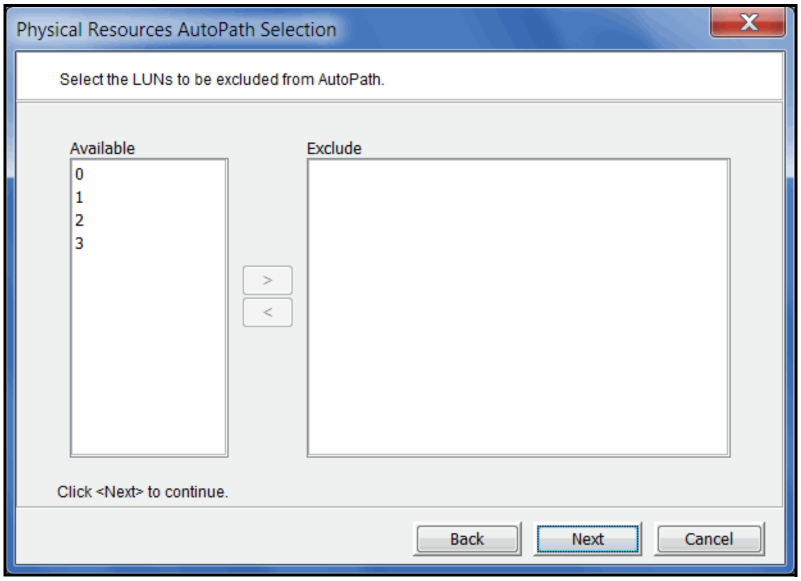

Select the LUNs to be excluded from AutoPath.

Click <Next> to continue.

6. Confirm all information and click *Finish*.

   The path configuration becomes effective immediately, but is not saved permanently. If you are satisfied with the results, you should save them so that they can be reused during startup and rescan so they can be restored. To do this, right-click *Storage Devices* (under *Physical Resources*) and select *System Preferred Path --> Save.*

   If you ever need to restore your settings back to the last version that was saved, right-click *Storage Devices* (under *Physical Resources*) and select *System*

Preferred Path --> Restore. This is useful if someone manually changes the settings and you want to revert to the saved version.

Note: Setting autopathing and saving/restoring the system preferred path will impact the I/O path that may be set for devices with multiple paths. Refer to 'Load balance the path for each downstream storage LUN' if you need to reset your settings.

Monitor space usage

On a VTL server, select *Dashboard Summary* --> *VTL/NAS Space Usage* to display information about space used by VTL and NAS file systems.

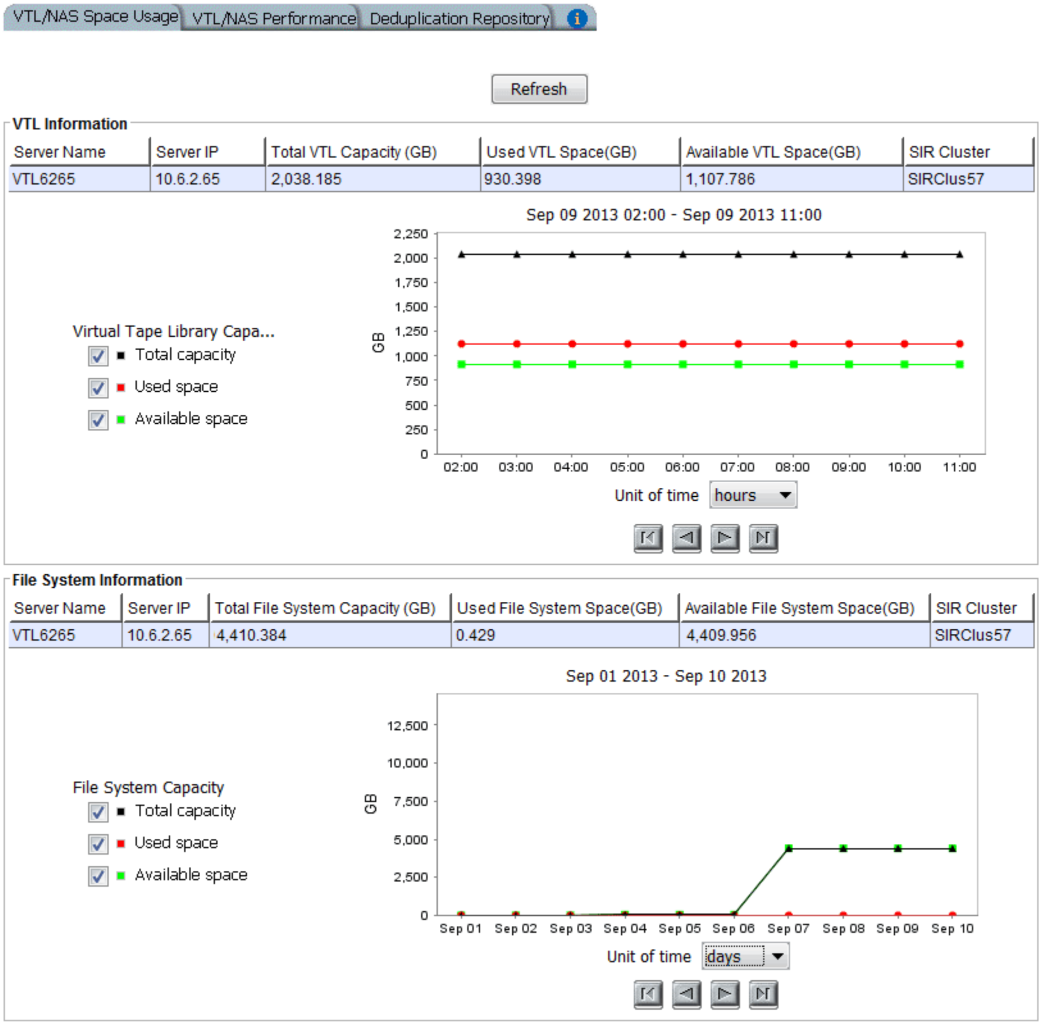

VTL Information

Server Name	Server IP	Total VTL Capacity (GB)	Used VTL Space(GB)	Available VTL Space(GB)	SIR Cluster
VTL6265	10.6.2.65	2,038.185	930.398	1,107.786	SIRClus57

Unit of time: hours

File System Information

Server Name	Server IP	Total File System Capacity (GB)	Used File System Space(GB)	Available File System Space(GB)	SIR Cluster
VTL6265	10.6.2.65	4,410.384	0.429	4,409.956	SIRClus57

Unit of time: days

The top section displays information about VTL resources, including capacity information and the associated deduplication cluster.

If NAS is enabled, the bottom section displays information about NAS resources, including capacity information and the associated deduplication cluster.

Data on the dashboard is recorded every minute.

Select a *Unit of time* (hours, days, weeks, or months) from the drop-down list to adjust the granularity of the graph. The data points in the graph will match the starting point for that unit. For example, if you select *Months*, the data point for March will show statistics for just after midnight on March 1. If you select hours, all data read/written between 7:00-8:00 will be displayed at the 7:00 data point. Use the arrow buttons to scan through accumulated data.

You can put your cursor on a data point to see detailed information.

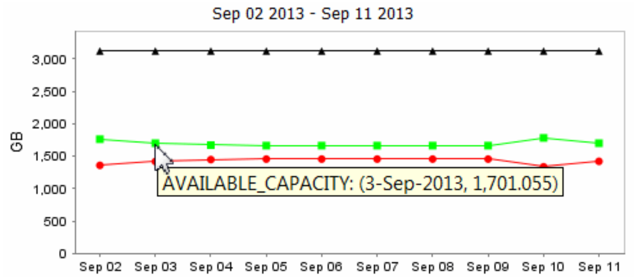

If you want to zoom into the chart to enlarge it, drag your cursor from left to right over the area you want to expand.

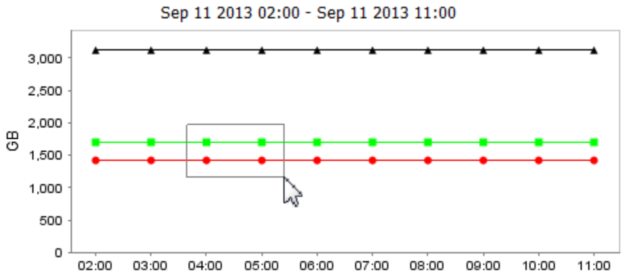

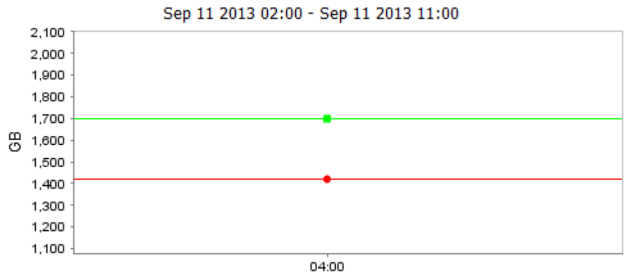

When you are finished, drag your cursor from right to left anywhere in the chart and the display will zoom out, back to a normal view.

Note: If the server has failed over, you will not see any information in the dashboard summary. Once failback occurs, dashboard summary information will be displayed.

Manage user accounts

You must add an account for each person who will have administrative rights on a VTL server. Only the root user can add or delete a VTL administrator or change an administrator’s password.

There are three types of user accounts, each with a different set of permissions:

- *VTL Administrators* are authorized for full console access (except that only the root user can add or delete a VTL administrator, change an administrator’s password, or access the system maintenance options).
- *VTL Read-Only Users* are only permitted to view information in the console. They are not authorized to make changes and they are not authorized for client authentication.
- *VTL iSCSI Users* are used for iSCSI protocol login authentication (from iSCSI initiator machines). They do not have console access. You will only be able to add this type of administrator if iSCSI is enabled.

Strong passwords

For security purposes, you can require that strong, complex passwords be used by:

- Console users - to log in to a server
- CIFS (in user mode), NDMP, OST, and iSCSI users - to log in to a server
- Encryption keys - for virtual tape encryption and physical tape export encryption

With the *Strong Passwords* option, passwords must contain a minimum of eight characters, including at least one lower-case letter, one upper-case letter, and one digit, plus at least one space or special character: ~!@#$%^&*()-_=+\|[{}];:'",<.>/?

The password cannot be the same as the administrator name or its reverse order. It also cannot contain part of the administrator name.

The first time the user logs in via the console, he or she will be required to change the password following the rules above.

When the *Strong Passwords* option is enabled, strong passwords are required for all users except the root user and existing users with passwords that were set prior to enabling the option. If the password is changed for an existing user, the strong password requirement will become valid and the user will be required to change the password the first time it is used to log in.

The *Strong Passwords* option includes a locking mechanism for user accounts that repeatedly fail to log in because of an incorrect password. The default number of times the wrong password can be entered is three. If a user is locked out, a system administrator can unlock the account or the user can wait for the lockout period to expire and try again. The default duration for the lockout period is five minutes. If you want to modify a default, contact Technical Support.

Enable strong passwords

To enable the *Strong Passwords* option, right-click your VTL server and select *Options --> Enable Strong Passwords*.

> **Note:** Once the *Strong Passwords* option is enabled, it cannot be disabled.

Password authentication

By default, all system passwords are hashed with the MD5 algorithm. For added security, you can switch to the SHA512 hashing algorithm.

To enable authentication with the SHA512 hashing algorithm, run the following command on the server:

```
authconfig --passalgo=sha512 --kickstart
```

> **Notes:**
>
> - Existing users must change their password after enabling this feature.
> - Once SHA512 authentication is enabled, you cannot go back to MD5 authentication.

Add or modify users

1. Right-click the server (or group) and select *Accounts*.

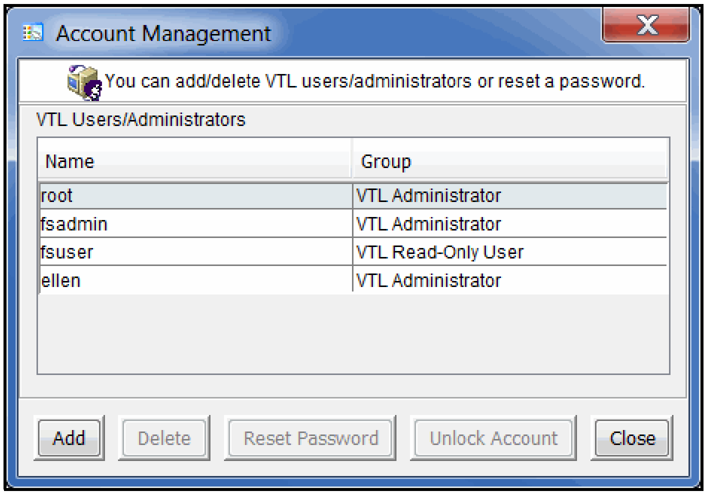

Account Management
You can add/delete VTL users/administrators or reset a password.
VTL Users/Administrators

Name	Group
root	VTL Administrator
fsadmin	VTL Administrator
fsuser	VTL Read-Only User
ellen	VTL Administrator

If you accessed *Administrators* from the group level, you can add an administrator, modify a password, or delete a user for all servers in the group.

2. Select the appropriate option.

   When you add an administrator, the name must adhere to the naming convention of the operating system running on your server. Refer to your operating system’s documentation for naming restrictions.

   If you are using the *Strong Password* option, the password must adhere to the rules for strong passwords, discussed earlier.

You cannot delete the root user or change the root user’s password from this screen. Use the *Change Password* option instead.

Change password

After initial setup, it is recommended that you change the default password.

1. Right-click the server name and select *Change Password*.
2. Enter the original password (*IPStor101*, on DSI appliances), new password, confirm the new password, then click *OK*.

Unlock an account

If the *Strong Passwords* option is enabled, users will be locked out if their log in attempt fails three times. If this happens, the user can wait for the five minute lockout period to expire and try again or the system administrator can unlock the account. To unlock an account:

1. Right-click the server (or group) and select *Accounts*.
2. Highlight the account and click *Unlock Account*.

Event Log

The Event Log details significant occurrences during server operation. You can view the Event Log in the console when you highlight a server or group in the tree and select the *Event Log* tab in the right pane.

Information displayed in the Event Log comes from the */var/log/messages* file on the server. A maximum of 10,000 records will be displayed in the Event Log.

The columns displayed in the Event Log are:

Type	**I**: This is an informational message. No action is required. **W**: This is a warning message that states that something occurred that may require maintenance or corrective action, although the system is still operational. **E**: This is an error that indicates a failure has occurred such that a device is not available, an operation has failed, or a licensing violation. Corrective action should be taken to resolve the cause of the error. **C**: These are critical errors that stop the system from operating properly.
Server	The server that this message is about. You will only see this column if you are viewing the Event Log at the group level.
Date & Time	The date and time on which the event occurred. Events are listed in chronological order. If you have servers from different time zones in a group, the events will be sorted using coordinated universal time (UTC).
ID	This is the message number.
Event Message	This is a text description of the event describing what has occurred.

The Event Log is refreshed every three seconds, meaning that new events are added on a regular basis. If you are at the top of the Event Log when new events are added, the screen will automatically scroll down to accommodate the new events. If you are anywhere else in the Event Log, your current view will not change when new events are added. This allows you to read messages without the screen scrolling.

To see more information about a warning, error, or critical error in the Event Log, double-click on the event message to see possible causes and suggested actions to take to correct the issue.

Sort the Event Log

When you initially view the Event Log, all information is displayed in chronological order (most recent at the top). If you want to reverse the order (oldest at top) or change the way the information is displayed, you can click on a column heading to re-sort the information. For example, if you click on the *ID* heading, you can sort the events numerically. This can help you identify how often a particular event occurs.

Filter the Event Log

By default, all informational system messages, warnings, and errors are displayed. To filter the information that is displayed:

1. Click the *Filter* button.

2. Specify your search criteria.

   You can search for specific message types, records that contain/do not contain specific text, category types, and/or time or date range for messages. You can also specify the number of lines to display.

Export data from the Event Log

You can save the data from the Event Log in one of the following formats: comma delimited (.csv) or tab delimited (.txt) text. Click the *Export* button to export information.

Print the Event Log

Click the *Print* button to print the Event Log to a printer.

Clear the Event Log

You can purge the messages from the Event Log. You will have the option of saving the existing messages to a file before purging them. Click the *Purge* button to clear the Event Log.

Attention Required tab

The *Attention Required* tab displays information that may require your attention, such as:

- Physical library failures
- Hardware appliance errors
- Replication errors
- Import/export job status

It also notifies you when an import/export job has completed.

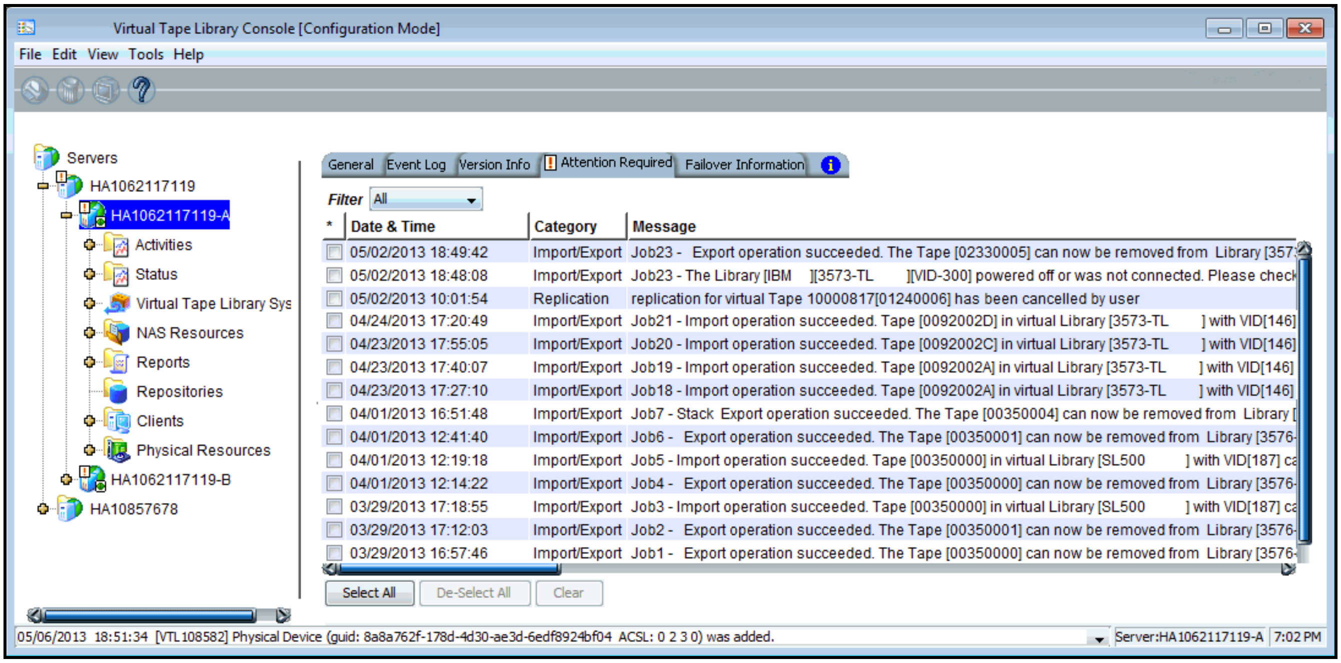

Filter: All

*	Date & Time	Category	Message
□	05/02/2013 18:49:42	Import/Export	Job23 - Export operation succeeded. The Tape [02330005] can now be removed from Library [357
□	05/02/2013 18:48:08	Import/Export	Job23 - The Library [IBM][3573-TL][VID-300] powered off or was not connected. Please check
□	05/02/2013 10:01:54	Replication	replication for virtual Tape 10000817[01240006] has been cancelled by user
□	04/24/2013 17:20:49	Import/Export	Job21 - Import operation succeeded. Tape [0092002D] in virtual Library [3573-TL] with VID[146]
□	04/23/2013 17:55:05	Import/Export	Job20 - Import operation succeeded. Tape [0092002C] in virtual Library [3573-TL] with VID[146]
□	04/23/2013 17:40:07	Import/Export	Job19 - Import operation succeeded. Tape [0092002A] in virtual Library [3573-TL] with VID[146]
□	04/23/2013 17:27:10	Import/Export	Job18 - Import operation succeeded. Tape [0092002A] in virtual Library [3573-TL] with VID[146]
□	04/01/2013 16:51:48	Import/Export	Job7 - Stack Export operation succeeded. The Tape [00350004] can now be removed from Library [
□	04/01/2013 12:41:40	Import/Export	Job6 - Export operation succeeded. The Tape [00350001] can now be removed from Library [3576-
□	04/01/2013 12:19:18	Import/Export	Job5 - Import operation succeeded. Tape [00350000] in virtual Library [SL500] with VID[187] ca
□	04/01/2013 12:14:22	Import/Export	Job4 - Export operation succeeded. The Tape [00350000] can now be removed from Library [3576-
□	03/29/2013 17:18:55	Import/Export	Job3 - Import operation succeeded. Tape [00350000] in virtual Library [SL500] with VID[187] ca
□	03/29/2013 17:12:03	Import/Export	Job2 - Export operation succeeded. The Tape [00350001] can now be removed from Library [3576-
□	03/29/2013 16:57:46	Import/Export	Job1 - Export operation succeeded. The Tape [00350000] can now be removed from Library [3576-

05/06/2013 18:51:34 [VTL108582] Physical Device (guid: 8a8a762f-178d-4d30-ae3d-6edf8924bf04 ACSL: 0 2 3 0) was added. | Server:HA1062117119-A | 7:02 PM

The *Attention Required* tab only appears for a server (or at the group level) when an error/notification occurs; it will not appear at other times. When the tab does appear, you will see an exclamation icon on the server. .

If you check the *Attention Required* tab at the group level, it will display events from all servers in the group, listed in chronological order. The server name will be included for each event to identify the source of the event.

If you have servers from different time zones in a group, the events will be sorted using coordinated universal time (UTC).

To view only a specific category of events, select the category from the *Filter* drop-down box.

Clear issues from the list

After you have resolved an issue, you can click the check box next to it and click the *Clear* button. You can clear individual issues or you can clear all listed issues by clicking *Select All* and then *Clear*.

Monitor performance

System performance

On a backup server, select *Dashboard Summary* --> *VTL/NAS Performance* to display information about VTL and NAS performance.

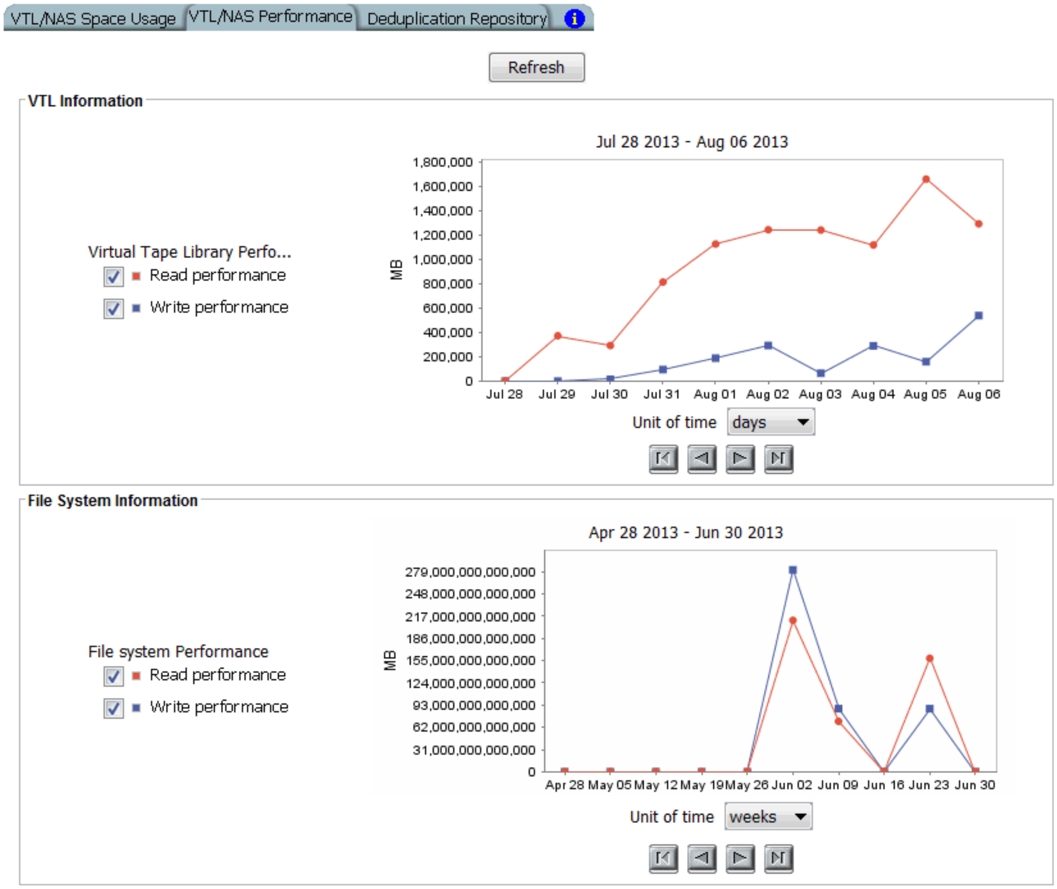

Each section displays read and write throughput for all related resources. Performance statistics are acquired from all adapters (configured in initiator mode and dual mode) and from local storage. If compression is enabled, the write values will be the compressed values. The statistics include all of the I/O data that is transferred regardless of the activity type.

Select a *Unit of time* (hours, days, weeks, or months) from the drop-down list to adjust the granularity of the graph. The data points in the graph will match the starting point for that unit. For example, if you select *Months*, the data point for March will show statistics for just after midnight on March 1. If you select hours, all data read/written between 7:00-8:00 will be displayed at the 7:00 data point. Use the arrow buttons to scan through accumulated data.

You can put your cursor on a data point to see detailed information.

If you want to zoom into the chart to enlarge it, drag your cursor from left to right over the area you want to expand.

When you are finished, drag your cursor from right to left anywhere in the chart and the display will zoom out, back to a normal view.

Note: If the server has failed over, you will not see any information in the dashboard summary. Once failback occurs, dashboard summary information will be displayed.

Object performance

Performance statistics are available for each virtual tape library, tape drive, tape, adapter, LUN, physical tape library/drive, and Fibre Channel client. They are also available at the *Virtual Tape Libraries* level.

At the *Virtual Tape Libraries* level, the *Performance Statistics* tab shows the aggregate throughput of all I/O activity on *all* virtual libraries.

Each *Performance Statistics* tab displays a chart showing read and write throughput for the last 60 minutes. Current performance is also displayed. All information is displayed in MB per second.

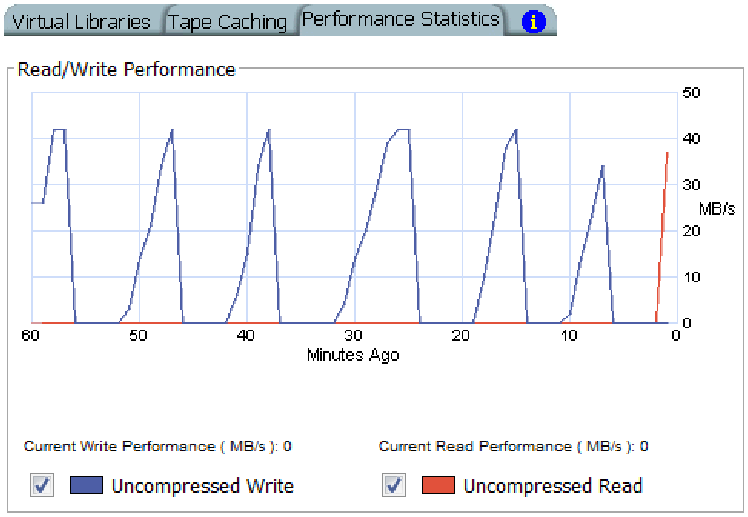

Current Write Performance (MB/s): 0

Current Read Performance (MB/s): 0

To hide a read or write performance chart, click the appropriate checkbox.

Server properties

To set properties for a specific server or group:

1. Right-click the server/group and select *Properties*.
2. On the *Activity Database Maintenance* tab, indicate how often VTL activity data should be purged.

   The Activity Log is a database that tracks all system activity, including all data read, data written, number of read commands, write commands, number of errors etc. This information is used to generate information for the VTL reports. The default values are 50 MB and 365 days.

3. On the *SNMP Maintenance* tab, indicate the system information that should be available in your SNMP manager and the types of event log messages that should be sent as traps to your SNMP manager.

   SysLocation - Enter the location of your system.

   SysContact - Enter contact information. This could be a name or an email address.

   Trap Level - By default, event log messages are *not* sent, but you may want to configure VTL to send certain types of messages. Five levels of messages are available:

   - None – (Default) No messages will be sent.
   - Critical – Only critical errors that stop the system from operating properly will be sent.
   - Error – Errors (failure such as a resource is not available or an operation has failed) and critical errors will be sent.
   - Warning – Warnings (something occurred that may require maintenance or corrective action), errors, and critical errors will be sent.
   - Informational – Informational messages, errors, warnings, and critical error messages will be sent.

   Once you have selected a trap level, the bottom of the dialog will display a table where you can click *Add* to enter information about your SNMP manager.

   If you are configuring SNMP version 1 or 2, you will see the following dialog:

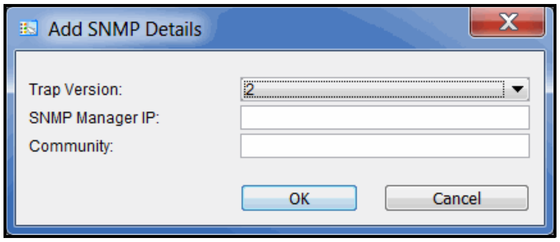

Trap Version: 2

   SNMP Manager IP - IP address of your SNMP server.

   Community - Community name used for SNMP traps (not MIB browsing).

If you are configuring SNMP version 3, you will see the following dialog:

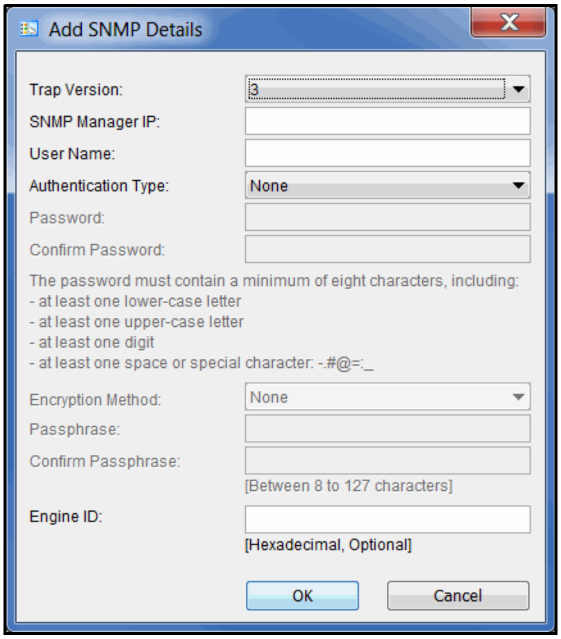

Add SNMP Details

Trap Version: 3

SNMP Manager IP:

User Name:

Authentication Type: None

Password:

Confirm Password:

The password must contain a minimum of eight characters, including:
- at least one lower-case letter
- at least one upper-case letter
- at least one digit
- at least one space or special character: -.#@=:_

Encryption Method: None

Passphrase:

Confirm Passphrase:

[Between 8 to 127 characters]

Engine ID:

[Hexadecimal, Optional]

SNMP Manager IP - IP address of your SNMP server.

User Name - SNMP user.

Authentication Type/Password - If you select the MD5 or SHA algorithm for user authentication, you must enter and confirm the password to use, including at least one lower-case letter, one upper-case letter, and one digit, plus at least one special character: -.#@=:_

Encryption Method/Passphrase - If you select AES or DES for encryption of data sent over the network, you must enter and confirm the passphrase (8-127 characters) to use.

Engine ID - Optional. Enter only if your SNMP manager requires a fixed Engine ID.

4. On the *Performance* tab, indicate if you want to enable replication throttling and then enter the maximum number of KBs per second that should be used for replication bandwidth.

   You can limit the amount of available network bandwidth that is used for replication (of VITs and non-deduplicated virtual tapes) on the source server side. Transmission will not exceed the set value. This is a global server parameter and affects all resources.

   Once enabled, the default is 10 KBs per second. If throttling is not used, replication will use the maximum bandwidth that is available. Besides 0, valid input is 10-1,000,000 KB/s (1G). For example, if throttling is set to 2,000 KB/s, this equates to 15.6Mb/s: (2000/ 1024) * 8 = 15.6Mb/s

5. On the *Auto Save Config* tab, set your system to automatically replicate your system configuration to an FTP server on a regular basis.

Auto Save takes a point-in-time snapshot of the server configuration prior to replication.

Select the *Enable Auto Save Configuration File* option and enter the appropriate information into the fields.

The target server you specify in the *Server Name* field must have FTP server installed and enabled.

The *Target Directory* is an existing directory on the FTP server where the files will be stored. The directory name you enter here (such as `vtlconfig`) is a directory on the FTP server (for example `ftp\vtlconfig`). You should not enter an absolute path like `c:\vtlconfig.`

The *Username/Password* will be the user that the system will log in as. You must create this user on the FTP site. This user must have read/write access to the directory named here.

In the *Interval* field, determine how often to replicate the configuration. Depending upon how frequently you make configuration changes to your system, set the interval accordingly.

In the *Number of Copies* field, enter the maximum copies to keep. The oldest copy will be deleted as each new copy is added.

6. On the *Storage Monitoring* tab, enter the maximum percentage of storage that can be used by VTL before you should be alerted.

   When the utilization percentage is reached, a warning message will be sent to the Event Log. If you have an SNMP manager, the current status can be monitored from there.

7. On the *NAS Activity Database Maintenance* tab, indicate how often NAS activity data should be purged.

   This tab is only visible when NAS is enabled.

   The NAS activity database tracks file system disk usage, share disk usage, SMB activity, file system integrity check activity, NAS deduplication, and replication activity. This information is used to generate information for the NAS reports. The default values are 50 MB and 30 days.

   If you are planning to create reports including more than 30 days of activity, you must increase the number of days to keep data.

8. On the *Location* tab, enter information about the location of this server and who is responsible for maintaining it.

   You can also include a .JPG/.JPEG format photograph of the appliance or its location.

Apply software patch updates

Server patches

The *Version Info* tab displays the current version of the server and console.

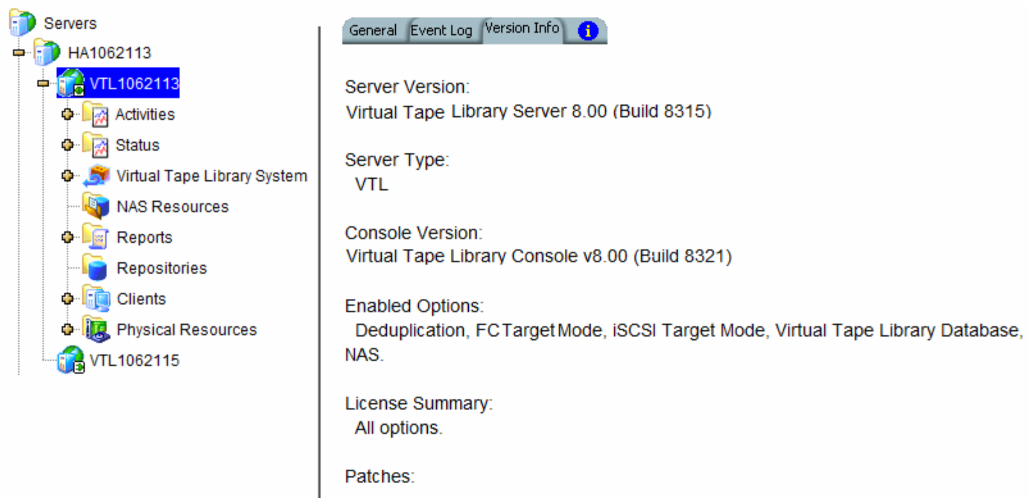

Server Version:
Virtual Tape Library Server 8.00 (Build 8315)

Server Type:
VTL

Console Version:
Virtual Tape Library Console v8.00 (Build 8321)

Enabled Options:
Deduplication, FC Target Mode, iSCSI Target Mode, Virtual Tape Library Database, NAS.

License Summary:
All options.

Patches:

With this information, you can apply maintenance patches to your VTL server through the console.

> **Note:** Server upgrades must be applied directly on the server and cannot be applied or rolled back via the console.

Apply patch - standalone server

To apply a patch on a standalone server:

1. Download the patch onto the computer where the console is installed or a location accessible from that machine.

   Patches can be downloaded from the DSI customer support portal (support.dynamicsolutions.com).

   If you are using the utility that automatically downloads server patches, the patches are located in the `$ISHOME/newpatches` directory.

2. Highlight a server in the tree.
3. Select *Tools* menu --> *Add Patch.*
4. Confirm that you want to continue.
5. Locate the patch file and click *Open*.

   The patch will be copied to the server (if not already there) and installed.

6. Check the Event Log to confirm that the patch installed successfully.

Apply patch - VTL failover configuration

To apply a patch on servers in a failover configuration and avoid unnecessary failover:

1. Make sure both servers are healthy and are not in a failover state.
2. From the console, suspend failover on both servers.
3. Apply the patch on one of the servers.

   Refer to the section above about applying a patch on a standalone server for details.
4. Check the Event Log to confirm that the patch installed successfully.
5. Repeat steps 3 and 4 on the other server.
6. From the console, resume failover on both servers.

Apply patch - SIR redundant node configuration

To apply a patch on servers in a redundant node configuration and avoid unnecessary failover:

1. Stop VTL services on the standby server.

   Before continuing, check the server messages to confirm that the services are stopped.
2. Apply the patch on each active server.

   Refer to the section above about applying a patch on a standalone server for details.
3. Check the Event Log to confirm that the patch installed successfully on each active server.
4. Restart VTL services on the standby server.
5. Apply the patch to the standby server.
6. Check the Event Log to confirm that the patch installed successfully on the standby server.

Roll back patch

To remove (uninstall) a patch and restore the original files:

1. Highlight a server in the tree.
2. Select *Tools* menu --> *Rollback Patch.*
3. Confirm that you want to continue.

   If this server is part of a failover configuration, you must suspend failover before continuing.
4. Select the patch and click *OK*.
5. Check the Event Log to confirm that the patch uninstalled successfully.

Download server patches automatically

A utility is available to automatically download server patches from the DSI support portal. This utility runs on the server and can be run manually or via a schedule (i.e., as a cron job).

In order to use this utility, you must have a valid maintenance contract for the server and all of its options. Additionally, the server must have Internet access. Downloaded patches are stored on the server in the `$ISHOME/newpatches` directory.

To run this utility manually:

1. Execute the following utility on each VTL server:

   ```
   dwlpatch
   ```

2. Navigate to the `$ISHOME/newpatches` directory and apply the newly discovered patches.

To run this utility via schedule:

1. Define a cron job on each VTL server to execute the dwlpatch.pl
2. Enable Email Alerts and enable the `chknewpatch.pl` trigger.

   This trigger checks for new patches in the `$ISHOME/newpatches` directory once a day (by default) and sends an email alert when new patches are detected.

3. Navigate to the `$ISHOME/newpatches` directory and apply the newly discovered patches.

Console patches

Windows console

You need an account with administrator privileges to install the full Windows console package.

1. Close any console that is running.
2. Run the Windows executable file to uninstall the current version of the console.

   You might need to select the *Run as administrator* option to launch the program based on your login account.

3. Re-run the Windows executable file to install the new version.

Java console

1. Close any console that is running.
2. Go to the Bin sub-directory of the console installation folder.
3. Copy the existing console *jar* file to another folder and add the date to the name so that the file can be used as a backup.
4. Copy the new *jar* file to the Bin directory, making sure it has the same name as the existing *jar* file.

Mirror repository disks to protect your configuration

You can mirror the *repository* disks in order to protect your configuration in the event of a hardware failure.

While the data on your tapes will be maintained even if you lose your server, **Mirroring your repository disks is the only way to protect your configuration** if a disk is lost. Mirroring is highly recommended.

When you mirror a disk, each time data is written to the disk, the same data is simultaneously written to the mirrored copy. This disk maintains an exact copy of the original primary disk. In the event that the primary is unusable, VTL seamlessly swaps to the mirrored copy.

For maximum redundancy, the mirror should be on a separate physical device from the primary (preferably on different controllers). The mirror can be defined with disks that are not necessarily identical to each other in terms of vendor, type, or even interface (SCSI, FC, iSCSI).

To set mirroring:

1. Prepare a physical device to use for the mirror.

   Be sure to select the appropriate reservation type, *Configuration Repository* or *Deduplication Repository* (for index, folder, and data disks).

   Refer to ‘Prepare physical storage devices’ for details.

2. Select *Repositories* in the tree, right-click the appropriate object in the right pane, and select *Mirror* --> *Add*.

3. Select the physical device you prepared to use for the mirror.

4. Confirm that all information is correct and then click *Finish* to create the mirroring configuration.

Check mirroring status

You can see the current status of your mirroring configuration by checking the *General* tab of the database.

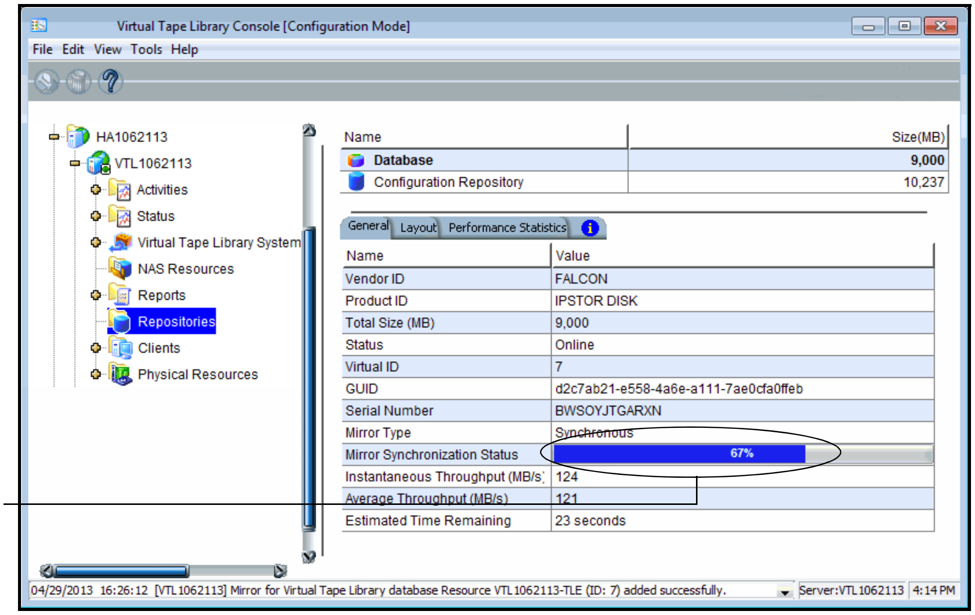

Name	Size(MB)
Database	**9,000**
Configuration Repository	10,237

Name	Value
Vendor ID	FALCON
Product ID	IPSTOR DISK
Total Size (MB)	9,000
Status	Online
Virtual ID	7
GUID	d2c7ab21-e558-4a6e-a111-7ae0cfa0ffeb
Serial Number	BWSOYJTGARXN
Mirror Type	Synchronous
Mirror Synchronization Status	67%
Instantaneous Throughput (MB/s)	124
Average Throughput (MB/s)	121
Estimated Time Remaining	23 seconds

04/29/2013 16:26:12 [VTL1062113] Mirror for Virtual Tape Library database Resource VTL1062113-TLE (ID: 7) added successfully. Server:VTL1062113 4:14 PM

Current status of mirroring configuration.

- *Synchronized* - Both disks are synchronized. This is the normal state.
- *Not synchronized* - A failure in one of the disks has occurred or synchronization has not yet started. If there is a failure in the primary database, VTL swaps to the mirrored copy.
- If the synchronization is occurring, you will see a progress bar along with the percentage that is completed.

Replace a failed disk

If a mirrored disk has failed and needs to be replaced:

1. Select *Repositories*, then right-click the appropriate database object and select *Mirror* --> *Remove* to remove the mirroring configuration.

2. Physically replace the failed disk.

   The failed disk is always the mirrored copy because if the primary database disk fails, VTL swaps the primary with the mirrored copy.

   Important: To replace the disk without having to reboot the server, refer to ‘Replace a failed physical disk without rebooting your server’.

3. Right-click the database object and select *Mirror* --> *Add* to create a new mirroring configuration.

Fix a minor disk failure

If one of the mirrored disks has a minor failure, such as a power loss:

1. Fix the problem (turn the power back on, plug the drive in, etc.).
2. Select *Repositories*, then right-click the appropriate database object and select *Mirror* --> *Synchronize*.

   This re-synchronizes the disks and re-starts the mirroring.

Replace a disk that is part of an active mirror configuration

If you need to replace a disk that is part of an active mirror configuration:

1. If you need to replace the primary database's disk, Select *Repositories*, then right-click the appropriate database object and select *Mirror* --> *Swap* to reverse the roles of the disks and make it a mirrored copy.
2. Select *Mirror* --> *Remove* to cancel mirroring.
3. Replace the disk.

   Important: To replace the disk without having to reboot the server, refer to 'Replace a failed physical disk without rebooting your server'.
4. Right-click the *Database* object and select *Mirror* --> *Add* to create a new mirroring configuration.

Swap the primary disk with the mirrored copy

Select *Repositories*, then right-click the *Database* object and select *Mirror* --> *Swap* to reverse the roles of the primary database disk and the mirrored copy. You will need to do this if you are going to perform maintenance on the primary database disk or if you need to remove the primary database disk.

Replace a failed physical disk without rebooting your server

Do the following if you need to replace a failed physical disk without rebooting your server.

1. If you are not sure which physical disk to remove, execute the following to access the drive and cause the disk's light to blink:

   ```
   ipstorhdparm x x x x
   ```

where x x x x stands for A C S L numbers: Adapter, Channel, SCSI, and LUN number, which you can find in the console.

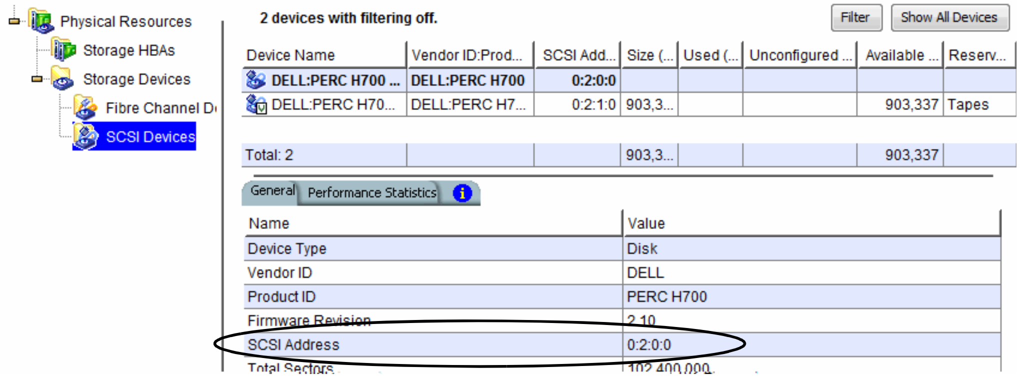

2 devices with filtering off.

Device Name	Vendor ID:Prod...	SCSI Add...	Size (...	Used (...	Unconfigured ...	Available ...	Reserv...
DELL:PERC H700 ...	**DELL:PERC H700**	**0:2:0:0**					
DELL:PERC H70...	DELL:PERC H7...	0:2:1:0	903,3...			903,337	Tapes
Total: 2			903,3...			903,337	

Name	Value
Device Type	Disk
Vendor ID	DELL
Product ID	PERC H700
Firmware Revision	2.10
SCSI Address	0:2:0:0
Total Sectors	102,400,000

2. You MUST remove the SCSI device from the Linux operating system by executing:

   ```
   echo "scsi remove-single-device x x x x">/proc/scsi/scsi
   ```

   where x x x x stands for A C S L numbers: Adapter, Channel, SCSI, and LUN number.

3. Execute the following to re-add the device so that Linux can recognize the drive:

   ```
   echo "scsi add-single-device x x x x">/proc/scsi/scsi
   ```

   where x x x x stands for A C S L numbers: Adapter, Channel, SCSI, and LUN number.

4. Rescan the adapter to which the device has been added.

   In the console, right-click *AdaptecSCSI Adapter.x* and select *Rescan*, where *x* is the adapter number the device is on.

Remove a mirror configuration

Select *Repositories*, then right-click the appropriate database object and select *Mirror* --> *Remove* to delete the mirrored copy and cancel mirroring. You will not be able to access the mirrored copy afterwards.

Mirroring and VTL Failover

If mirroring is in progress during failover/recovery, after the failover/recovery the mirroring will restart from where it left off.

If the mirror is synchronized but there is a Fibre disconnection between the server and storage, the mirror may become unsynchronized. It will resynchronize automatically after failover/recovery.

A synchronized mirror will always remain synchronized during a recovery process.

Manually save/restore the Virtual Tape Library configuration

VTL includes a utility (*vtlrecover*) that enables you to protect your databases and recover your backup server or deduplication server configuration in case of the following:

- The Linux boot disk of the appliance is lost or corrupted.
- The file system where the Virtual Tape Library software is installed is lost.

Information and requirements

- In order to use this utility, all appliance hardware, including FC HBAs and network adapters, storage, physical libraries/drives, and connectivity must be intact and functioning properly.
- The utility does not back up or restore software patches. Patches need to be saved and restored prior to running the restore process.
- The utility does not restore your database mirroring configuration. If you had database mirroring configured before recovery, you will need to reconfigure it after recovery.
- The tape import/export queue and scheduled reporting will not be saved and cannot be restored.
- In a redundant node (N+1) failover environment, we do not recommend using the *vtlrecover* tool to recover the redundant node because it is simpler to just remove the redundant node and reconfigure redundant node failover after the problem is fixed.

After restoring your configuration

- If you deleted any tapes after saving your configuration, those tapes will show up with a red dot (incomplete) after restoring the configuration.
- If you created any tapes after saving your configuration, those tapes will go to the vault.
- If you reclaimed any direct link tapes after saving your configuration, those tapes will show up with red dots.
- After the restore is completed, any expansions or shrinking done after saving your configuration will be adjusted after the restore is completed as part of the normal tape consistency checking done at during startup.

Save your configuration

> **Note:** You should run this utility after you make any major Virtual Tape Library configuration changes.

Use the following procedure to save configuration information:

1. Run the following command: `$ISHOME/bin/vtlrecover save archive.tar`

   This generates an output file which includes all configuration information needed for recovery.

2. If the system is a failover configuration, repeat the first step on each node.

This must be done when the failover system is in a normal mode of operation (i.e., both nodes are functioning correctly).

3. Copy the output file(s) to a safe remote location outside the appliance.

   The file(s) will be needed to restore the configuration.

Restore configuration - standalone system

Use the following procedure to restore Virtual Tape Library configuration on a standalone (non-failover) system.

1. If the Linux operating system is lost, reinstall Linux using an approved procedure.

   For a DSI appliance, use the DSI installation procedure to install the operating system.

2. Configure the hostname, network IP addresses, and other network settings as before.

   > **Note:** The hostname MUST match that of the server that was previously saved.

3. Install Virtual Tape Library software using the recommended installation procedure.

   Be sure to apply the same level of patches as the previous system had.

4. Copy the saved configuration file from the remote location to $ISHOME/bin.

5. Run the following command: `$ISHOME/bin/vtlrecover restore archive.tar.bz2`

   NOTE: Depending on how many tapes are present on a backup server, it may take up to 10 minutes to restore the system.

6. Connect from the console and verify that all configuration information has been restored.

   All virtual tapes, including direct link tapes, will be automatically moved to the appropriate libraries.

Restore a VTL configuration - failover environment

There are two scenarios for restoring a configuration in a failover environment:

- Takeover was successful - Follow the instructions below to restore your configuration.
- Takeover was not successful because the failed server's configuration and database were corrupted - Contact Technical Support for help restoring the configuration.

Takeover was successful

Use the following procedure to restore the Virtual Tape Library configuration on a failed server when the surviving server has taken over successfully.

1. If the Linux operating system is lost, reinstall Linux using an approved procedure.

   For a DSI appliance, use the DSI installation procedure to install the operating system.

2. Configure the hostname, network IP addresses, and other network settings as before.

   > **Note:** The hostname MUST match that of the server that was previously saved.

3. Install Virtual Tape Library software using the recommended installation procedure.

   Be sure to apply the same level of patches as the previous system had.

4. Copy the saved configuration file from the remote location to $ISHOME/bin.

5. Run the following command: `$ISHOME/bin/vtlrecover restore archive.tar.bz2`

   > **Note:** Depending on how many tapes are present on a backup server, it may take up to 10 minutes to restore the system.

6. In the console, right-click the secondary server, then select *Failover* --> *Stop takeover* to perform a failback.

7. Connect from the console and verify that all configuration information has been restored.

   All virtual tapes, including direct link tapes, will be automatically moved to the appropriate libraries.

Multi-Node Groups

If you have multiple VTL tape backup servers, you can create multi-node groups in the console, allowing all servers in the group to be managed together.

Each multi-node group can contain up to eight VTL or eight VTL-S tape backup servers and can include a combination of single servers and/or failover pairs (but failover must be configured before adding them to a group). NAS-only servers and SIR servers cannot be added to a group. VTL servers must have the virtual tape library database created before being added to a group.

A multi-node group can be built by simply connecting all nodes through switches.

The following diagram shows how a multi-node group can be built containing failover pairs and standalone servers.

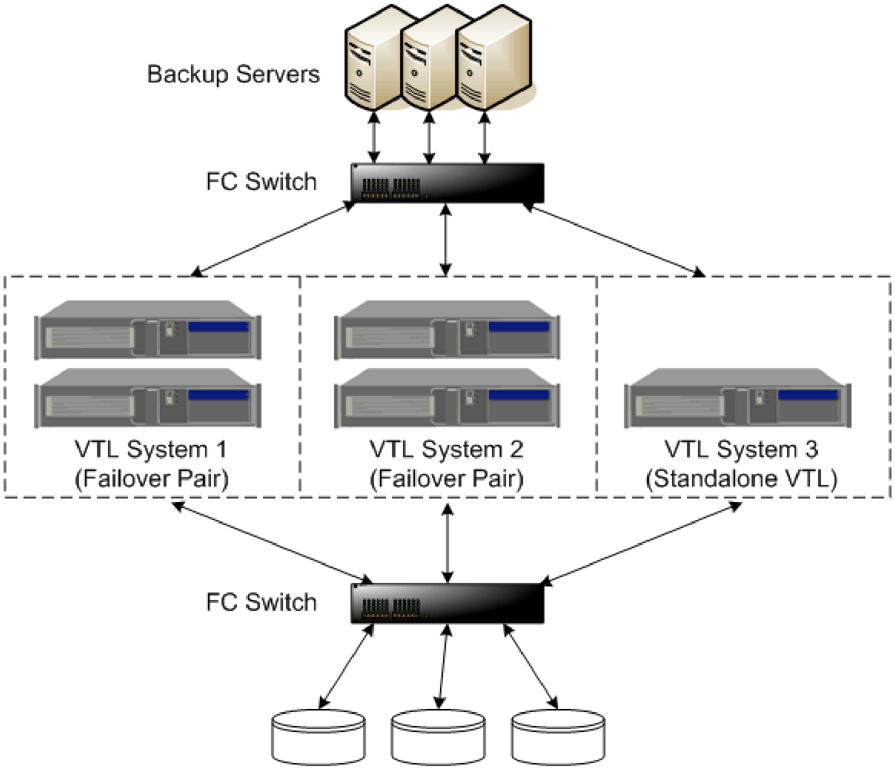

All of the servers in a group can be managed together. The following management functions are available at the group level:

- Single sign-on - Log in to all of the servers in the group at the same time with a single user name and password that exists for all servers in the groups.
- Add/remove members

- Consolidated reporting - VTL tape reports can be generated for all servers in a group. This includes standard reports that are generated on each server in the group and contain data specific to that server. You can also run a consolidated *Group Disk Space Allocation for Virtual Tapes in Libraries Report* that includes every server in the group in one single report.
- Consolidated Event log/Attention required monitoring - The Event log displays events from all servers in chronological order.
- Common configuration settings - Including compression, SNMP, storage monitoring triggers, tape caching thresholds, and X-ray creation.
- Consolidated user management - System users and administrators can be added/deleted at the group level.

Note that if any server in a group is offline, you will not be able to change global properties. In such cases, you will need to remove the offline server from the group before any global properties can be changed.

Create a group

To create a group:

1. Right-click the *Servers* object and select *Create Group*.

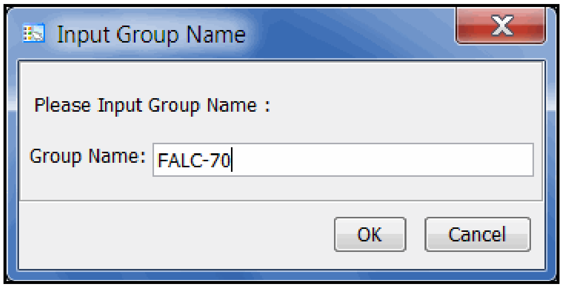

Input Group Name

Please Input Group Name :

Group Name: FALC-70

2. You can also right-click a server and select *Join Multi-Node Group*. If the group name you enter does not already exist, a new group will be created for that server.

3. Enter a name for the group.

   You can enter letters, numbers, a dash, or underscore. Spaces and other characters are not allowed.

Add servers to a group

Notes:

- There is a maximum of eight VTL or eight VTL-S tape backup servers per group.
- VTL and VTL-S servers cannot be in the same group.
- Only VTL (tape) servers can be added to a group; NAS-only servers and SIR servers cannot be added to a group.
- Each server can only be part of one multi-node group.
- You do not need to connect to a server before adding it to a group.
- If you want to add failover servers to a group, failover must be configured first.
- In order to join a group, the new server should have the same user name and password as the servers that are already in group because this is not changed when the server is added to a group.
- Common configuration settings, including hardware/software compression, reporting configuration, SNMP, storage monitoring triggers, and tape caching thresholds, are not automatically applied to servers that are added to the group. If the servers are not all configured the same way, you must manually update each one after adding it to the group.

To add a server, you can do either of the following:

- If you are already connected to a server, right-click the server and select *Join Multi-Node Group*. You will then need to type the group name *exactly* as it appears (names are case sensitive).
- If you are not connected to a server, right-click a group and select *Add Member.* You will then need to enter an IP address and a valid user name and password. When you add subsequent servers, you will only have to enter the IP address. The system will use the user name and password from the first server that you added.

When you are done, all of the servers in a group will be listed in alphabetical order beneath the group in the console. Failover pairs will be displayed together, one below the other.

Your console will now look similar to the following:

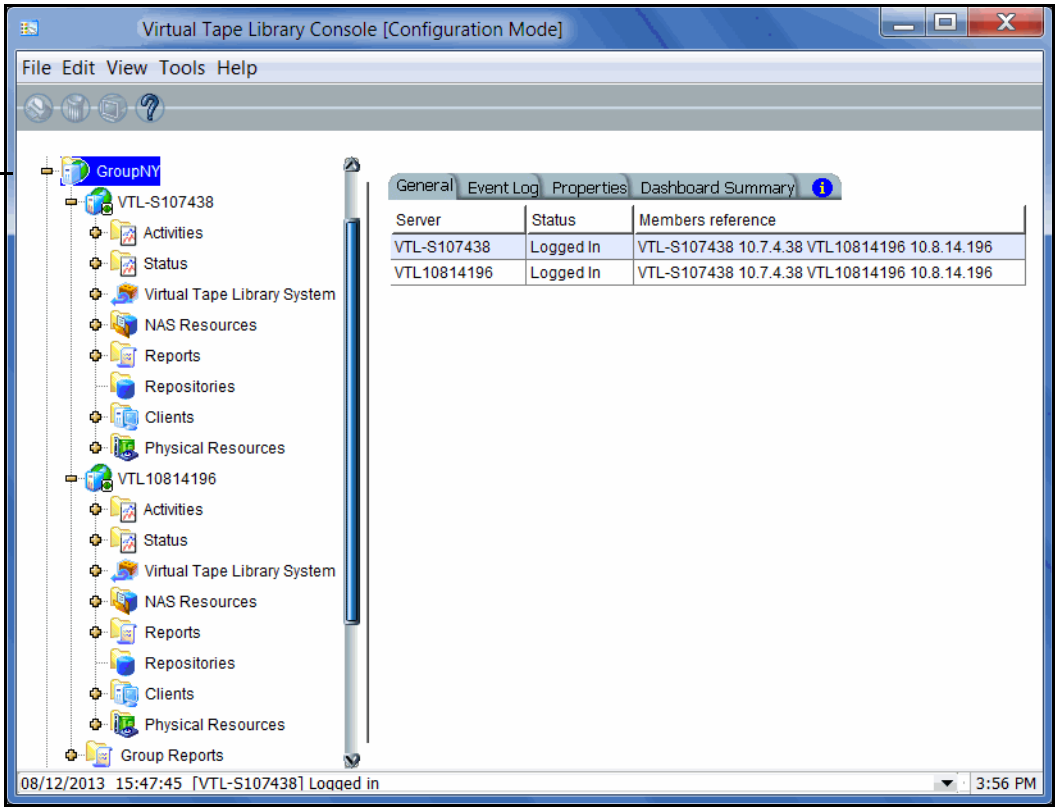

This is a group.

These servers are members of this group.

Server	Status	Members reference
VTL-S107438	Logged In	VTL-S107438 10.7.4.38 VTL10814196 10.8.14.196
VTL10814196	Logged In	VTL-S107438 10.7.4.38 VTL10814196 10.8.14.196

08/12/2013 15:47:45 [VTL-S107438] Logged in

Remove a server from a group

Both online and offline servers can be removed.

Notes:

- If you delete the only server in a group, the group itself will be deleted.
- When a server leaves a group, all administrator accounts that were added at the group level remain with the server.

To remove a server from a group:

1. Right-click the server you want to remove and select *Leave Multi-Node Group*.
2. Answer *Yes* to confirm.

Tape Libraries, Tape Drives, and Tapes

Create virtual tape libraries

You can create virtual tape libraries that emulate your physical tape libraries. If you have a preconfigured VTL appliance with a default virtual tape library that does not emulate your backup server, you can replace the default library.

There are two ways to create a virtual tape library:

- Use the configuration wizard - If you have already configured your system, you can launch the wizard by right-clicking on the *Virtual Tape Library System* object in the console and selecting *Configuration Wizard.*
- Right-click the *Virtual Tape Libraries* object and select *New*.

Note: If you have recently added additional storage to your VTL system, before you can use it to create a virtual tape library, you must reserve it for virtual use. To do this: Right-click *Physical Resources* and select *Prepare Devices.* Set hard drives to *Reserved for Virtual Device*.

1. Select the physical tape library that you are emulating.

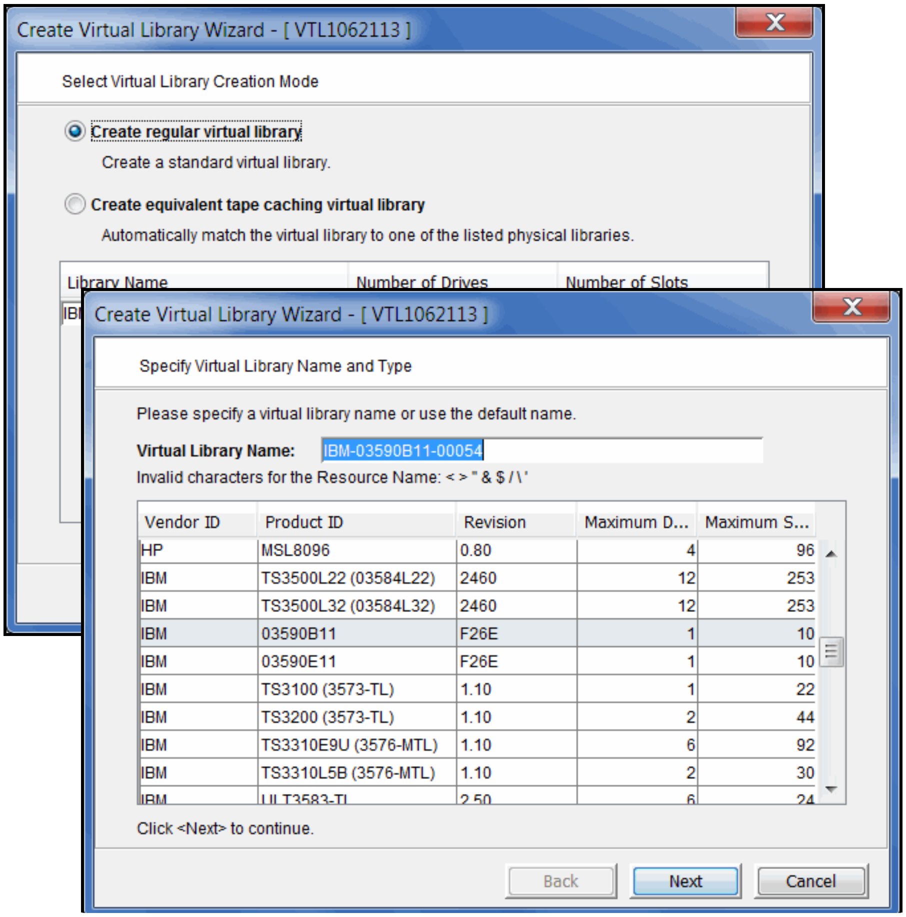

This is the dialog you will see if you have already assigned a physical tape library to your VTL.

Create Virtual Library Wizard - [VTL1062113]

Select Virtual Library Creation Mode

● Create regular virtual library

Create a standard virtual library.

○ Create equivalent tape caching virtual library

Automatically match the virtual library to one of the listed physical libraries.

Library Name	Number of Drives	Number of Slots
IBI		

This is the dialog you will see if you are not using a physical tape library (or have not yet assigned it to your VTL).

Create Virtual Library Wizard - [VTL1062113]

Specify Virtual Library Name and Type

Please specify a virtual library name or use the default name.

Virtual Library Name: IBM-03590B11-00054

Invalid characters for the Resource Name: < > " & $ / \ '

Vendor ID	Product ID	Revision	Maximum D...	Maximum S...
HP	MSL8096	0.80	4	96
IBM	TS3500L22 (03584L22)	2460	12	253
IBM	TS3500L32 (03584L32)	2460	12	253
IBM	03590B11	F26E	1	10
IBM	03590E11	F26E	1	10
IBM	TS3100 (3573-TL)	1.10	1	22
IBM	TS3200 (3573-TL)	1.10	2	44
IBM	TS3310E9U (3576-MTL)	1.10	6	92
IBM	TS3310L5B (3576-MTL)	1.10	2	30
IBM	ULT3583-TL	2.50	6	24

Click <Next> to continue.

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.

> **Note:** For IBM iSeries clients, you must select the IBM 03590E11, IBM TS3500L32(03584L32), or the IBM ULT3583-TL tape library.

Creating an equivalent tape caching library will use VTL *Automated Tape Caching*, which enhances the functionality of VTL by acting as a cache to your physical tape library, providing transparent access to data regardless of its location. If you choose this option, you will only see your available (not already configured) physical tape libraries listed. Select the check box and the system will automatically match your virtual library to the physical library.

> **Note:** The automatic matching of virtual libraries to physical libraries is not available for ACSLS-managed libraries.

2. (Equivalent tape caching library only) Specify a name for the virtual library.

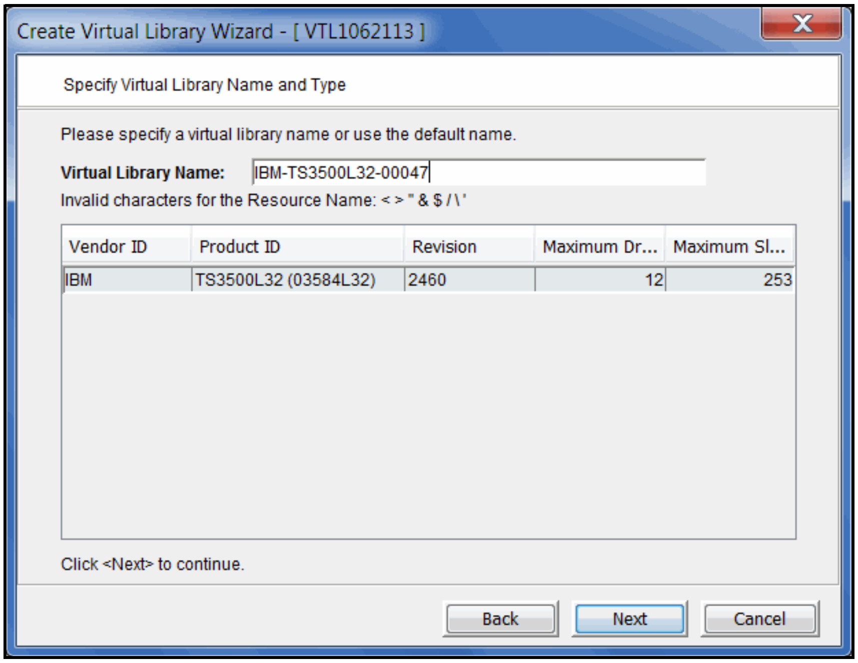

Create Virtual Library Wizard - [VTL1062113]
Specify Virtual Library Name and Type
Back
Next
Cancel

Please specify a virtual library name or use the default name.

Virtual Library Name: IBM-TS3500L32-00047

Invalid characters for the Resource Name: < > " & $ / \ '

Vendor ID	Product ID	Revision	Maximum Dr...	Maximum Sl...
IBM	TS3500L32 (03584L32)	2460	12	253

Click <Next> to continue.

3. Enter information about the tape drives in your library.

This is an example of what you will see if you are creating an equivalent tape caching library.

This is an example of what you will see if you are creating a standard tape library.

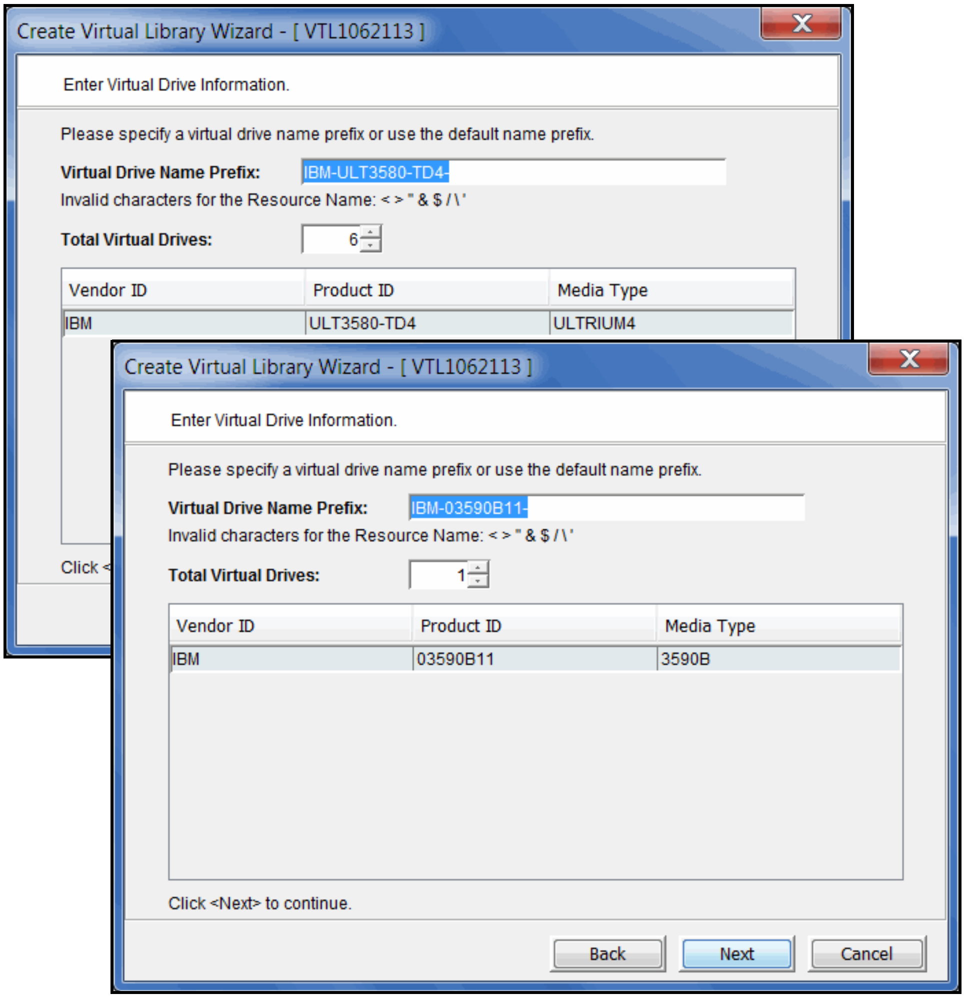

Create Virtual Library Wizard - [VTL1062113]

Enter Virtual Drive Information.

Please specify a virtual drive name prefix or use the default name prefix.

Virtual Drive Name Prefix: IBM-ULT3580-TD4-

Invalid characters for the Resource Name: < > " & $ / \ '

Total Virtual Drives: 6

Vendor ID	Product ID	Media Type
IBM	ULT3580-TD4	ULTRIUM4

Click <

Create Virtual Library Wizard - [VTL1062113]

Enter Virtual Drive Information.

Please specify a virtual drive name prefix or use the default name prefix.

Virtual Drive Name Prefix: IBM-03590B11-

Invalid characters for the Resource Name: < > " & $ / \ '

Total Virtual Drives: 1

Vendor ID	Product ID	Media Type
IBM	03590B11	3590B

Click <Next> to continue.

If you are creating an equivalent tape caching library, the appropriate drives are selected for you.

Virtual Drive Name Prefix - The prefix is combined with a number to form the name of the virtual drive.

Total Virtual Drives - 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.

Note: After creating this virtual tape library, if you need to add drives to it, right-click your library and select *New Drive(s)*.

4. Indicate if you want to use encryption for this virtual tape library.

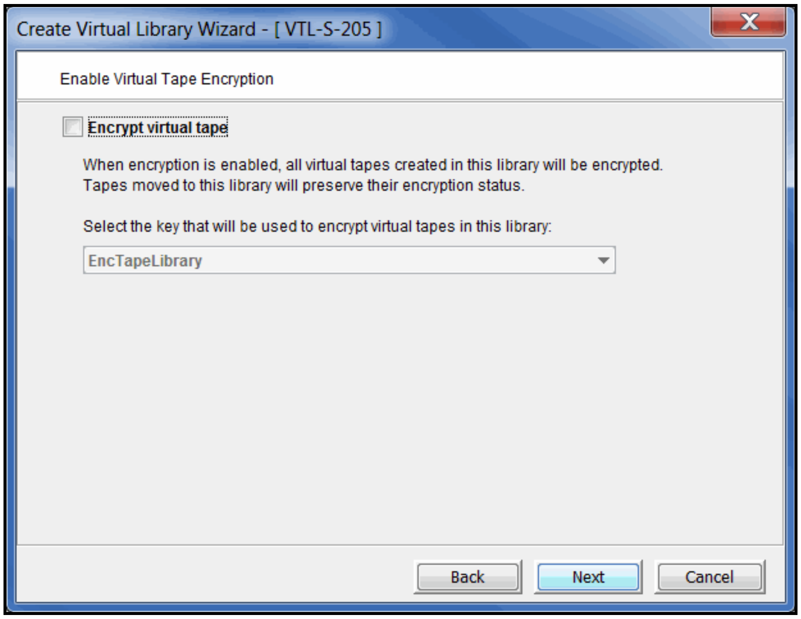

Create Virtual Library Wizard - [VTL-S-205]

Enable Virtual Tape Encryption

□ Encrypt virtual tape

When encryption is enabled, all virtual tapes created in this library will be encrypted.
Tapes moved to this library will preserve their encryption status.

Select the key that will be used to encrypt virtual tapes in this library:

EncTapeLibrary

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. An E icon is displayed on each virtual tape that is encrypted. 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.

5. If this virtual library was not automatically matched to a physical library, indicate if you want to use Automated Tape Caching for this library.

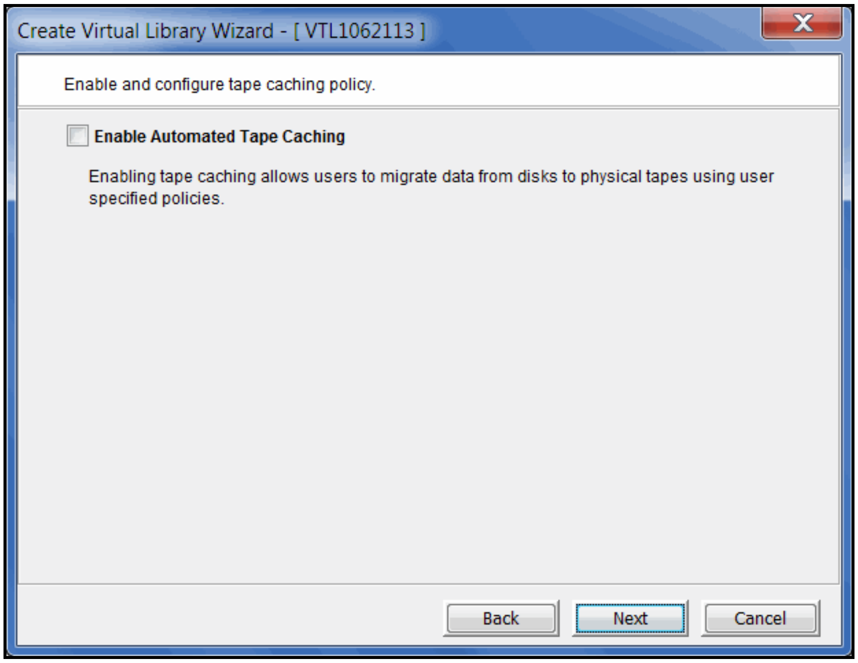

Create Virtual Library Wizard - [VTL1062113]

Enable and configure tape caching policy.

□ **Enable Automated Tape Caching**

Enabling tape caching allows users to migrate data from disks to physical tapes using user specified policies.

Selecting this option here means that the virtual library will act as a cache to your physical tape library, providing transparent access to data regardless of its location.

However, selecting this option here will *not* automatically match this virtual library to a physical library. This means that you do not have to maintain a 1:1 mapping between virtual and physical tape drives.

6. (Automated Tape Caching only) Specify your migration triggers.

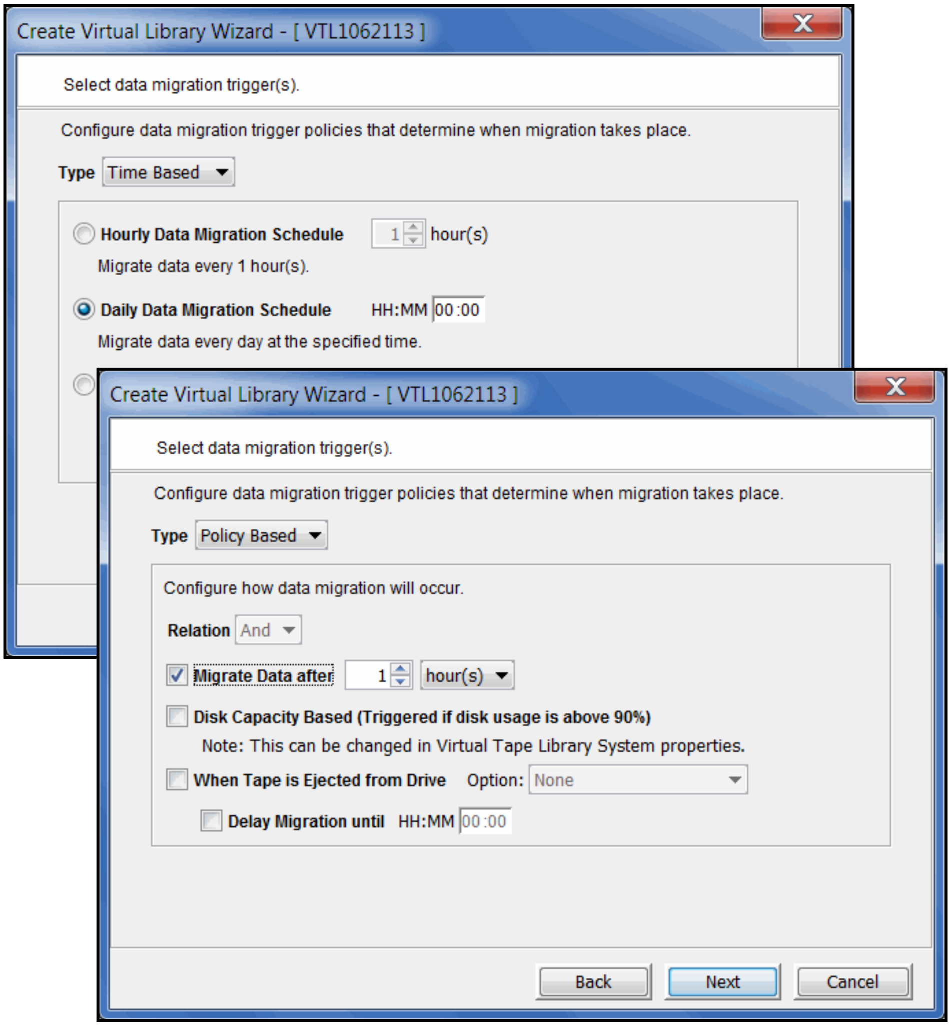

Create Virtual Library Wizard - [VTL1062113]

Select data migration trigger(s).

Configure data migration trigger policies that determine when migration takes place.

Type: Time Based

○ Hourly Data Migration Schedule 1 hour(s)

Migrate data every 1 hour(s).

● Daily Data Migration Schedule HH:MM 00:00

Migrate data every day at the specified time.

Create Virtual Library Wizard - [VTL1062113]

Select data migration trigger(s).

Configure data migration trigger policies that determine when migration takes place.

Type: Policy Based

Configure how data migration will occur.

Relation: And

☑ Migrate Data after 1 hour(s)

□ Disk Capacity Based (Triggered if disk usage is above 90%)

Note: This can be changed in Virtual Tape Library System properties.

□ When Tape is Ejected from Drive Option: None

□ Delay Migration until HH:MM 00:00

Select *Time Based* or *Policy Based* to toggle between the two types of triggers.

Data migration triggers control when data in the cache will be copied to physical tape.

For detailed information about these settings, refer to ‘Automated Tape Caching’.

7. (Automated Tape Caching only) Specify reclamation triggers.

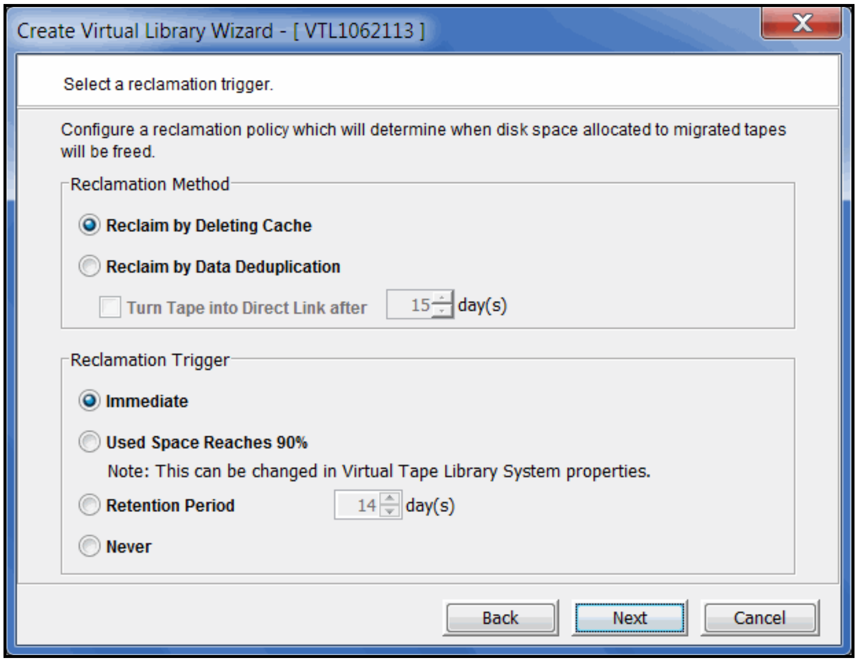

Create Virtual Library Wizard - [VTL1062113]

Select a reclamation trigger.

Configure a reclamation policy which will determine when disk space allocated to migrated tapes will be freed.

Reclamation Method

● Reclaim by Deleting Cache

○ Reclaim by Data Deduplication

□ Turn Tape into Direct Link after 15 day(s)

Reclamation Trigger

● Immediate

○ Used Space Reaches 90%

Note: This can be changed in Virtual Tape Library System properties.

○ Retention Period 14 day(s)

○ Never

Reclamation triggers control when the data that has been migrated to physical tape can be deleted to free up cache disk space.

For detailed information about these settings, refer to ‘Automated Tape Caching’.

8. (Non-Tape Caching environments) Determine if you want to use auto archive/ replication for this virtual library.

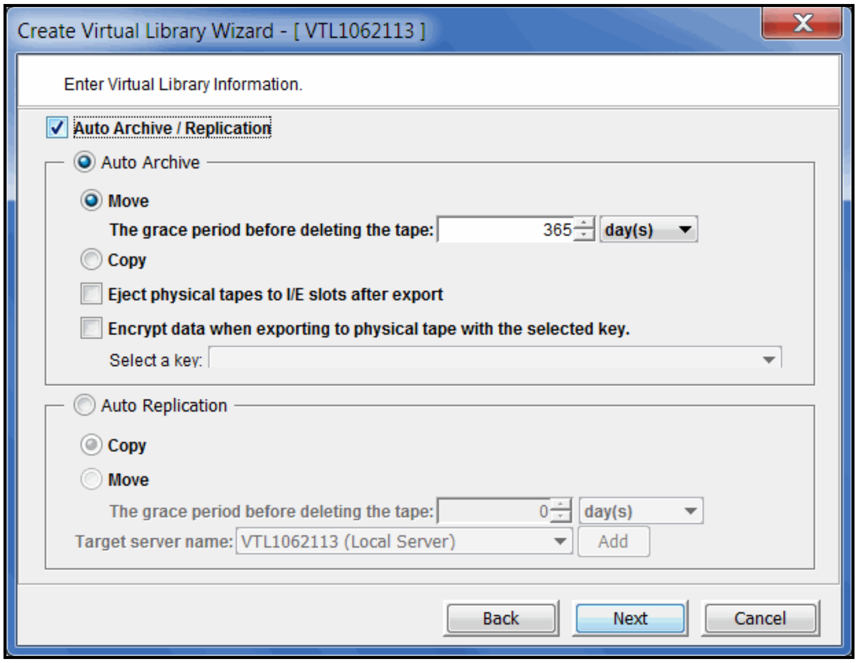

Create Virtual Library Wizard - [VTL1062113]

Enter Virtual Library Information.

☑ Auto Archive / Replication

● Auto Archive

● Move

The grace period before deleting the tape: 365 day(s)

○ Copy

□ Eject physical tapes to I/E slots after export

□ Encrypt data when exporting to physical tape with the selected key.

Select a key:

○ Auto Replication

● Copy

○ Move

The grace period before deleting the tape: 0 day(s)

Target server name: VTL1062113 (Local Server)

You can select either *Auto Archive* or *Auto Replication* for a virtual library, but not both.

Auto Archive writes data to physical tape whenever a virtual tape is moved to an Import/Export (IE) slot from a virtual library by a backup application or other utility after a backup. (You will see the tape in the virtual vault.) In order for the *Auto Archive* function to work*,* the physical tape library must support barcodes because when VTL attempts to export to a physical tape, it must find a matching barcode in a physical library (you do not need to specify which physical library).

- Determine if you want the virtual tape copied (retained) or moved (removed) after the data is transferred. If you select *Move*, indicate how long to wait before deleting it.
- Indicate if you want to eject your *physical* tapes to the library's import/ export slots after exporting.
- You can encrypt the data while exporting as long as you have created at least one key. Refer to 'Create a key' for more information.

Auto Replication replicates data to another VTL server whenever a virtual tape is moved to an IE slot from a virtual library (such as from a backup application or other utility). If selected, determine whether you want the virtual tape copied (retained) or moved (removed) after the data is replicated. If you select *Move*, indicate how long to wait before deleting it. Also, select the remote server from the list of existing target servers. You can also click *Add* to add another VTL server.

9. If you have at least two physical tape libraries (same model, same number of drives, same tapes with the same barcodes) connected to your system, indicate if you want to enable Tape Duplication for this library.

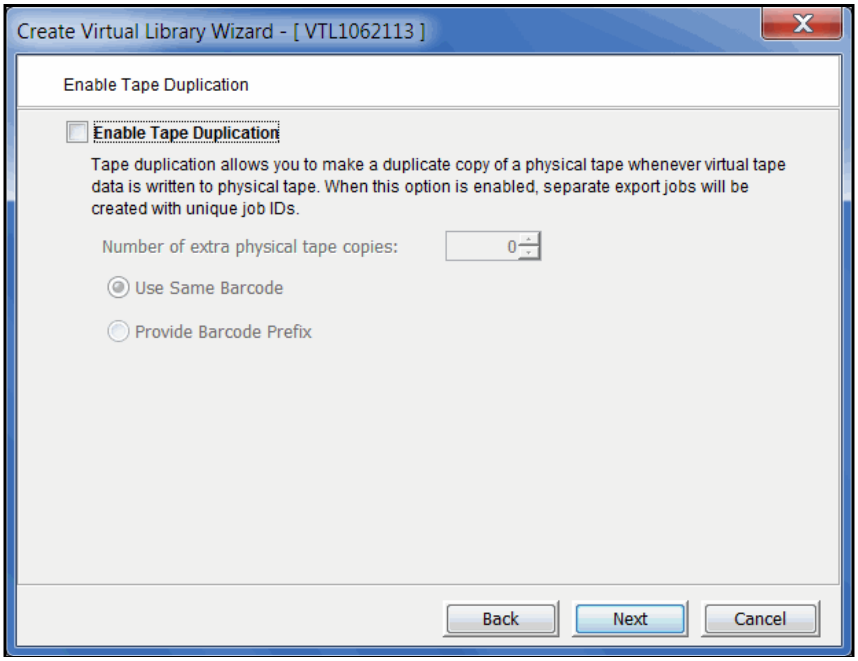

Create Virtual Library Wizard - [VTL1062113]

Enable Tape Duplication

□ Enable Tape Duplication

Tape duplication allows you to make a duplicate copy of a physical tape whenever virtual tape data is written to physical tape. When this option is enabled, separate export jobs will be created with unique job IDs.

Number of extra physical tape copies: 0

● Use Same Barcode

○ Provide Barcode Prefix

Tape duplication allows you to make up to five duplicate copies of a physical tape whenever virtual tape data is exported to physical tape. The number of physical libraries controls the number of duplicate copies; you need six physical libraries to make five duplicates.

You can select whether the physical tape will have the same barcode as the virtual tape or will use a specific prefix that replaces the first character of the barcode. For example, if your virtual tape barcode is “123456” and you specify that the prefix is “A”, the system will look for a physical tape with the barcode “A23456”.

The duplication job will look for a tape with the correct barcode in one of the physical libraries. If one is found, the data is duplicated to that physical library. If an appropriate tape is not found, but there are additional physical libraries, the system will continue to look for a match.

When data is exported, separate export jobs will be created for each physical library and each job will have a unique job ID. Multiple duplication jobs will run concurrently.

If this library is using Automated Tape Caching or if you selected the *Move* option for Auto Archive on the previous dialog, the virtual tape data will not be deleted until the duplication job finishes successfully.

Notes:

- You should not have duplicate physical tape barcodes in your system *unless* you are using tape duplication.
- Once you configure Tape Duplication, you should not unassign any of the physical libraries from VTL.
- If this library is using Automated Tape Caching, once the virtual tape has been reclaimed and a direct link tape exists, tape duplication will no longer occur and only the primary physical tape can be updated.

10. Enter barcode information for the virtual library.

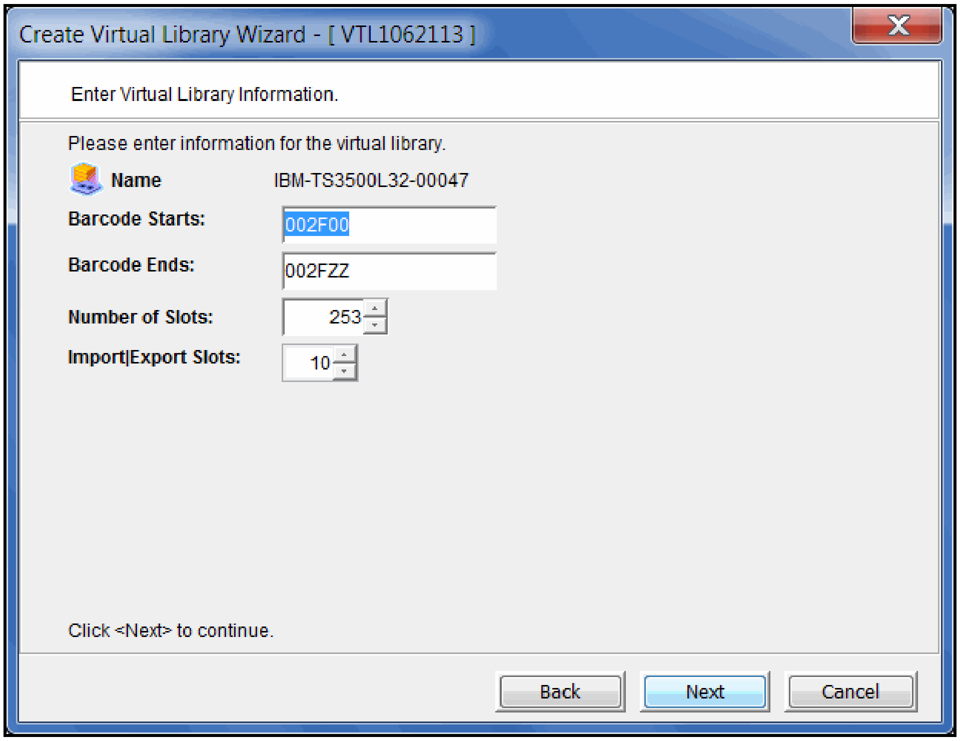

Create Virtual Library Wizard - [VTL1062113]
Enter Virtual Library Information.

Please enter information for the virtual library.

Name IBM-TS3500L32-00047

Barcode Starts: 002F00

Barcode Ends: 002FZZ

Number of Slots: 253

Import|Export Slots: 10

Click <Next> to continue.

Barcode Starts/Ends - 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 *Barcode Ends* field to **999**; for example, **XXX0999**

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

Slot - Maximum number of tape slots in your tape library.

Import/Export Slots - Number of slots used to take tapes in and out of the bin.

Note: If you are using an HP EML E-Series library with LTO drives with IBM® Tivoli® Storage Manager, you need to change the default library barcode to six digits.

11. Enter the guidelines for expanding virtual tape capacity.

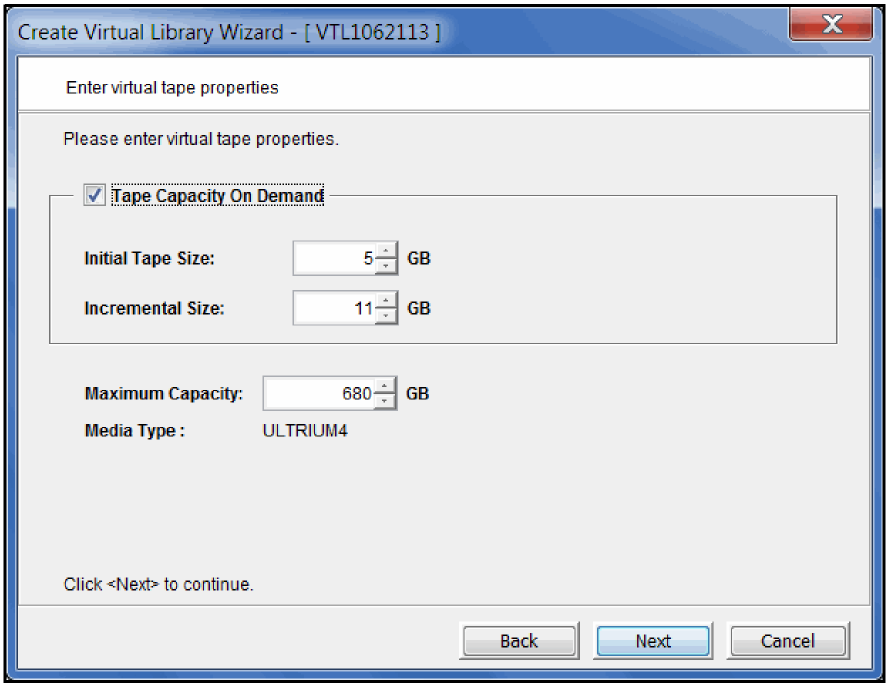

Create Virtual Library Wizard - [VTL1062113]

Enter virtual tape properties

Please enter virtual tape properties.

☑ Tape Capacity On Demand

Initial Tape Size: 5 GB

Incremental Size: 11 GB

Maximum Capacity: 680 GB

Media Type : ULTRIUM4

Click <Next> to continue.

You will only see this dialog if you have enabled the *Advanced Tape Creation* method (set in *Tools --> Console Options).* If *Advanced Tape Creation* is not enabled, Tape Capacity On Demand will automatically be set for you.

Tape Capacity On Demand - Allows you 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. If you do not select this option, VTL will allocate each virtual tape at the full size of the tape you are emulating.

If Tape Capacity on Demand is used, when a tape is overwritten, all disk segments beyond the segment being written to are freed up and the tape is reset to its initial size. Space allocated for a replica resource will be adjusted to match the primary tape allocation before the replication starts, optimizing the disk space used by replica resources.

Initial Tape Size/Incremental Size - Enter the initial size of each resource and the amount by which it will be incremented.

Maximum Capacity - Indicate the maximum size for each tape.

- 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 VTL’s 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 this virtual library will use tape caching, you should not resize the virtual tapes because migration will fail when the amount of data exceeds the size of the physical tape.

- If you *will* be exporting data to physical tape *and you will* be using VTL's 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 VTL's compression algorithm can vary depending upon the dataset; certain file types (ZIP, PDF, GIF, RAR, etc.) are already compressed and cannot be compressed further.

12. Verify all information and then click *Finish* to create the virtual tape library.

    If you are not using Automated Tape Caching, you will be prompted to create virtual tapes. Answer *Yes* to continue. Refer to the following section for more information about creating virtual tapes.

    If you are using Automated Tape Caching and you selected the *Reclaim by data deduplication* reclamation trigger, the system will automatically create an associated deduplication policy. After synchronizing the virtual library to your physical library, you will be able to edit this deduplication policy, at which point you can configure replication, if desired. Refer to 'Create tape deduplication policies' for more information about configuring replication for a deduplication policy.

    If you are using Automated Tape Caching, you will be prompted to synchronize your virtual library to your physical library. Refer to 'Create a cache for your physical tapes' for more information.

Create standalone virtual tape drives

You can create standalone virtual tape drives that emulate your physical tape drives.

> **Note:** This procedure is for *standalone* virtual tape drives. If you want to add virtual tape drives to an existing virtual tape library, right-click your library and select *New Drive(s)*.

1. Right-click the *Virtual Tape Drives* object and select *New*.

2. Select the physical tape drive you are emulating.

3. Enter the guidelines for expanding virtual tape capacity.

   You will only see this dialog if you have enabled the *Advanced Tape Creation* method (set in *Tools --> Console Options).* If *Advanced Tape Creation* is not enabled, Tape Capacity On Demand will automatically be set for you.

   Tape Capacity On Demand - Allows you 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. If you do not select this option, VTL will allocate each virtual tape at the full size of the tape you are emulating.

   If Tape Capacity on Demand is used, when a tape is overwritten, all disk segments beyond the segment being written to are freed up and the tape is reset to its initial size. Space allocated for a replica resource will be adjusted to match the primary tape allocation before the replication starts, optimizing the disk space used by replica resources.

   Initial Tape Size/Incremental Size - Enter the initial size of each resource and the amount by which it will be incremented.

   Maximum Capacity - Indicate the maximum size for each tape.

   - 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 VTL's 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 VTL's 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 VTL's compression algorithm can vary depending upon the dataset; certain file types (ZIP, PDF, GIF, RAR, etc.) are already compressed and cannot be compressed further.

4. Verify all information and then click *Finish* to create the virtual tape drive.

Create virtual tapes

You can create virtual tapes in the following ways:

- After you create a virtual tape library, you will be prompted to create tapes for it.
- Use the configuration wizard - If you have already configured your system, you can launch the wizard by right-clicking on the *Virtual Tape Library System* object in the console and selecting *Configuration Wizard.* Skip to Step 2, which lets you create a virtual library and tapes for that library.
- Right-click a virtual tape library or the *Tapes* object and select *New Tape(s)*.

The *Create Virtual Tape wizard* will vary depending on whether or not you have enabled the *Advanced Tape Creation* method (set in *Tools --> Console Options).*

1. (*Advanced Tape Creation* only) Select how you want to create the virtual tape(s).

   Custom lets you select which physical device(s) to use and lets you designate how much space to allocate from each.

   Express automatically creates the resource(s) for you using available device(s). If you select *Express*, you can create multiple virtual tapes at the same time.

2. Indicate if you want to enable remote export and, if you do, select a target server.

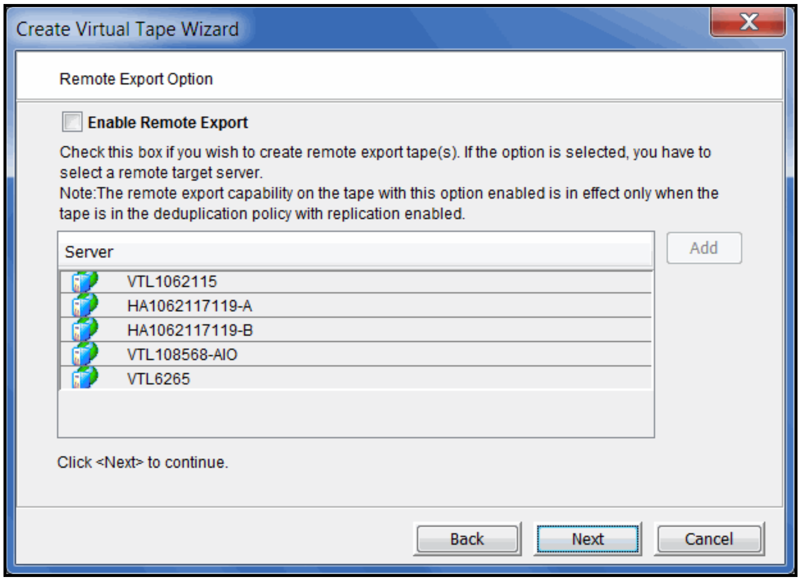

Create Virtual Tape Wizard

Remote Export Option

☐ **Enable Remote Export**

Check this box if you wish to create remote export tape(s). If the option is selected, you have to select a remote target server.
Note:The remote export capability on the tape with this option enabled is in effect only when the tape is in the deduplication policy with replication enabled.

Server
VTL1062115
HA1062117119-A
HA1062117119-B
VTL108568-AIO
VTL6265

Click <Next> to continue.

Remote export automatically exports the contents of a replicated virtual index tape (VIT) to a physical tape at the remote site, providing a copy of a physical tape at the remote site.

The target server must have an available physical tape library and the media type must match that of the virtual tape library.

If remote export is enabled for a virtual tape, the virtual tape must be added into a deduplication policy.

Once replication is successfully completed, an export job is submitted at the replica site.

If a deduplication policy includes cascaded replication, the tape will be added to the policy only if the target server is the same as one of the two target servers in the policy. This tape will be remotely exported to physical tape(s) on the target servers only if physical tapes with the same barcodes exist on the target servers.

3. (*Remote export* only) Select a physical library on the target server that should be used to synchronize barcodes.

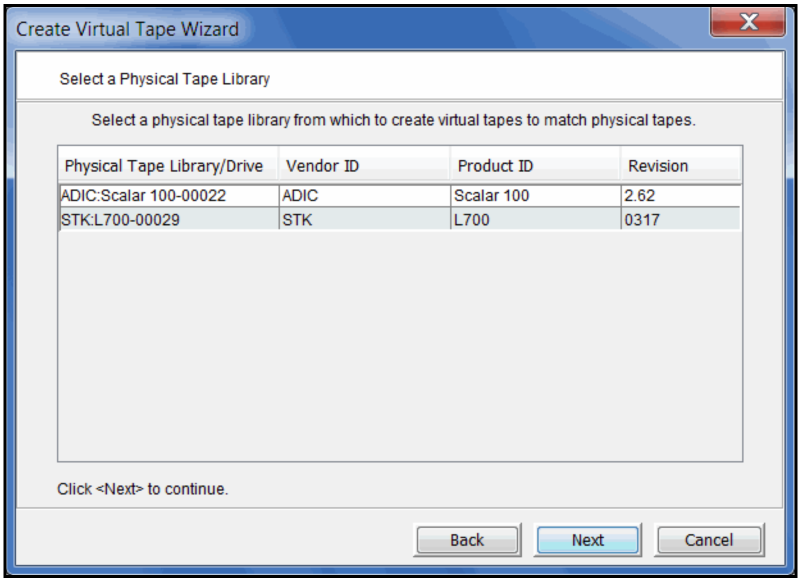

Select a Physical Tape Library

Select a physical tape library from which to create virtual tapes to match physical tapes.

Physical Tape Library/Drive	Vendor ID	Product ID	Revision
ADIC:Scalar 100-00022	ADIC	Scalar 100	2.62
STK:L700-00029	STK	L700	0317

Click <Next> to continue.

4. (*Remote export* only) Select the physical tapes you want to synchronize with at the remote site.

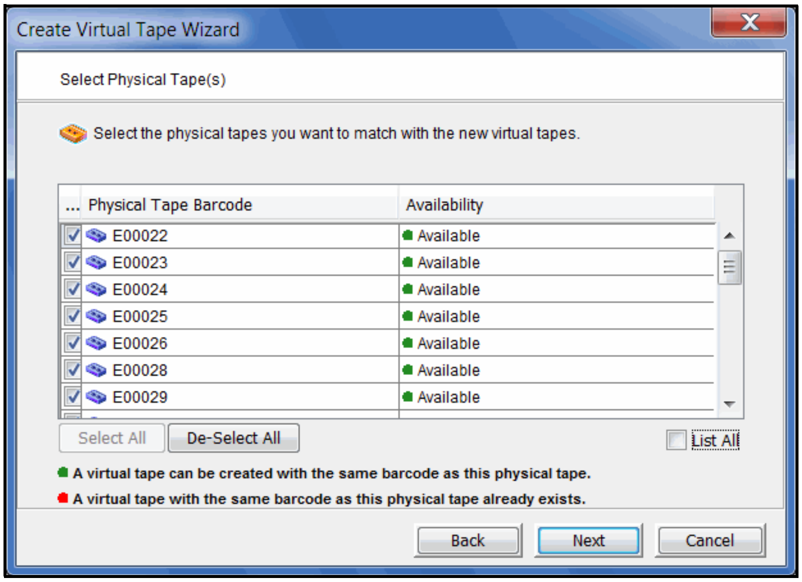

Select Physical Tape(s)

Select the physical tapes you want to match with the new virtual tapes.

...	Physical Tape Barcode	Availability
☑	E00022	🟩 Available
☑	E00023	🟩 Available
☑	E00024	🟩 Available
☑	E00025	🟩 Available
☑	E00026	🟩 Available
☑	E00028	🟩 Available
☑	E00029	🟩 Available

🟩 A virtual tape can be created with the same barcode as this physical tape.

🟥 A virtual tape with the same barcode as this physical tape already exists.

5. If *Auto Archive* is enabled for the virtual library, select the physical tape library you want to match with virtual tapes.

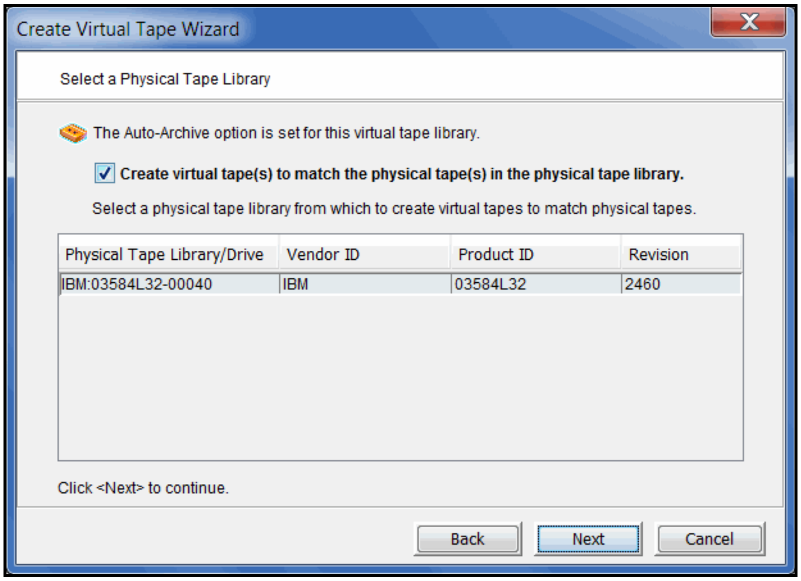

Select a Physical Tape Library

The Auto-Archive option is set for this virtual tape library.

☑ **Create virtual tape(s) to match the physical tape(s) in the physical tape library.**

Select a physical tape library from which to create virtual tapes to match physical tapes.

Physical Tape Library/Drive	Vendor ID	Product ID	Revision
IBM:03584L32-00040	IBM	03584L32	2460

Click <Next> to continue.

This enables you to have a physical tape with a barcode that matches your virtual tape. This is important for exporting functions.

If you selected the *Auto Archive* option but have not yet connected a physical tape library to the VTL appliance, you can either create virtual tapes without setting a matching barcode or discontinue creating virtual tapes at this time and exit the dialog.

6. If *Auto Archive* is enabled for the virtual library and you are matching physical tapes, select the physical tapes for which you want to create matching virtual tapes.

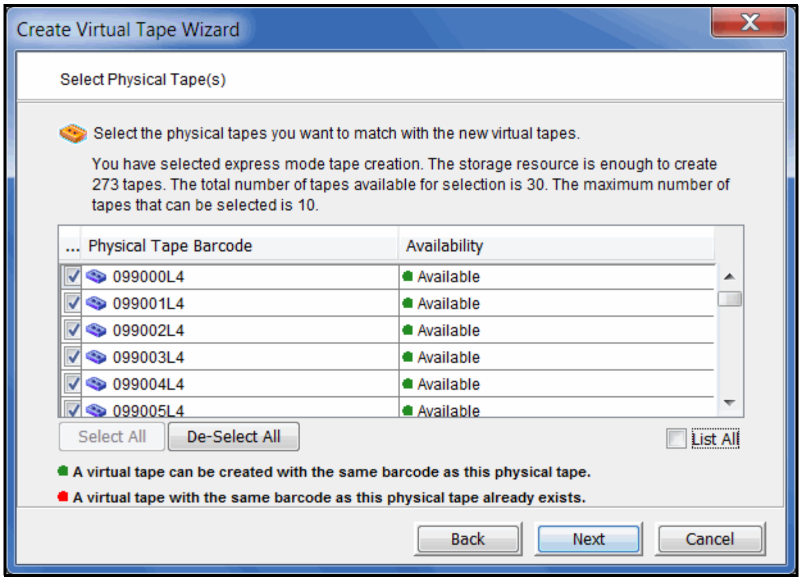

Select Physical Tape(s)

Select the physical tapes you want to match with the new virtual tapes.

You have selected express mode tape creation. The storage resource is enough to create 273 tapes. The total number of tapes available for selection is 30. The maximum number of tapes that can be selected is 10.

...	Physical Tape Barcode	Availability
☑	099000L4	🟢 Available
☑	099001L4	🟢 Available
☑	099002L4	🟢 Available
☑	099003L4	🟢 Available
☑	099004L4	🟢 Available
☑	099005L4	🟢 Available

🟢 **A virtual tape can be created with the same barcode as this physical tape.**

🔴 **A virtual tape with the same barcode as this physical tape already exists.**

7. (*Advanced Tape Creation* only) If you selected the *Custom* creation method, specify which physical device should be used to create the virtual tapes.

   Storage space is allocated from the local server even if this server is part of a multi-node group.

8. Depending upon which method you selected, specify the tape prefix, tape size, and, if applicable, the number of tapes to create.

You will be able to specify the tape name if the *Advanced Tape Creation* method is enabled.

You will see this dialog if the *Advanced Tape Creation* method is not enabled.

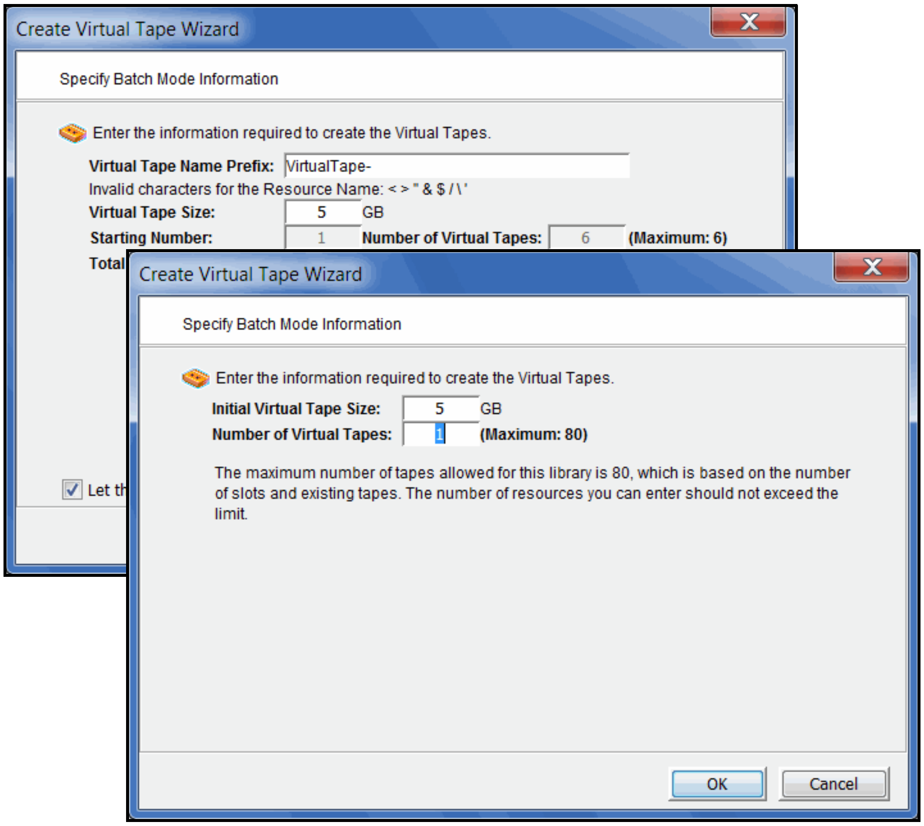

Create Virtual Tape Wizard

Specify Batch Mode Information

Enter the information required to create the Virtual Tapes.

Virtual Tape Name Prefix: VirtualTape-

Invalid characters for the Resource Name: < > " & $ / \ '

Virtual Tape Size: 5 GB

Starting Number: 1 Number of Virtual Tapes: 6 (Maximum: 6)

Total

☑ Let th

Create Virtual Tape Wizard

Specify Batch Mode Information

Enter the information required to create the Virtual Tapes.

Initial Virtual Tape Size: 5 GB

Number of Virtual Tapes: 1 (Maximum: 80)

The maximum number of tapes allowed for this library is 80, which is based on the number of slots and existing tapes. The number of resources you can enter should not exceed the limit.

If you enabled *remote export* in the previous dialogs, you cannot specify the number of tapes to create; the number of virtual tapes will be fixed by the number of physical tapes you want to synchronize.

9. If *Auto Replication* is enabled for the virtual library and you want it enabled for this/these tapes, select the target server.

   You will be asked to confirm the hostname/IP address and indicate how long the system should attempt to replicate data before timing out and how often it should attempt to retry before skipping a scheduled replication.

   Then, indicate if you want to use the *Compression* and/or *Encryption* options. The *Compression* option provides enhanced throughput during replication by compressing the data stream. The *Encryption* option secures data transmission over the network during replication.

Notes:

- Do not enable auto replication for tapes for which you will be defining a deduplication policy. This feature is not supported for virtual index tapes (VITs).
- Encryption must be enabled on the target server; all keys used by the source tapes must exist on both servers and be identical. This means that the keys have the same name and were created using the same secret phrase. If the secret phrase is not the same, you can export a key from the source server and import it to the target.
- Compression/encryption for transmission over a network should not be set if the source tapes are already encrypted.

10. (*Advanced Tape Creation* only) If you are not matching physical tapes, you can set a barcode range for the virtual tapes you are creating.

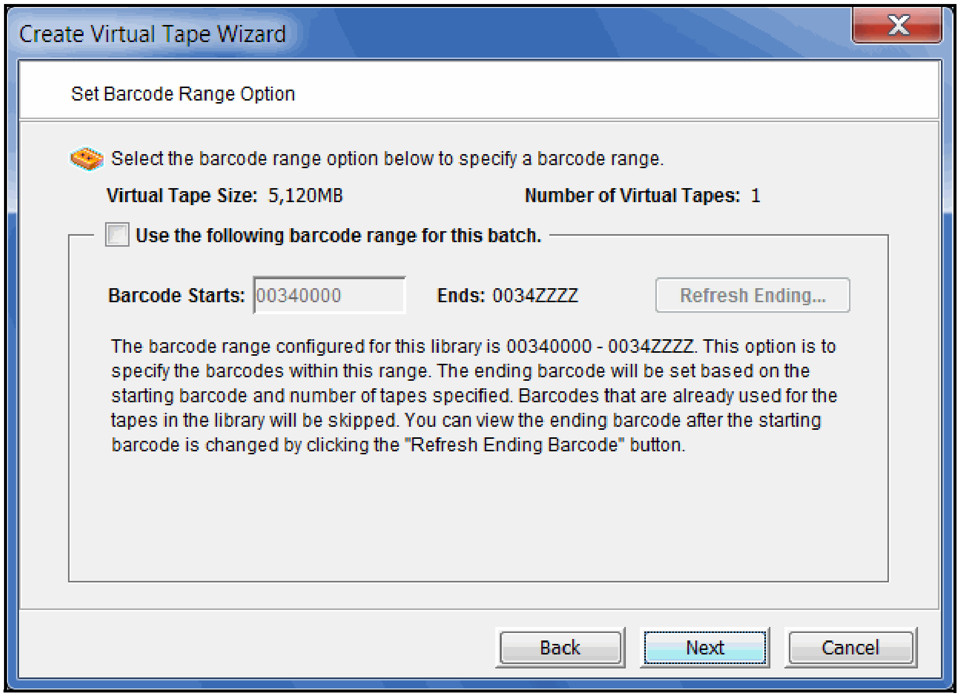

Create Virtual Tape Wizard

Set Barcode Range Option

Select the barcode range option below to specify a barcode range.

Virtual Tape Size: 5,120MB **Number of Virtual Tapes:** 1

□ **Use the following barcode range for this batch.**

Barcode Starts: 00340000 **Ends:** 0034ZZZZ

The barcode range configured for this library is 00340000 - 0034ZZZZ. This option is to specify the barcodes within this range. The ending barcode will be set based on the starting barcode and number of tapes specified. Barcodes that are already used for the tapes in the library will be skipped. You can view the ending barcode after the starting barcode is changed by clicking the "Refresh Ending Barcode" button.

11. Verify all information and then click *Finish* to create the virtual tape(s).

How virtual tapes are allocated

VTL uses a sophisticated methodology to determine which LUN to use when allocating space for virtual tapes.

Using two algorithms, *Dynamic LUN Allocation* and *Round Robin*, virtual tapes being expanded or created in *Express* mode are allocated from the LUN that is currently experiencing the least amount of I/O. (Virtual tapes created in *Custom* mode use the LUN that is specified.)

When virtual space is needed, the system looks at the available LUNs. A scoring method is used to determine how *busy* each LUN is. If the scores for all LUNs are equal (i.e. all LUNs are *free* or all are equally busy), Round Robin logic is used to select the next available LUN in the rotation queue.

Once a LUN is selected, VTL looks to see if there is enough continuous space available on the LUN to match what is needed. If there is enough, that space is allocated. Afterward, the LUN is pushed to the “back” of the queue, ensuring that tapes are evenly distributed across all LUNs.

If there is not enough continuous space on a single LUN, VTL allocates the biggest chunk of continuous space available. Smaller chunks on the same LUN are then allocated (but never chunks less than 1 GB) to reach the total amount needed. If there is not enough space available, VTL continues allocating from another LUN.

When Tape Capacity on Demand is used and tape expansion is needed, VTL will attempt to expand the tape on the current LUN, provided there is enough space available. Once that LUN is filled, Round Robin logic will select the next LUN to allocate space from.

By default, a single allocation pool is used for all available storage. All available LUNs are assigned to this pool. For enhanced performance, multiple allocation pools can be defined to further distribute I/O between multiple controllers and RAID units. Because configuration is different for every environment, contact DSI Professional Services if you would like to configure multiple allocation pools for LUN allocation.

Locate and display virtual tapes in the Console

Because it is possible to have a large number of virtual tapes, we have included tools to help you locate just the tape(s) you are looking for.

Search by barcode

To search by barcode for a specific virtual tape:

1. Highlight any object on the server where the tape resides.
2. Select *Edit* menu --> *Find.*

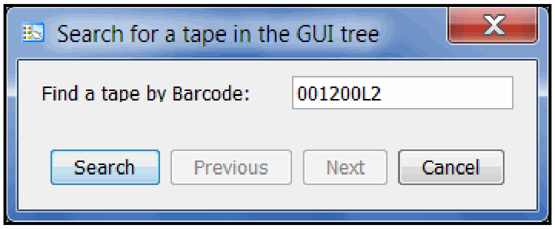

Find a tape by Barcode: 001200L2

3. Enter the full barcode.

   Note that the search is case sensitive. Once you click *Search*, you will be taken directly to that tape in the right pane.

Display virtual tapes

When you highlight the *Tapes* object in the tree, a list of all tapes in that virtual library is displayed in the right-hand pane. When you highlight the *Virtual Vault* object, a list of all tapes in the vault is displayed in the right-hand pane. The icon next to the tape name indicates the status of the last operation performed (“D” for deduplication and “R” for deduplication with replication).

Icon	Color	Source Server	Target Server
D	Green	The last deduplication process was successful (pure VIT)	NA
D	Yellow	Deduplication is pending or in-progress (not a pure VIT)	NA
D	Red	The last deduplication process failed (not a pure VIT)	NA
R	Green	The last replication process was successful	The tape has been successfully resolved
R	Yellow	The replication process is pending or in-progress on the source server	NA

Icon	Color	Source Server	Target Server
R	Red	The last replication process was unsuccessful	The last attempt at resolving the tape failed or the tape is currently being resolved or has not been resolved

While the right pane is usually just for informational purposes, you can perform tape functions directly from the right pane by highlighting one or more tapes and using the right-click context menu. You can also highlight any tape to see detailed tape information in the lower part of the pane.

For single tapes, the right-click menu allows you to create a remote copy; rename the tape; delete the tape; move the tape to the virtual vault, slot, or drive; configure replication, and display/set tape properties (barcode, tape capacity on demand, write protection, auto archive, auto replication, tape duplication, and remote export).

For multiple selected tapes, the right-click menu allows you to delete the tapes, move them to the virtual vault, and configure replication.

To load tapes into all empty virtual tape drives or to dismount tapes from all virtual tape drives, right-click the virtual tape library object in the tree and select *Auto Load Tapes* or *Auto Unload Tapes*.

Sort all tapes

You can sort the tapes displayed in the right-hand pane. To do this:

1. Select the appropriate heading in the drop-down box next to *Sort*.
2. Indicate whether they should be sorted in *Ascending* or *Descending* order.

Filter the display of tapes

Because it is possible to have a large number of tapes in the right-hand pane, you may want to filter the tapes and display only specific tapes. To do this:

1. Click the *Filter* button.

2. On the *General* tab, you can indicate the type of tape(s) you are looking for.

You will see this dialog if you started from the *Tapes* object.

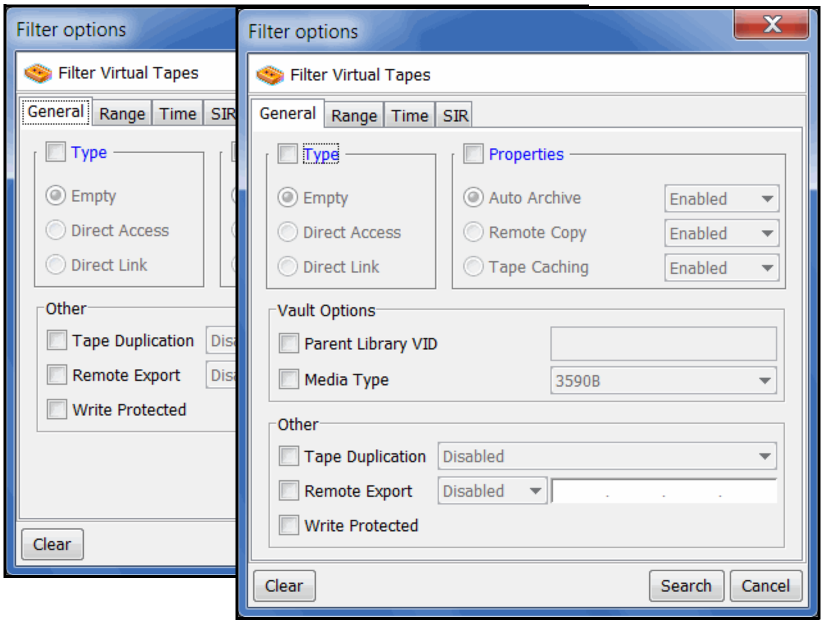

Filter options — Filter Virtual Tapes (General tab)

- ☐ Type: ● Empty ○ Direct Access ○ Direct Link
- ☐ Properties: ● Auto Archive: Enabled; ○ Remote Copy: Enabled; ○ Tape Caching: Enabled
- Vault Options: ☐ Parent Library VID; ☐ Media Type: 3590B
- Other: ☐ Tape Duplication: Disabled; ☐ Remote Export: Disabled; ☐ Write Protected

You will see this dialog if you started from the *Virtual Vault* object.

The dialog will offer different options depending upon whether you are in the virtual vault or not.

3. On the *Range* tab, you can enter a range of barcodes and/or sizes.

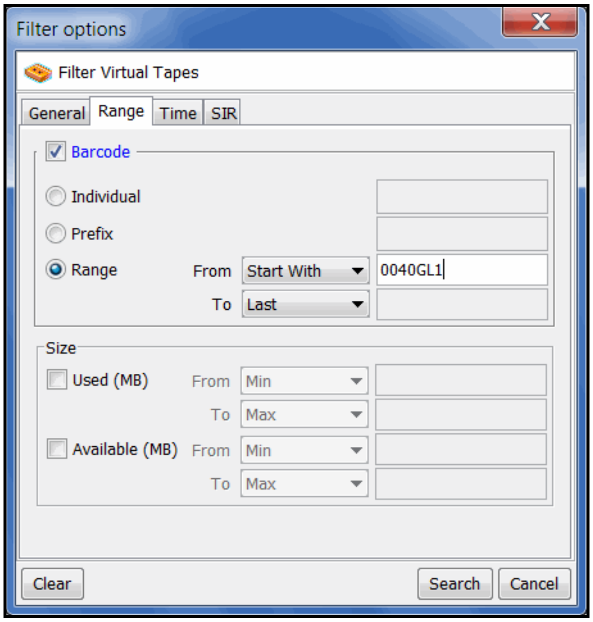

Filter options — Filter Virtual Tapes (Range tab)

- ☑ Barcode: ○ Individual ○ Prefix ● Range — From: Start With 0040GL1; To: Last
- Size: ☐ Used (MB) From: Min To: Max; ☐ Available (MB) From: Min To: Max

If you want to specify a particular number, select *Start With* or *End With* in the *From/To* fields. You can then type the number in the box to the right.

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.

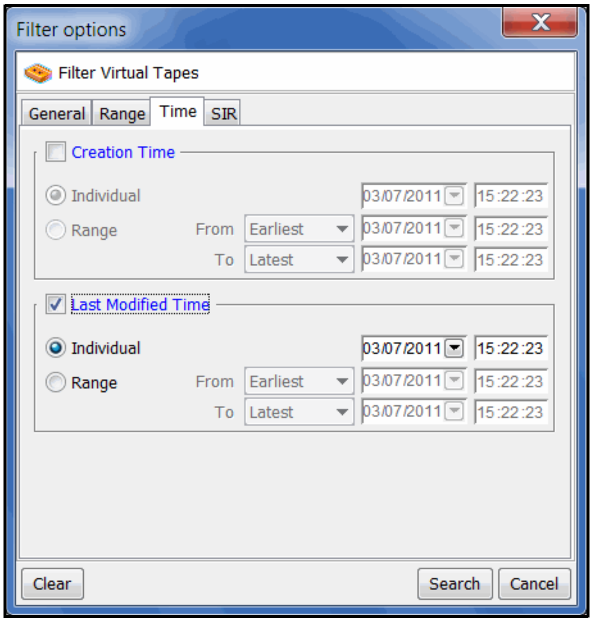

Filter options

Filter Virtual Tapes

□ Creation Time

◉ Individual 03/07/2011 15:22:23

○ Range From Earliest 03/07/2011 15:22:23

To Latest 03/07/2011 15:22:23

☑ Last Modified Time

● Individual 03/07/2011 15:22:23

○ Range From Earliest 03/07/2011 15:22:23

To Latest 03/07/2011 15:22:23

If you want to specify a particular date/time, select *Start At* or *End At* in the *From/ To* fields. You can then change the number in the box to the right.

5. On the *SIR* tab, you can look for tapes associated with a deduplication policy or from a specific source server.

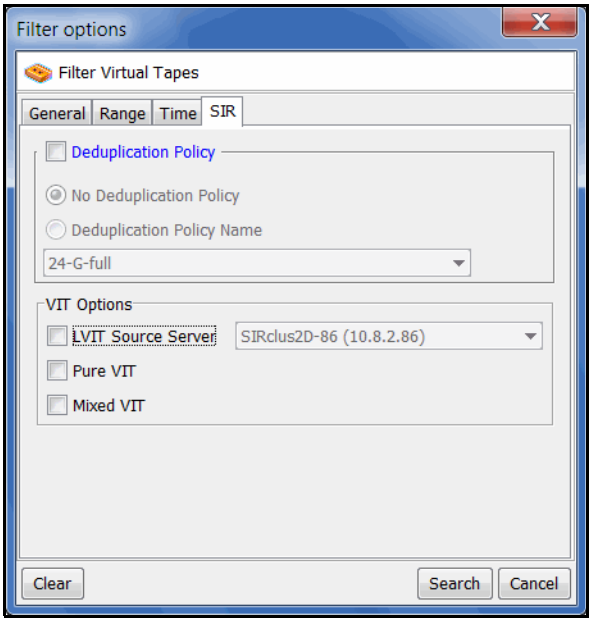

Filter options

Filter Virtual Tapes

□ Deduplication Policy

◉ No Deduplication Policy

○ Deduplication Policy Name

24-G-full

VIT Options

□ LVIT Source Server SIRclus2D-86 (10.8.2.86)

□ Pure VIT

□ Mixed VIT

6. Click *Search*.

Afterwards, *just* the tapes that match the selected criteria will be displayed in the right pane. You can click the *Show All Tapes* button when you are done.

Assign virtual tape libraries and drives to backup servers

You can assign a virtual tape library or drive to the target of a backup server listed in the VTL console under the *Clients* object. The backup server can then access the assigned virtual tape library/drive(s).

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

There are three ways to assign a library or drive to a client (backup server):

- Use the configuration wizard - If your system is already configured, you can launch the wizard by right-clicking on the *Virtual Tape Library System* object in the console and selecting *Configuration Wizard.* After adding a virtual tape library, you can assign it to a backup server.
- Begin with a client object and select a virtual tape library or drive.
- Begin with a virtual tape library or drive object and select the backup server to assign it to.

Configuration wizard or client object

If you started from the configuration wizard or a client object, follow these steps to continue:

1. Select a virtual tape library or drive.

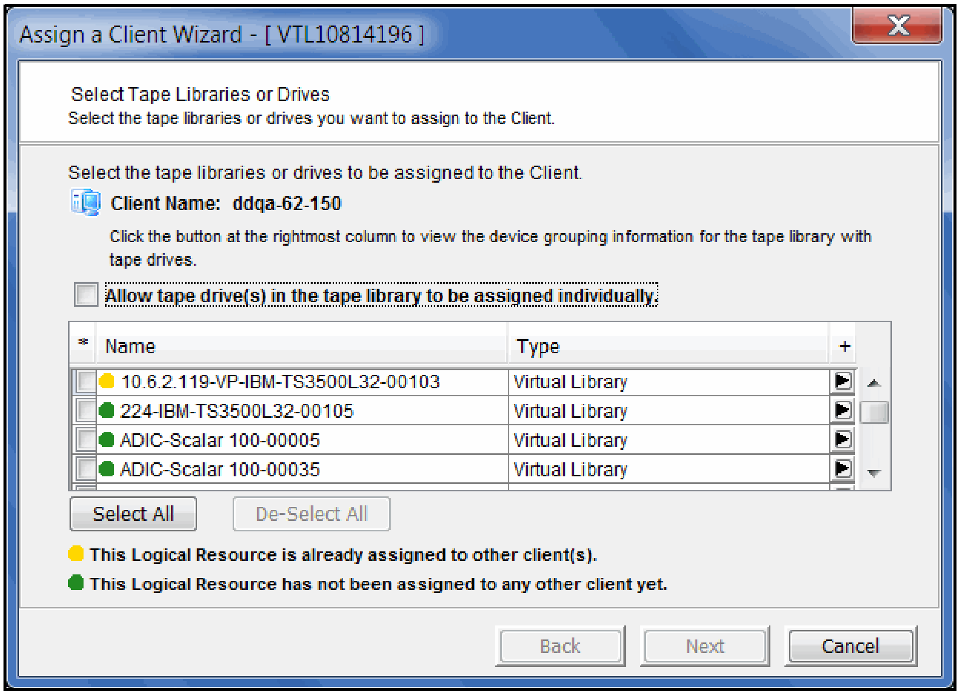

Assign a Client Wizard - [VTL10814196]

Select Tape Libraries or Drives

Select the tape libraries or drives you want to assign to the Client.

Select the tape libraries or drives to be assigned to the Client.

Client Name: ddqa-62-150

Click the button at the rightmost column to view the device grouping information for the tape library with tape drives.

☐ **Allow tape drive(s) in the tape library to be assigned individually.**

*	Name	Type	+
☐	🟡 10.6.2.119-VP-IBM-TS3500L32-00103	Virtual Library	▶
☐	🟢 224-IBM-TS3500L32-00105	Virtual Library	▶
☐	🟢 ADIC-Scalar 100-00005	Virtual Library	▶
☐	🟢 ADIC-Scalar 100-00035	Virtual Library	▶

🟡 **This Logical Resource is already assigned to other client(s).**

🟢 **This Logical Resource has not been assigned to any other client yet.**

All tape drives in a library will be assigned to the selected client.

If you want to assign tape drives in the library individually, select the checkbox for that option. The VTL server and backup server will treat each individually assigned drive as if it were a standalone tape drive.

2. Click *Finish* when you are done.
3. Use the backup server's operating system to discover the VTL server.

   The steps to do this vary according to the backup server's operating system.

   For Fibre Channel environments, if your zoning has been correctly configured, and devices have been properly assigned to clients, a simple bus rescan performed on the client should show the new backup devices. Of course, this procedure varies depending on the OS.

   For Windows, *Control Panel* --> *Computer Management* --> *Device Manager* --> right-click the device in the right pane --> *Scan for hardware changes*.
4. Use your backup software to discover the library.

   The steps to do this vary according to your backup software.

Virtual tape library or drive

If you started from a virtual tape library or drive, follow these steps to continue:

1. Select the appropriate protocol for the backup server to which you want to assign the library or drive.
2. Select a backup server.

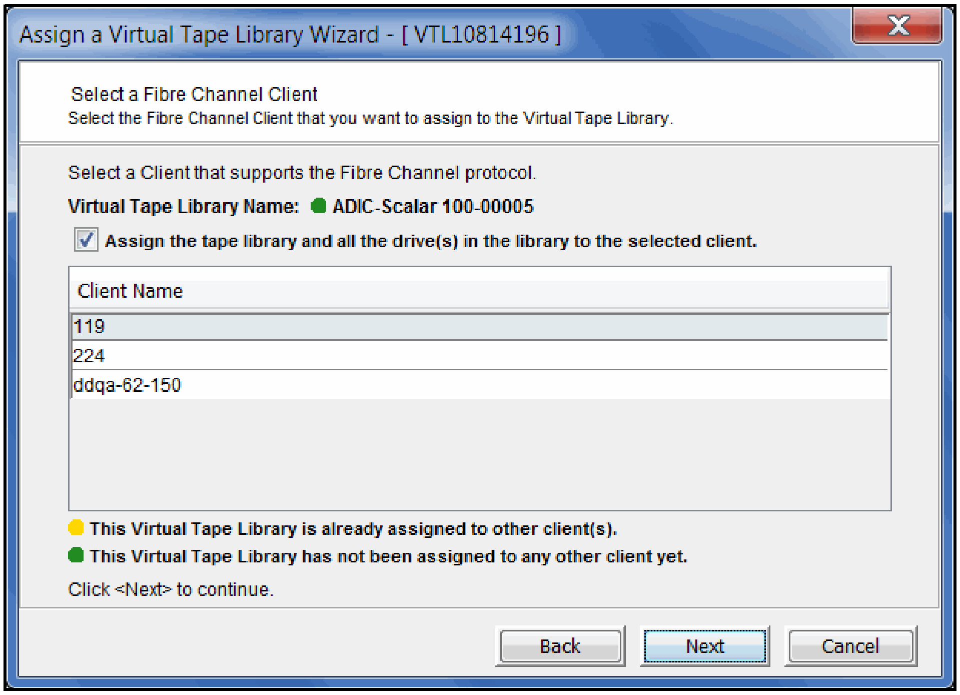

Assign a Virtual Tape Library Wizard - [VTL10814196]

Select a Fibre Channel Client

Select the Fibre Channel Client that you want to assign to the Virtual Tape Library.

Select a Client that supports the Fibre Channel protocol.

Virtual Tape Library Name: ● **ADIC-Scalar 100-00005**

☑ **Assign the tape library and all the drive(s) in the library to the selected client.**

Client Name
119
224
ddqa-62-150

● **This Virtual Tape Library is already assigned to other client(s).**

● **This Virtual Tape Library has not been assigned to any other client yet.**

Click <Next> to continue.

3. Click *Next* and then click *Finish* when you are done.

Physical tape libraries

You can import data from physical tapes into your virtual tape library or export data from virtual tapes or deduplicated virtual tapes to physical tapes. Before you can perform these operations, do the following:

1. Connect the physical tape library to the VTL system.

   This can be done via a Fibre Channel connection.

2. Rescan devices to make sure the server can see the new physical library/drives.

   You can now identify these new devices in the navigation tree. Expand the *Physical Resources* object and then the *Fibre Channel Devices* object.

   On the *General* tab of the console information pane, a physical library has a Device Type of *Medium Changer* and a physical drive has a Device Type of *Tape Device*. New devices have a Category value of *Unassigned*.

3. Assign the physical library/drives to the virtual tape library (refer to 'Assign physical libraries/drives to VTL').

4. For export purposes, create virtual tapes that correspond to barcoded physical tapes (refer to 'Create virtual tapes').

Assign physical libraries/drives to VTL

If you will be importing data from physical tapes into your virtual tape library or exporting virtual tapes to physical tapes, you must assign your physical tape libraries/drives to VTL. This process also inventories the physical tapes in your library/drive so that you can create virtual tapes that match your physical tapes.

Notes:

- A physical tape library can only be assigned to one VTL server at a time, unless the ACSLS or IBM 3494 option is being used.
- VTL does not support physical libraries when tape drive numbering does not start with 0 or is not sequential.

1. Assign physical libraries/drives to VTL in one of the following two ways:

   - Use the configuration wizard - You can launch the wizard by right-clicking on the *Virtual Tape Library System* object in the console and selecting *Configuration Wizard.* If you haven't already prepared your physical library/ drive, you can do that as well*.*
   - Right-click the *Physical Tape Libraries* object or the *Physical Tape Drives* object and select *Assign*.

2. Select the physical libraries/drives to be assigned to VTL.

3. Click *Finish/Assign* to assign.

Inventory physical tapes

You can perform an inventory of the physical tapes in a physical tape library. This allows you to create virtual tapes that match your physical tapes. To do this, right-click a physical tape library and select *Inventory*.

Designate a physical library or drive as disabled

For maintenance purposes, a physical library or drive can be marked as disabled. While it is designated as disabled, the library/drive cannot be used for any import/export functions. To do this, right-click a physical tape library or drive and select *Disable*. Afterwards, when the library/drive is available, you can enable it.

Reset physical tapes in a library

Resetting a physical library reinitializes the library and puts tapes back in the appropriate slots. This function is useful after a problem (such as a physical tape stuck in a slot) is resolved. To reset a library, right-click a physical tape library and select *Reset.*

Import data from tapes

One of the advantages of using a virtual tape library is that you can protect data on your existing physical tapes by importing them into your virtual tape system.

You can also import data from virtual tapes being used by another DSI virtual tape library.

If you need to recover files from a physical tape, you can use the import function to directly access the physical tape for immediate recovery.

Import data from a physical tape

The import function allows you to:

- Copy the contents of a physical tape to a virtual tape
- Directly access a physical tape without copying the entire tape
- Recycle a physical tape
- Import data from virtual tapes used by another DSI virtual tape library

Before you import a tape

You must do the following before you can import tapes:

- Verify that the drive type (i.e. IBM:ULTRIUM-TD3) and the media type (i.e. Ultrium3) of the library you are importing from match the library to which you are importing.
- If you are connecting through a switch, verify that the VTL server is in the same zone as the physical tape library or other VTL server from which you are importing.
- If you are importing from a standalone physical tape drive, you must right-click the drive and select *Retrieve Barcode* before importing.

Import a tape

1. Right-click your physical tape library or physical tape drive and select *Import Tape*.

   If the tape was exported with stacking, select *Import Stacked Tapes*.

2. Select which virtual library to import into.

   Be sure to pick a drive/library with the same tape size capacity.

   Note that stacked tapes can only be imported to virtual libraries that are not enabled with Automated Tape Caching.

3. Select how you want the data copied.

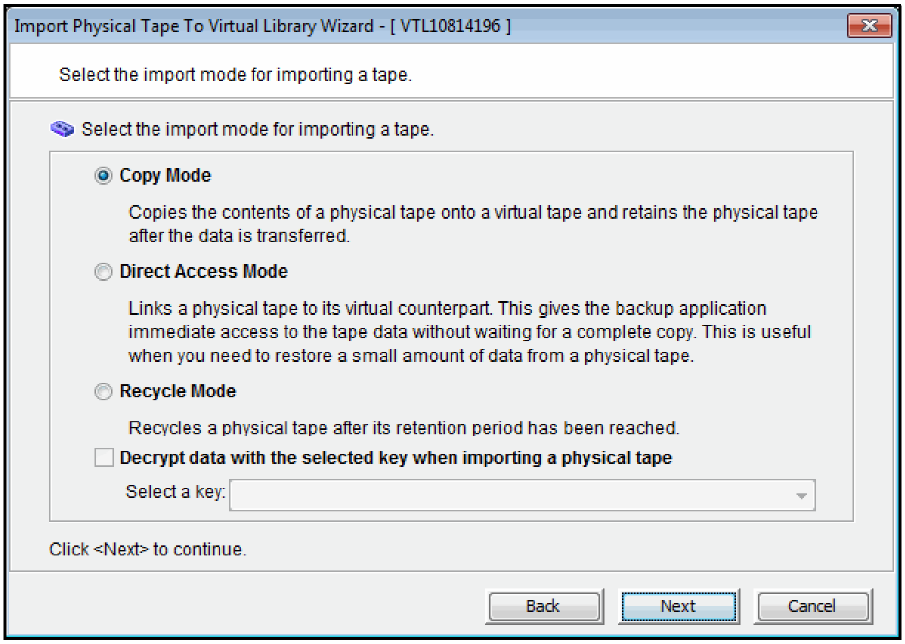
Import Physical Tape To Virtual Library Wizard - [VTL10814196]

Select the import mode for importing a tape.

Select the import mode for importing a tape.

● Copy Mode

Copies the contents of a physical tape onto a virtual tape and retains the physical tape after the data is transferred.

○ Direct Access Mode

Links a physical tape to its virtual counterpart. This gives the backup application immediate access to the tape data without waiting for a complete copy. This is useful when you need to restore a small amount of data from a physical tape.

○ Recycle Mode

Recycles a physical tape after its retention period has been reached.

□ Decrypt data with the selected key when importing a physical tape

Select a key:

Click <Next> to continue.

Copy Mode - Copies the entire contents of a physical tape onto a virtual tape and leaves the physical tape unchanged. This is the only mode available if you are importing a stacked tape.

Direct Access Mode - Links a physical tape to its virtual counterpart. This gives the backup application immediate access to the tape data without waiting for a complete copy. This is useful when you need to restore a small amount of data from a physical tape. Direct access tapes are write protected. Therefore, you can only read from the tape and not write to it.

Recycle Mode - Recycles a physical tape after its retention period has been reached. If you import a tape in recycle mode and the virtual tape is subsequently initialized, the physical tape is now considered recycled and can be used for future export operations.

4. Specify whether or not to decrypt the data on the tape.

You can select this option only if at least one key exists. (For more information, refer to 'Manage encryption keys'.) If you select this option, you must select the key to use.

> **Note:** Selecting this option if the data was not previously encrypted, or an incorrect key is selected, or an invalid password is provided, will import data that is indecipherable. Clearing this option will not decrypt data on the tape.

5. Select the physical tape you want to import.

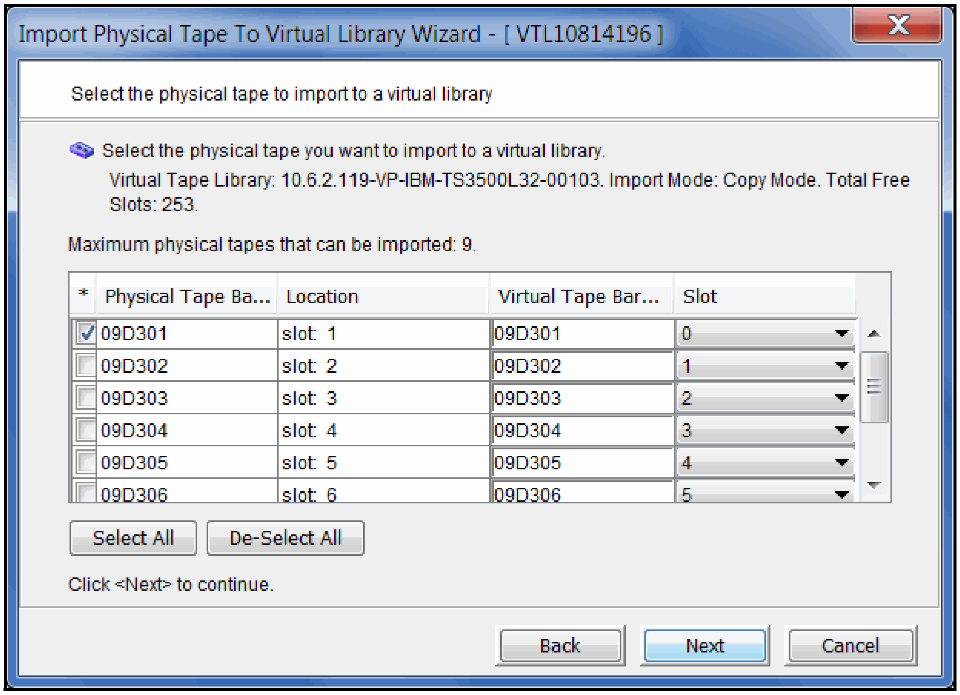

Import Physical Tape To Virtual Library Wizard - [VTL10814196]

Select the physical tape to import to a virtual library

Select the physical tape you want to import to a virtual library.

Virtual Tape Library: 10.6.2.119-VP-IBM-TS3500L32-00103. Import Mode: Copy Mode. Total Free Slots: 253.

Maximum physical tapes that can be imported: 9.

*	Physical Tape Ba...	Location	Virtual Tape Bar...	Slot
☑	09D301	slot: 1	09D301	0
□	09D302	slot: 2	09D302	1
□	09D303	slot: 3	09D303	2
□	09D304	slot: 4	09D304	3
□	09D305	slot: 5	09D305	4
□	09D306	slot: 6	09D306	5

Click <Next> to continue.

You can select a tape based on its barcode or slot location. You can then use the same barcode for the virtual tape or you can enter a new barcode. You can also select a slot for the virtual tape.

You can import whichever tapes you need; you are not required to import all tapes in a library.

6. Verify the information and then click *Finish* to import the tape.

   You can check the *Tape Import/Export Queue* (under the *Activities* object) to watch the progress of the job.

   When import is completed, the virtual tape is automatically moved to the virtual vault.

Import data from a tape in another virtual tape library

In order to import data from a virtual tape in another DSI virtual tape library, the two VTL servers must be connected or zoned together so that they can communicate with each other. To import data:

1. On the source VTL server (where your original tape is located), create a client to represent the target VTL server (the server to which the tape will be imported).

2. On the VTL server to which you are importing tapes, highlight *Physical Resources --> Storage HBAs* and select *Rescan*.

   Select the *Discover New Devices* option.

3. After re-scanning, verify that the tape library and drives you are importing from appear under *Storage Devices*.
4. Right-click the *Physical Tape Libraries* object under *Virtual Tape Library System* and select *Assign*.
5. Select the physical tape library and click *Assign*.
6. Select all tape drives that belong to the physical tape library created previously and click *Assign*.
7. Right-click the physical tape library and select *Import Tape*.
8. Select the virtual library to which you want to import the tape.
9. Select *Copy Mode* to copy the data.
10. Select the tape(s) you want to import.

    You can select a tape based on its barcode or slot location. You can use the same barcode for the virtual tape or you can enter a new barcode. A free destination slot in the virtual library has been automatically assigned to store the virtual tape. You can choose another empty destination slot manually by clicking on the down arrow button.
11. Verify the information and then click *Finish* to import the tape(s).

    You can view status by highlighting the job in the *Tape Import/Export Queue* (under the *Activities* object) . You will see a *Percent Complete* progress bar in the right-hand frame.

    When completed, the virtual tape will be automatically moved to the virtual tape library.

Export data to physical tape

You can export data from a virtual tape to a physical tape in a physical tape library. This process is useful for offsite data archiving.

For deduplicated tapes that have been replicated, exporting a tape is useful for disaster recovery purposes because data can be accessed from the replicated tape. Whenever a deduplicated tape is exported, a complete physical tape is created. This means that the entire, original data set (prior to deduplication) is written to the physical tape. This applies to all VITs, pure or mixed.

With VTL’s built-in incremental export functionality for automatic export, only the modified data is exported and appended to physical tape, if the same tape was exported previously.

On the other hand, if a tape has been reformatted or rewritten, the export job will overwrite the physical tape with all of the information on the virtual tape. If a different tape is exported to that same physical tape, the data will be overwritten on the physical tape.

Notes:

- You cannot use the VTL export function if you are using the Automated Tape Caching option.
- Because some third-party backup applications alter what they write to the tape depending on the type of cartridge used, VTL only exports tapes to *like* media. You cannot export to a dissimilar physical tape. This guarantees that the backup application will accept the tape as valid; from the backup application’s point of view, there is no difference between the virtual and physical tape.

Moving data from virtual tape to physical tape can be accomplished in several ways:

- From your backup software using the software’s own *Tape Copy* function
- Using VTL’s *Export* function, either manually or automatically after each backup using the *Auto Archive* function
- Using VTL’s *Export with stacking* function to export multiple virtual tapes to a single physical tape

As an alternative to exporting data to a physical tape, the VTL Automated Tape Caching option provides a cache to your physical tape library, providing transparent access to data regardless of its location. The Automated Tape Caching option provides advanced flexibility that allows you to set up policies that automatically trigger data migration to physical tapes.

The VTL export methods are explained below. Refer to ‘Automated Tape Caching’ for more information about Automated Tape Caching.

Export manually

Manual exporting data from a virtual tape builds a full physical tape. Incremental export functionality does not apply to manual export.

To manually export data:

1. Right-click a virtual tape and select *Move to Vault*.

   If you are exporting a deduplicated tapes that has been replicated, it will already be in the virtual vault.

2. If you have not already done so, inventory the physical tapes in your library by right-clicking on the physical library and selecting *Inventory*.
3. Right-click the virtual tape under *Virtual Vault* and select *Export Tape*.
4. Select the physical tape library/drive to which you want to export.
5. Select how you want the data exported and if you want the physical tape exported to the import/export slots.

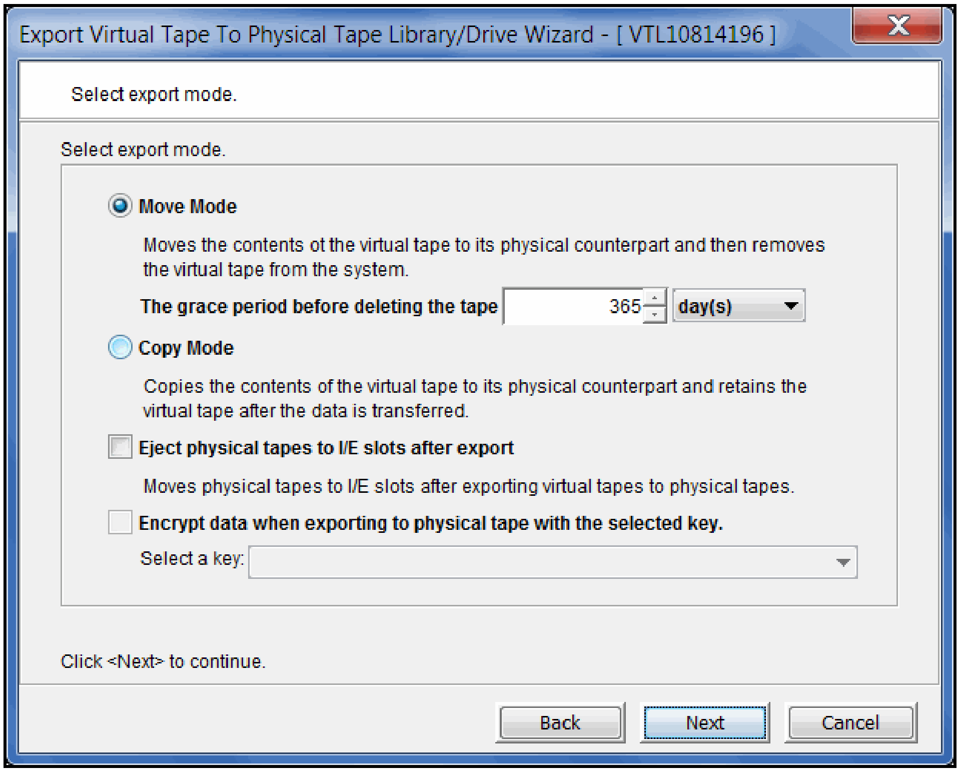

Export Virtual Tape To Physical Tape Library/Drive Wizard - [VTL10814196]

Select export mode.

Select export mode.

● **Move Mode**

Moves the contents ot the virtual tape to its physical counterpart and then removes the virtual tape from the system.

The grace period before deleting the tape 365 **day(s)**

○ **Copy Mode**

Copies the contents of the virtual tape to its physical counterpart and retains the virtual tape after the data is transferred.

□ **Eject physical tapes to I/E slots after export**

Moves physical tapes to I/E slots after exporting virtual tapes to physical tapes.

□ **Encrypt data when exporting to physical tape with the selected key.**

Select a key:

Click <Next> to continue.

Move Mode - Copies the contents of the virtual tape to its physical counterpart and then removes the virtual tape from the system. Specify a grace period if you want to keep the virtual tape for a time before deleting it. If you select *Enable Tape Duplication* below, the virtual tape data will not be deleted until the duplication job finishes successfully.

Copy Mode - Copies the contents of the virtual tape to its physical counterpart and retains the virtual tape after the data is transferred.

Eject physical tapes to I/E slots after export - Move physical tapes to I/E slots after exporting.

Encrypt data when exporting to physical tape with the selected key - Select if you want to encrypt the data on the tape. You can select this option only if at least one key has been created. If you select this option, you must select the key to use. All data on the tape will be indecipherable until it is imported back to a virtual tape library and decrypted using the same key.

You can use a single key to all virtual tapes when you export them or you can create a unique key for each one. Creating multiple keys provides more security; in the unlikely event that a key is compromised, only the tapes that use that key would be affected. However, if you use multiple keys, you must keep track of which key applies to each tape so that you use the correct key to decrypt the data when you import the physical tape back to virtual tape. For more information about encryption, refer to 'Manage encryption keys'.

6. Determine if you want to use Tape Duplication.

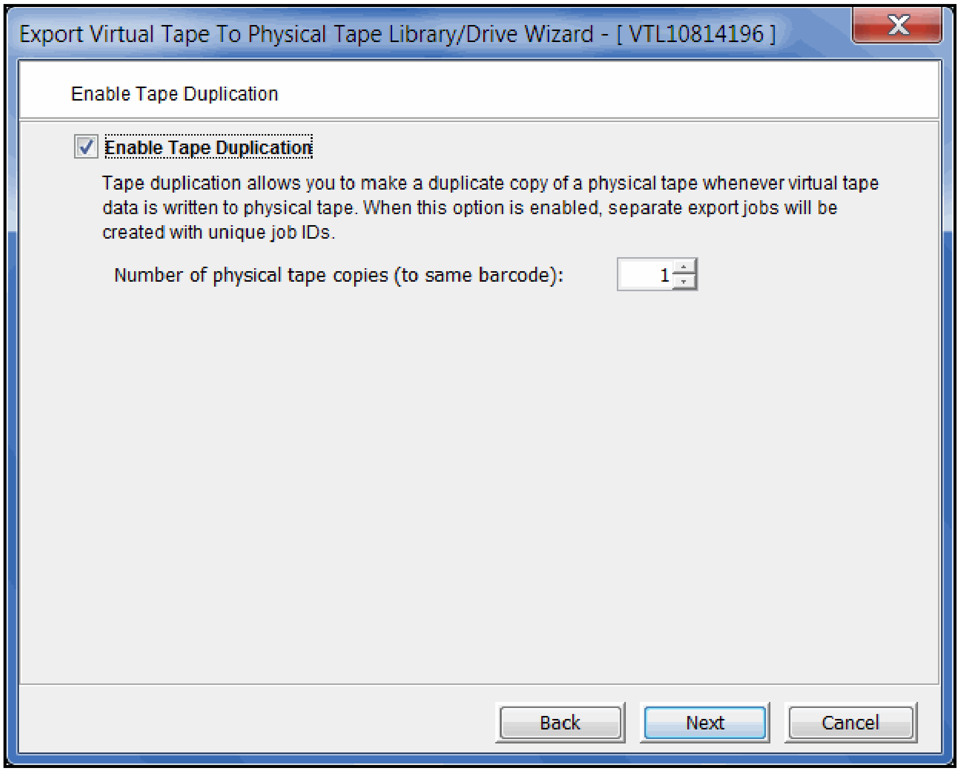

Export Virtual Tape To Physical Tape Library/Drive Wizard - [VTL10814196]

Enable Tape Duplication

☑ Enable Tape Duplication

Tape duplication allows you to make a duplicate copy of a physical tape whenever virtual tape data is written to physical tape. When this option is enabled, separate export jobs will be created with unique job IDs.

Number of physical tape copies (to same barcode): 1

Tape Duplication makes a duplicate copy of the physical tape when data is exported. You must have at least two identical physical libraries (same model, same number of drives, same tapes with the same barcodes). When data is exported, separate export jobs will be created for each physical library and each job will have a unique job ID. Refer to 'Tape Duplication' for more information about Tape Duplication.

7. Select the virtual tape(s) you want to export.

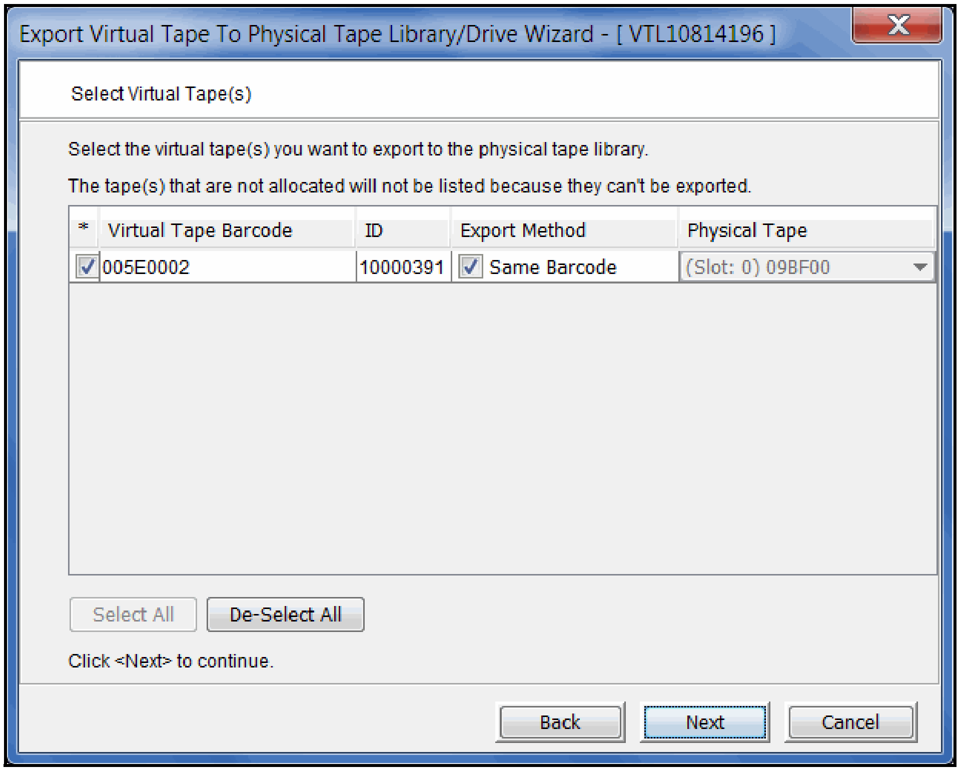

Export Virtual Tape To Physical Tape Library/Drive Wizard - [VTL10814196]

Select Virtual Tape(s)

Select the virtual tape(s) you want to export to the physical tape library.

The tape(s) that are not allocated will not be listed because they can't be exported.

*	Virtual Tape Barcode	ID	Export Method	Physical Tape
☑	005E0002	10000391	☑ Same Barcode	(Slot: 0) 09BF00

Click <Next> to continue.

If your physical tape drive/library uses barcodes, we highly recommend that you select to use the same barcode as the physical tape.

If your physical tape drive/library does not use barcodes, you can then select which slot to use for the physical tape.

8. Verify the information and then click *Finish* to export the tape.

   You can check the *Tape Import/Export Queue* (under the *Activities* object) to watch the progress of the job.

Auto Archive

Auto Archive writes data to physical tape whenever a virtual tape is moved to an Import/Export slot by a backup application or other utility after a backup (you will see the tape in the virtual vault). In order to use *Auto Archive*, the physical tape library must support barcodes because when VTL attempts to export to physical tape it must find a matching barcode in a physical library (you do not need to specify which physical library).

> **Note:** You can only use Auto Archive if you are not currently using the Automated Tape Caching option or the Auto Replication feature on this virtual tape library.

You can configure Auto Archive when you create the library or afterward, as described below:

1. Right-click a virtual tape library and select *Properties*.
2. Select the *Auto Archive* checkbox.
3. Select export options as described for manual export (refer to 'Export manually').

Export with stacking

Tape stacking allows multiple virtual tapes to be exported to a single physical tape in a physical tape drive or library, maximizing physical tape usage and allowing the conversion of virtual media with a smaller capacity to physical media with a higher capacity (i.e. DLT to LTO).

Tape stacking is a tool for archival purposes only. Backup applications will not have direct access to stacked tapes. While VTL's direct access mode links a physical tape to its virtual counterpart, restoring from a stacked tape will require importing data from a stacked tape back to virtual tape. You cannot use VTL's direct access mode for stacked tapes.

> **Notes:**
>
> - You can use tape stacking only if you are not currently using Automated Tape Caching on the selected virtual tape library.
> - When exporting to a physical tape library, if you want tape stacking to append data to a physical tape (rather than overwrite the data on that physical tape), you must scan the tape before exporting.
> - When exporting to a standalone physical tape drive, you must prepare the tape to assign it a barcode before exporting (right-click the drive and select *Prepare Stacked Tape*). If you want to append data to a physical tape in the drive, you must retrieve the barcode before exporting (right-click the drive and select *Retrieve Barcode*).

To stack virtual tapes:

1. Move the tape(s) you want to stack to the virtual vault.
2. Right-click the tape in the virtual vault and select *Export with stacking*.
3. Select the physical tape drive or library to which you want to export.
4. Select export options as described for manual export (refer to 'Export manually').
5. Select the virtual tape(s) you want to export.
6. Select the physical tape to which you want to export and specify if you want to append or overwrite data on the physical tape.

   The default is to append data to the physical tape unless you select the *Overwrite the physical tape* checkbox.

7. Verify the information and then click *Finish* to export the tape.

   You can check the *Tape Import/Export Queue* (under the *Activities* object) to watch the progress of the job.

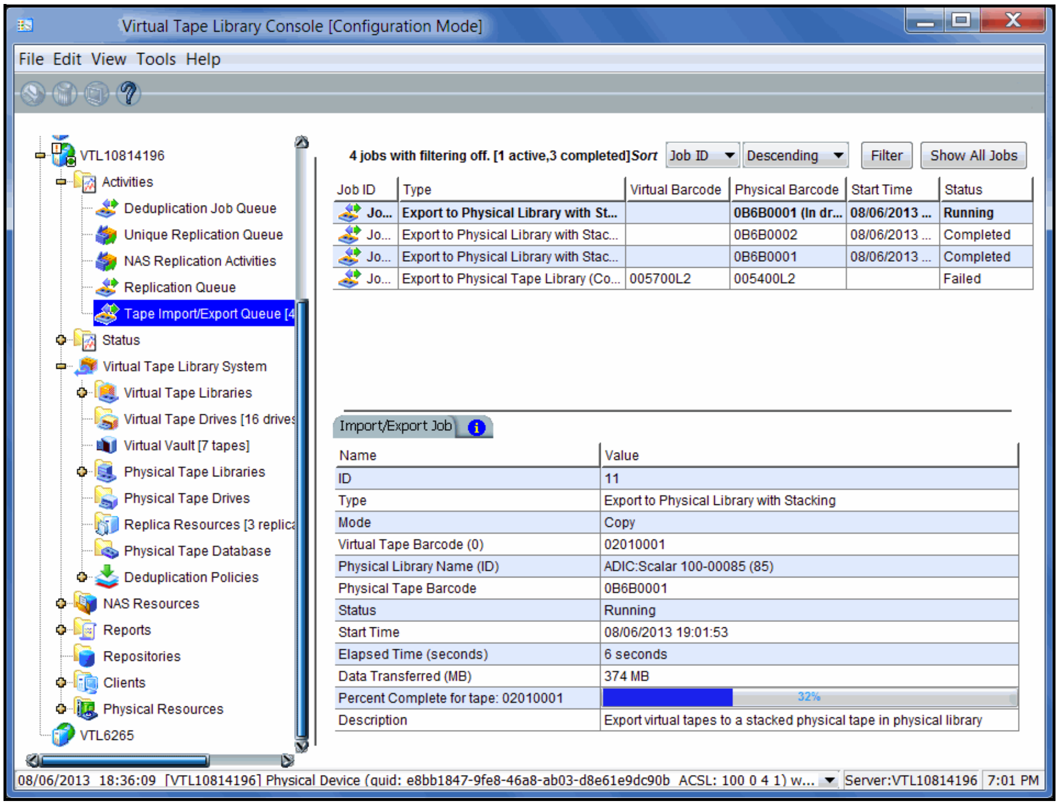

4 jobs with filtering off. [1 active,3 completed]

Job ID	Type	Virtual Barcode	Physical Barcode	Start Time	Status
Jo...	**Export to Physical Library with St...**		**0B6B0001 (In dr...**	**08/06/2013 ...**	**Running**
Jo...	Export to Physical Library with Stac...		0B6B0002	08/06/2013 ...	Completed
Jo...	Export to Physical Library with Stac...		0B6B0001	08/06/2013 ...	Completed
Jo...	Export to Physical Tape Library (Co...	005700L2	005400L2		Failed

Name	Value
ID	11
Type	Export to Physical Library with Stacking
Mode	Copy
Virtual Tape Barcode (0)	02010001
Physical Library Name (ID)	ADIC:Scalar 100-00085 (85)
Physical Tape Barcode	0B6B0001
Status	Running
Start Time	08/06/2013 19:01:53
Elapsed Time (seconds)	6 seconds
Data Transferred (MB)	374 MB
Percent Complete for tape: 02010001	32%
Description	Export virtual tapes to a stacked physical tape in physical library

```
08/06/2013  18:36:09  [VTL10814196] Physical Device (guid: e8bb1847-9fe8-46a8-ab03-d8e61e9dc90b  ACSL: 100 0 4 1) w...   Server:VTL10814196   7:01 PM
```

When the job has completed, you can look on the *Physical Tapes* tab for the physical tape library to see which tapes are stacked. Highlight a tape to see which virtual tapes are stacked on each physical tape.

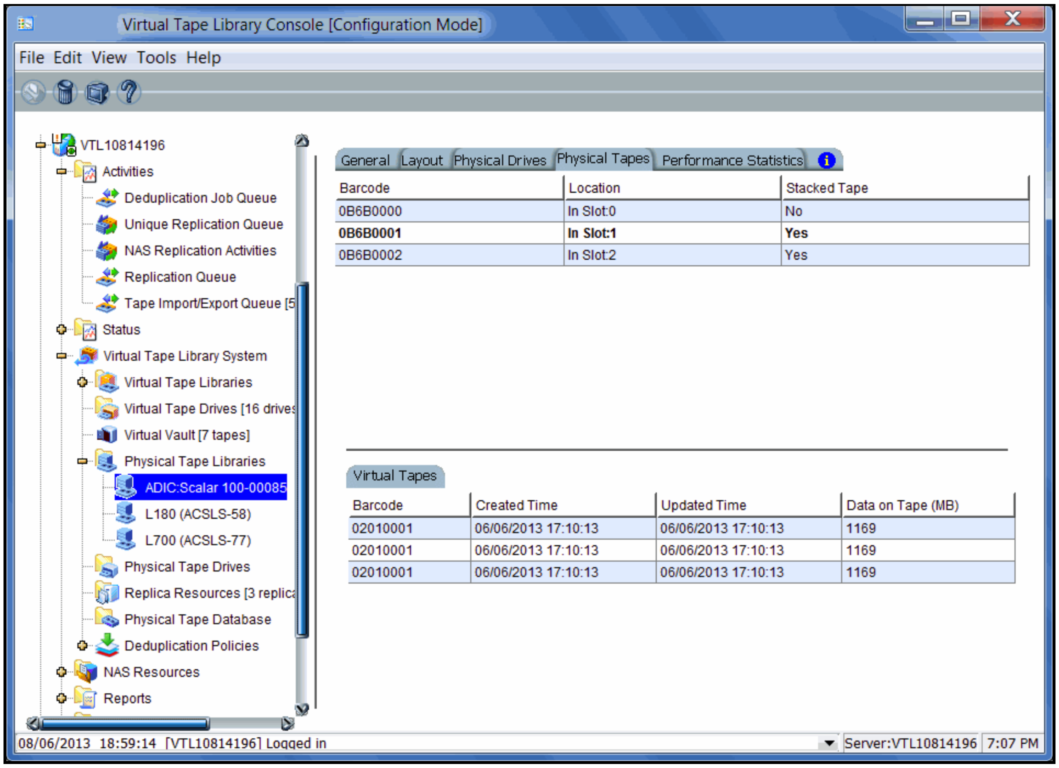

Barcode	Location	Stacked Tape
0B6B0000	In Slot:0	No
0B6B0001	**In Slot:1**	**Yes**
0B6B0002	In Slot:2	Yes

Barcode	Created Time	Updated Time	Data on Tape (MB)
02010001	06/06/2013 17:10:13	06/06/2013 17:10:13	1169
02010001	06/06/2013 17:10:13	06/06/2013 17:10:13	1169
02010001	06/06/2013 17:10:13	06/06/2013 17:10:13	1169

08/06/2013 18:59:14 [VTL10814196] Logged in Server:VTL10814196 7:07 PM

For stacked tapes on a standalone tape drive, you can look at the *Stacked Tapes* tab for the physical tape drive to see which tapes are stacked.

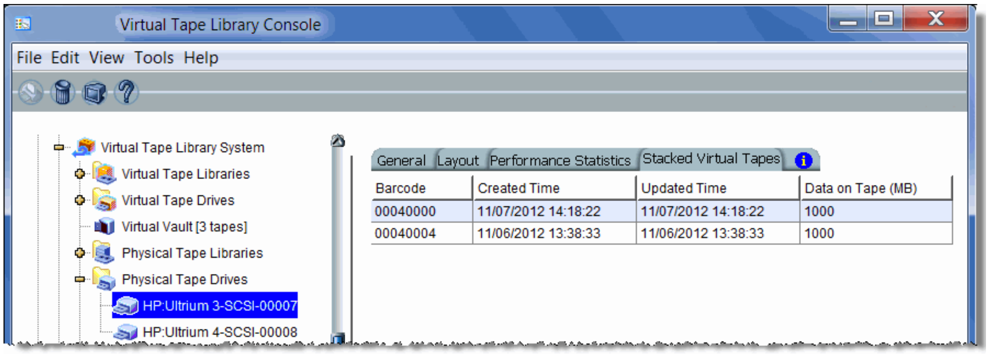

Barcode	Created Time	Updated Time	Data on Tape (MB)
00040000	11/07/2012 14:18:22	11/07/2012 14:18:22	1000
00040004	11/06/2012 13:38:33	11/06/2012 13:38:33	1000

You can also see the stacked tapes from the *Physical Tape Database* object. When you highlight a tape, a list of virtual tapes that are stacked on the physical tape appears.

In addition, the Physical Tape Usage Report lists information about each stacked physical tape in the physical tape database.

Scan stacked physical tapes

You can scan your stacked physical tapes. Scan is used to update the Physical Tape Database about stacked tape information on physical tapes currently in the system.

Scanning physical tapes is useful if you have loaded or ejected a tape from the drive or library or if you have different tapes with the same barcode and they are used in the system at different times.

If a scanned tape is a stacked tape, it will be added to the Physical Tape Database.

To scan a tape:

1. Right-click your physical tape drive or library and select *Scan Tapes*.
2. For physical tape libraries, select which physical tape(s) to scan.

   Afterwards, you can see the stacked tapes that were scanned by going to the *Physical Tape Database* object. Here, you can highlight a tape to see a list of virtual tapes that are stacked on the physical tape.

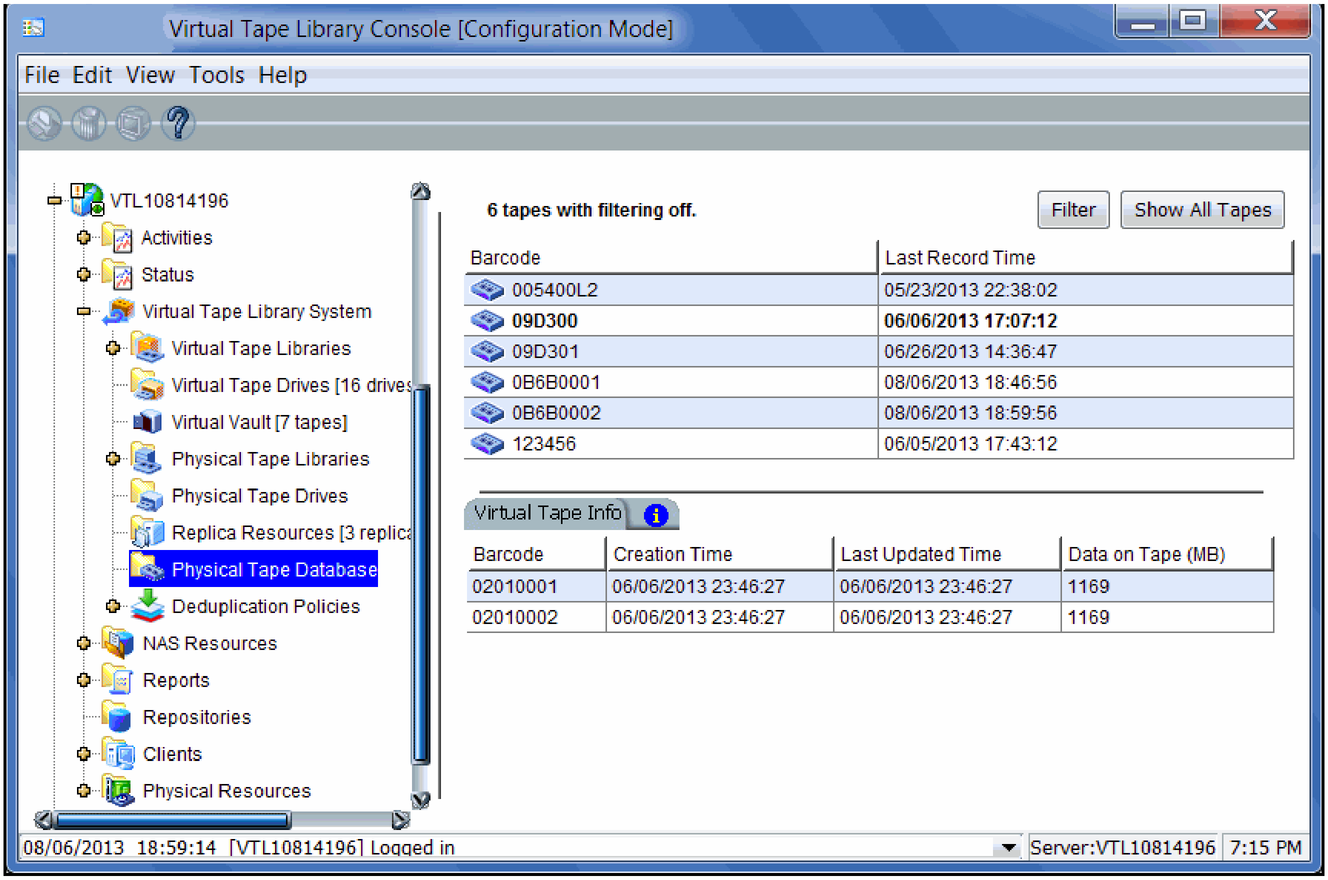

6 tapes with filtering off.

Barcode	Last Record Time
005400L2	05/23/2013 22:38:02
09D300	**06/06/2013 17:07:12**
09D301	06/26/2013 14:36:47
0B6B0001	08/06/2013 18:46:56
0B6B0002	08/06/2013 18:59:56
123456	06/05/2013 17:43:12

Barcode	Creation Time	Last Updated Time	Data on Tape (MB)
02010001	06/06/2013 23:46:27	06/06/2013 23:46:27	1169
02010002	06/06/2013 23:46:27	06/06/2013 23:46:27	1169

08/06/2013 18:59:14 [VTL10814196] Logged in

Server:VTL10814196 7:15 PM

Manage jobs in the Tape Import/Export Queue

When you highlight the *Tape Import/Export Queue* (under the *Activities* object) in the tree, a list of all import and export jobs and Automated Tape Caching jobs that have been submitted is displayed in the right-hand pane.

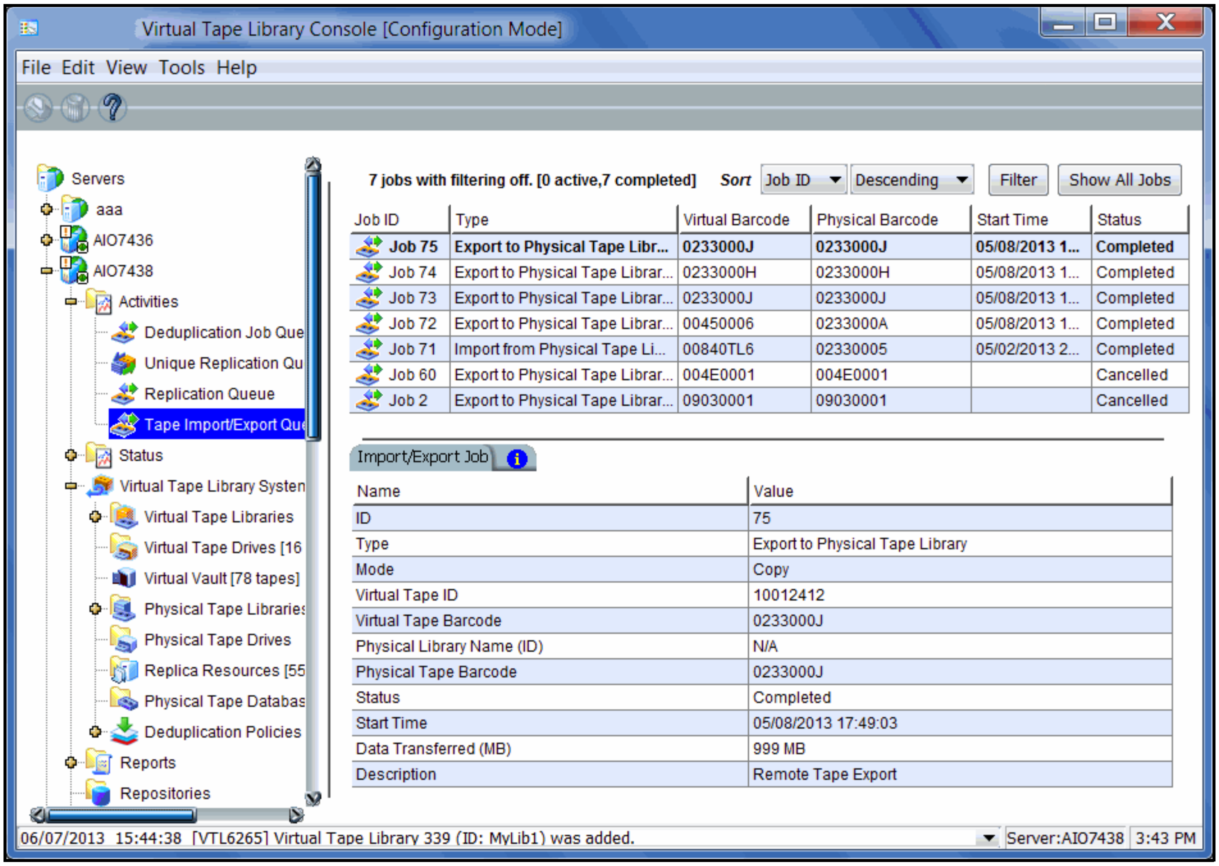

7 jobs with filtering off. [0 active,7 completed]

Job ID	Type	Virtual Barcode	Physical Barcode	Start Time	Status
Job 75	**Export to Physical Tape Libr...**	**0233000J**	**0233000J**	**05/08/2013 1...**	**Completed**
Job 74	Export to Physical Tape Librar...	0233000H	0233000H	05/08/2013 1...	Completed
Job 73	Export to Physical Tape Librar...	0233000J	0233000J	05/08/2013 1...	Completed
Job 72	Export to Physical Tape Librar...	00450006	0233000A	05/08/2013 1...	Completed
Job 71	Import from Physical Tape Li...	00840TL6	02330005	05/02/2013 2...	Completed
Job 60	Export to Physical Tape Librar...	004E0001	004E0001		Cancelled
Job 2	Export to Physical Tape Librar...	09030001	09030001		Cancelled

Name	Value
ID	75
Type	Export to Physical Tape Library
Mode	Copy
Virtual Tape ID	10012412
Virtual Tape Barcode	0233000J
Physical Library Name (ID)	N/A
Physical Tape Barcode	0233000J
Status	Completed
Start Time	05/08/2013 17:49:03
Data Transferred (MB)	999 MB
Description	Remote Tape Export

06/07/2013 15:44:38 [VTL6265] Virtual Tape Library 339 (ID: MyLib1) was added.

Server:AIO7438 3:43 PM

While the right pane is usually just for informational purposes, you can cancel, put on hold, resume, restart a failed job, or delete a job directly from the right pane by highlighting one or more jobs and using the right-click context menu. You can also highlight any job to see detailed job information.

For tape stacking, a job will be created for every 16 virtual tapes stacked to a physical tape. Therefore, if you select 80 tapes to be stacked, five jobs will run.

Completed jobs will be purged after 30 days. You can also delete jobs manually. To delete jobs, highlight one or more jobs, right-click, and select *Delete*.

Filter the display of jobs

Because it is possible to have a large number of jobs in the right-hand pane, you may want to filter the jobs and display only specific ones. To do this:

1. Click the *Filter* button.

2. On the *ID* tab, you can specify a job, physical library, or virtual library ID.

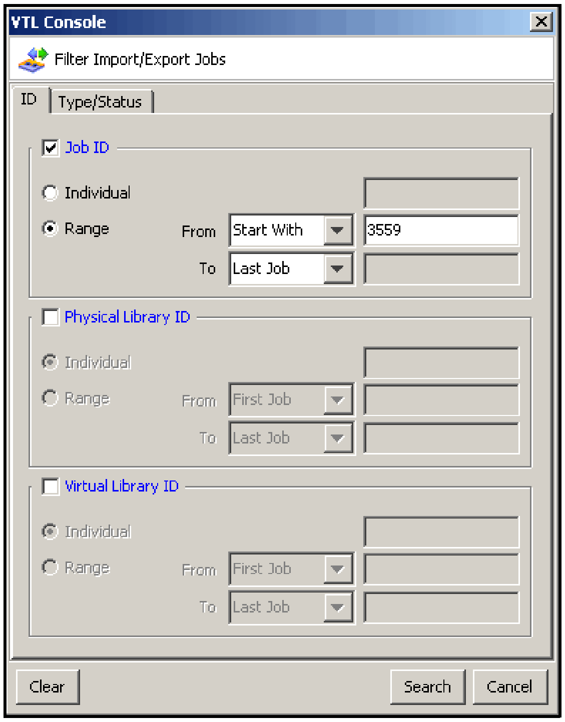

VTL Console

Filter Import/Export Jobs

ID | Type/Status

☑ Job ID

○ Individual

● Range From: Start With 3559

To: Last Job

□ Physical Library ID

● Individual

○ Range From: First Job

To: Last Job

□ Virtual Library ID

● Individual

○ Range From: First Job

To: Last Job

If you want to specify a particular number, select *Start With* or *End With* in the *From/To* fields. You can then type the number in the box to the right.

3. On the *Type/Status* tab, you can specify a job type or status.

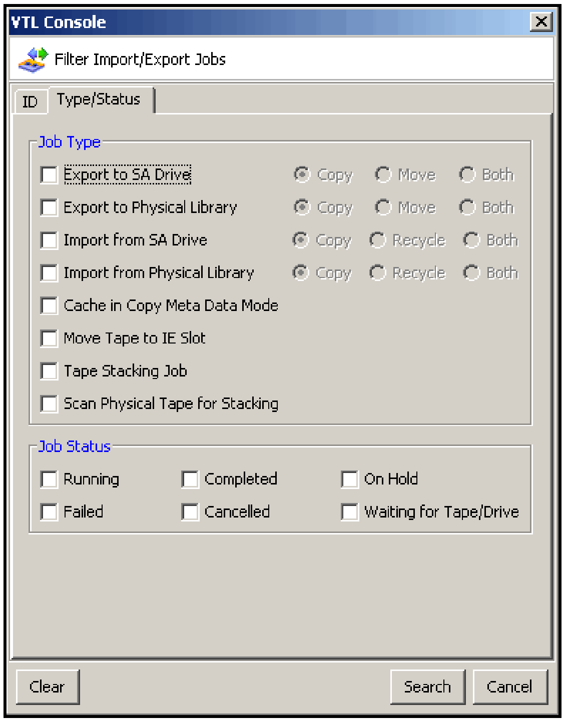

VTL Console

Filter Import/Export Jobs

ID | Type/Status

Job Type

□ Export to SA Drive ● Copy ○ Move ○ Both
□ Export to Physical Library ● Copy ○ Move ○ Both
□ Import from SA Drive ● Copy ○ Recycle ○ Both
□ Import from Physical Library ● Copy ○ Recycle ○ Both
□ Cache in Copy Meta Data Mode
□ Move Tape to IE Slot
□ Tape Stacking Job
□ Scan Physical Tape for Stacking

Job Status

□ Running □ Completed □ On Hold
□ Failed □ Cancelled □ Waiting for Tape/Drive

4. Click *Search*.

   Afterwards, *just* the jobs that match the selected criteria will be displayed in the right pane. You can click the *Show All Jobs* button when you are done.

Set import/export job properties for failed jobs

You can set properties that determine the number of times a failed job should be retried and the duration between each retry. To do this:

1. Right-click the *Tape Import/Export Queue* object and select *Properties*.

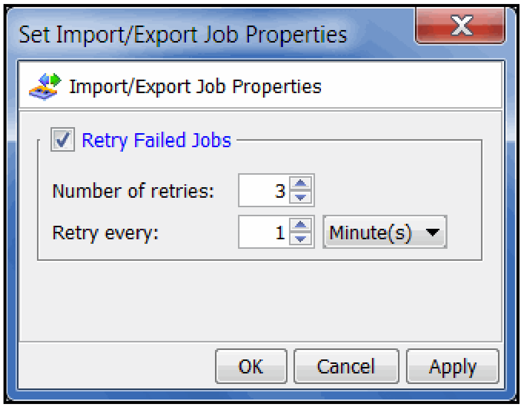

Set Import/Export Job Properties

Import/Export Job Properties

☑ Retry Failed Jobs

Number of retries: 3

Retry every: 1 Minute(s)

2. Indicate if you want failed jobs to be retried and the specifications for retrying.

Set virtual tape library system properties

You can set global options for all virtual tape libraries. To do this:

1. In the VTL console, right-click *Virtual Tape Library System* and select *Properties*.

   If the server is a member of a group, right-click the group and select *VTL Properties*.

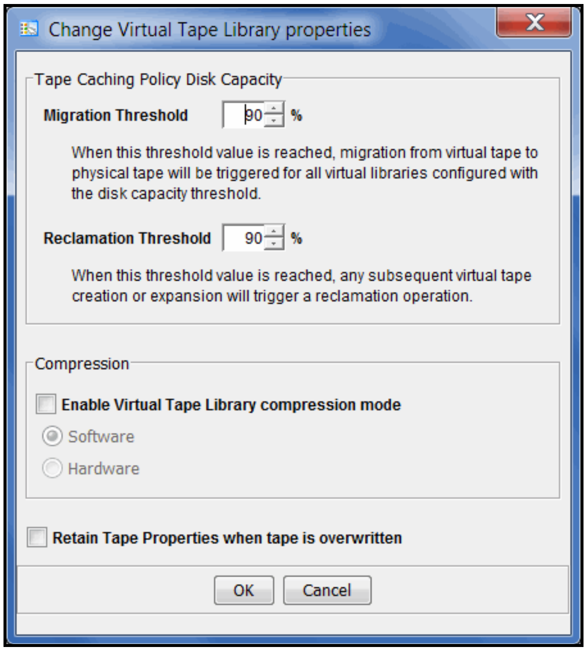

Change Virtual Tape Library properties

Tape Caching Policy Disk Capacity

Migration Threshold 90 %

When this threshold value is reached, migration from virtual tape to physical tape will be triggered for all virtual libraries configured with the disk capacity threshold.

Reclamation Threshold 90 %

When this threshold value is reached, any subsequent virtual tape creation or expansion will trigger a reclamation operation.

Compression

□ **Enable Virtual Tape Library compression mode**

● Software

○ Hardware

□ **Retain Tape Properties when tape is overwritten**

2. Select the options you want to use.

   Tape Caching Policy Disk Capacity Migration Threshold - Migration will occur when the used disk space exceeds the specified disk capacity.

   Tape Caching Policy Disk Capacity Reclamation Threshold - Cache disk space is freed up when the used space reaches this threshold.

   Enable Virtual Tape Library compression mode - VTL's compression saves disk space by compressing files so that more data can be stored by a virtual tape drive. Refer to 'Use virtual tape drive compression' for more information.

   Retain Tape Properties when tape is overwritten - Allows tape properties to persist even when the tape is overwritten.

Use virtual tape drive compression

VTL's software compression uses an LZO algorithm to save disk space by compressing files so that more data can be stored by a virtual tape drive. The increase in capacity is directly related to the compressibility of the data being backed up. If you can compress the data being backed up by a factor of up to 2:1, you can store up to twice as much information on the virtual tape. Disk compression can vary depending upon the dataset; certain file types (ZIP, PDF, GIF, RAR, etc.) are already compressed and cannot be compressed further.

In order to use compression, you must enable tape drive compression from your backup server.

> **Note:** If you are already using software compression that is performed by your backup application, you should not use VTL's compression. Using both types of compression will cause VTL to try to compress already-compressed data and this can slow down your backups.

Enable/disable compression

To enable or disable compression:

1. Enable tape drive compression in your backup application.

2. In the VTL console, right-click *Virtual Tape Library System* and select *Properties*.

   If the server is a member of a group, right-click the group and select *VTL Properties*.

3. Select the *Enable Virtual Tape Library compression mode* checkbox and select *Software* compression.

   Hardware compression is a legacy item that is only supported on older VTL appliances.

   Compression is a global setting, which means that it will apply to all tapes in your system. However, if compression is enabled on the VTL server, you can still disable or enable compression on each individual virtual tape drive in the same manner as real tape drives -- via your backup application or via SCSI commands which are sent by the operating system. Depending on your operating system, do one of the following:

   - UNIX — On backup servers that run Solaris or other UNIX operating systems, specify a compressed tape device file such as /dev/rmt/0cbn to enable compression or /dev/rmt/0ubn to disable compression.
   - Windows — On Windows servers, select the option in your backup software to enable or disable tape drive compression. If global VTL compression is disabled, it is possible to enable individual drive compression, but it will have no effect.

You will see a compression icon next to each virtual tape drive with compression enabled.

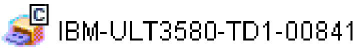

IBM-ULT3580-TD1-00841

Note: CRC checking can be enabled for hardware compression cards to help detect data corruption and identify the source of the corruption. This feature is disabled by default. Contact DSI Technical Support if you need to use this feature.

Change firmware of a virtual library or drive

You can change the firmware of a virtual library or drive to match that of the physical library/drive. To do this:

1. Right-click a virtual tape library or drive and select *Change Firmware*.
2. Enter the new firmware and click *OK*.

Shred a virtual tape

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.

Notes:

- Tape shredding may adversely affect backup performance. We recommend that you perform tape shredding when there are no backups running.
- When you shred a VIT, the index information will be erased and data in the repository will be cleaned up during reclamation.

To shred tapes:

1. Move the tape(s) you want to shred to the virtual vault.
2. Select the tape(s) you want to shred.

   For a single tape, right-click the tape in the virtual vault and select *Tape Shredding* --> *Shred Tape*.

   For multiple tapes, highlight all of the tapes you want, right-click, and select *Tape Shredding* --> *Shred Tapes.*
3. If desired, select the option to delete the tape after shredding it.
4. Type *YES* to confirm and click *OK*.

   You can view the status by highlighting the virtual tape in the vault. The status bar displays the progress.

   If you want to cancel the shredding process, right-click the tape or the *Virtual Vault* object and select *Tape Shredding* --> *Cancel*.

Deduplication

DSI’s data deduplication solution eliminates redundant data without impacting your established backup window. DSI data deduplication offers as much as a 30:1 reduction of backup data, minimizing replication time and storage requirements.

The system offers deduplication for tapes and NAS files.

How tape deduplication works

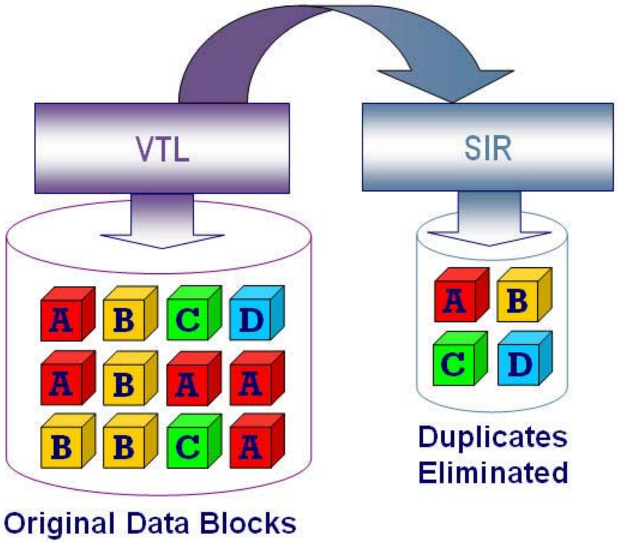

During deduplication, an intelligent, content-aware "Tape Scanner" process analyzes the data and determines whether data is unique or has already been copied to the repository. The process then passes only single instances of unique data to the repository; data is compressed automatically. The original virtual tape is replaced with a virtual index tape (VIT) that contains pointers to the data in the repository, freeing considerable space for more data.

Deduplication is triggered by policies managed in VTL. You can set policies for all tapes in a library, groups or ranges of tapes, or just an individual tape. Deduplication is performed in the background without user intervention. During normal use, deduplication is transparent to the backup operation.

Deduplication jobs are automatically suspended when the tape being deduplicated is needed for backup or restore; when the backup application finishes using that particular tape, the deduplication job automatically resumes from where it left off.

Depending upon your version of VTL, you may have a separate deduplication appliance. If you do, connectivity between the backup and deduplication appliances can be via a FC switch or it can be direct. Your deduplication cluster configuration, repository configuration, and storage are managed in the same console as your backup server. The Virtual Tape Library common console allows you to view real-time deduplication activity, as well as historical statistics, using one efficient interface.

In a cluster configuration, deduplication can operate as a single node (one active deduplication server) or as a multi-node cluster (two or four active deduplication servers).

VTL with deduplication provides replication capability. If replication is configured, deduplication replicates its repository and metadata, effectively performing global data deduplication. Any data duplicated across remote sites is deduplicated at the central site, enabling only globally unique data to be replicated to the disaster recovery site.

How NAS deduplication works

VTL uses standard network protocols such as Common Internet File System (CIFS) or Network File System (NFS) to present a simple, network-based file share as the target for backed-up data. Each NAS file share holds incoming data, acting as a "disk" for disk-to-disk (D2D) backup.

During deduplication, the system analyzes blocks of data and determines whether the data is unique or has already been copied to the repository (virtualized disks that hold deduplicated data). The process then passes only single instances of unique data to the repository and replaces each deduplicated file with a small file (called a *stub* file), that points to the repository and is used to retrieve stored data.

Even though the user interface is file-based, deduplication is performed at the block level, not at a file level. Block-level deduplication examines small sub-blocks, making it far more effective at reducing storage consumption than file-based deduplication.

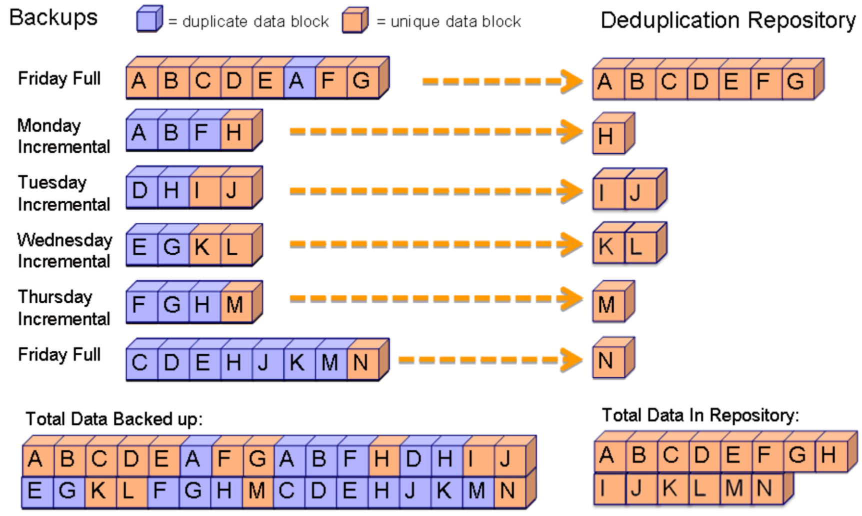

Because network-based file shares are used for backed-up data, restoring data is fast and easy. The administrator has direct access to all files without the need for a restore job. Even after deduplication occurs, pointers (*stub* files) on the share point to the full file in the repository. Restoring data copied by archiving software is as simple as copying the necessary files from a share back to the appropriate location.

Deduplication methods - at a glance

Typically, backup jobs to the VTL system are performed during the night-time “backup window”. When deduplication is performed depends upon several factors, including your environment and requirements, as well as the data type.

To maximize deduplication performance and minimize storage needs, DSI offers the following deduplication methods:

- Inline deduplication - available for tape deduplication; used for NAS deduplication
- Turbo deduplication - available for tape deduplication
- Post-processing - available for tape deduplication

NAS file backups use inline deduplication, which processes and deduplicates data during backup.

For tape backups, deduplication can be performed at any time, while backup is running or after each backup job completes, and can be scheduled or can be run on demand. This is set as part of each tape deduplication policy.

Turbo deduplication pre-processes data during backup but completes the deduplication process at a later time. Before deduplication occurs, the pre-processed, hashed data is stored in a temporary area (which requires approximately 1%-2% of the size of the backed-up data). When deduplication is triggered, the system processes the hashed data stored in the temporary area and the unique data blocks are stored in the repository. After a tape has finished deduplication, the space used by the temporary area is released and the original virtual tape is replaced with a VIT. Turbo deduplication minimizes disk contention by reducing I/O to the disk because data is pre-processed; less has to be read by the deduplication process.

Inline deduplication processes and deduplicates data during backup and then stores the unique data blocks in the repository. Inline deduplication uses less storage than Turbo deduplication because it does not require backup landing storage.

The following table compares the different deduplication methods:

	Inline Deduplication	Turbo Deduplication	Post-Processing (Without Inline or Turbo)
Tape or NAS Deduplication	• Used for NAS deduplication • Available for tape deduplication	• Available for tape deduplication	• Available for tape deduplication
Backup Landing Storage and Performance	• Requires the least amount of storage. • As the deduplication ratio increases over time, overall backup/ deduplication performance will increase because fewer writes will need to be made to the repository. The first few backup/deduplication sessions on a new system require additional processing time.	• Requires appropriate capacity for your daily backups plus additional space for VITs (about 2% of the size of the pre-deduplication capacity). • Pre-processing can significantly improve deduplication performance.	Requires appropriate capacity for your daily backups plus additional space for VITs (about 2% of the size of the pre-deduplication capacity) .
CPU Processing Power	Requires the most CPU power in the VTL server so that incoming backup data can be processed (12 cores of processing power or above recommended).	The more CPUs in the VTL server, the less impact to backup performance. With 12 cores of processing power or above, impact will be minimal.	No special requirements.
Compression	Software compression is performed. No hardware compression card is required.	Uses a hardware compression card if hardware compression is configured.	Uses a hardware compression card if hardware compression is configured.
Can be used with Tape Caching	No. Inline deduplication and Tape Caching are mutually exclusive. To reduce disk contention, data written to cached tapes will not be deduplicated in real time and will follow the schedule of the tape caching policy.	Yes	Yes

When the server deduplicates data for both VTL and NAS resources concurrently, deduplication occurs based on the order that data arrives without distinguishing between deduplication jobs.

Tape deduplication policies

Create tape deduplication policies

Deduplication policies specify which virtual tapes need to have deduplication and when deduplication should occur. You must have at least one virtual tape library in order to create a policy.

When you create a deduplication policy, you can configure replication for the tapes in the policy. If you intend to do this, you must first configure replication as described in 'Overview of steps to configure replication for deduplicated tapes'. 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.

> **Note:** Once you set your deduplication policies, you should not change the IP address of either appliance. If you need to change the IP address, do it BEFORE setting your policies.

1. Right-click the *Deduplication Policies* object and select *New*.

   To modify an existing policy, right-click the policy name and select *Edit*.

2. Enter a name for the policy.

   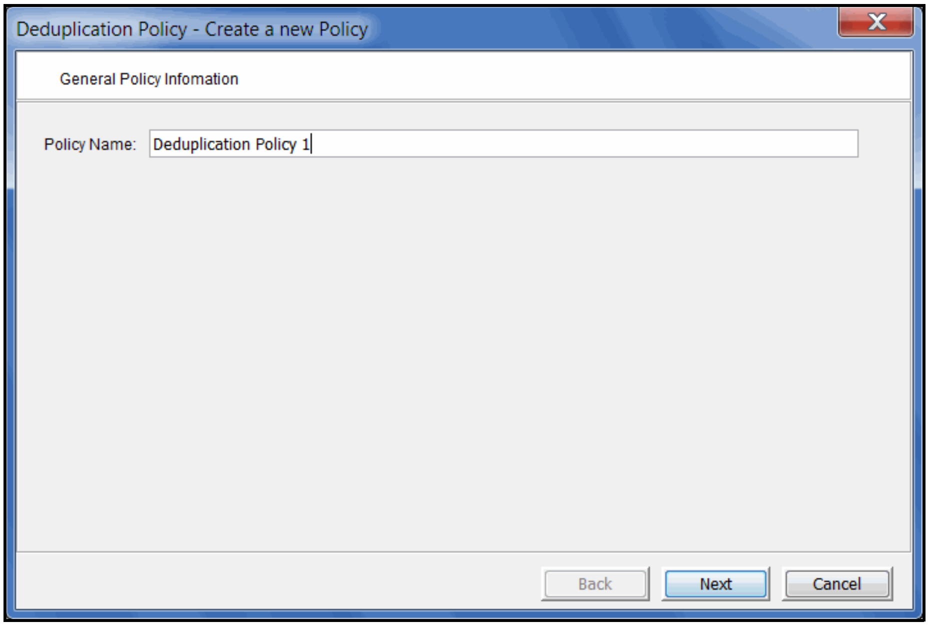
   Deduplication Policy - Create a new Policy
General Policy Infomation
Back
Next
Cancel

   Policy Name: Deduplication Policy 1

   Use standard characters. Unicode characters are not permitted.

3. Select your deduplication cluster.

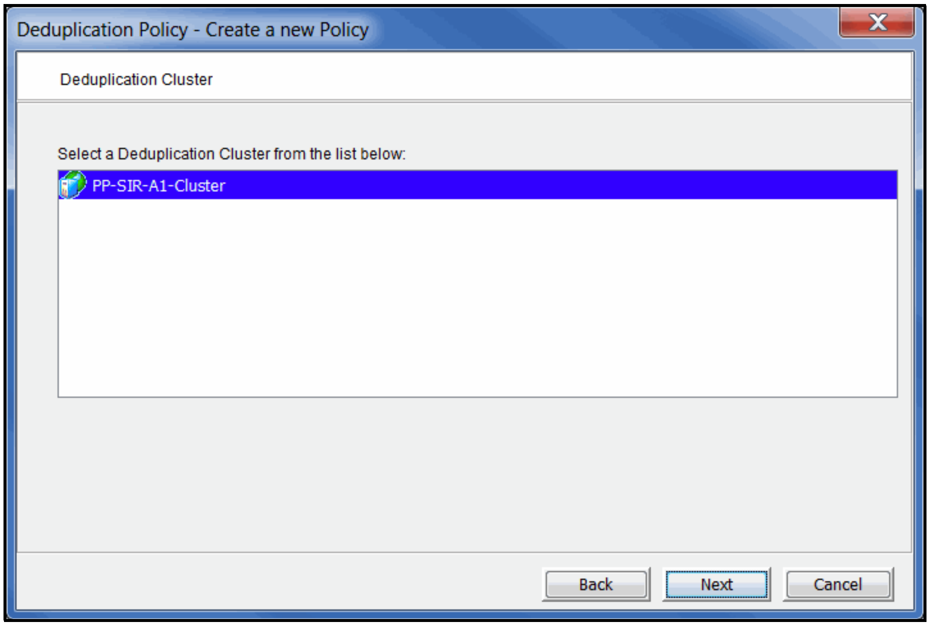

Select a Deduplication Cluster from the list below:

PP-SIR-A1-Cluster

The associated cluster is displayed for you. *Localcluster* will be displayed for VTL-S systems.

4. Specify when deduplication should occur.

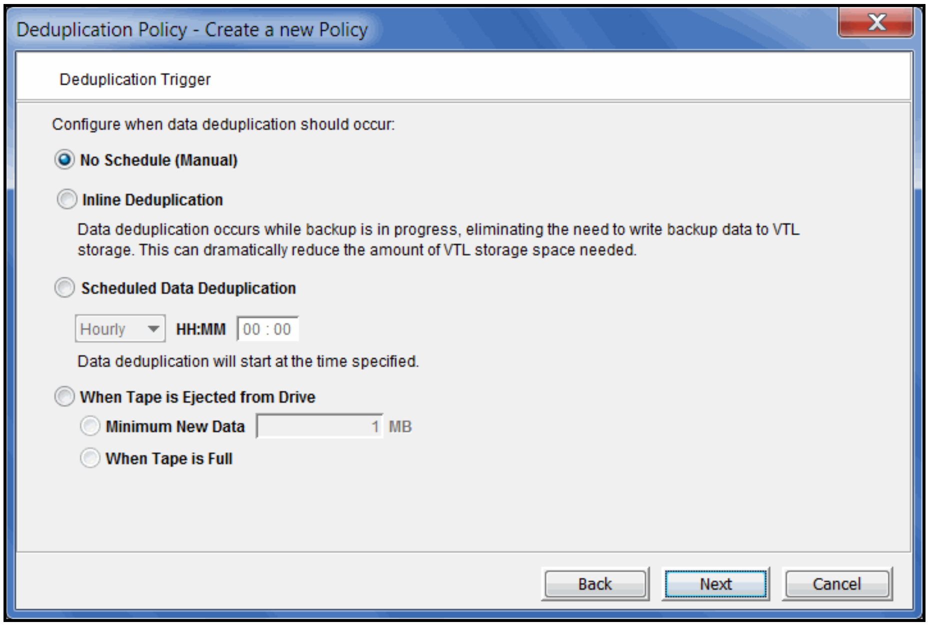

Configure when data deduplication should occur:

● **No Schedule (Manual)**

○ **Inline Deduplication**

Data deduplication occurs while backup is in progress, eliminating the need to write backup data to VTL storage. This can dramatically reduce the amount of VTL storage space needed.

○ **Scheduled Data Deduplication**

Hourly **HH:MM** 00 : 00

Data deduplication will start at the time specified.

○ **When Tape is Ejected from Drive**

○ **Minimum New Data** 1 MB

○ **When Tape is Full**

No Schedule (Manual) - Deduplication must be manually initiated.

Inline Deduplication - Data deduplication occurs while backup is in progress. If inline deduplication fails at the very beginning (or at the first write from the backup application), backup data will be written to VTL storage without being deduplicated in real time but will be deduplicated immediately once the backup is completed and the tape is ejected to a slot. If inline deduplication fails in the middle of a backup session, the current backup session will fail with a medium error status returned to the backup application. Most backup applications will then re-run the backup session to a new tape. Refer to 'Deduplication methods - at a glance' for more information about Inline deduplication.

Scheduled Data Deduplication - Deduplication will occur based on the schedule specified.

- *Hourly* - Deduplication will occur every hour at the specified time.
- *Daily* - Deduplication will occur at a specific time of day.
- *Weekly* - Deduplication will occur on a specific day of the week at a specific time.

When Tape is Ejected from Drive - Deduplication starts when a virtual tape has been written and is ejected from a drive to a slot, thereby running concurrently with backup.

- *Minimum New Data* - Specify the minimum amount of new data that must have been backed up in order for deduplication to occur.
- *When Tape is Full* - Deduplication will only occur if the tape is full.

5. If you did not select *Inline Deduplication*, specify if you want to use *Turbo Deduplication.*

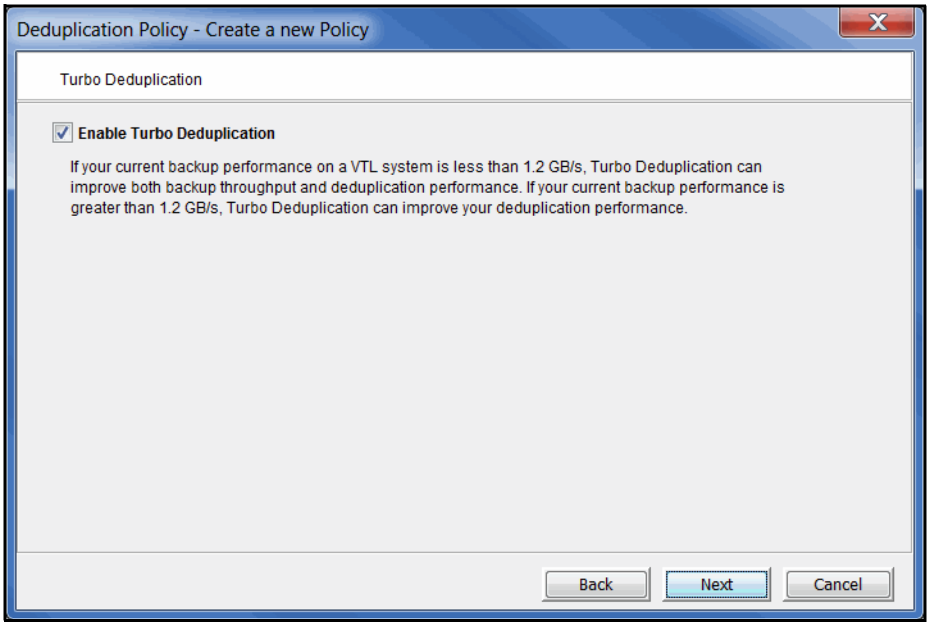

☑ Enable Turbo Deduplication

If your current backup performance on a VTL system is less than 1.2 GB/s, Turbo Deduplication can improve both backup throughput and deduplication performance. If your current backup performance is greater than 1.2 GB/s, Turbo Deduplication can improve your deduplication performance.

With Turbo deduplication, VTL will pre-process tapes during backup. Refer to 'Deduplication methods - at a glance' for more information about Turbo deduplication.

Many backup applications report backup performance at the end of each backup job. You can also find your performance by displaying *Write Performance* on the *Dashboard Summary* tab or by looking at the *Data Throughput* chart in the *VTL Performance Report*.

6. To prioritize the order in which deduplication jobs will be launched, select *Enable Policy Priority* and set a priority level for this policy. Also, set the retry parameters for this policy.

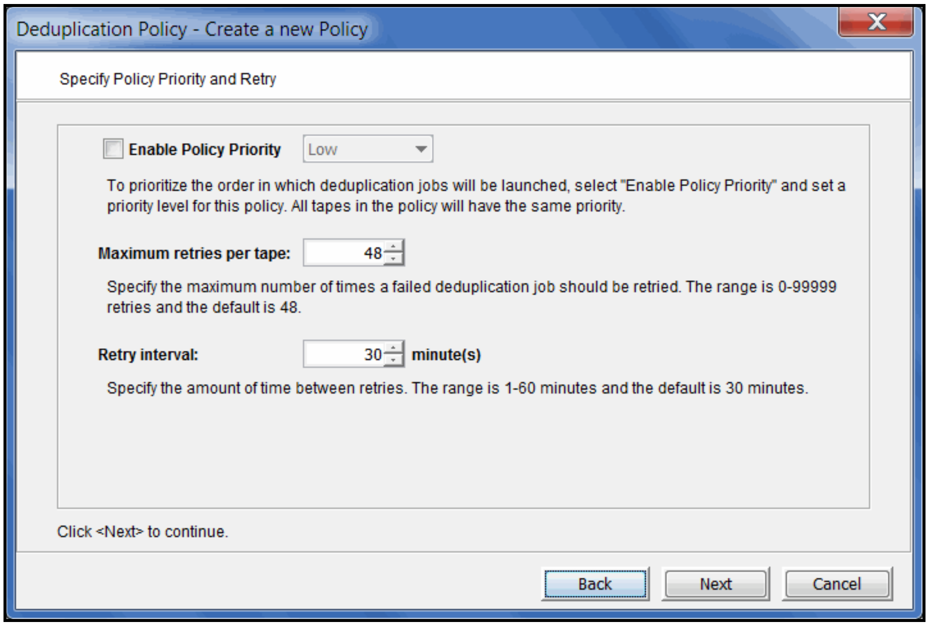

Deduplication Policy - Create a new Policy

Specify Policy Priority and Retry

□ **Enable Policy Priority** Low

To prioritize the order in which deduplication jobs will be launched, select "Enable Policy Priority" and set a priority level for this policy. All tapes in the policy will have the same priority.

Maximum retries per tape: 48

Specify the maximum number of times a failed deduplication job should be retried. The range is 0-99999 retries and the default is 48.

Retry interval: 30 **minute(s)**

Specify the amount of time between retries. The range is 1-60 minutes and the default is 30 minutes.

Click <Next> to continue.

Setting a priority is valid for policies configured for Turbo deduplication, Post-processing deduplication, and Inline deduplication with replication. It also becomes valid for Inline deduplication policies if the job fails and the deduplication becomes a post-processing job.

You can select a *Low, Medium*, or *High* priority. All tapes in the policy will have the same priority.

If you do not specify a priority, the order in the Deduplication Job Queue is determined by the last time the tape was written. The older the tape, the higher its position in the job queue and the sooner it will get processed. However, tapes without a priority will always have a lower standing than tapes with a priority.

If you specify a priority, the tapes in this policy will get a higher placement in the job queue than the tapes without any priority.

At execution time, if multiple policies have the same priority, the system will alternate jobs among all of the policies with the same priority. However, in order to better utilize replication bandwidth, tapes with replication may run along with or ahead of higher priority deduplication jobs. For example:

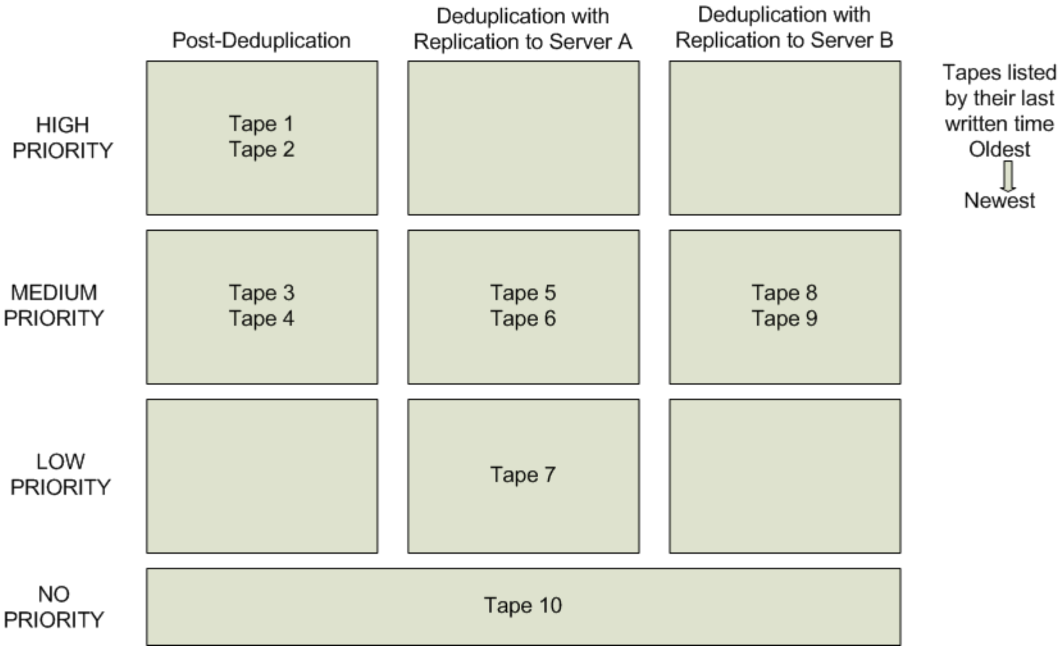

	Post-Deduplication	Deduplication with Replication to Server A	Deduplication with Replication to Server B
HIGH PRIORITY	Tape 1 Tape 2		
MEDIUM PRIORITY	Tape 3 Tape 4	Tape 5 Tape 6	Tape 8 Tape 9
LOW PRIORITY		Tape 7	
NO PRIORITY	Tape 10		

In this example, the tapes will be processed in the following order:

- Tape 1
- Tape 5
- Tape 8
- Tape 2
- Tape 6
- Tape 9
- Tape 3
- Tape 7
- Tape 4
- Tape 10

If you change the priority of tapes in the Deduplication Job Queue and select *Run Next*, the *Run Next* priority is higher than any existing policy priority that is set.

If a deduplication, VIT replication, or resolver process fails, the tape will be returned to the job queue and the entire job will be retried up to the maximum number of times specified in the *Maximum retries per tape* field. The range is 0-99999 retries and the default is 48.

You can specify the amount of time between retries in the *Retry Interval* field. The range is 1-60 minutes and the default is 30 minutes.

The number of times a tape has been retried and failed will be logged in the Event Log.

The retry policy settings apply to all tapes, regardless of whether or not a priority policy is set.

7. To define a replication policy for the VITs in this policy, select *Enable Replication*.

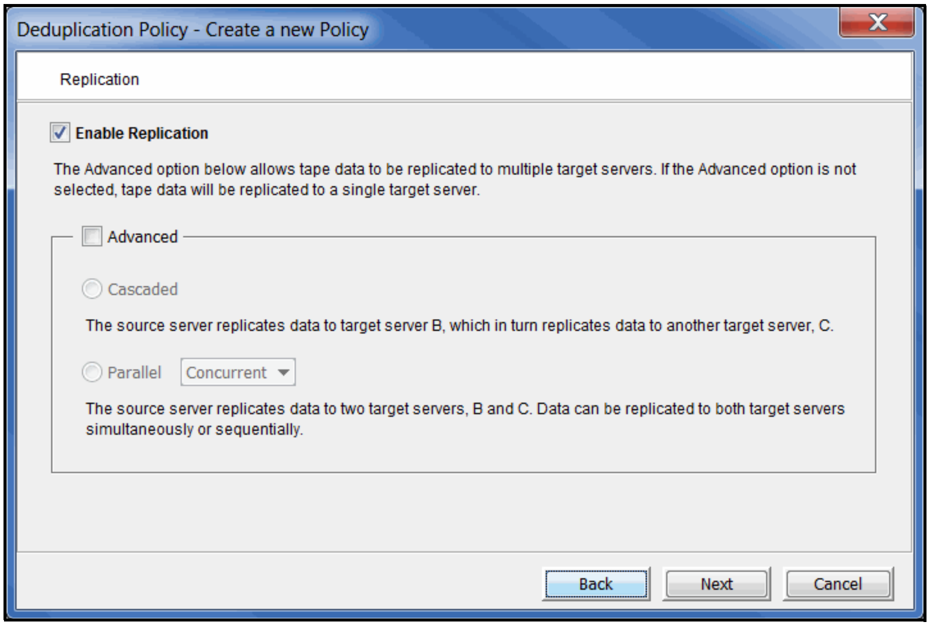

Deduplication Policy - Create a new Policy

Replication

☑ Enable Replication

The Advanced option below allows tape data to be replicated to multiple target servers. If the Advanced option is not selected, tape data will be replicated to a single target server.

□ Advanced

○ Cascaded

The source server replicates data to target server B, which in turn replicates data to another target server, C.

○ Parallel Concurrent

The source server replicates data to two target servers, B and C. Data can be replicated to both target servers simultaneously or sequentially.

Single mode is the default if you do not choose an *Advanced* option. *Single* mode replicates tape data to a single target server.

Advanced replication includes several options.

- In *Cascaded* mode, the source server replicates data to Replication Target 1, which in turn replicates data to Replication Target 2.

  Configuring a deduplication policy with replication in *Cascaded* mode automatically creates the same policy on Replication Target 1, where the name of the policy is the server name followed by the policy name.
- In *Parallel* mode, source server A replicates data to two target servers, 1 and 2. You can choose either *Concurrent* (replication to both target servers at the same time), or *Serial* (replication first to Replication Target 1 and then to Replication Target 2).

Note: All servers using advanced replication must be VTL 8.00 or higher.

Data encryption is determined by the target server, regardless of whether or not encryption is enabled on the source:

- If the source deduplication repository is encrypted but the target is not, data will not be encrypted in the deduplication repository of the target server.
- If the source deduplication repository is not encrypted but the target is, data will be encrypted in the deduplication repository of the target server.

If both the source and target deduplication repositories are encrypted, they can have different encryption keys.

8. Confirm/enter the IP address of the source server.

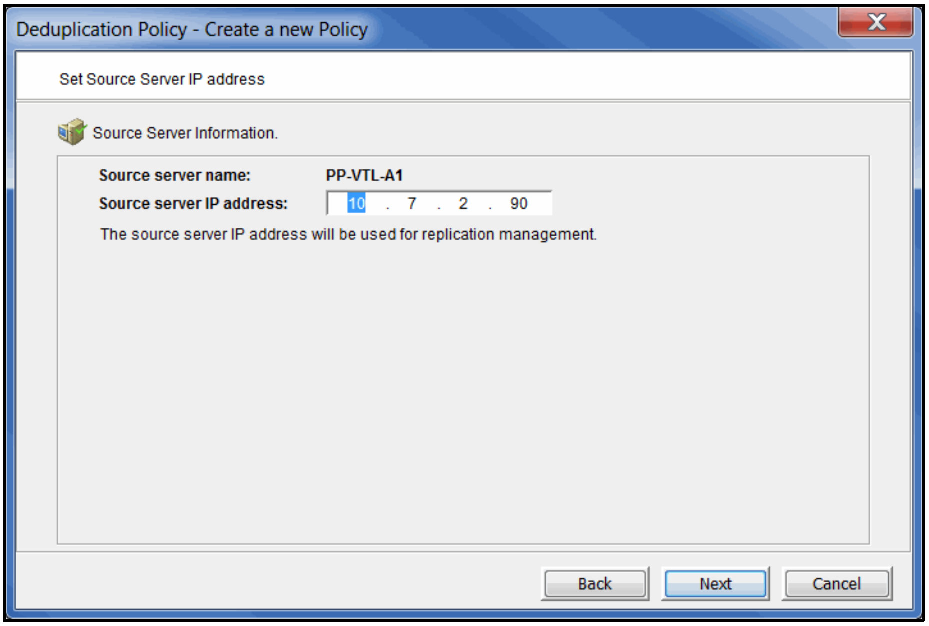

Set Source Server IP address

Source Server Information.

Source server name: PP-VTL-A1

Source server IP address: 10 . 7 . 2 . 90

The source server IP address will be used for replication management.

9. If you selected *Single* mode, select a target server.

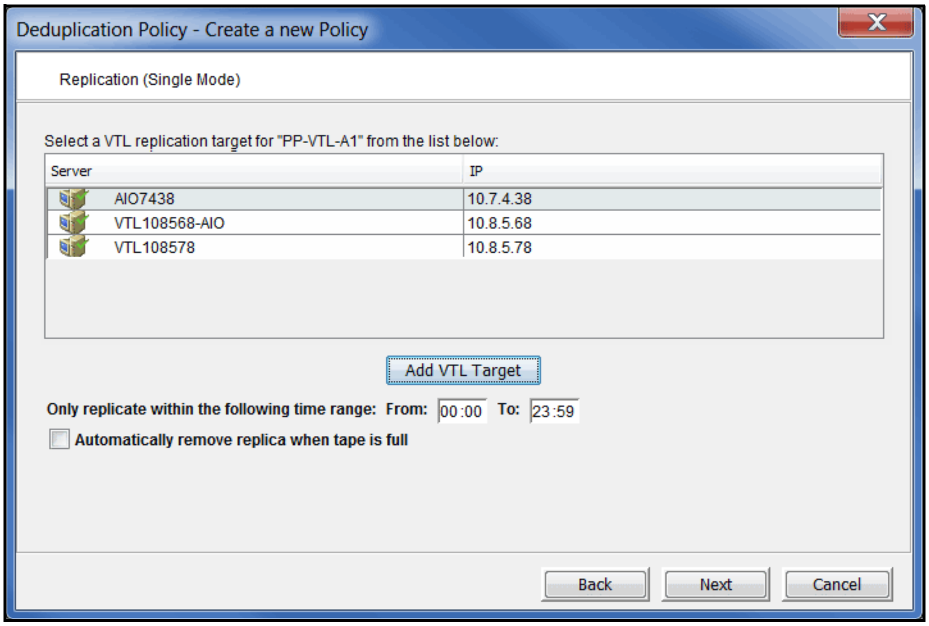

Replication (Single Mode)

Select a VTL replication target for "PP-VTL-A1" from the list below:

Server	IP
AIO7438	10.7.4.38
VTL108568-AIO	10.8.5.68
VTL108578	10.8.5.78

Only replicate within the following time range: From: 00:00 **To:** 23:59

☐ **Automatically remove replica when tape is full**

This is either the VTL target server that is associated with the *target* deduplication cluster or it is a VTL-S server. The dialog lists backup servers that have already been configured as replication targets for the source server.

> **Notes:**
>
> - The VTL target server must already be associated with the *target* deduplication cluster. Also, the deduplication source cluster must already have been configured to the same deduplication target cluster. Refer to ‘Overview of steps to configure replication for deduplicated tapes’ for more information.
> - If N+1 SIR failover is configured, only protected IP addresses can be selected; if failover is not configured, any IP address can be selected.

If the target server you want is not listed, click *Add VTL Target* and specify the IP address, user name, and password of the target VTL server.

Once you select a replication target for a policy, you cannot edit the policy to change the target.

If desired, specify a period of time within which replication can occur.

Select *Automatically remove replica when tape is full* to delete a tape’s virtual index tape from the target server when the tape is full, thereby reducing the total tape count.

10. In *Single* mode, if replication to the target server is enabled, specify if the local virtual index tape (LVIT) on the target server should be moved to a virtual library or should stay in the virtual vault after the tape is resolved.

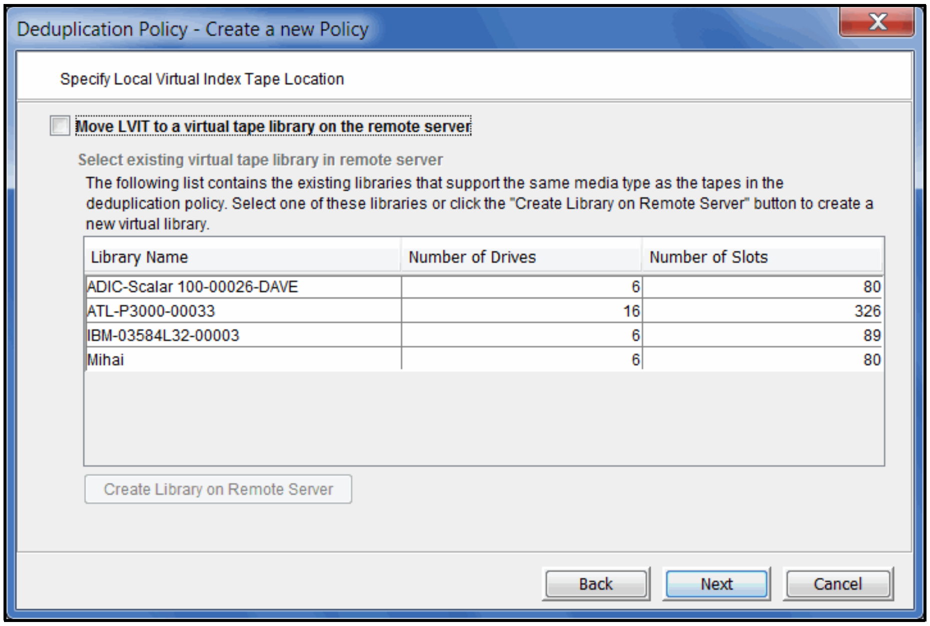

Specify Local Virtual Index Tape Location

☐ Move LVIT to a virtual tape library on the remote server

Select existing virtual tape library in remote server

The following list contains the existing libraries that support the same media type as the tapes in the deduplication policy. Select one of these libraries or click the "Create Library on Remote Server" button to create a new virtual library.

Library Name	Number of Drives	Number of Slots
ADIC-Scalar 100-00026-DAVE	6	80
ATL-P3000-00033	16	326
IBM-03584L32-00003	6	89
Mihai	6	80

You can create a library for this purpose on the target server.

11. If you selected *Cascaded* mode, choose two replication targets and an LVIT location for each target in the next four dialogs.

a. The first dialog lists backup servers that have already been identified as replication targets, just as it does in *Single* mode. The same guidelines apply here as well.

   Select a replication target for the source server (the server on which you are creating this policy). This will be displayed as *Replication Target 1* in the console.

b. In the next dialog, select a replication target for Replication Target 1. This will be displayed as *Replication Target 2* in the console.

c. Next, select an LVIT location for Replication Target 1. The same guidelines and options apply as for *Single* mode.

d. Next, select an LVIT location for Replication Target 2.

12. If you selected *Parallel* mode, choose two replication targets and an LVIT location for each target server in the next four dialogs.

    The procedure is the same for either the *Concurrent* or *Serial* option. When the deduplication policy runs, replication will be performed according to the option you selected.

    a. The first dialog lists backup servers that have already been identified as replication targets, just as it does in *Single* mode. The same guidelines apply here as well.

       Select Replication Target 1. If you chose the *Serial* option, data will be replicated to this target first.

    b. In the next dialog, select Replication Target 2. If you chose the *Serial* option, data will be replicated to this target after replication to the first target is complete.

    c. Next, select an LVIT location for Replication Target 1. The same guidelines and options apply as for *Single* mode.

    d. Next, select an LVIT location for Replication Target 2.

13. If replication is enabled, indicate how often the system should retry if a scheduled replication attempt fails.

This step applies to all replication modes.

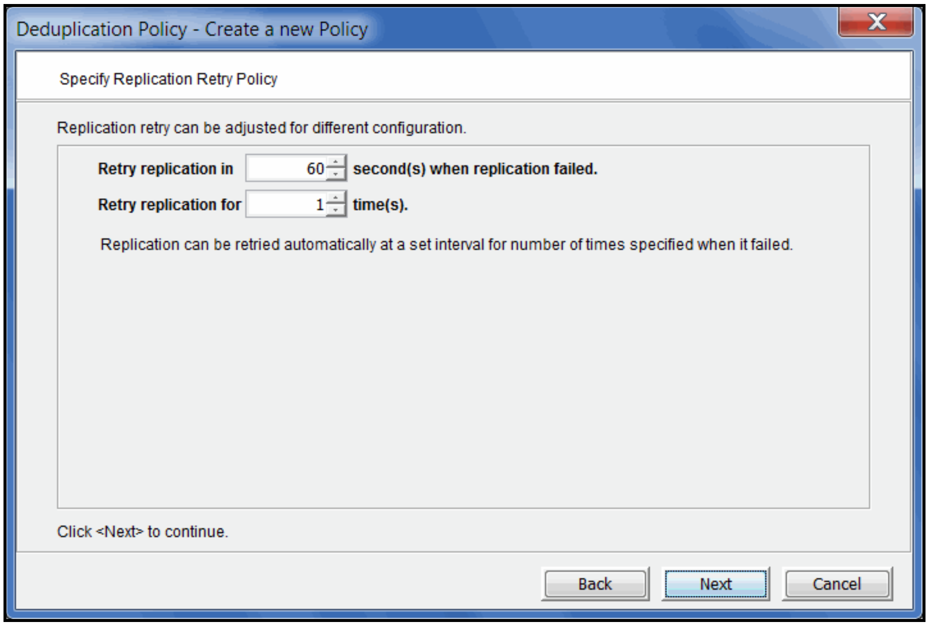

Specify Replication Retry Policy

Replication retry can be adjusted for different configuration.

Retry replication in 60 **second(s) when replication failed.**

Retry replication for 1 **time(s).**

Replication can be retried automatically at a set interval for number of times specified when it failed.

Click <Next> to continue.

The replication retry policy is specific to index replication failures and will only retry index replication-specific functions.

14. If replication is enabled, indicate if you want to use *Compression* and/or *Encryption*.

    This step applies to all replication modes.

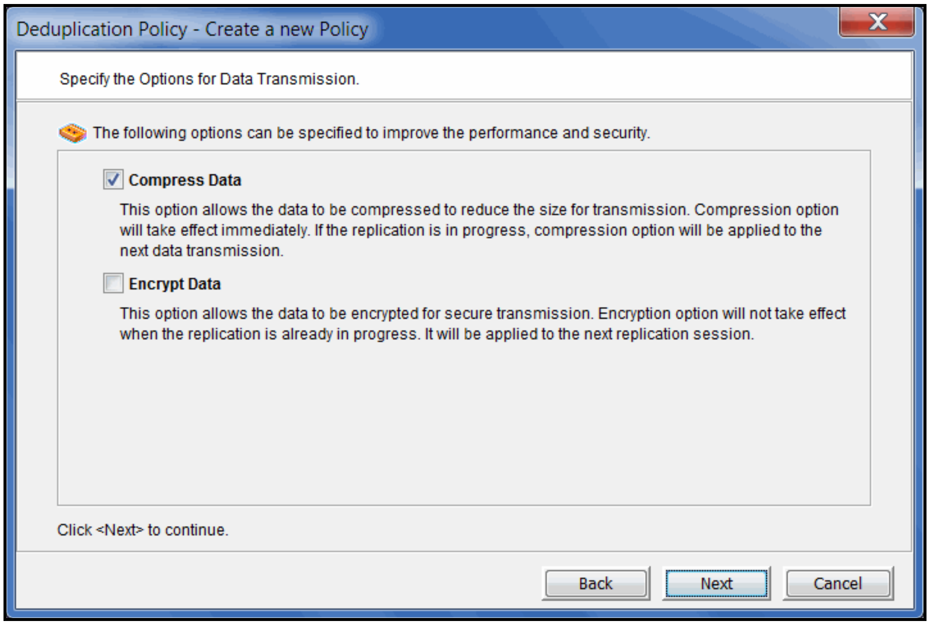

Specify the Options for Data Transmission.

The following options can be specified to improve the performance and security.

☑ **Compress Data**

This option allows the data to be compressed to reduce the size for transmission. Compression option will take effect immediately. If the replication is in progress, compression option will be applied to the next data transmission.

☐ **Encrypt Data**

This option allows the data to be encrypted for secure transmission. Encryption option will not take effect when the replication is already in progress. It will be applied to the next replication session.

Click <Next> to continue.

The *Compression* option provides enhanced throughput during replication by compressing the data stream.

The *Encryption* option 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.

15. Click *Finish* to finalize the policy.

    Once the policy is created, you will be prompted to add tapes to the policy.

16. Select the virtual tape(s) that you want to include in this policy.

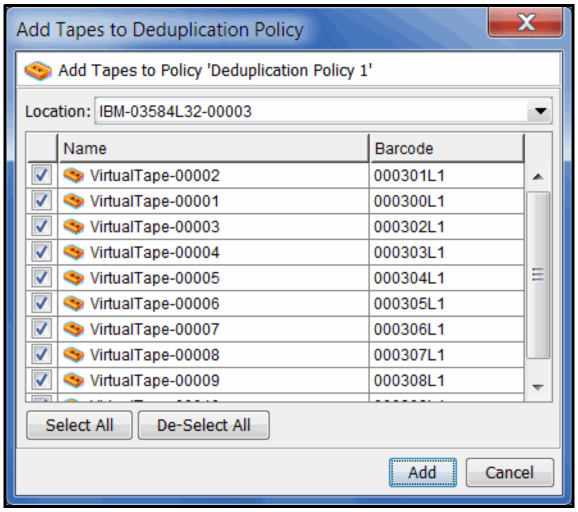
Add Tapes to Deduplication Policy
Add Tapes to Policy 'Deduplication Policy 1'
Select All
De-Select All
Add
Cancel

Location: IBM-03584L32-00003

	Name	Barcode
☑	VirtualTape-00002	000301L1
☑	VirtualTape-00001	000300L1
☑	VirtualTape-00003	000302L1
☑	VirtualTape-00004	000303L1
☑	VirtualTape-00005	000304L1
☑	VirtualTape-00006	000305L1
☑	VirtualTape-00007	000306L1
☑	VirtualTape-00008	000307L1
☑	VirtualTape-00009	000308L1

A virtual tape can be part of only one deduplication policy at a time.

Use the *Location* drop-down box to select a virtual tape library. Then, select one or more tapes. You can select tapes from multiple virtual tape libraries for the same policy.

Notes:

- You can add virtual tapes that have already been configured for VTL replication only if the target server is the same as the one specified in this policy; you will not see virtual tapes that have been configured for VTL replication to a different target server.
- Write-protected tapes will not be deduplicated.
- If inline deduplication is configured for this policy, you cannot add cached tapes. You are able to add uncached tapes, even if the library is configured to use tape caching.
- If you add a tape for which Remote Export has been configured, the tape will be added as long as it has the same target server as one of the two possible target servers in the policy.
- When *Cascaded* mode is configured, standard tapes and LVIT tapes created by replicating and resolving source tapes are automatically added to the policy on target server B.
- You should not add tapes to a policy that will use Auto Replication because Auto Replication is not supported for VITs.

Modify a tape deduplication policy

To modify the properties of a policy, right-click the policy and select *Edit.* You can change deduplication triggers, enable/disable Turbo deduplication, enable/disable replication in some cases, and change replication properties. You cannot change the replication target for a policy and you cannot change the deduplication policy name.

> **Note:** You should not use the VTL console to modify deduplication policies that were created for FSOST (via the FSOST `cfglsu` command).

Modify replication in a deduplication policy according to the following guidelines:

- You can modify a policy to enable or disable replication.
- If you enable cascaded replication in the policy on the source server and the server you choose as Replication Target 1 has a pre-existing deduplication policy with the same name, the policy on Replication Target 1 will be overwritten.
- If you disable cascaded replication in a policy on the source server, replication is disabled in the policy on Replication Target 1 automatically.

  You cannot disable cascaded replication in the policy on Replication Target 1.
- You can modify a policy with cascaded replication on the source server as well as on Replication Target 1. However, consider the following:
  - If you modify the policy on the source server, updates are made in the policy on Replication Target 1 automatically.
  - If you modify the policy on Replication Target 1, updates are made only in the policy on Replication Target 1.
- You cannot change the name of a deduplication policy.
- You cannot change replication mode from *Single* to *Advanced* or vice versa.
- You cannot change the replication target server(s).

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 deduplication policy
- Is a pure VIT
- Is in the vault

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

Add/remove tapes from a tape deduplication policy

To add tapes to a policy, right-click the policy and select *Add Tapes.*

To remove tapes from a policy, right-click the policy and select *Remove Tapes.*

If the tapes you are removing are configured for replication, the replication configuration will be maintained and the replicas for the tapes will not be deleted from the target server. If you later add the same tapes to another deduplication policy, replication will be intact. They will not have to be re-replicated, assuming the new policy has the same replication configuration (i.e., same target server). Each tape's icon will be yellow until the tape is processed again. If everything is the same, the icon turns green.

Manually run a tape deduplication policy

To execute a policy right now, right-click the policy and select *Run*. This executes deduplication/replication for all tapes in the policy.

If you want to deduplicate/replicate one or more individual tapes in the policy, highlight the policy, select the *Tapes* tab in the right column, right-click the tape(s) and select *Run*.

Delete a tape deduplication policy

To completely remove a policy, right-click the policy and select *Delete*.

If you want to delete multiple policies, right-click the *Deduplication Policies* object and select *Delete.*

If this policy contains virtual tapes that are configured for replication, the replication configuration will be maintained and the replicas for the tapes will not be deleted from the target server.

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

Suspend replication on a target in a tape deduplication policy

To suspend replication on any replication target configured for a policy, right-click the policy and select *Suspend Replication*.

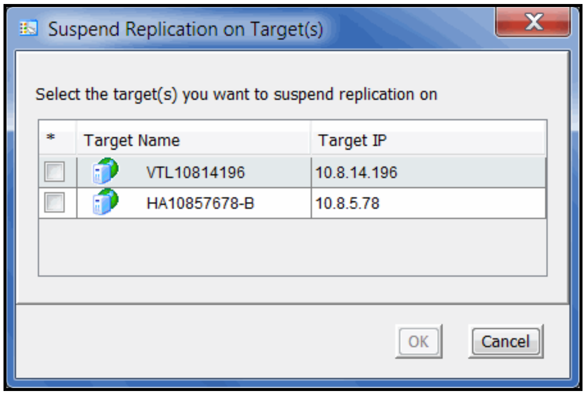
Suspend Replication on Target(s)

Select the target(s) you want to suspend replication on

*	Target Name	Target IP
□	VTL10814196	10.8.14.196
□	HA10857678-B	10.8.5.78

The dialog lists all replication targets. If the policy includes cascaded replication, only Replication Target 1 will be listed, whether you suspend replication from the source server or from Replication Target 1.

Select the checkbox next to each target for which you want to suspend replication and click *OK*.

Resume replication on a target in a tape deduplication policy

After you suspend replication on any replication target configured for a policy, *Resume Replication* is available in the right-click menu for the policy. To resume replication on a replication target, right-click the policy and select *Resume Replication*.

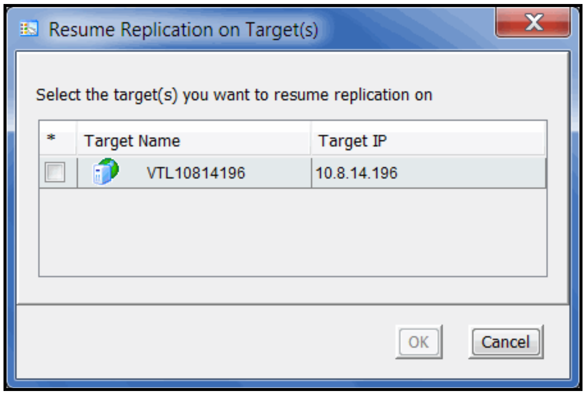
Resume Replication on Target(s)

Select the target(s) you want to resume replication on

*	Target Name	Target IP
□	VTL10814196	10.8.14.196

Select the checkbox next to each target for which you want to resume replication and click *OK*.

Manage active tape deduplication policies

Backup server

If a policy is running and you want it to stop, right-click the policy and select *Stop*.

If you want to stop multiple policies that are running, right-click the *Deduplication Policies* object and select *Stop.*

Select *Activities --> Deduplication Job Queue* to view a list of all of the tapes that are being processed. From here, you can change the priority of tapes in the queue and cancel processing for a specific tape in the queue.

This is useful for disaster recovery (DR) purposes when a tape has replication configured and is needed at the DR site. You can also cancel processing for a tape by selecting *Remove* from the right-click menu.

To change the priority of tapes, right-click a queued tape and select *Run Next* or *Run Later*.

To cancel processing for a specific tape, select *Remove* from the right-click menu.

Deduplication server

You can also monitor the Deduplication Job Queue from the command line.

Type `sirnodeqcli` at the command line of a deduplication server to display a list of commands. The commands are:

- `sirnodeqcli init` - Initialize the database for the queue. (This option should only be used in conjunction with DSI Technical Support.)
- `sirnodeqcli reset` - Reset the queue, deleting all tasks in the database.
- `sirnodeqcli suspend` - Suspend the queue.
- `sirnodeqcli resume` - Resume the queue.
- `sirnodeqcli list` - List jobs in the queue. (Sample output is shown below.)
- `sirnodeqcli find -d driveSN` - Specify the virtual drive’s serial number to find a job in the queue. You can find the serial number of the virtual drive on the *Virtual Drives* tab of your virtual library in the VTL console.
- `sirnodeqcli remove -d driveSN` - Specify the virtual drive’s serial number to remove a job from the queue.

Sample output for the `sirnodeqcli list` command is:

```
ID|DRIVE SN|SOURCE IP|CONTROL|TIMESTAMP|PAYLOAD TYPE|OP CODE
50,76Q6M0020K,,RUNNING,2011-03-03 15:30:48,SCANTAPE,1
53,76Q6M0020G,,RUNNING,2011-03-03 15:36:07,SCANTAPE,14
```

Failover of tape deduplication policies

When failover occurs, deduplication policies will continue to run according to their existing schedule. However, you cannot edit or create policies while the server is in failover state.

If failover occurs on the backup server while deduplication is occurring, the original virtual tapes will be returned to their original location, and the temporary VIT tapes will be moved to the virtual vault. It is safe to delete these temporary tapes from the virtual vault after you have confirmed that the original tapes are in a good state.

NAS deduplication

Check integrity of deduplicated NAS data

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.

Enable integrity checking

When you enable integrity checking, all of the configuration options that you set are part of a policy. If you do not select to schedule integrity checks, the configuration options will still be used when you manually run an integrity check.

1. Right-click your server and select *Deduplication* --> *NAS Integrity Check* --> *Enable*.

2. Specify when the integrity check should be run.

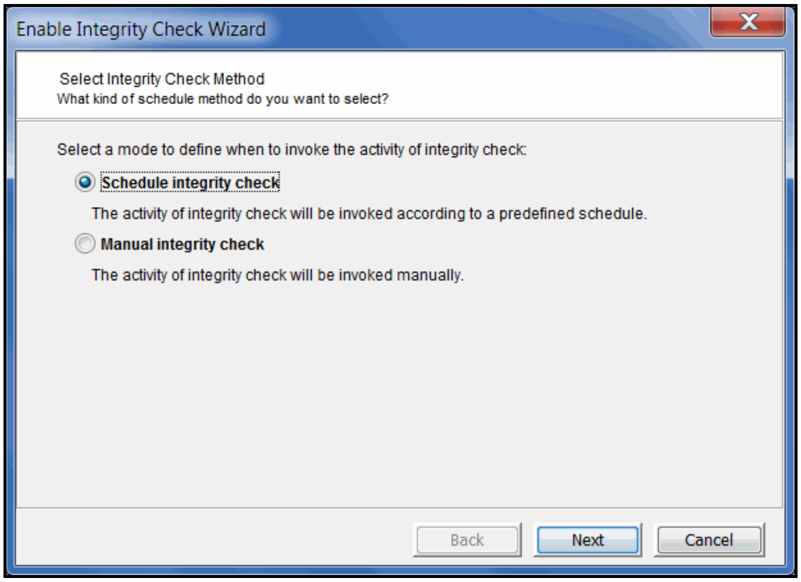

Select Integrity Check Method

What kind of schedule method do you want to select?

Select a mode to define when to invoke the activity of integrity check:

● **Schedule integrity check**

The activity of integrity check will be invoked according to a predefined schedule.

○ **Manual integrity check**

The activity of integrity check will be invoked manually.

3. If you selected to schedule the integrity check, set the schedule.

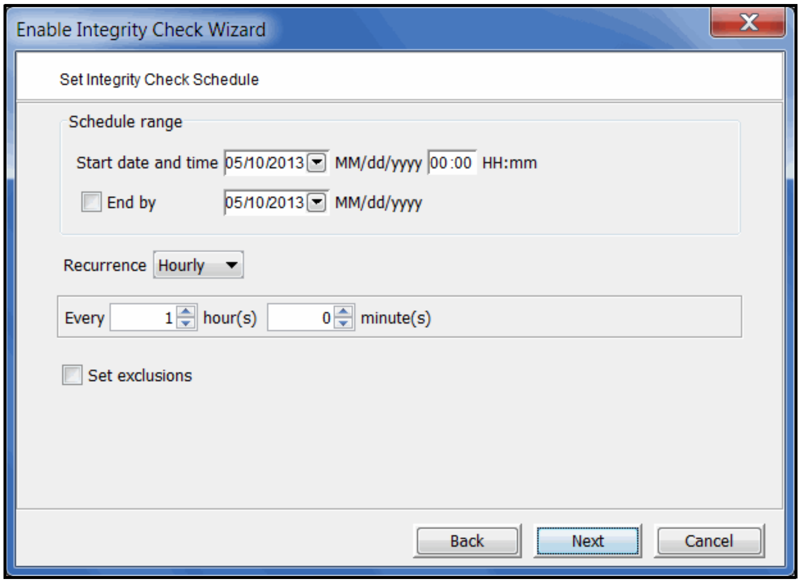

Set Integrity Check Schedule

Schedule range

Start date and time 05/10/2013 MM/dd/yyyy 00:00 HH:mm

□ End by 05/10/2013 MM/dd/yyyy

Recurrence Hourly

Every 1 hour(s) 0 minute(s)

□ Set exclusions

Specify when the integrity check should begin and, optionally, when it should end. Also specify the frequency. If you want to exclude specific days/hours/months, select *Set exclusions*.

4. If you selected *Set exclusions,* select the hours, days, or months during which the integrity check should not run.

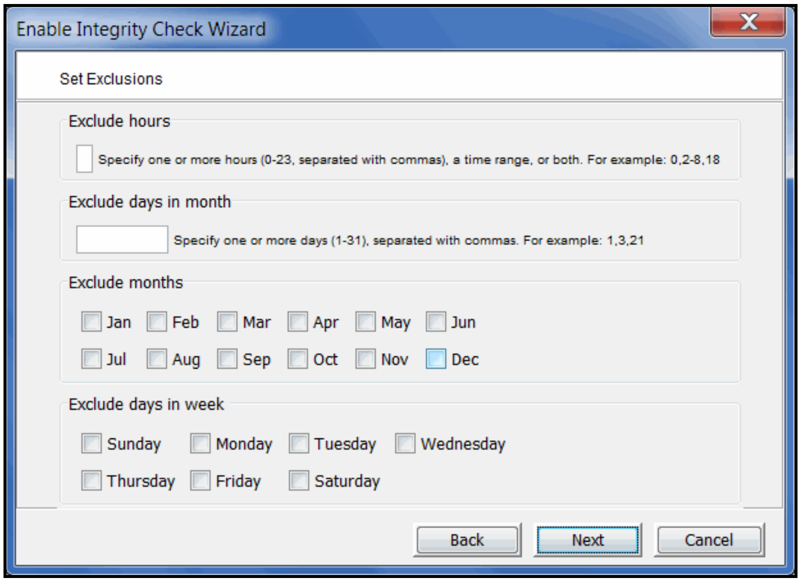

Enable Integrity Check Wizard

Set Exclusions

Exclude hours

Specify one or more hours (0-23, separated with commas), a time range, or both. For example: 0,2-8,18

Exclude days in month

Specify one or more days (1-31), separated with commas. For example: 1,3,21

Exclude months

□ Jan □ Feb □ Mar □ Apr □ May □ Jun
□ Jul □ Aug □ Sep □ Oct □ Nov □ Dec

Exclude days in week

□ Sunday □ Monday □ Tuesday □ Wednesday
□ Thursday □ Friday □ Saturday

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

5. Specify if you want to include all paths in the integrity check or if you want to select specific paths to include.

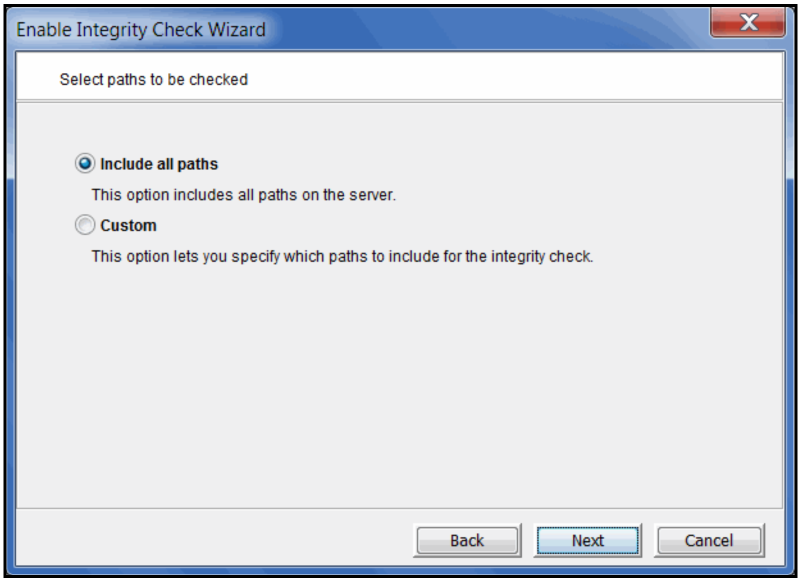

Enable Integrity Check Wizard

Select paths to be checked

● **Include all paths**

This option includes all paths on the server.

○ **Custom**

This option lets you specify which paths to include for the integrity check.

6. If you selected *Custom*, select what paths to include/exclude from the integrity check.

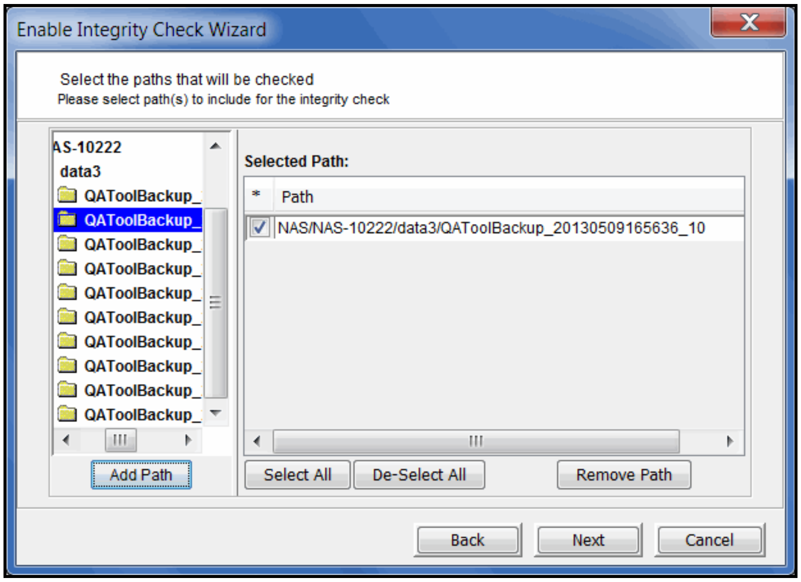

Enable Integrity Check Wizard

Select the paths that will be checked

Please select path(s) to include for the integrity check

Selected Path:

*	Path
☑	NAS/NAS-10222/data3/QAToolBackup_20130509165636_10

The left pane lists all NAS resources. Select a folder that you want to include and click the *Add Path* button. The newly added row in the table shows the detailed path.

7. Specify the validation criteria.

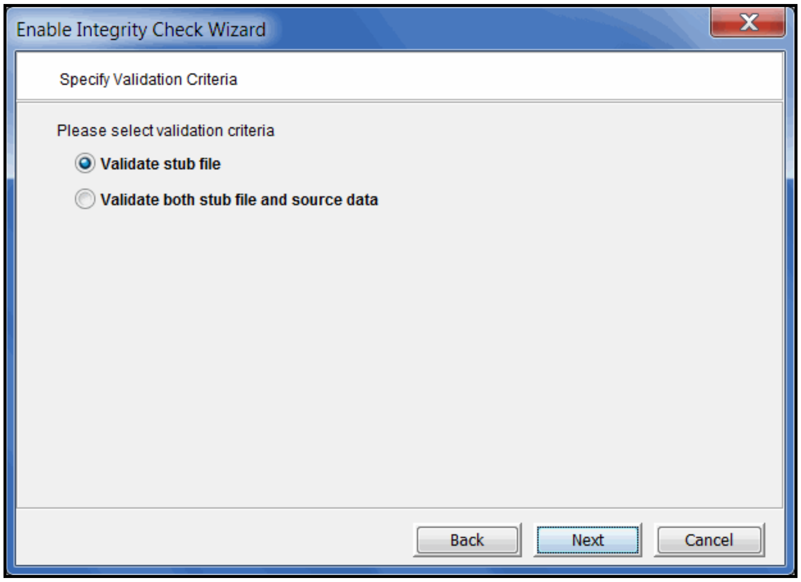

Enable Integrity Check Wizard

Specify Validation Criteria

Please select validation criteria

● Validate stub file

○ Validate both stub file and source data

You can select to validate stub pointers only or include source data validation as well.

8. Confirm all information and click *Finish*.

Manually run an integrity check

To manually run an integrity check, right-click your server and select *Deduplication --> NAS Integrity Check* --> *Start*.

If you created an integrity check policy, all of the configuration options that you set will be used when you manually run an integrity check.

View integrity check statistics

You can see the status by selecting the *NAS Integrity Check Statistics* object (under *Status*).

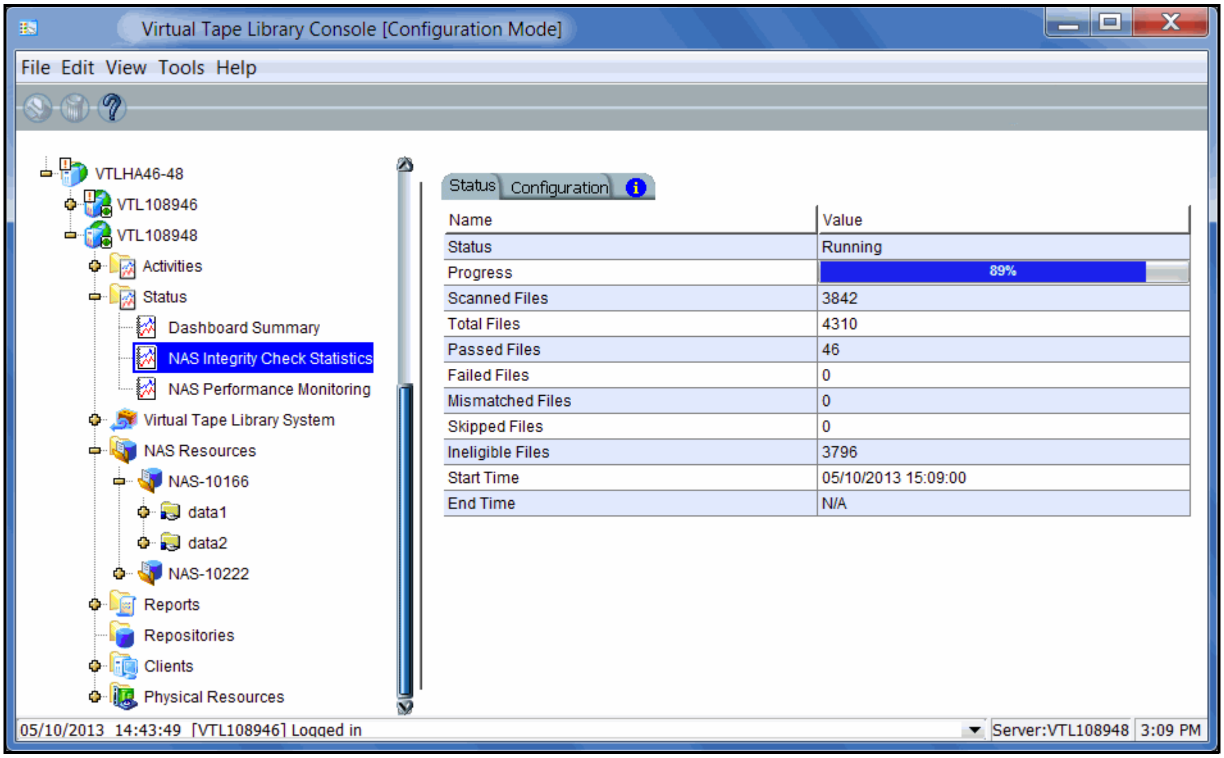

Name	Value
Status	Running
Progress	89%
Scanned Files	3842
Total Files	4310
Passed Files	46
Failed Files	0
Mismatched Files	0
Skipped Files	0
Ineligible Files	3796
Start Time	05/10/2013 15:09:00
End Time	N/A

05/10/2013 14:43:49 [VTL108946] Logged in Server:VTL108948 3:09 PM

Scanned Files - Number of files scanned so far. If the integrity check is finished, this will be the same as the number of total files.

Total Files - Number of files in the system.

Passed Files - Number of files that have successfully passed the integrity check.

Failed Files - Number of files for which the integrity checking operation failed and the system cannot tell if integrity of the file is good.

Mismatched Files - Files that were checked but did pass the integrity check due to a mismatch. There may be a problem retrieving data from these files.

Skipped Files - Files that were skipped during the integrity check. This can be caused when a high-priority process or application tries to access the file. It can also mean that the file was deleted before the integrity check process started.

Ineligible Files - Files that cannot be verified because integrity checking had not been enabled when the file was deduplicated.

Stop/suspend an integrity check

To stop an integrity check that is running, right-click your server and select *Deduplication* --> *NAS Integrity Check* --> *Stop*.

To temporarily suspend an integrity check that is running, right-click your server and select *Deduplication* --> *NAS Integrity Check* --> *Suspend.* When you want to resume the job, right-click your server and select *Deduplication* --> *NAS Integrity Check* --> *Resume.*

Change integrity check properties

To change the configuration, right-click your server and select *Deduplication* --> *NAS Integrity Check* --> *Configure*.

Disable integrity checking

To disable integrity checking, right-click your server and select *Deduplication* --> *NAS Integrity Check* -->*Disable.* You will need to type “Yes” to confirm.

Include/exclude NAS resources and shares for deduplication

By default, all NAS files are 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. Therefore, there is no benefit to deduplicating files smaller than 8 KB.

Exclude

If you need to exclude specific NAS resources or shares from deduplication, right-click the resource or share and select *Deduplication* --> *Exclude from Deduplication.*

You will see an X icon on every resource/share that is excluded.

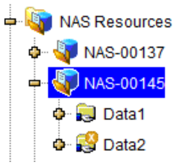

You can see a list of all excluded resources and shares on the *Excluded Paths* tab of the *NAS Resources* object.

> **Note:** NAS OST resources should not be excluded from deduplication; excluding OST resources from being deduplicated can cause duplication jobs to hang in Symantec Backup Exec.

Include

To include a folder that was previously excluded, right-click the folder and select *Deduplication* --> *Include in Deduplication*.

Copy NAS files and directories using Fast Copy

At any time, you can map/mount a NAS share and copy files and directories.

You can also use the Fast Copy feature to copy files and directories between NAS resources within the same server. With Fast Copy, if the source files have been deduplicated, the stub file will be copied so that the file does 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.

1. Right-click the *NAS Resources* object or an share/folder and select *Fast Copy*.

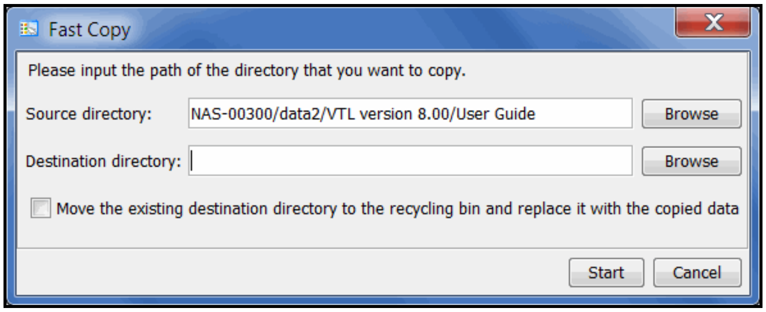

Fast Copy

Please input the path of the directory that you want to copy.

Source directory: NAS-00300/data2/VTL version 8.00/User Guide

Destination directory:

☐ Move the existing destination directory to the recycling bin and replace it with the copied data

2. If needed, select the source directory.

   If you started Fast Copy from a share/folder, that path is automatically filled. Otherwise, you can click *Browse* to find the source.

3. Select the destination directory.

   You can click *Browse* to find the directory. Click *New Folder* to create a new folder.

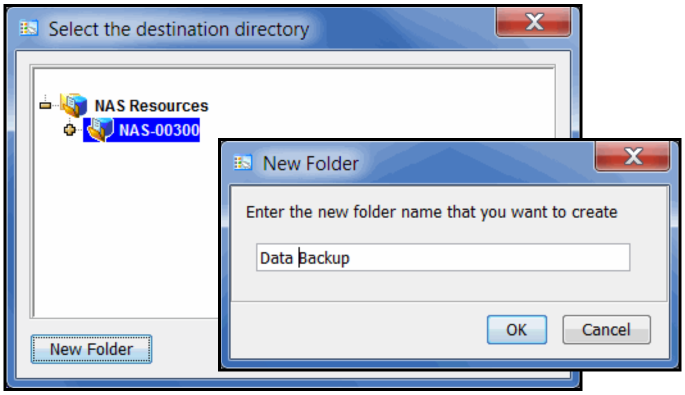

Select the destination directory

New Folder

Enter the new folder name that you want to create

Data Backup

4. If you selected an existing destination directory, you **must** select the *Move the existing destination directory to the recycling bin...* checkbox.

   If you do not select that checkbox, the job will fail.

   The fastcopy-recycle.bin holds the content of the directory that was replaced. If necessary, you can share the bin and then map/mount it in order to recover files and directories.

5. Click *Start* to begin copying.

Monitor performance

NAS Performance Monitoring (under the *Status* object) displays charts showing real-time performance statistics for file systems, outgoing replication, and incoming replication.

NAS Resources

The *NAS Resources* tab lists all of the NAS resources on this server. You can see performance information for all NAS resources or a single NAS resource. Highlight a NAS resource and the *Performance* section displays current read and write throughput for the NAS resource as well as average read and write throughput for the period of time for which status is currently displayed. The graph shows the throughput information for up to the past 30 minutes.

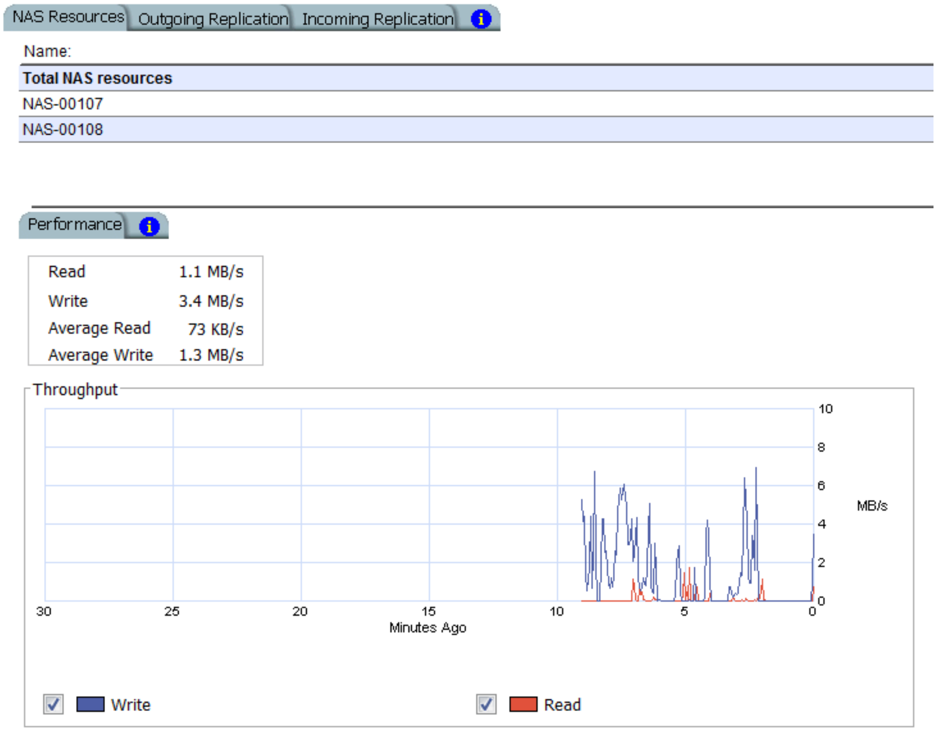

Name:
Total NAS resources
NAS-00107
NAS-00108

Read	1.1 MB/s
Write	3.4 MB/s
Average Read	73 KB/s
Average Write	1.3 MB/s

Click the appropriate checkboxes to display write and/or read performance.

Outgoing replication

The *Outgoing Replication* tab displays the current replication throughput and indicates if throttling is on. The graph shows the throughput information for up to the past 30 minutes.

Incoming replication

The *Incoming Replication* tab lists all of the source replication servers. Highlight a source server and the *Performance* section displays current throughput for that server as well as average throughput for up to the past 30 minutes.

Monitor deduplication and view statistics

In order to monitor deduplication, including replication of deduplicated information, it is important to be able to identify all of the servers that interact with the cluster. Having this information allows you to quickly access information on current activities and status. When you highlight a cluster in the tree, you will see the following tabs:

- *Associated Servers* - This tab lists all of the backup servers that are associated with the cluster, including the IP address and type for each. If the backup servers are configured in an Active-Passive failover configuration, only the active node will be listed.
- *Replication* - This tab lists all *source* and *target* relationships that include this cluster and the communications protocol for each.
- *Redundant Node* - This tab is displayed if failover is configured for this cluster.

Information available from the deduplication server

You can view the following on the deduplication server:

- Status of running policies
- Scanner history
- Repository statistics for the cluster

Status of running tape policies

To view information for tape policies that are currently running, select the *Activities* object for the server and then select *Scanner processes*.

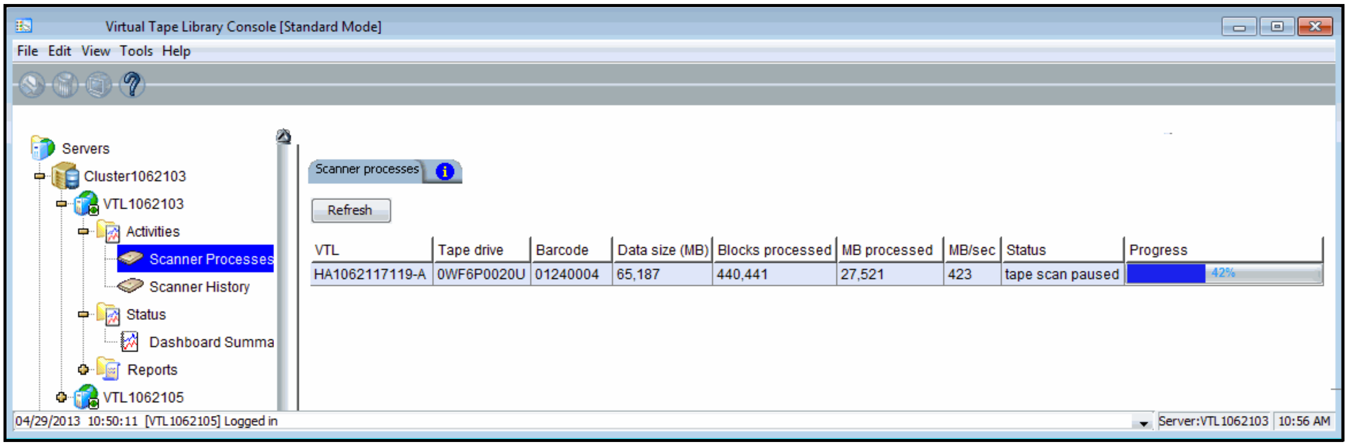

VTL	Tape drive	Barcode	Data size (MB)	Blocks processed	MB processed	MB/sec	Status	Progress
HA1062117119-A	0WF6P0020U	01240004	65,187	440,441	27,521	423	tape scan paused	42%

04/29/2013 10:50:11 [VTL1062105] Logged in

Server:VTL1062103 10:56 AM

Here you can view the progress of each job. You can see which virtual tape is being deduplicated, how large the tape is, how much has been processed, and the performance, measured in MB per second.

Tape scanner history

To view details of past activity for a specific deduplication server, select the *Activities* object for the server and then select *Scanner History*.

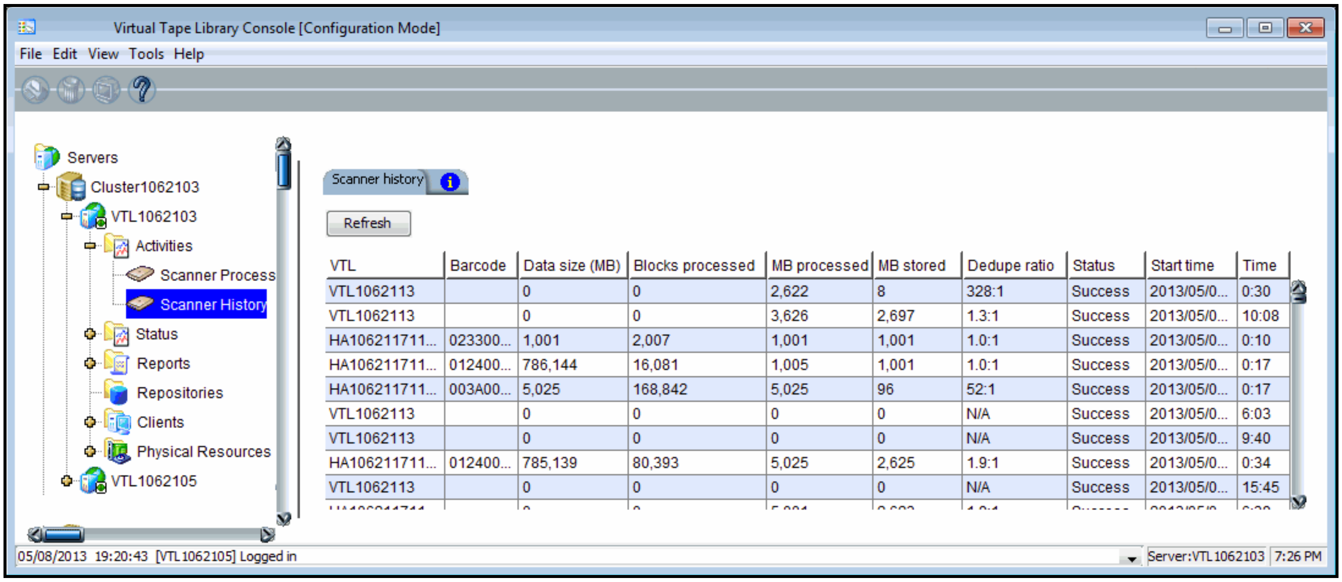

VTL	Barcode	Data size (MB)	Blocks processed	MB processed	MB stored	Dedupe ratio	Status	Start time	Time
VTL1062113		0	0	2,622	8	328:1	Success	2013/05/0...	0:30
VTL1062113		0	0	3,626	2,697	1.3:1	Success	2013/05/0...	10:08
HA106211711...	023300...	1,001	2,007	1,001	1,001	1.0:1	Success	2013/05/0...	0:10
HA106211711...	012400...	786,144	16,081	1,005	1,001	1.0:1	Success	2013/05/0...	0:17
HA106211711...	003A00...	5,025	168,842	5,025	96	52:1	Success	2013/05/0...	0:17
VTL1062113		0	0	0	0	N/A	Success	2013/05/0...	6:03
VTL1062113		0	0	0	0	N/A	Success	2013/05/0...	9:40
HA106211711...	012400...	785,139	80,393	5,025	2,625	1.9:1	Success	2013/05/0...	0:34
VTL1062113		0	0	0	0	N/A	Success	2013/05/0...	15:45

05/08/2013 19:20:43 [VTL1062105] Logged in | Server:VTL1062103 | 7:26 PM

The display includes the history of the virtual tapes that have been deduplicated.

Under *MB processed* you will see the amount of original data sent from the tape, and under *MB stored* you will see the amount of unique data on that tape that has been stored in the deduplication repository.

You will also see the amount of time required for each tape's deduplication and the start time, from which you can determine the time when deduplication completed.

Repository statistics

To view repository statistics for the entire cluster, highlight the cluster object and select the *Dashboard Summary* tab in the right pane. For a VTL-S server, this information is available from the *Status* object --> *Dashboard Summary* --> *Deduplication Repository* tab.

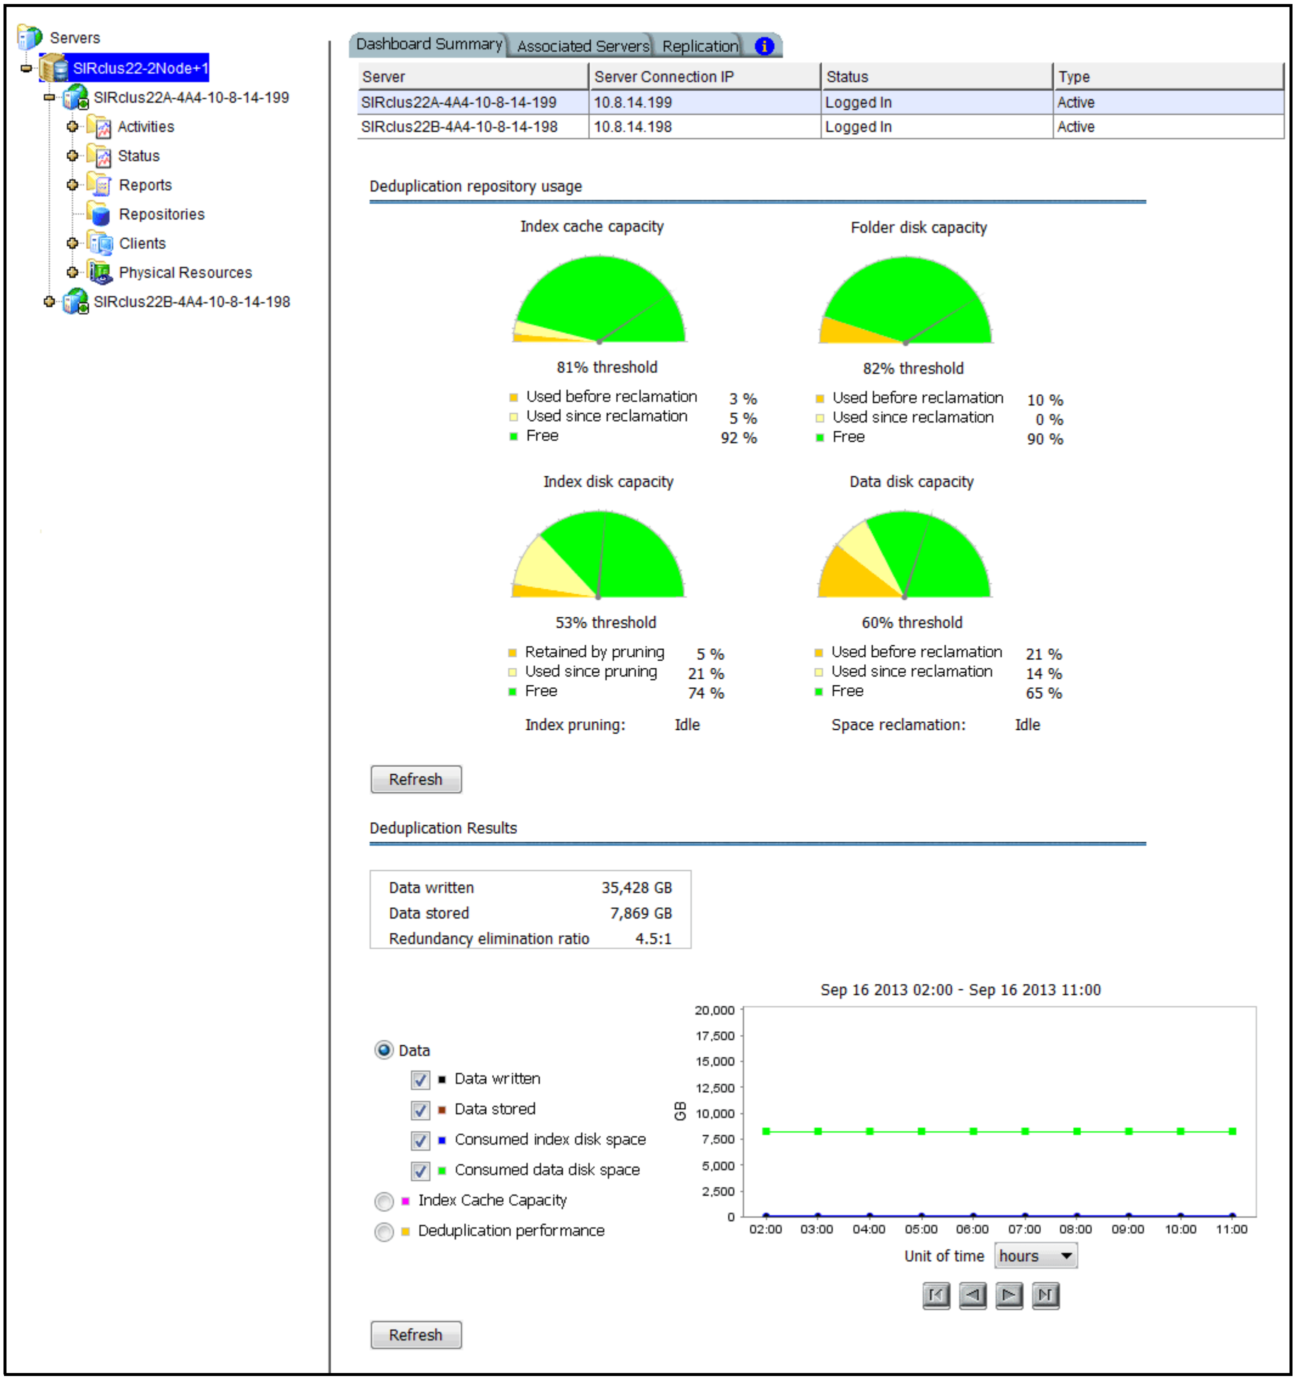

Server	Server Connection IP	Status	Type
SIRclus22A-4A4-10-8-14-199	10.8.14.199	Logged In	Active
SIRclus22B-4A4-10-8-14-198	10.8.14.198	Logged In	Active

Deduplication repository usage

Index cache capacity

81% threshold

🟧 Used before reclamation	3 %
🟨 Used since reclamation	5 %
🟩 Free	92 %

Folder disk capacity

82% threshold

🟧 Used before reclamation	10 %
🟨 Used since reclamation	0 %
🟩 Free	90 %

Index disk capacity

53% threshold

🟧 Retained by pruning	5 %
🟨 Used since pruning	21 %
🟩 Free	74 %

Index pruning: Idle

Data disk capacity

60% threshold

🟧 Used before reclamation	21 %
🟨 Used since reclamation	14 %
🟩 Free	65 %

Space reclamation: Idle

Refresh

Deduplication Results

Data written	35,428 GB
Data stored	7,869 GB
Redundancy elimination ratio	4.5:1

● Data
- ☑ ■ Data written
- ☑ ■ Data stored
- ☑ ■ Consumed index disk space
- ☑ ■ Consumed data disk space

○ ■ Index Cache Capacity
○ ■ Deduplication performance

Sep 16 2013 02:00 - Sep 16 2013 11:00

Unit of time: hours

Refresh

Deduplication Repository usage

This section graphically displays the current state of deduplication storage for the cluster. Values are based on all scans performed during the life span of the selected server. When reclamation is enabled, the graphs in this section resemble a dashboard. Green represents free space, while the shades of yellow represent space used as of last reclamation and used since reclamation. The needle shows where the current threshold is. After each reclamation, the system calculates a new trigger.

Index cache capacity shows the amount of total deduplication index cache that is used and available.

- *Used before reclamation* - This number indicates index cache usage immediately after the last successful reclamation. The value will be zero before reclamation is run for the first time.
- *Used since reclamation* - This number indicates the amount of index cache capacity used after the last successful reclamation.
- *Free* - This number indicates memory that is free.

Folder disk capacity shows the capacity of the folder disk, how much space has been used, and how much space is available.

- *Used before reclamation* - This number indicates folder disk space usage immediately after the last successful reclamation. The value will be zero before reclamation is run for the first time.
- *Used since reclamation* - This number indicates the amount of folder capacity used after the last successful pruning.
- *Free* - This number indicates available folder space.

Index disk capacity shows the capacity of the index disk, how much space has been used, and how much space is available. If index pruning is taking place, a status bar displays the progress.

- *Retained by pruning* - This number indicates index disk space usage immediately after the last successful pruning. The value will be zero before reclamation is run for the first time.
- *Used since pruning* - This number indicates the amount of index space used after the last successful pruning.
- *Free* - This number indicates available index space.

Data disk capacity shows the capacity of the deduplication data disk(s), how much space has been used, and how much space is available. If space reclamation is taking place, a status bar displays the progress.

- *Used before reclamation* - This number indicates deduplication data usage immediately after the last successful reclamation. The value will be zero before reclamation is run for the first time.
- *Used since reclamation* - This number indicates the amount of deduplicated data added after the last successful reclamation.
- *Free* - This number indicates space that is available for deduplicated data.

Select *Refresh* to update the display.

Deduplication results

This section displays overall deduplication statistics and statistics for a user-defined period of time. Note that NAS statistics are updated every hour.

The box above the graph displays data written, data stored, and the redundancy elimination ratio.

Data written represents data scanned, updated upon completion of each deduplication process and replication job (on the target server). *Data written* is calculated after folders are closed, such as when a backup is complete. Therefore, updated values may be delayed if an application keeps a folder open.

Data stored is the amount of unique data stored in the repository during each time interval, updated upon completion of each deduplication process and target replication job.

The *Redundancy elimination ratio* (frequently referred to in the industry as the *Deduplication Ratio*) is an approximate ratio that represents this formula: [(data scanned) / (data stored)].

This bottom section allows you to display your choice of deduplication statistics for a specific period of time:

- *Data written* - Data scanned during each time interval (i.e., each hour).
- *Data stored* - The amount of unique data stored in the repository during each time interval.
- *Consumed index disk space* - Space used by index. Each point represents the amount used as of the end of that time interval. Data is affected by reclamation.
- *Consumed data disk space* - Space used by data. Each point represents the amount used as of the end of that time interval. Data is affected by reclamation.
- *Index cache capacity* - Space used by index cache. Each point represents the percentage used as of the end of that time interval. Data is affected by reclamation.
- *Deduplication performance* - Cumulative amount of data deduplicated for that unit of time. The statistics are not dependent upon the completion of a job.

These statistics show deduplication activity over time. Viewing data in this way allows you to calculate the redundancy elimination ratio for any period of time.

Reviewing deduplication operations for successive weeks of full backup reveals the approximate redundancy ratios of week-to-week data evolution and can be used to accurately forecast repository requirements. You can identify how quickly you are using your repository disk space and when you are likely to need to add more.

Select a *Unit of time* (hours, days, weeks, or months) from the drop-down list to adjust the granularity of the graph. The data points in the graph will match the starting point for that unit. For example, if you select *Months*, the data point for March will show statistics for just after midnight on March 1. Use the arrow buttons to scan through accumulated data.

Click *Refresh* to include data for deduplication activity that has occurred since the last refresh.

You can put your cursor on a data point to see detailed information.

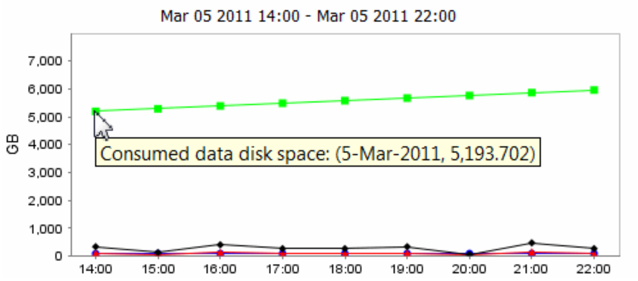
Mar 05 2011 14:00 - Mar 05 2011 22:00
GB
7,000
6,000
5,000
4,000
3,000
2,000
1,000
0
Consumed data disk space: (5-Mar-2011, 5,193.702)
14:00
15:00
16:00
17:00
18:00
19:00
20:00
21:00
22:00

If you want to zoom into the chart to enlarge it, drag your cursor from left to right over the area you want to expand.

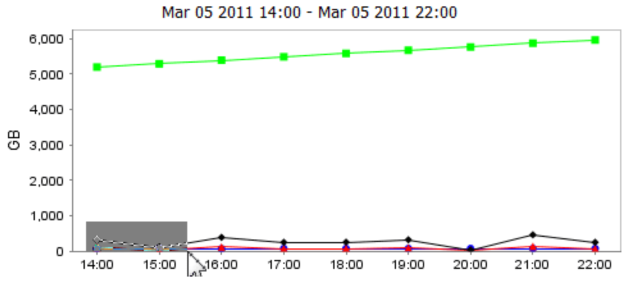

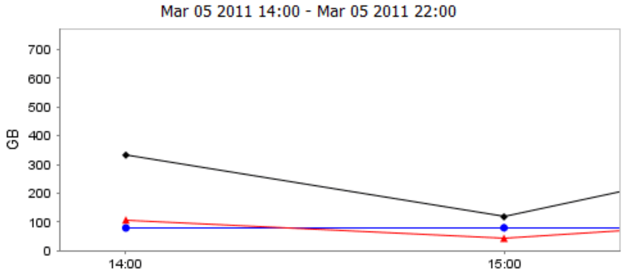

When you are finished, drag your cursor from right to left anywhere in the chart and the display will zoom out, back to a normal view.

Tape information available from the backup server

The management console offers the following tape deduplication-related information for your backup server:

- Information about each policy
- Information about each virtual tape in a policy
- Status of running policies
- Virtual tape history
- Event log entries pertaining to each policy
- Virtual index tape status

Deduplication Policies object

When you highlight the *Deduplication Policies* object, the right-hand pane lists all of the policies that are configured for deduplication on this server.

For each policy, you can see the number of tapes included, schedule information (such as status, history, and next run time), and the deduplication cluster to which this policy belongs.

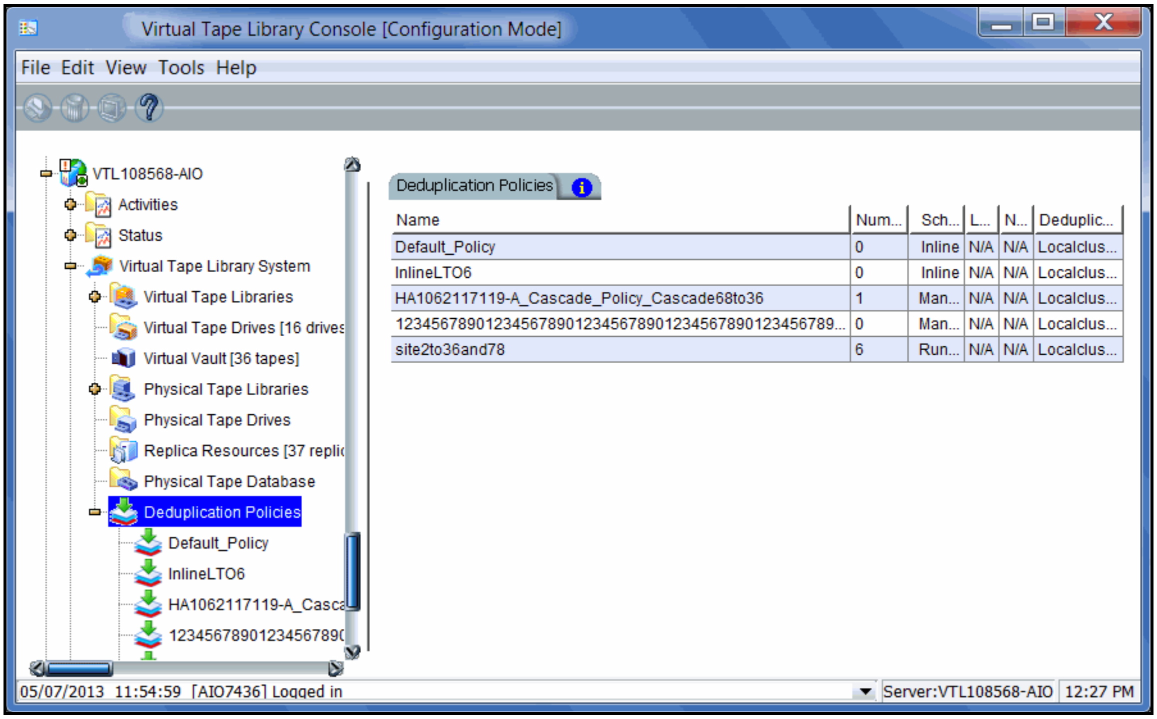

Name	Num...	Sch...	L...	N...	Deduplic...
Default_Policy	0	Inline	N/A	N/A	Localclus...
InlineLTO6	0	Inline	N/A	N/A	Localclus...
HA1062117119-A_Cascade_Policy_Cascade68to36	1	Man...	N/A	N/A	Localclus...
1234567890123456789012345678901234567890123456789...	0	Man...	N/A	N/A	Localclus...
site2to36and78	6	Run...	N/A	N/A	Localclus...

05/07/2013 11:54:59 [AIO7436] Logged in Server:VTL108568-AIO 12:27 PM

Individual tape deduplication policies

When you highlight a policy in the tree, you can view information about that policy.

General Info tab

The *General Info* tab shows how many tapes are included in this policy, deduplication server information, and schedule and replication information. If *Advanced Replication* is enabled, information is provided for Replication Target 1 and Replication Target 2.

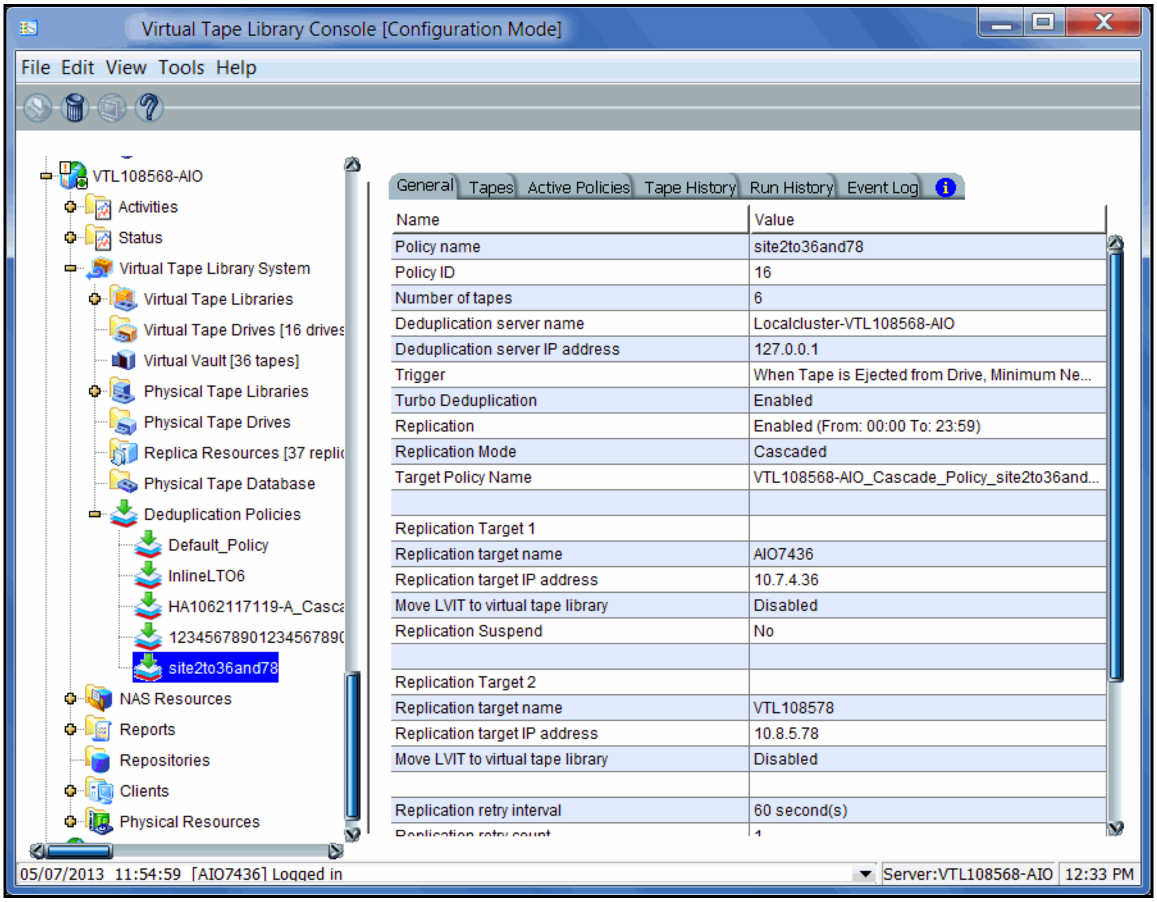

Name	Value
Policy name	site2to36and78
Policy ID	16
Number of tapes	6
Deduplication server name	Localcluster-VTL108568-AIO
Deduplication server IP address	127.0.0.1
Trigger	When Tape is Ejected from Drive, Minimum Ne...
Turbo Deduplication	Enabled
Replication	Enabled (From: 00:00 To: 23:59)
Replication Mode	Cascaded
Target Policy Name	VTL108568-AIO_Cascade_Policy_site2to36and...
Replication Target 1	
Replication target name	AIO7436
Replication target IP address	10.7.4.36
Move LVIT to virtual tape library	Disabled
Replication Suspend	No
Replication Target 2	
Replication target name	VTL108578
Replication target IP address	10.8.5.78
Move LVIT to virtual tape library	Disabled
Replication retry interval	60 second(s)
Replication retry count	1

05/07/2013 11:54:59 [AIO7436] Logged in | Server:VTL108568-AIO | 12:33 PM

Tapes tab

The *Tapes* tab lists information about each virtual tape in the policy. The icon next to the tape name indicates the status of the last operation performed (“D” for deduplication or “R” for deduplication with replication). Refer to ‘Display virtual tapes’ for information about what each status icon means. Note that all values are rounded.

Maximum Capacity - Maximum uncompressed storage capacity of the tape. This is determined when the tape was created.

Written - The amount of data (before compression) that is written to tape by backup applications. This amount can be greater than the tape size if the data is compressed.

New - The amount of data (before compression) that has not yet been deduplicated, including newly appended data to a tape.

In SIR - The amount of data (before compression) written that has now been moved to the deduplication repository. This is basically the difference between the data written and the data not yet deduplicated.

Unique - The actual physical storage used to store tape data. This includes the effect of deduplication compression.

Dedupe ratio - The ratio between the data moved to the deduplication repository and the unique data. A ratio of >10000:1 indicates that extremely little data has changed and the amount of unique data is very small or zero.

Last run Dedupe - The last time the tape was deduplicated.

Last run Replicated - The last time the tape was replicated.

Next run - The next time the tape will be deduplicated.

When you highlight a tape in the top section, the *Policy Tape Info* tab in the bottom section displays additional details about the tape.

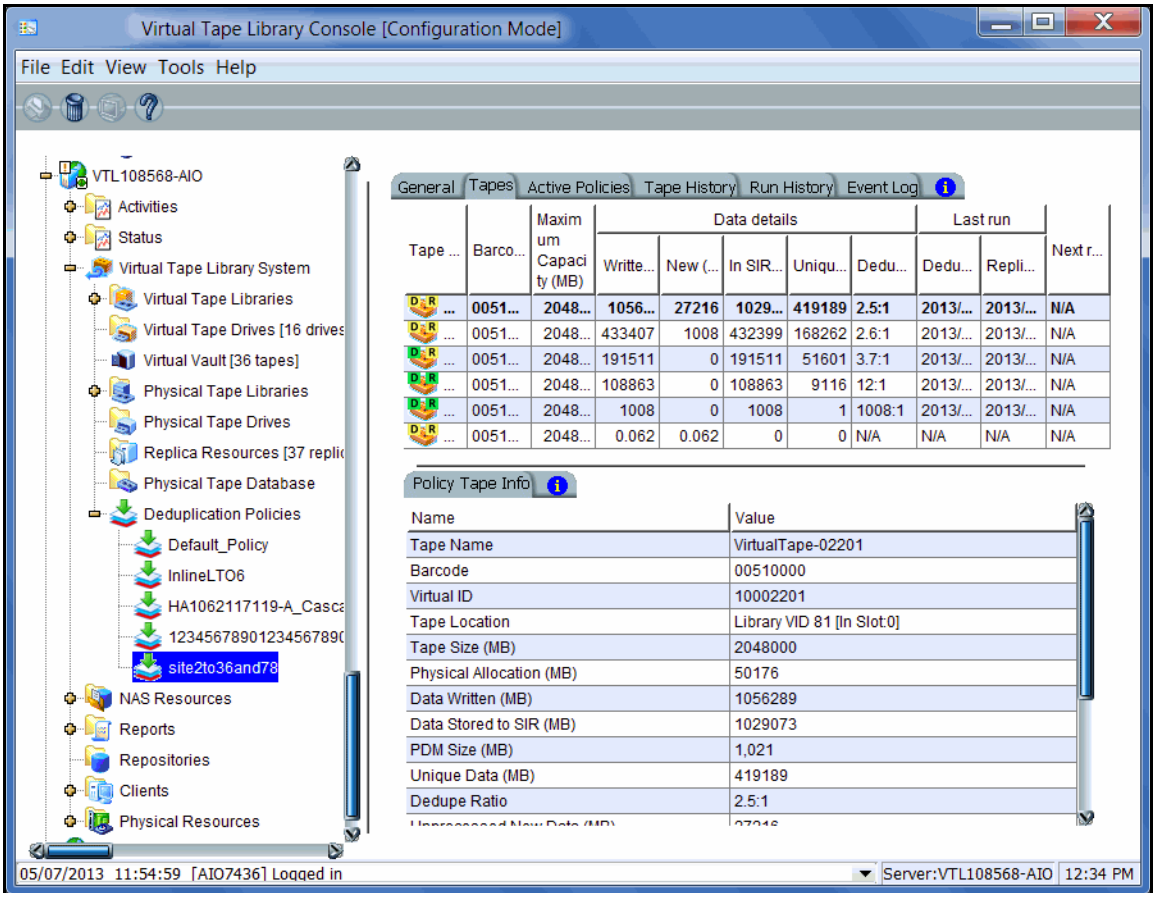

Tape ...	Barco...	Maximum Capacity (MB)	Data details					Last run		Next r...
			Writte...	New (...	In SIR...	Uniqu...	Dedu...	Dedu...	Repli...	
D R ...	**0051...**	**2048...**	**1056...**	**27216**	**1029...**	**419189**	**2.5:1**	**2013/...**	**2013/...**	**N/A**
D R ...	0051...	2048...	433407	1008	432399	168262	2.6:1	2013/...	2013/...	N/A
D R ...	0051...	2048...	191511	0	191511	51601	3.7:1	2013/...	2013/...	N/A
D R ...	0051...	2048...	108863	0	108863	9116	12:1	2013/...	2013/...	N/A
D R ...	0051...	2048...	1008	0	1008	1	1008:1	2013/...	2013/...	N/A
D R ...	0051...	2048...	0.062	0.062	0	0	N/A	N/A	N/A	N/A

Name	Value
Tape Name	VirtualTape-02201
Barcode	00510000
Virtual ID	10002201
Tape Location	Library VID 81 [In Slot:0]
Tape Size (MB)	2048000
Physical Allocation (MB)	50176
Data Written (MB)	1056289
Data Stored to SIR (MB)	1029073
PDM Size (MB)	1,021
Unique Data (MB)	419189
Dedupe Ratio	2.5:1
Unprocessed New Data (MB)	27216

Active Policies tab

The *Active Policies* tab lists information about currently running policies and replication jobs. The data is automatically refreshed. All values are rounded.

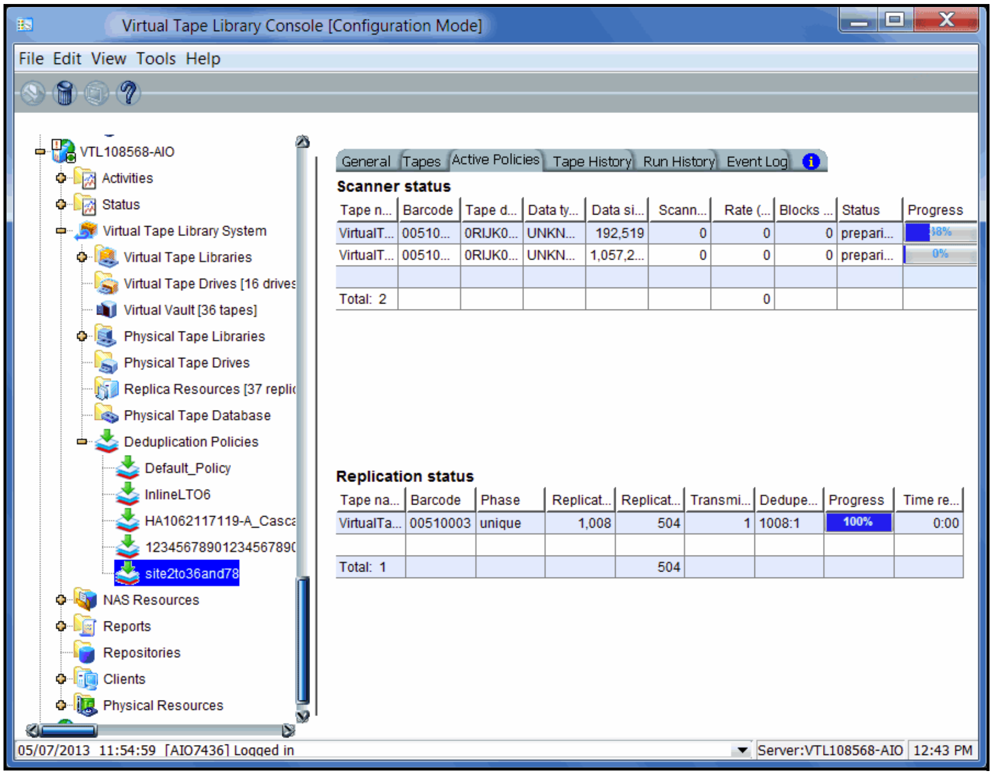

Scanner status

Tape n...	Barcode	Tape d...	Data ty...	Data si...	Scann...	Rate (...	Blocks ...	Status	Progress
VirtualT...	00510...	0RIJK0...	UNKN...	192,519	0	0	0	prepari...	38%
VirtualT...	00510...	0RIJK0...	UNKN...	1,057,2...	0	0	0	prepari...	0%
Total: 2						0			

Replication status

Tape na...	Barcode	Phase	Replicat...	Replicat...	Transmi...	Dedupe...	Progress	Time re...
VirtualTa...	00510003	unique	1,008	504	1	1008:1	100%	0:00
Total: 1				504				

05/07/2013 11:54:59 [AIO7436] Logged in Server:VTL108568-AIO 12:43 PM

Tape History tab

The *Tape History* tab lists all of the deduplication and replication jobs that have run and provides statistics for each. All values are rounded.

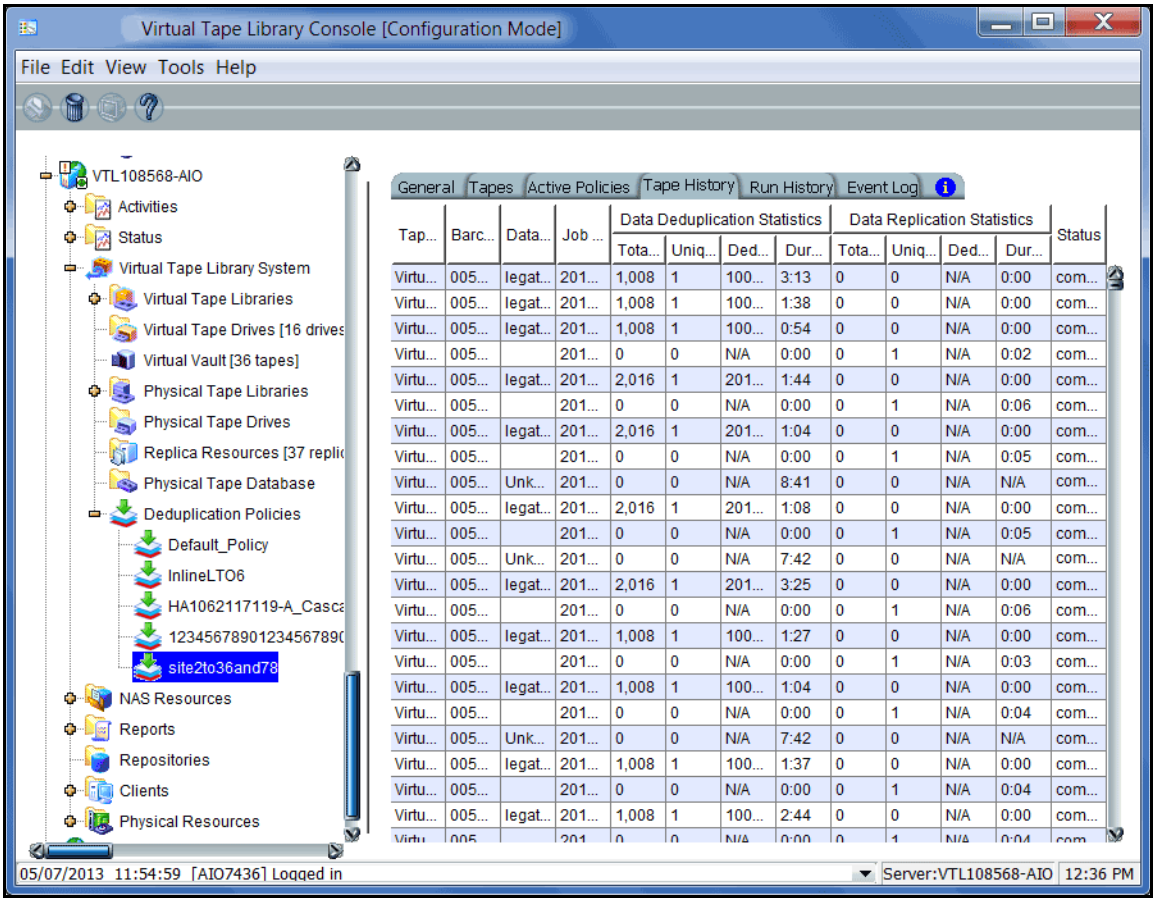

Tap...	Barc...	Data...	Job ...	Data Deduplication Statistics				Data Replication Statistics				Status
				Tota...	Uniq...	Ded...	Dur...	Tota...	Uniq...	Ded...	Dur...	
Virtu...	005...	legat...	201...	1,008	1	100...	3:13	0	0	N/A	0:00	com...
Virtu...	005...	legat...	201...	1,008	1	100...	1:38	0	0	N/A	0:00	com...
Virtu...	005...	legat...	201...	1,008	1	100...	0:54	0	0	N/A	0:00	com...
Virtu...	005...		201...	0	0	N/A	0:00	0	1	N/A	0:02	com...
Virtu...	005...	legat...	201...	2,016	1	201...	1:44	0	0	N/A	0:00	com...
Virtu...	005...		201...	0	0	N/A	0:00	0	1	N/A	0:06	com...
Virtu...	005...	legat...	201...	2,016	1	201...	1:04	0	0	N/A	0:00	com...
Virtu...	005...		201...	0	0	N/A	0:00	0	1	N/A	0:05	com...
Virtu...	005...	Unk...	201...	0	0	N/A	8:41	0	0	N/A	N/A	com...
Virtu...	005...	legat...	201...	2,016	1	201...	1:08	0	0	N/A	0:00	com...
Virtu...	005...		201...	0	0	N/A	0:00	0	1	N/A	0:05	com...
Virtu...	005...	Unk...	201...	0	0	N/A	7:42	0	0	N/A	N/A	com...
Virtu...	005...	legat...	201...	2,016	1	201...	3:25	0	0	N/A	0:00	com...
Virtu...	005...		201...	0	0	N/A	0:00	0	1	N/A	0:06	com...
Virtu...	005...	legat...	201...	1,008	1	100...	1:27	0	0	N/A	0:00	com...
Virtu...	005...		201...	0	0	N/A	0:00	0	1	N/A	0:03	com...
Virtu...	005...	legat...	201...	1,008	1	100...	1:04	0	0	N/A	0:00	com...
Virtu...	005...		201...	0	0	N/A	0:00	0	1	N/A	0:04	com...
Virtu...	005...	Unk...	201...	0	0	N/A	7:42	0	0	N/A	N/A	com...
Virtu...	005...	legat...	201...	1,008	1	100...	1:37	0	0	N/A	0:00	com...
Virtu...	005...		201...	0	0	N/A	0:00	0	1	N/A	0:04	com...
Virtu...	005...	legat...	201...	1,008	1	100...	2:44	0	0	N/A	0:00	com...
Virtu	005		201	0	0	N/A	0:00	0	1	N/A	0:04	com

05/07/2013 11:54:59 [AIO7436] Logged in

Server:VTL108568-AIO 12:36 PM

Run History tab

The *Run History* tab displays policy history, including when and why the policy was run, number of tapes, total amount of data scanned, total amount of unique data written to the repository, the deduplication ratio and the total run time for the policy.

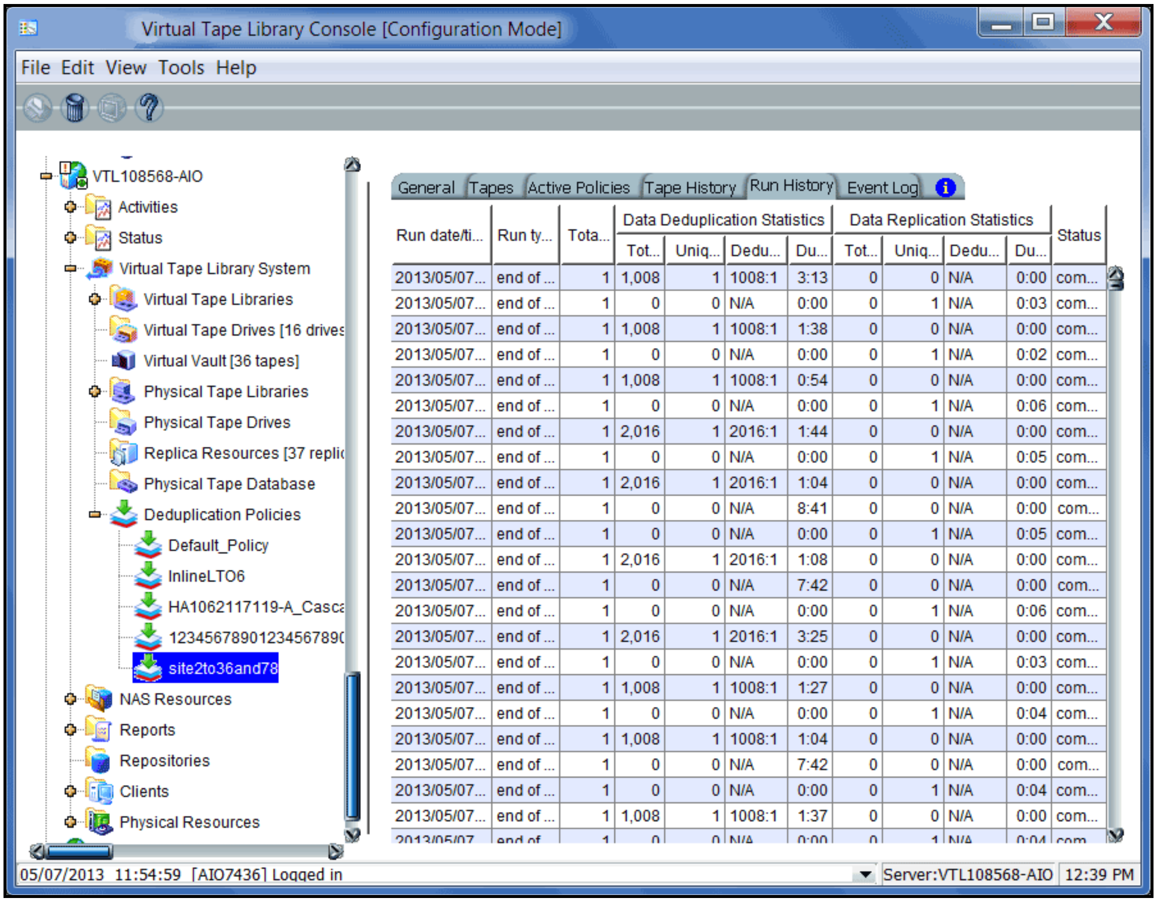

Run date/ti...	Run ty...	Tota...	Data Deduplication Statistics				Data Replication Statistics				Status
			Tot...	Uniq...	Dedu...	Du...	Tot...	Uniq...	Dedu...	Du...	
2013/05/07...	end of ...	1	1,008	1	1008:1	3:13	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	0:00	0	1	N/A	0:03	com...
2013/05/07...	end of ...	1	1,008	1	1008:1	1:38	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	0:00	0	1	N/A	0:02	com...
2013/05/07...	end of ...	1	1,008	1	1008:1	0:54	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	0:00	0	1	N/A	0:06	com...
2013/05/07...	end of ...	1	2,016	1	2016:1	1:44	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	0:00	0	1	N/A	0:05	com...
2013/05/07...	end of ...	1	2,016	1	2016:1	1:04	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	8:41	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	0:00	0	1	N/A	0:05	com...
2013/05/07...	end of ...	1	2,016	1	2016:1	1:08	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	7:42	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	0:00	0	1	N/A	0:06	com...
2013/05/07...	end of ...	1	2,016	1	2016:1	3:25	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	0:00	0	1	N/A	0:03	com...
2013/05/07...	end of ...	1	1,008	1	1008:1	1:27	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	0:00	0	1	N/A	0:04	com...
2013/05/07...	end of ...	1	1,008	1	1008:1	1:04	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	7:42	0	0	N/A	0:00	com...
2013/05/07...	end of ...	1	0	0	N/A	0:00	0	1	N/A	0:04	com...
2013/05/07...	end of ...	1	1,008	1	1008:1	1:37	0	0	N/A	0:00	com...
2013/05/07	end of	1	0	0	N/A	0:00	0	1	N/A	0:04	com

05/07/2013 11:54:59 [AIO7436] Logged in

Event Log tab

The *Event Log* tab displays informational events and errors pertaining to this policy.

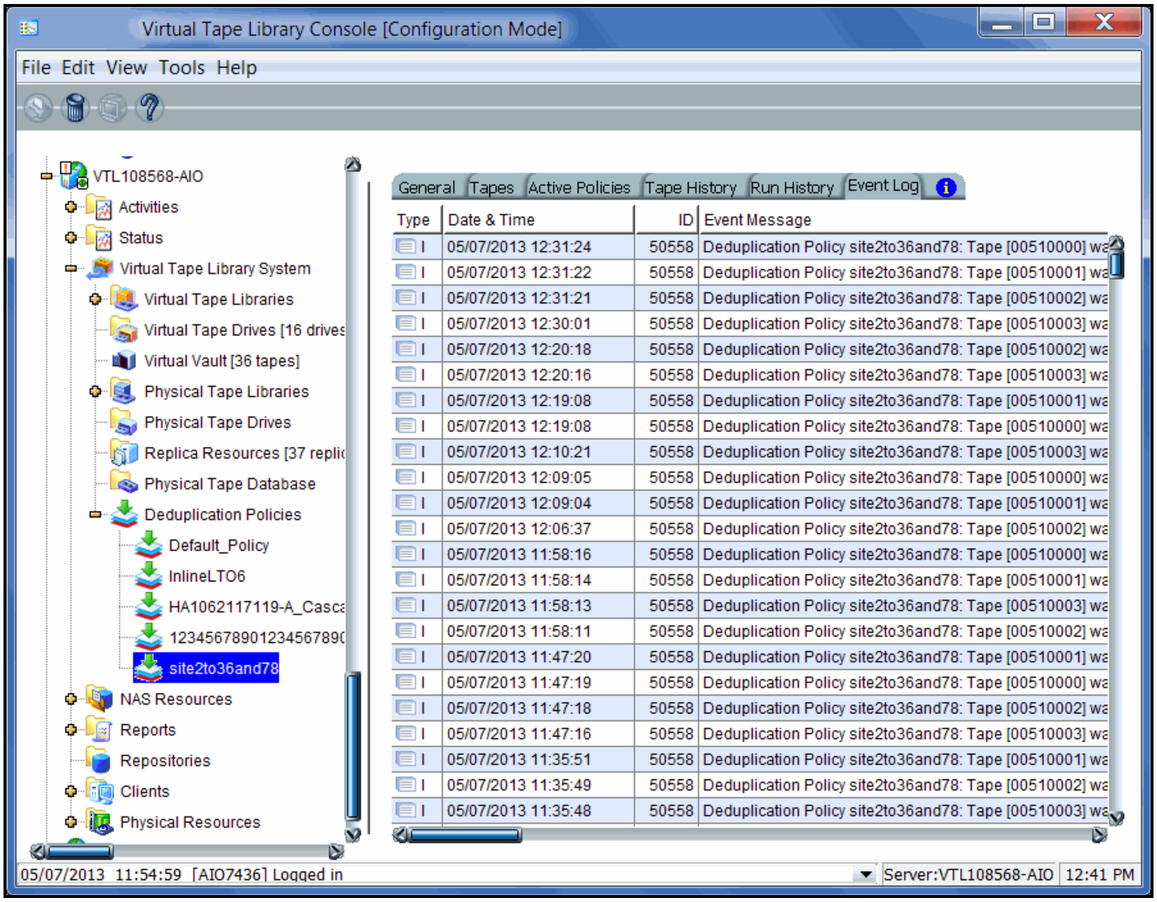

Type	Date & Time	ID	Event Message
I	05/07/2013 12:31:24	50558	Deduplication Policy site2to36and78: Tape [00510000] wa
I	05/07/2013 12:31:22	50558	Deduplication Policy site2to36and78: Tape [00510001] wa
I	05/07/2013 12:31:21	50558	Deduplication Policy site2to36and78: Tape [00510002] wa
I	05/07/2013 12:30:01	50558	Deduplication Policy site2to36and78: Tape [00510003] wa
I	05/07/2013 12:20:18	50558	Deduplication Policy site2to36and78: Tape [00510002] wa
I	05/07/2013 12:20:16	50558	Deduplication Policy site2to36and78: Tape [00510003] wa
I	05/07/2013 12:19:08	50558	Deduplication Policy site2to36and78: Tape [00510001] wa
I	05/07/2013 12:19:08	50558	Deduplication Policy site2to36and78: Tape [00510000] wa
I	05/07/2013 12:10:21	50558	Deduplication Policy site2to36and78: Tape [00510003] wa
I	05/07/2013 12:09:05	50558	Deduplication Policy site2to36and78: Tape [00510000] wa
I	05/07/2013 12:09:04	50558	Deduplication Policy site2to36and78: Tape [00510001] wa
I	05/07/2013 12:06:37	50558	Deduplication Policy site2to36and78: Tape [00510002] wa
I	05/07/2013 11:58:16	50558	Deduplication Policy site2to36and78: Tape [00510000] wa
I	05/07/2013 11:58:14	50558	Deduplication Policy site2to36and78: Tape [00510001] wa
I	05/07/2013 11:58:13	50558	Deduplication Policy site2to36and78: Tape [00510003] wa
I	05/07/2013 11:58:11	50558	Deduplication Policy site2to36and78: Tape [00510002] wa
I	05/07/2013 11:47:20	50558	Deduplication Policy site2to36and78: Tape [00510001] wa
I	05/07/2013 11:47:19	50558	Deduplication Policy site2to36and78: Tape [00510000] wa
I	05/07/2013 11:47:18	50558	Deduplication Policy site2to36and78: Tape [00510002] wa
I	05/07/2013 11:47:16	50558	Deduplication Policy site2to36and78: Tape [00510003] wa
I	05/07/2013 11:35:51	50558	Deduplication Policy site2to36and78: Tape [00510001] wa
I	05/07/2013 11:35:49	50558	Deduplication Policy site2to36and78: Tape [00510002] wa
I	05/07/2013 11:35:48	50558	Deduplication Policy site2to36and78: Tape [00510003] wa

05/07/2013 11:54:59 [AIO7436] Logged in

Server:VTL108568-AIO 12:41 PM

Deduplication Job Queue

Select *Activities --> Deduplication Job Queue* to see the current active deduplication job as well as all jobs in the queue. When you select an individual job, its details are shown at the bottom of the screen, as well as progress bar that updates automatically when a job is running.

If you want to change the priority of tapes in the queue, right-click a tape and select *Run Next* or *Run Later*. This is useful for disaster recovery (DR) purposes when a tape has replication configured and is needed at the DR site. You can also cancel processing for a tape by selecting *Remove* from the right-click menu.

During failover, deduplication and replication jobs running on a failed server will fail. Based on your retry policy, each job should restart after the failed server has been taken over. At this point, you can continue to view job activity in the Deduplication Job Queue.

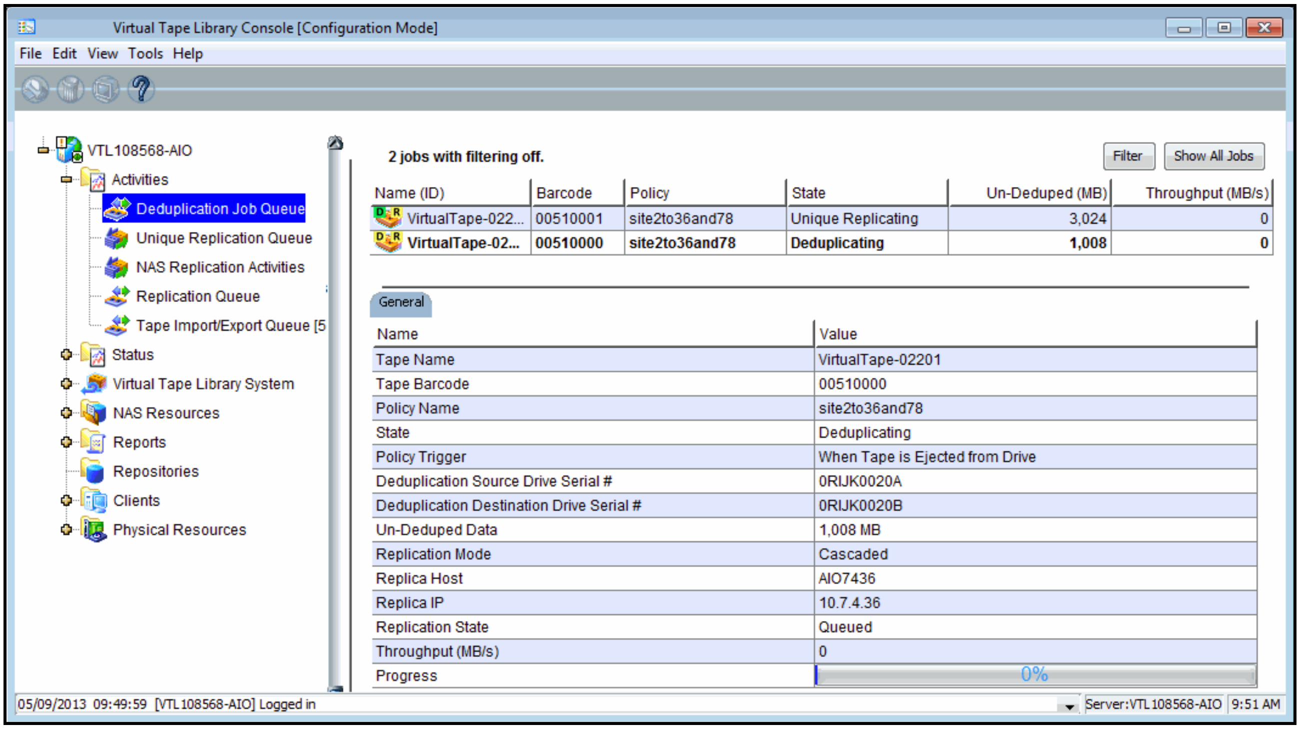

2 jobs with filtering off.

Name (ID)	Barcode	Policy	State	Un-Deduped (MB)	Throughput (MB/s)
VirtualTape-022...	00510001	site2to36and78	Unique Replicating	3,024	0
VirtualTape-02...	**00510000**	**site2to36and78**	**Deduplicating**	**1,008**	**0**

Name	Value
Tape Name	VirtualTape-02201
Tape Barcode	00510000
Policy Name	site2to36and78
State	Deduplicating
Policy Trigger	When Tape is Ejected from Drive
Deduplication Source Drive Serial #	0RIJK0020A
Deduplication Destination Drive Serial #	0RIJK0020B
Un-Deduped Data	1,008 MB
Replication Mode	Cascaded
Replica Host	AIO7436
Replica IP	10.7.4.36
Replication State	Queued
Throughput (MB/s)	0
Progress	0%

05/09/2013 09:49:59 [VTL108568-AIO] Logged in

Server:VTL108568-AIO 9:51 AM

Virtual index tape status

VITs replace virtual tapes after they have been deduplicated. You can review the status of these virtual tapes from backup server objects in the console:

1. Expand the *Virtual Tape Libraries* object.

2. Expand the library and select the *Tapes* object.

3. Select a tape in the right pane and locate the *Allocation Type* field in the *General* tab in the lower portion of the display.

   This field includes the tape status.

   If the field includes *Pure VIT*, all data on this VIT has been deduplicated. The tape contains only pointers to its data in the deduplication repository.

   If the field includes *Mixed*, this tape contains pointers to its data in the deduplication repository as well as data that has not yet been deduplicated. This status can occur due to deduplication policy settings or to power factors that may affect the number of deduplication jobs that can run simultaneously.

   If the field displays *Virtual Tape*, the data on the tape has not been deduplicated.

4. Select the *Layout* tab to display information about the physical devices that were used to create this VIT.

5. If replication has been configured for the tape, the *Replication* tab includes information about the replica target and policies.

NAS information available from the backup server

The management console offers the following NAS deduplication-related information for your backup server:

- Information about each NAS resource
- Deduplication and replication statistics

Capacity information

When you highlight the *NAS Resources* object, the *General* tab displays information about capacity and usage for all NAS resources. The *NAS Resources* tab breaks this information down for each NAS resource.

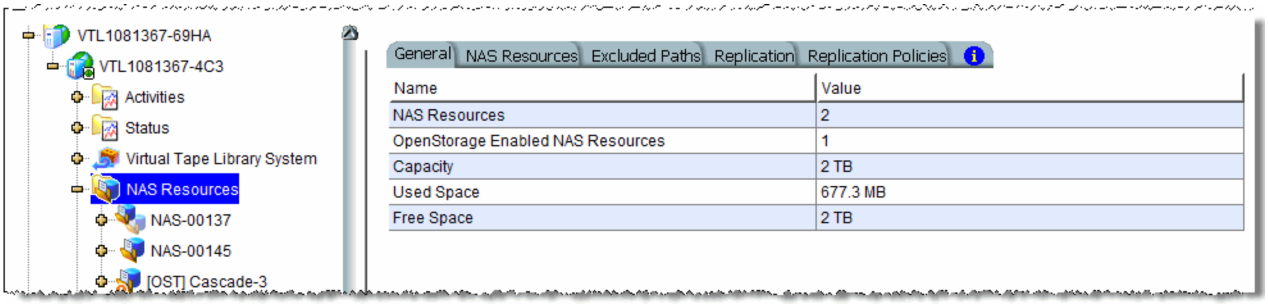

Name	Value
NAS Resources	2
OpenStorage Enabled NAS Resources	1
Capacity	2 TB
Used Space	677.3 MB
Free Space	2 TB

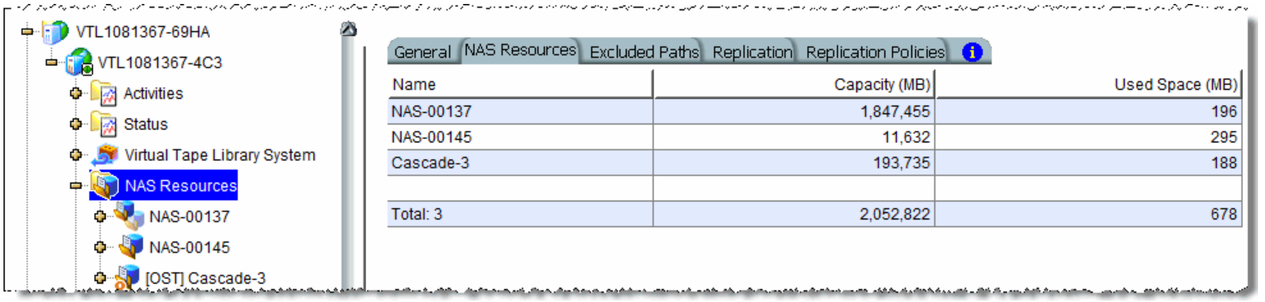

Name	Capacity (MB)	Used Space (MB)
NAS-00137	1,847,455	196
NAS-00145	11,632	295
Cascade-3	193,735	188
Total: 3	2,052,822	678

NAS deduplication/replication statistics

NAS resource device statistics

Highlight a NAS resource and select the *Device Statistics* tab to see deduplication and replication statistics for the resource.

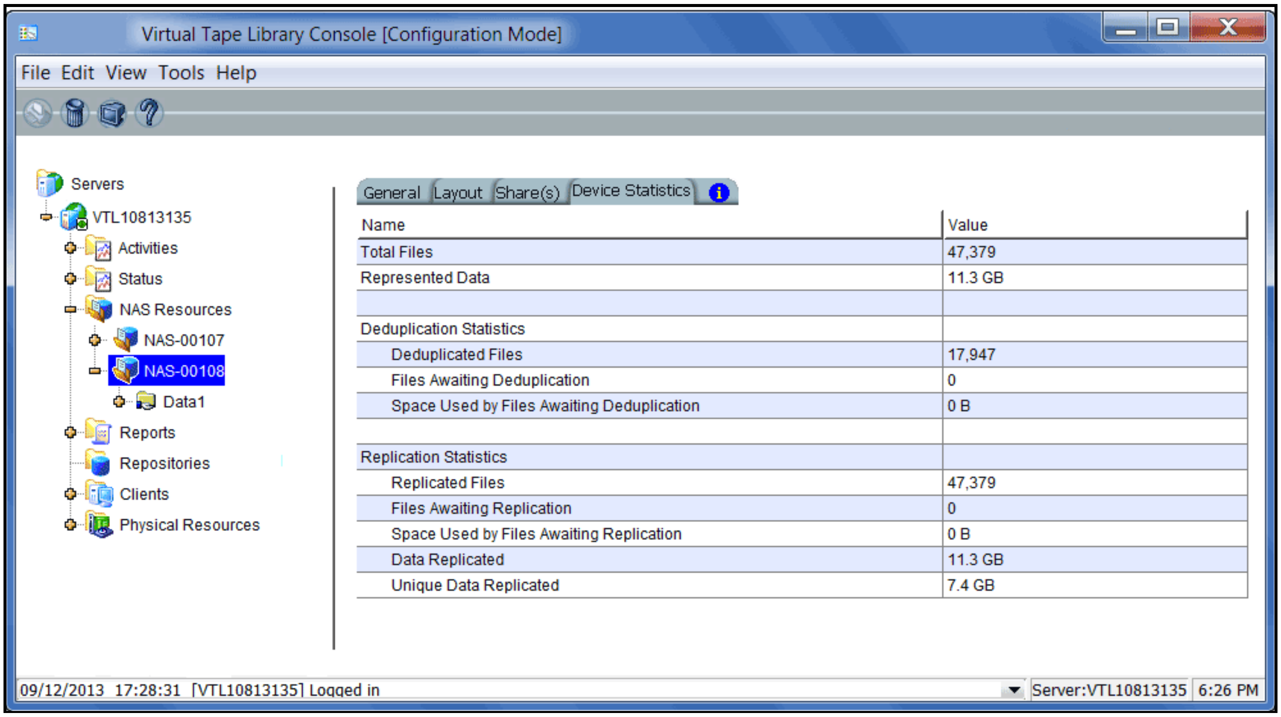

Name	Value
Total Files	47,379
Represented Data	11.3 GB
Deduplication Statistics	
Deduplicated Files	17,947
Files Awaiting Deduplication	0
Space Used by Files Awaiting Deduplication	0 B
Replication Statistics	
Replicated Files	47,379
Files Awaiting Replication	0
Space Used by Files Awaiting Replication	0 B
Data Replicated	11.3 GB
Unique Data Replicated	7.4 GB

09/12/2013 17:28:31 [VTL10813135] Logged in

Server:VTL10813135 6:26 PM

Total Files - Total number of files. This includes files that have been deduplicated, files that were excluded, and files awaiting deduplication.

Represented Data - Total size of files copied to all file systems (including those that have not yet been deduplicated).

Deduplicated Files - Total number of files that have been deduplicated. This does not include excluded files (smaller than 8 KB or encrypted/compressed files with known file extensions).

Files Awaiting Deduplication - Total number of files that have not yet been deduplicated.

Space Used by Files Awaiting Deduplication - Total size of files that have not yet been deduplicated.

Replicated Files - Total number of files replicated. Depending upon how replication was configured, this may not include excluded files.

Files Awaiting Replication - Total number of files that have not yet been replicated.

Space Used by Files Awaiting Replication - Total size of files that have not yet been replicated.

Data Replicated - Total amount of data replicated.

Unique Data Replicated - Total amount of unique data replicated.

Share and folder-level statistics

General tab

When you highlight a folder or share, the *General* tab displays information about that share/folder. Here, you will find the path for the share, which can be useful for mapping/mounting shares. Additional information is displayed if the folder has been shared, including share settings.

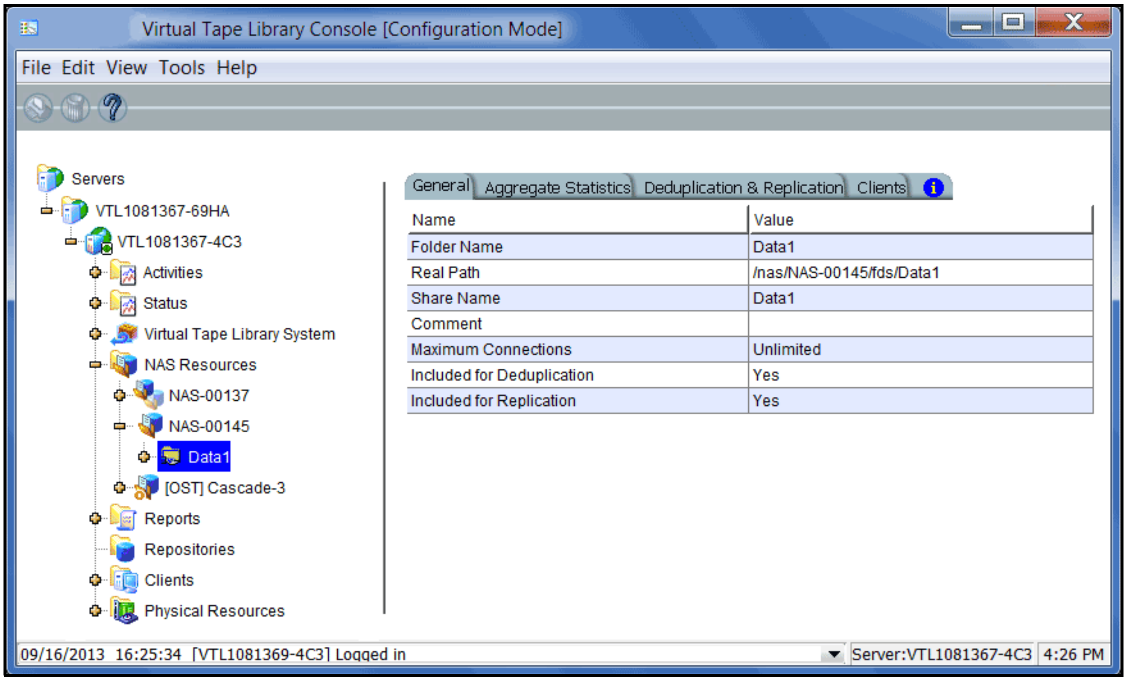

Name	Value
Folder Name	Data1
Real Path	/nas/NAS-00145/fds/Data1
Share Name	Data1
Comment	
Maximum Connections	Unlimited
Included for Deduplication	Yes
Included for Replication	Yes

Aggregate Statistics tab

When you highlight a folder, share, or directory, the *Aggregate Statistics* tab displays deduplication and replication statistics for the folder/share/directory and all subdirectories below that share/folder. The information here is similar to the information that is presented on the *Device Statistics* tab for the NAS resource. Refer to 'NAS resource device statistics' for a description of the fields listed on the *Aggregate Statistics* tab.

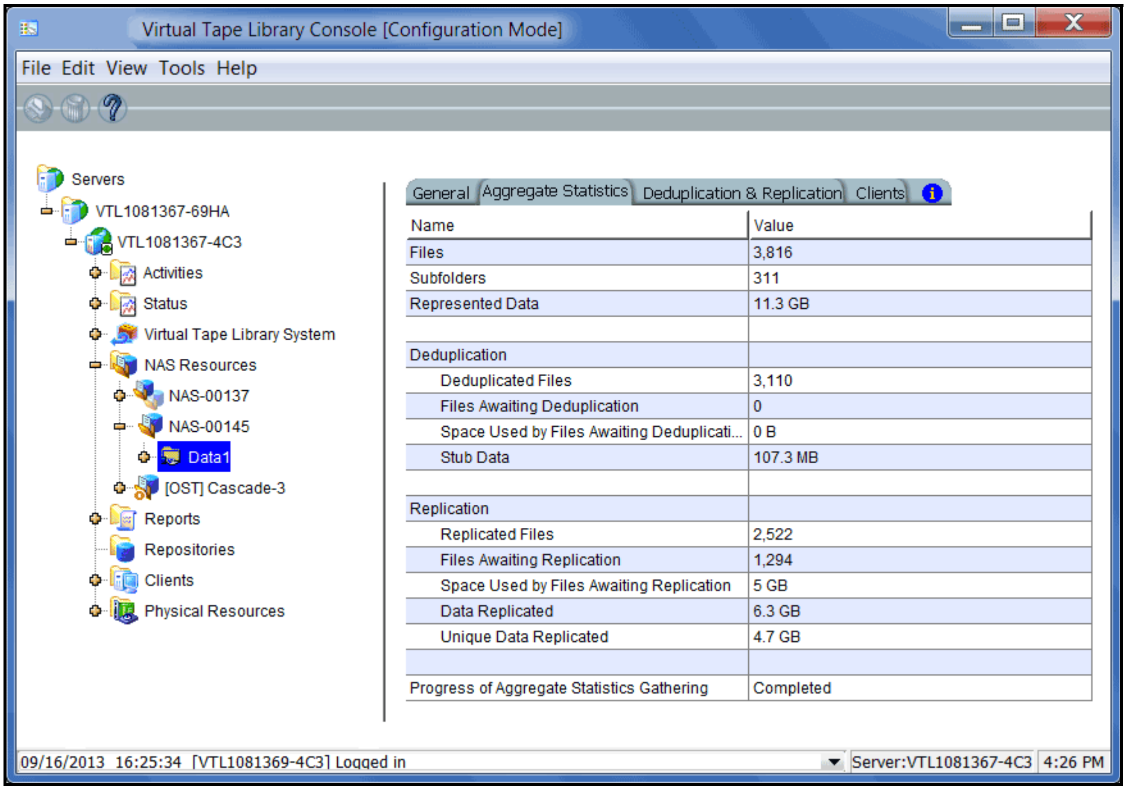

Name	Value
Files	3,816
Subfolders	311
Represented Data	11.3 GB
Deduplication	
Deduplicated Files	3,110
Files Awaiting Deduplication	0
Space Used by Files Awaiting Deduplicati...	0 B
Stub Data	107.3 MB
Replication	
Replicated Files	2,522
Files Awaiting Replication	1,294
Space Used by Files Awaiting Replication	5 GB
Data Replicated	6.3 GB
Unique Data Replicated	4.7 GB
Progress of Aggregate Statistics Gathering	Completed

09/16/2013 16:25:34 [VTL1081369-4C3] Logged in

Server:VTL1081367-4C3 4:26 PM

Deduplication & Replication tab

When you highlight a folder or share, the *Deduplication & Replication* tab displays a list of all files and folders in that folder.

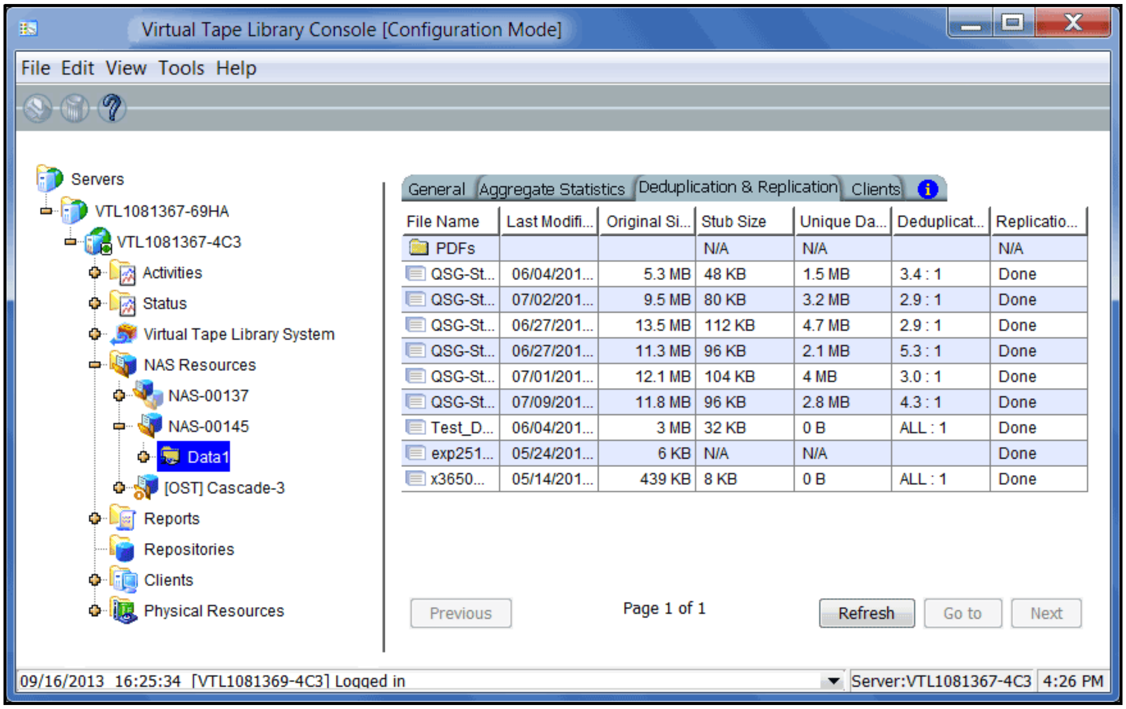

File Name	Last Modifi...	Original Si...	Stub Size	Unique Da...	Deduplicat...	Replicatio...
PDFs			N/A	N/A		N/A
QSG-St...	06/04/201...	5.3 MB	48 KB	1.5 MB	3.4 : 1	Done
QSG-St...	07/02/201...	9.5 MB	80 KB	3.2 MB	2.9 : 1	Done
QSG-St...	06/27/201...	13.5 MB	112 KB	4.7 MB	2.9 : 1	Done
QSG-St...	06/27/201...	11.3 MB	96 KB	2.1 MB	5.3 : 1	Done
QSG-St...	07/01/201...	12.1 MB	104 KB	4 MB	3.0 : 1	Done
QSG-St...	07/09/201...	11.8 MB	96 KB	2.8 MB	4.3 : 1	Done
Test_D...	06/04/201...	3 MB	32 KB	0 B	ALL : 1	Done
exp251...	05/24/201...	6 KB	N/A	N/A		Done
x3650...	05/14/201...	439 KB	8 KB	0 B	ALL : 1	Done

09/16/2013 16:25:34 [VTL1081369-4C3] Logged in | Server:VTL1081367-4C3 | 4:26 PM

File Name - The name of the file or folder.

Last Modified Time - The date and time the file was last modified. This field is only valid for files.

Original Size - The size of the file before deduplication.

Stub size - A stub file is a small file that replaces the deduplicated file and points to the stored data in the repository. The minimum stub file size is 8 KB.

Unique data - The amount of unique data that has been moved to the repository.

Deduplication ratio - The ratio between the original file size and the amount of unique data moved to the repository. If you see a ratio of *ALL:1*, the file’s data was already in the repository and no new unique data was stored for this file.

Replication Status - Indicates whether the file has been replicated yet.

Reclaim disk space

During the deduplication process, only single instances of unique data are passed to the deduplication repository. For tapes, the original virtual tape is replaced with a VIT pointing to deduplication storage. For NAS, the original file is replaced with a stub file pointing to deduplication storage.

Over time, VITs and files can be erased, formatted, updated, or overwritten by your backup application (such as when a tape has expired). It is also possible that you have manually deleted stub files or a VIT from the console.

When a VIT or stub file is eliminated, the pointers to deduplication storage are deleted but the actual deduplicated data is not. Reclamation eliminates the unneeded data in the repository, thereby reclaiming storage space on the index and data disks.

There are two types of reclamation:

- *Space Reclamation* - Reclaims index cache capacity, data disk space, and folder disk space
- *Index Pruning* - Reclaims space on the index disk(s)

Note: Putting a VIT in a scratch pool of a backup application does not mean that the storage used by that VIT can be reclaimed. Storage can be reclaimed only when a VIT is deleted from the console or erased/formatted/overwritten by the backup application.

Space reclamation

Space Reclamation is composed of two processes, *Index Reclamation* (reclaims index cache and folder disk space) and *Storage Reclamation* (reclaims data disk space).

Index Reclamation

Index Reclamation reads through all VITs and NAS stub files to determine which hashes need to be kept. Hashes that are no longer needed are removed from memory and the index disk is marked to indicate where deletion can occur. Folder space that is no longer needed is also freed up.

Storage Reclamation

After *Index Reclamation* completes, *Storage Reclamation* looks at the areas of the data disk(s) that were referenced by the hashes that have been removed and marks this space as re-usable.

Index pruning

Index Pruning uses the marks made to the index disk(s) during *Index Reclamation* and removes the hash records that are no longer needed as well as the deletion notes themselves.

Index pruning is CPU intensive and does not need to be run frequently.

Reclamation requirements

Reclamation requires:

- At least one VIT on any of the backup servers associated with the deduplication cluster or at least one NAS resource
- The backup server to have access to deduplication data drives

Reclamation thresholds

The system will run a reclamation process automatically whenever a threshold is met. There are four independent thresholds and the reclamation processes run independently for each.

After successful reclamation, the system recalculates the threshold values using the following formula:

(Initial Trigger) (Free Space [After Reclamation]) + (Used before reclamation) = New Threshold

For example, if there was 80% free after reclamation, the new threshold would be 60%. You can see how this is calculated:

(50% threshold) (80% free space) + (20%) = 60% threshold

To see the thresholds on a VTL-S system, expand the *Status* object, select the *Dashboard Summary* object, and select the *Deduplication Repository* tab.

To see the thresholds for a cluster, highlight your deduplication cluster in the VTL console and select the *Dashboard Summary* tab.

The thresholds are displayed below the dashboard charts. You cannot modify these values:

- *Index cache capacity* - Index reclamation
- *Folder capacity* - Folder space reclamation
- *Index capacity* - Index pruning
- *Deduplication data capacity* - Data space reclamation

Run reclamation

Reclamation can be run in the following ways:

- Automatic – Reclamation runs based on system usage, on a scheduled basis, or both.
- Manual – Reclamation is initiated from the console.
- Command Line – Reclamation is run from the command line.

Automatic reclamation

Reclamation can be run automatically based on system usage, on a scheduled basis, or both. To automatically run reclamation:

1. Right-click your deduplication cluster or VTL-S server and select *Deduplication --> Reclamation --> Reclamation Properties*.

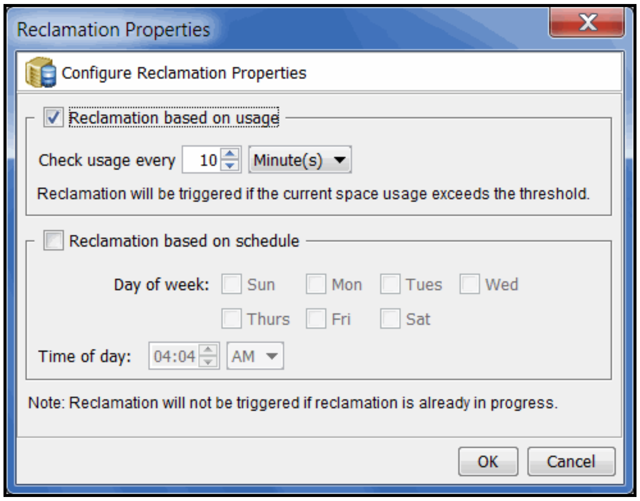

Reclamation Properties

Configure Reclamation Properties

☑ Reclamation based on usage

Check usage every 10 Minute(s)

Reclamation will be triggered if the current space usage exceeds the threshold.

□ Reclamation based on schedule

Day of week: □ Sun □ Mon □ Tues □ Wed □ Thurs □ Fri □ Sat

Time of day: 04:04 AM

Note: Reclamation will not be triggered if reclamation is already in progress.

2. Select when reclamation should run.
   - *Reclamation based on usage* - Specify how often the system should check usage. Reclamation will run automatically whenever a threshold is met.
   - *Reclamation based on schedule* - Reclamation can be scheduled to run at a set time on selected days. Specify the days and the time. Index pruning will be triggered immediately after reclamation.

   Note that if you do not check *Reclamation based on usage* or *Reclamation based on schedule*, reclamation will not run automatically; you will need to run it manually.

Manual reclamation

To manually run reclamation from the console, right-click your deduplication cluster or VTL-S server and select *Deduplication --> Reclamation --> Start Space Reclamation* (or *Start Index Pruning*). A message is displayed when the process begins and again to confirm that the process is complete. Click *OK* to close the confirmation box.

Command Line

Use the `startsirreclamation` command to initiate reclamation. Refer to 'Start reclamation' in the Command Line chapter for more information.

Switch to a flexible storage configuration

If you configured your deduplication cluster via the *Classic* method (which requires organizing your disks into columns/rows), you can change your storage configuration to use the *Flexible* method. This method allows you to add one or more LUNs, rather than needing to add an equal amount per column.

Notes:

- Once you switch to the *Flexible* method, it is not reversible.
- You will need to restart all nodes in the cluster in order for the change to take effect.

To switch to the *Flexible* method:

1. Right-click your cluster and select *Deduplication* --> *Switch to Flexible Storage Configuration*.
2. Restart all servers in the cluster.

Add data disks

If your deduplication appliance has additional available disk space, you can create additional logical resources for storage of deduplication data, index, or folders.

If you have physical disks available, prepare them, connect them to your server, and re-scan before following the steps below.

1. Right-click the server object and select *Deduplication --> Add Deduplication Data disks*.

   For index and folder storage, select *Add Index and Folder Disks*.

2. If your deduplication data disks were configured using the *Classic* method, select whether you want to use physical or virtual devices for the additional space.

   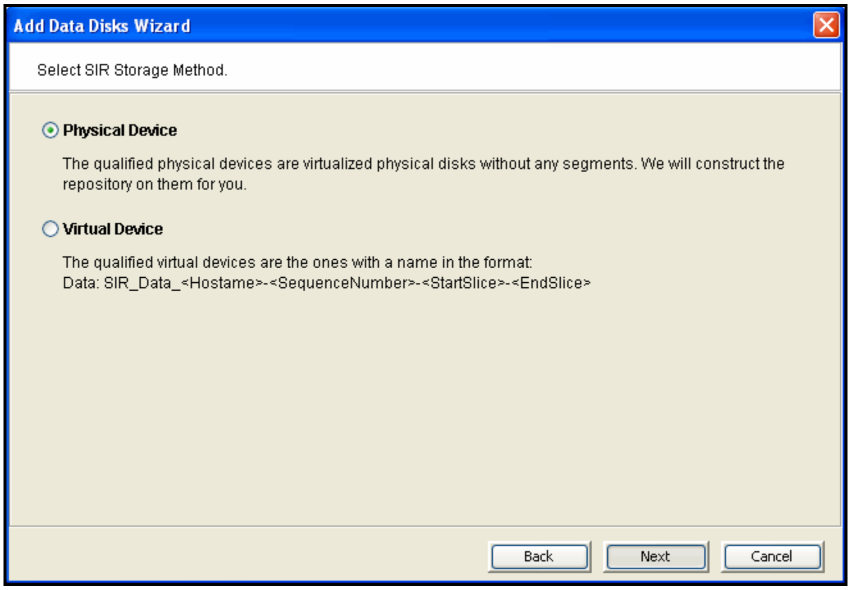

   Add Data Disks Wizard

   Select SIR Storage Method.

   ● **Physical Device**

   The qualified physical devices are virtualized physical disks without any segments. We will construct the repository on them for you.

   ○ **Virtual Device**

   The qualified virtual devices are the ones with a name in the format:
   Data: SIR_Data_<Hostame>-<SequenceNumber>-<StartSlice>-<EndSlice>

   You can choose either, regardless of whether your classic deduplication configuration was enabled using physical or pre-configured virtual resources.

3. Select an available device.

4. Confirm information about your selection and click *Finish*.

   The newly added disks will not be usable until all nodes in the deduplication cluster have been restarted.

   > **Note:** If your deduplication server is an active node in a redundant node failover configuration, you must suspend failover before restarting the deduplication server to prevent the standby node from taking over. To do this, right-click the cluster and select *Deduplication --> Redundant Node --> Suspend*.

5. On each backup server associated with this cluster, right-click the server object and select *Deduplication --> Deduplication Cluster - Check for new deduplication data disks*.

   You do not need to do anything for index or folder disks.

Encryption

Encryption can be used to ensure that tape data is confidential and secure. Encryption can protect:

- Deduplicated data stored in the deduplication repository
- Data backed up on virtual tapes
- Data exported to physical tapes

Encryption utilizes the Advanced Encryption Standard (AES) 256-bit key CBC algorithms (Secure Tape) published by the National Institute of Standards and Technology (an agency of the U.S. government) and is FIPS-140-2-compliant.

Deduplication repository encryption

With deduplication repository encryption, deduplicated data from all cluster servers is encrypted in the deduplication repository.

Encryption requires that an activation password be created. In order for data on the repository to be accessible, the cluster servers must have encryption activated with the specified password each time the VTL services are started. All servers in the cluster use the same encryption activation password. Data in the repository will not be accessible unless the activation password has been entered.

Notes:

- Deduplication repository encryption uses a single internal encryption key for the cluster; you do not have to create separate encryption keys.
- NAS cannot be enabled for servers associated with a deduplication cluster repository encryption.
- In order to deactivate encryption for a cluster, you must restart services. However, if encryption is deactivated, backup and restore jobs will fail; for security reasons, encryption must be activated in order for data on the repository to be accessible.
- Once the deduplication repository is created, encryption cannot be enabled or disabled and data cannot be rolled back to an unencrypted state. In order to disable deduplication repository encryption, you will have to disable your deduplication cluster.
- When encryption is enabled, there will be an overall performance decrease for read and write operations. The actual impact will depend upon a number of factors, including the number of CPU cores and speed, number of concurrent IO operations, data compression ratio, and data deduplication ratio. Faster server processors with more cores can be used to minimize the impact.
- For VTL-S servers, if virtual tape encryption is enabled when you enable deduplication, deduplication repository encryption will be enabled by default. You will need to enter the secret phrase for the encryption key.

Enable encryption for a deduplication server

Deduplication repository encryption can be enabled as follows:

1. If you are enabling encryption for a new VTL 8.20 or higher installation, run `unlockdataencryptionoption` from the command line of your deduplication server.

   Refer to 'Unlock data encryption' for more information.

2. Create a deduplication cluster and enable deduplication repository encryption.

   Refer to 'Create a deduplication cluster' for more information.

Activate encryption for a cluster

To activate encryption, right-click the cluster and select *Encryption Management* --> *Activate Encryption*.

Change the encryption activation password for a cluster

To change the encryption activation password, right-click the cluster and select *Encryption Management* --> *Change Activation Password*. You will need to provide the current encryption activation password.

Virtual tape encryption

With virtual tape encryption, data backed up on virtual tapes is protected.

Encryption requires that an activation password be created. In order for data on encrypted virtual tapes to be accessible, the server must have encryption activated with the specified password each time the VTL services are started. Encryption can be activated from the DSI Management Console or via the command line interface. Encrypted virtual tapes will not be accessible for backup or restore unless the activation password in entered.

Virtual tape encryption is enabled at the server level and can then be enabled for your virtual tape libraries.

When encryption is enabled for a virtual tape library, each tape in that library gets the selected encryption key. If the tape is moved to another library, it retains the key, even if that library does not have encryption enabled or uses a different encryption key.

You cannot create an encrypted tape in a standalone virtual tape drive, but if a tape is already encrypted, it retains the encryption key and remains encrypted in the standalone drive.

Notes:

- Once encryption is enabled, it cannot be disabled at the server level and data cannot be rolled back to an unencrypted state.
- When encryption is enabled, there will be an overall performance decrease for read and write operations. The actual impact will depend upon a number of factors, including the number of CPU cores and speed, number of concurrent IO operations, data compression ratio, and data deduplication ratio. Faster server processors with more cores can be used to minimize the impact.

Enable encryption for a tape backup server

Virtual tape encryption can be enabled as follows:

1. If you are enabling encryption for a new VTL 8.20 or higher installation, run `unlockdataencryptionoption` from the command line of your virtual tape server.

   Refer to 'Unlock data encryption' for more information.

2. Right-click your VTL tape backup server and select *Options* --> *Enable Virtual Tape Encryption*.

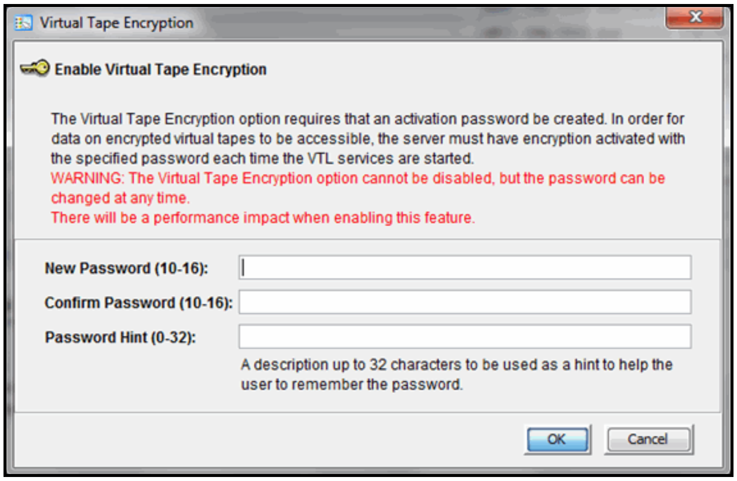

Virtual Tape Encryption

Enable Virtual Tape Encryption

The Virtual Tape Encryption option requires that an activation password be created. In order for data on encrypted virtual tapes to be accessible, the server must have encryption activated with the specified password each time the VTL services are started.
WARNING: The Virtual Tape Encryption option cannot be disabled, but the password can be changed at any time.
There will be a performance impact when enabling this feature.

New Password (10-16):

Confirm Password (10-16):

Password Hint (0-32):

A description up to 32 characters to be used as a hint to help the user to remember the password.

   Virtual tape encryption may have been enabled when you prepared your server via the configuration wizard. If it was enabled, continue with step 4 below.

3. Create the encryption activation password that will need to be entered each time the VTL services are started.

   You must also enter a hint (0–32 characters) to help you remember the password. This hint appears when you type an incorrect password and request a hint.

4. Create one or more encryption keys.

   Refer to 'Create a key' for more information.

5. Enable encryption for virtual tape libraries.

   You can enable encryption when you create a new virtual tape library. To enable it for an existing library, right-click the library and select *Properties*. Existing tapes will not be encrypted; new tapes will be encrypted when they are created.

   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 libraries preserve their encryption status. This means that unencrypted tapes moved to a library with encryption 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 a library, tapes created afterward will not be encrypted. Therefore, each library can have both encrypted and unencrypted tapes. An E icon is displayed on each virtual tape that is encrypted. 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.

Activate encryption for a server

To activate encryption, right-click the server and select *Encryption Management* --> *Activate Encryption*.

Change the encryption activation password for a server

To change the encryption activation password, right-click the server and select *Encryption Management* --> *Change Activation Password*. You will need to provide the current encryption activation password.

Physical tape export encryption

You can encrypt data when exporting from a virtual tape to a physical tape in a physical tape library. You need at least one encryption key. Refer to 'Create a key' for more information.

You can apply a single key to all virtual tapes when you export them to physical tape, or you can create a unique key for each one. Creating multiple keys provides more security; in the unlikely event that a key is compromised, only the tapes that use that key would be affected. However, if you use multiple keys, you must keep track of which key applies to each tape so that you use the correct key to decrypt the data when you import the physical tape back to virtual tape.

Once you have created one or more keys, you can export them to a separate file called a *key package*. If you send encrypted tapes to other locations that run VTL, you can also send the key package. By importing the key package, administrators at the other sites can then decrypt the tapes when they are imported back into virtual tape libraries managed by VTL. Without this key package, data will not be usable, providing protection against storage being stolen and used on a different VTL/SIR system.

You can enable encryption and specify which key to use when you:

- Manually import or export a tape
- Configure to use the auto archive feature
- Create a cache for physical tapes, if you are using Automated Tape Caching

Note: If you apply an incorrect key when importing a tape, the data imported from that tape will be indecipherable.

Manage encryption keys

Keys can be created and managed for virtual tape encryption and physical tape export encryption. Deduplication repository encryption uses an internal key that is created when encryption is enabled; it is not visible in *Key Management*.

Create a key

Each key can be used for both virtual and physical tapes.

Each key consists of a secret phrase. For additional security, each key is password-protected. You must provide this password in order to activate encryption, change the key name, password, or password hint, or to delete or export the key.

To create a key:

1. In the navigation tree, right-click the server name and click *Key Management*.
2. Click *New*.

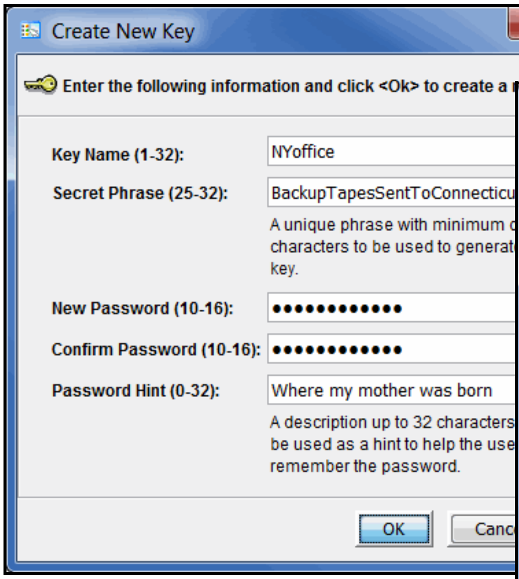

Enter the following information and click <Ok> to create a

Key Name (1-32): NYoffice

Secret Phrase (25-32): BackupTapesSentToConnecticu

A unique phrase with minimum c
characters to be used to generat
key.

New Password (10-16): ●●●●●●●●●●●●

Confirm Password (10-16): ●●●●●●●●●●●●

Password Hint (0-32): Where my mother was born

A description up to 32 characters
be used as a hint to help the use
remember the password.

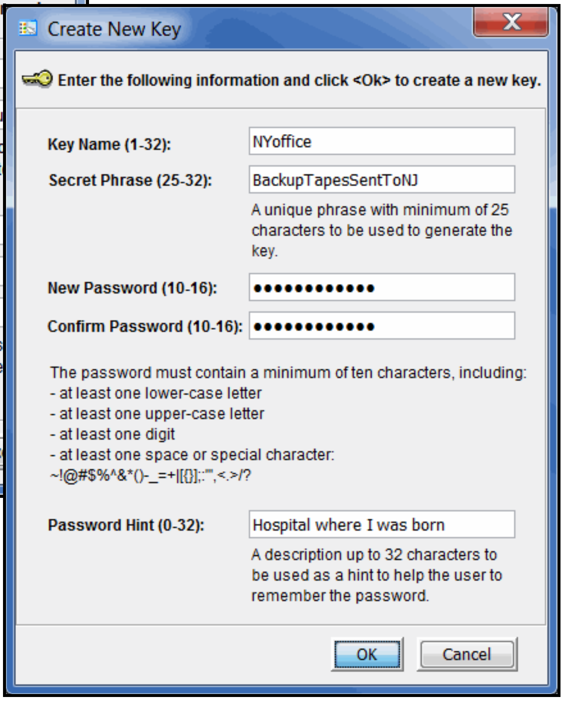

Enter the following information and click <Ok> to create a new key.

Key Name (1-32): NYoffice

Secret Phrase (25-32): BackupTapesSentToNJ

A unique phrase with minimum of 25 characters to be used to generate the key.

New Password (10-16): ●●●●●●●●●●●●

Confirm Password (10-16): ●●●●●●●●●●●●

The password must contain a minimum of ten characters, including:
- at least one lower-case letter
- at least one upper-case letter
- at least one digit
- at least one space or special character:
~!@#$%^&*()-_=+|[{}];:'",<.>/?

Password Hint (0-32): Hospital where I was born

A description up to 32 characters to be used as a hint to help the user to remember the password.

The dialog you will see when the *Strong Passwords* option is enabled.

3. In the *Key Name* text box, type a unique name for the key (1–32 characters).
4. In the *Secret Phrase* text box, type a phrase (25–32 characters, including numbers and spaces) that will be used to encrypt the data.

> **Note:** We recommend that you make a note of your secret phrase somewhere. This is important in case you need to recreate the key if it is ever accidentally deleted.

5. In the *New Password* and *Confirm Password* text boxes, type a password for accessing the key (10–16 characters).

   You will need to provide this password in order to change the key name, password, or password hint, or to delete or export the key.

   You do not have to provide a unique password for each key. In fact, if you use the same password for multiple keys, you have to provide the password only once when you export multiple keys that all use the same password.

   If you are using the *Strong Passwords* option, the password contain at least one lower-case letter, one upper-case letter, and one digit, plus at least one space or special character: ~!@#$%^&*()-_=+\|[{}];:'",<.>/?

6. In the *Password Hint* text box, type a hint (0–32 characters) to help you remember the password.

   This hint appears when you type an incorrect password and request a hint.

7. Click *OK*.

Change a key name or password

Once you have created a key, you cannot change the secret phrase associated with that key. However, you can change the password used to access the key and the hint associated with that password. For tape import purposes, you can also change the name of the key; you cannot rename a key used for virtual tape encryption.

If you rename a key, you can still use that key to decrypt data that was exported and encrypted using the old key name. For example, if you encrypt data using Key1, and you change its name to Key2, you can decrypt the data using Key2, since the secret phrase is the same.

To change a key name or password:

1. In the navigation tree, right-click the server name and click *Key Management*.
2. From the *Key Name* list, click the key you want to change.
3. Click *Edit*.
4. If you closed the *Key Management* dialog box after creating the key, type the current password for accessing this key in the *Password* text box.

   If you just created the key, did not close the *Key Management* dialog box, and subsequently decided to change the key, you are not prompted for the password.

5. Make the desired changes.
6. Click *OK*.

Delete a key

Caution: Once you delete a key, you can no longer decrypt tapes that were encrypted using that key unless you subsequently create a new key that uses the exact same secret phrase, or import the key from a key package.

To delete a key:

1. In the navigation tree, right-click the server name and click *Key Management*.
2. From the *Key Name* list, click the key that you want to delete.
3. Click *Delete*.
4. In the *Password* text box, type the password for accessing this key.
5. Type YES to confirm.
6. Click *OK*.

Export a key

When you export a key, you create a separate file called a *key package* that contains one or more keys. You can then send this file to another site that uses VTL, and administrators at that site can import the key package and use the associated keys to encrypt or decrypt data.

Creating a key package also provides you with a backup set of keys. If a particular key is accidentally deleted, you can import it from the key package so that you can continue to access the data encrypted using that key.

To export a key:

1. In the navigation tree, right-click the server name and click *Key Management*.
2. Click *Export*.
3. In the *Package Name* text box, type the file name to use for this key package (1–32 characters).
4. In the *Decryption Hint* text box, type a three-character hint.

   When you subsequently attempt to import a key from this key package, you are prompted for a password. If you provide the correct password, the decryption hint specified here appears correctly on the *Import Keys* dialog box. If you provide an incorrect password, a different decryption hint appears. You can import keys using an incorrect password, but you will not be able to decrypt any files using those keys.
5. From the *Select Keys to Export* list, select the key(s) that you want to include in the key package.

   When you select a key or click *Select All*, you are prompted to provide the password for each key. (If multiple selected keys use the same password, you are prompted for the password only once, when you select the first key that uses that password.)

After you type the password in the *Password* text box, that password appears in the *Password for All Keys in Package* area on the *Export Keys* dialog box. By default, the password is displayed as asterisks. To display the actual password, select the *Show clear text* check box.

If you selected a key and subsequently decide not to include it in the key package, you can clear the key. You can also clear all selected keys by clicking *De-Select All*.

6. Select *Prompt for new password for all keys in package* if you want to create a new password for the key package.

   If you select this option, you will be prompted to provide the new password when you click *OK* on the *Export Keys* dialog box. You will subsequently be prompted for this password when you try to import a key from this package. In addition, all keys imported from this package will use this new password rather than the password originally associated with each key.

   If you clear this option, this package will use the same password as the first selected key (which appears in the *Password for All Keys in Package* area), and you must provide this password when you try to import a key from this package. You must also provide this password when you subsequently change, delete, or export any key imported from this package.

7. In the *Save in this directory* text box, type the full path for the file.

   Alternatively, you can click 📁, select the desired directory, and click *Save*.

8. Click *OK*.

   If you selected the *Prompt for new password for all keys in package* check box, type the new password (10–16 characters) in the *New Password* and *Confirm Password* text boxes, type a hint for that password (0–32 characters) in the *Password Hint* text box.

   A file with the specified package name and the extension *.key* is created in the specified location.

Import a key

Once you have created a key package, you can open that package and specify which keys to import into VTL. Once you import a key, you can use that key to encrypt or decrypt data.

To import a key:

1. In the navigation tree, right-click the server name and click *Key Management*.
2. Click *Import*.
3. In the *Find Package* text box, type the full path to the key package.

Alternatively, you can click , select the file in the appropriate location, and click *Open*.

4. Click *View*.

5. Type the password for accessing the key package in the *Password* text box.

   Note: After you provide the password, make sure that the displayed *Decryption Hint* matches the decryption hint specified when the key package was created. If the hint is not correct, click *Password* and provide the correct password for accessing the key package. If you provide an incorrect password, you will still be able to import the keys in the package, but you will not be able to use them to decrypt any data that was previously encrypted using those keys.

6. From the *Select Keys to Import* list, select the keys that you want to import.

   You can select only those keys that have a green dot and the phrase *Ready for Import* in the *Status* column. A red dot and the phrase *Duplicate Key Name* indicates that a key of the same name already exists in this instance of VTL and cannot be imported.

   If you selected a key and subsequently decide not to import it, you can clear the key. You can also clear all selected keys by clicking *De-Select All*. (You can click this button only if the *Show All Keys* check box is cleared.)

   Note: A key of the same name might not necessarily have the same secret phrase. For example, you might have a key named Key1 with a secret phrase of ThisIsTheSecretPhraseForKey1. If the key package was created by another instance of VTL, it might also have a key named Key1, but its secret phrase might be ThisIsADifferentSecretPhrase. Since the key names are the same, you will not be able to import the key in the key package unless you rename the existing Key1. After you rename the key, you can continue to use it to decrypt tapes that were encrypted using that key, and you can also import the key named Key1 from the key package and use it to decrypt tapes that were encrypted using that key.

7. Click *OK*.

   The imported keys appear in the *Key Name* list on the *Key Management* dialog box. When you subsequently export or import a tape, these key names also appear in the *Select a Key* list.

VTL Server Failover

Overview

VTL server failover provides high availability for your VTL servers, protecting you from a wide variety of problems, including:

- Storage device path failure
- VTL server failure

The following illustrates a basic VTL configuration with potential points of failure and a high availability configuration, where VTL's high availability options work with redundant hardware to eliminate the points of failure:

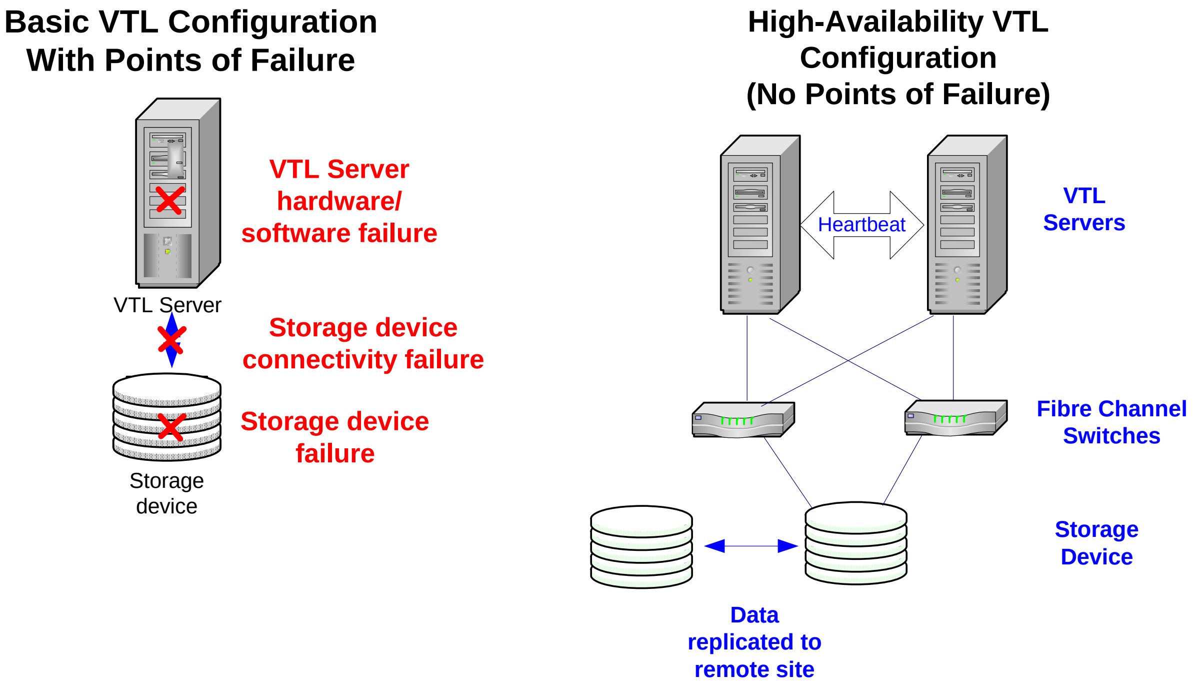

Basic VTL Configuration With Points of Failure

High-Availability VTL Configuration (No Points of Failure)

VTL Server hardware/ software failure
VTL Server
Storage device connectivity failure
Storage device failure
Storage device
Heartbeat
VTL Servers
Fibre Channel Switches
Storage Device
Data replicated to remote site

Storage device path failure

A storage device path failure can occur due to a cable or switch/router failure.

You can eliminate this potential point of failure by providing a multi-path configuration, using multiple Fibre Channel switches, and/or multiple adapters, and/ or storage devices with multiple controllers. In a multi-path configuration, all paths to the storage devices are automatically detected. If one path fails, there is an automatic switch to another path.

Note that Fibre Channel switches can demonstrate different behavior in a multi-path configuration. Before using this configuration with VTL, you must verify that the configuration can work on your server *without* the VTL software. To verify:

1. Use the hardware vendor’s utility to see the devices after the driver is loaded.

   You can also use Linux’s *cat /proc/scsi/scsi* command.
2. Use the hardware vendor’s utility to access the devices.

   You can also use Linux’s *hdparm* command.
3. Unplug the cable from one device and use the utilities listed above to verify that everything is working.
4. Repeat the test by reversing which device is unplugged and verify that everything is still working.

VTL server failure

A server failure can occur due to a software or hardware failure.

In the VTL failover design, a VTL server is configured to monitor another VTL server. In the event that the server being monitored fails to fulfill its responsibilities to the clients it is serving, the monitoring server will seamlessly take over its identity so that the clients will transparently fail over to the monitoring server.

A unique monitoring system is used to check the health of the VTL servers. This system includes a self-monitor and a partner heartbeat monitor module.

The *self-monitor* is part of all VTL servers, not just the servers configured for failover and provides continuous health status of the server. It is part of the process that provides operational status to any interested and authorized parties, including the console and supported network management applications through SNMP. The self-monitor checks the VTL processes and connectivity to the server’s storage devices.

In a failover configuration, the *heartbeat monitor* continuously monitors the partner server through the same network path that the server uses to serve its clients. The health monitoring results are evaluated to determine if any action is needed.

Failover triggers

Failover will occur under any of the following conditions:

- Any network-protected interface fails on a NAS server.
- All network-protected interfaces fail on a VTL tape server.
- A Fibre Channel target port reports a link down while a virtual library/drive is assigned to a client through that port.
- An attempt to access a storage device by the primary fails and the secondary server is able to access it.
- There is a server panic on the primary server.
- Any critical module fails on the primary server.
- The primary server is powered off.

VTL uses its *self-monitor* and *heartbeat monitor* to determine whether failover should be triggered when one of the above conditions is met.

- The self-monitor detects a critical server process has failed that is determined to be fatal, yet the error does not affect the network interface. In this case, the server will ask its partner server to take over.
- The self-monitor detects a *storage device connectivity failure* but cannot determine if the failure is local or applies to the partner server also. In this case, the server reports the device error condition to its partner server through the heartbeat. The partner server will check to see if it can successfully access the storage. If it can, it attempts to access all devices. If it can successfully access all devices, the partner server initiates a takeover. If it cannot successfully access all devices, no failover occurs.
- Because the heartbeat uses the same network path that the server uses to serve its clients, the partner server knows when the clients cannot access the partner server when network connectivity is incapacitated. This scenario is considered a *catastrophic failure*. In this case the server will initiate a takeover.
- The partner server must be able to be accessed and reset via a remote power control device (IPMI or HP iLO) in order for automatic failover to occur.

> **Note:** If you want to avoid triggering failover when you need to perform storage maintenance, such as updating the firmware, you should suspend failover. You can resume failover when you are done.

Failover terminology

Primary/ Secondary VTL servers

VTL *primary* and *secondary* servers are separate, independent VTL servers that each have their own assigned clients. In *Active-Passive Failover*, the primary VTL server is monitored by the secondary VTL server. In the event the primary server fails, the secondary takes over. *Active-Passive Failover* is available for VTL tape servers and for NAS-only servers.

In *Active-Active Failover (*or *Mutual Failover)*, both servers are configured to monitor each other. *Active-Active Failover* is available for VTL tape servers. The terms *Primary* and *Secondary* are purely from the client's perspective. Each server is *primary* to its own clients and *secondary* to the other's clients. Each server normally services its own clients. In the event one server fails, the other will take over and serve the failed server's clients.

Failover/ Takeover

Failover/takeover is the process that occurs when the secondary server takes over the identity of the primary.

Recovery/ Failback

Recovery/Failback is the process that occurs when the secondary server releases the identity of the primary to allow the primary to restore its operation. There is a three minute delay during recovery. During this time, no I/O is permitted.

Failover group

For ease of identification, the primary and secondary failover servers in a set are grouped together in the console beneath the failover group object.

Heartbeat IP address

A static IP address used by failover servers to monitor each other's health. The Heartbeat IP address remains with the server in the event of failure so that the server's health can be continually monitored.

Virtual IP address

An IP address used for console and client connections. The Virtual IP address is transferred to the partner server during failure so that console and client connections using this IP address remain active without any interruptions.

Power control management interface

A hardware-level interface that monitors various hardware functions on a server. IPMI and HP iLO are supported interfaces.

Failover behavior

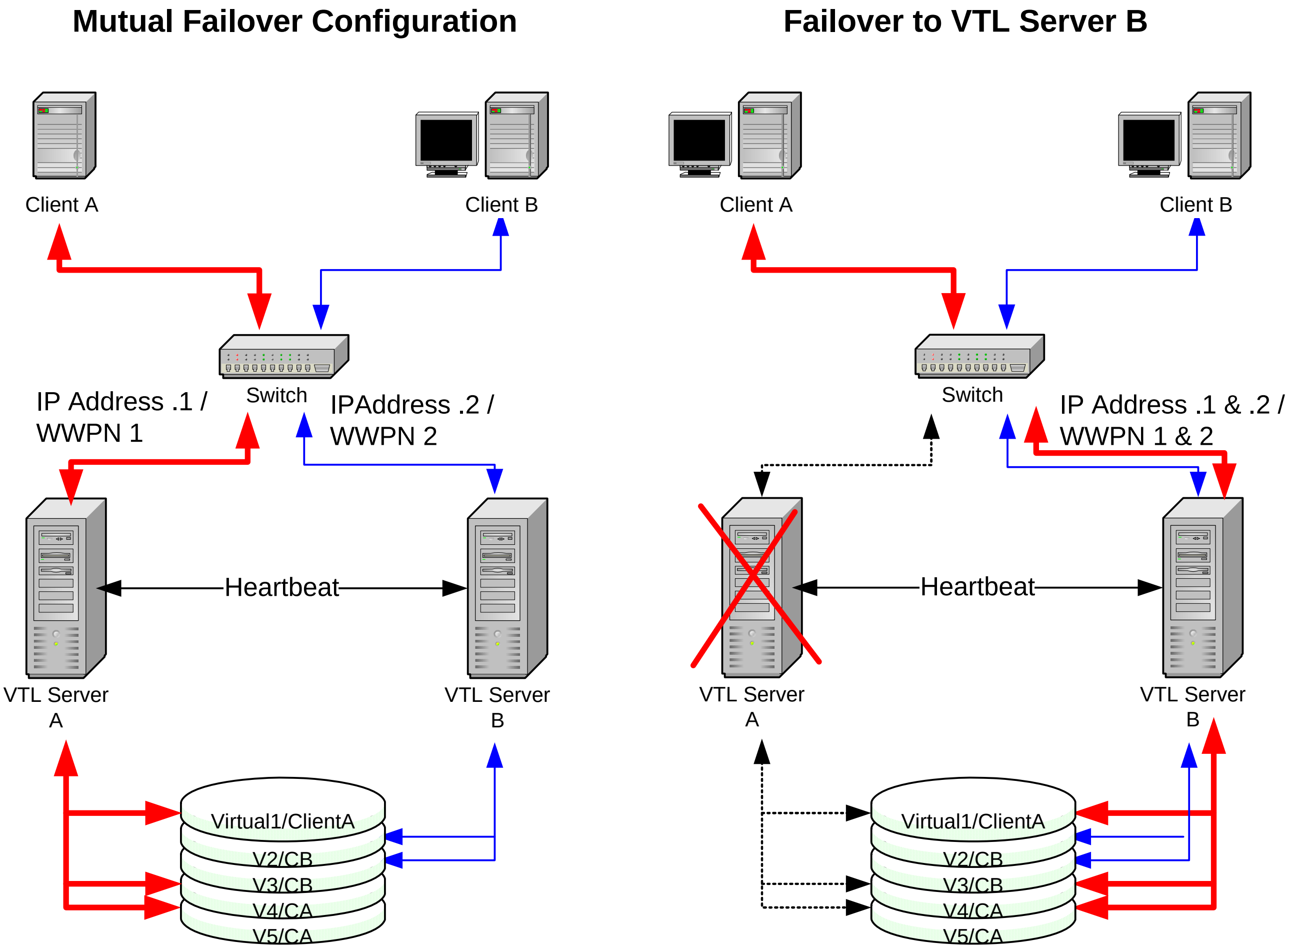

Mutual Failover Configuration

Failover to VTL Server B

This diagram illustrates a VTL failover configuration. When server A fails, server B takes over and serves the clients of server A in addition to its own clients.

When failover occurs, there is a delay of several minutes until the secondary server takes over. During this time, no I/O is permitted on the failed server. Therefore, all backup jobs (including import/export and replication) running on the primary server will fail.

When failover occurs, the following operations cannot be performed:

- Change the failover configuration on either server
- Select the system maintenance console option on the failed server
- Change server properties on the failed server
- Run reports on the failed server
- Change the storage hardware configuration on the failed server

IP connectivity

When a server fails over to its partner, its virtual IP address is transferred to the partner server so that console and client connections using this virtual IP remain active without any interruptions. The following example illustrates the steps.

Sample environment

- There are two servers, "A" and "B", which act as failover partners.
- Each server has two IP network ports.
- Each network port has one heartbeat IP address for failover monitoring: `AHIP1` and `AHIP2` for server A, `BHIP1` and `BHIP2` for server B.
- Each network port has one alias IP that is used as a virtual IP address for console and client connections: `AVIP1` and `AVIP2` for server A, `BVIP1` and `BVIP2` for server B.

Before failover

Four IP addresses show on each server:

- Server A: `AHIP1 AVIP1 AHIP2 AVIP2`
- Server B: `BHIP1 BVIP1 BHIP2 BVIP2`

When server A fails over to server B

1. Virtual IP addresses are unloaded from server A. The result is:
   - Server A: `AHIP1 AHIP2`
   - Server B: `BHIP1 BVIP1 BHIP2 BVIP2`
2. Virtual IP addresses of server A are added to server B. The result is:
   - Server A: `AHIP1 AHIP2`
   - Server B: `BHIP1 BVIP1` **`AVIP1`** `BHIP2 BVIP2` **`AVIP2`**

When server A recovers from failover

1. New virtual IP addresses are unloaded from server B. The result is:
   - Server A: `AHIP1 AHIP2`
   - Server B: `BHIP1 BVIP1 BHIP2 BVIP2`
2. Virtual IP addresses are added back to server A. The result is:
   - Server A: `AHIP1` **`AVIP1`** `AHIP2` **`AVIP2`**
   - Server B: `BHIP1 BVIP1 BHIP2 BVIP2`

Fibre Channel connectivity

When a server takes over its partner, it impersonates its initiator standby FC ports as target ports of the failed server by spoofing the initiator WWPNs with the partner’s target WWPNs. In this way the Fibre Channel client connections remain active without any service interruption. The following example indicates the spoofing steps.

Sample environment

- There are two servers, "A" and "B", which act as failover partners.
- Each server has two target ports: `AT1` and `AT2` for server A, and `BT1` and `BT2` for server B.
- Each server has two standby initiator ports: (`AS1` and `AS2` for server A, and `BS1` and `BS2` for server B.

Note that `AT1, AT2, BT1, BT2, AS1, AS2, BS1`, and `BS2` represent WWPNs.

Before failover

Four FC WWPN entries from each server appear on the switch:

- Server A: `AT1 AT2 AS1 AS2`
- Server B: `BT1 BT2 BS1 BS2`

When server A fails over to server B

1. Target ports are unloaded from server A and their WWPNs disappear from the switch. The result is:
   - Server A: **`00 00`** `AS1 AS2`
   - Server B: `BT1 BT2 BS1 BS2`
2. Target ports are spoofed on server A with temporary WWPN `AX1` and `AX2` and re-appear on the switch as initiator ports. The result is:
   - Server A: **`AX1 AX2`** `AS1 AS2`
   - Server B: `BT1 BT2 BS1 BS2`
3. Standby initiator ports are unloaded from server B and their WWPNs disappear from the switch. The result is:
   - Server A: `AX1 AX2 AS1 AS2`
   - Server B: `BT1 BT2` **`00 00`**
4. Standby initiator ports are spoofed on server B with the original WWPNs from server A, get reloaded as target ports, and re-appear on the switch as target ports. The result is:
   - Server A: `AX1 AX2 AS1 AS2`
   - Server B: `BT1 BT2` **`AT1 AT2`**

When server A recovers from failover

1. New target ports are unloaded from server B and their WWPNs disappear from the switch. The result is:
   - Server A: `AX1 AX2 AS1 AS2`
   - Server B: `BT1 BT2` **`00 00`**
2. New target ports on server B regain their original WWPN ports and appear on the switch as initiator ports. The result is:

- Server A: `AX1 AX2 AS1 AS2`
- Server B: `BT1 BT2` **`BS1 BS2`**

3. Temporary target ports are unloaded from server A and their WWPNs disappear from the switch. The result is:
   - Server A: **`00 00`** `AS1 AS2`
   - Server B: `BT1 BT2 BS1 BS2`

4. Target ports on server A regain their original WWPN ports and appear on the switch as target ports. The result is:
   - Server A: **`AT1 AT2`** `AS1 AS2`
   - Server B: `BT1 BT2 BS1 BS2`

Resume backups after failover/failback

When failover and recovery occur, there is a three minute delay. During this time, no I/O is permitted and all backup jobs (including import/export and replication) will fail.

While failover and failback are transparent for the VTL server, after failover/failback occurs, you may need to take some action in your backup application in order for it to work properly with VTL. The action you take varies by backup application and operating system. We have described some of the actions we have used. Your environment may differ. Refer to the documentation that came with your backup application for more details.

Dell NetVault™

- After failover/failback occurs, if the Windows Device Manager hangs, you must reboot your NetVault server. Once devices are visible by both the operating system and NetVault, if the NetVault job hangs while waiting to connect to the tape media, reboot your NetVault server.
- Sometimes after failover or failback, a drive may be online but tapes that were previously in the drive or in the slots are missing. This is caused by an interruption of I/O, which causes the software to lose barcodes. To resolve this, the library, drives, and tapes need a complete rescan. Right-click the library (in the NetVault Device Manager) and select *Open Door*. Wait a few moments and select *Close Door*. You should now see all tapes properly labeled and back in their respective drives and/or slots.

CA ARCserve®

After failover/failback occurs, if the operating system and/or ARCserve losses devices, stop the ARCserve tape engine, rescan the operating system, and then restart the ARCserve tape engine. If you still cannot see devices, reboot your ARCserve server.

Once devices are visible by both the operating system and ARCserve, start ARCserve, eject all tapes from their drives and then re-inventory the library.

HP Data Protector

During failover/failback, backup jobs may fail. When this occurs, the tape will be marked as *poor* quality and it will stay in the tape drive. Manually move the tape back to the slot.

IBM® Tivoli® Storage Manager

After failover/failback occurs, reboot the Tivoli Storage Manager machine to get the devices back before submitting any jobs.

EMC NetWorker®

After failover/failback, NetWorker will mark the tape for the current backup job as *full* and will use a new blank tape to continue the backup job.

Symantec Backup Exec™

If Backup Exec has stalled jobs, reboot the Backup Exec server.

If there are no stalled jobs, but drives are down during failover or failback, everything should recover normally.

Symantec NetBackup™

On Windows, if NetBackup has stalled jobs, reboot the NetBackup server. After rebooting, check the NetBackup tape drive/library status and restart NetBackup services, if needed.

On Windows, if drives are failed or missing, bringing up the drives should be sufficient. However, if one or more of the drives cannot be brought up, reboot the NetBackup server.

On Solaris, jobs usually fail *gracefully* and subsequent jobs start without problem.

Failover requirements

General requirements

The following are the general requirements for setting up a failover configuration:

- You must have two VTL servers.
- The failover pair should be installed with identical operating system versions and must have identical storage configurations.
- Virtual storage devices must be accessible by both servers but devices cannot be owned by both servers. This means that storage devices must be attached in a multi-host SCSI configuration or attached on a Fibre loop or switched fabric. In this configuration, both servers can access the same devices at the same time (both read and write). You will see something similar to the following when you look at the physical storage in the console.

The same Fibre Channel device is shown from both servers (primary and secondary). The V icon indicates the disk is virtualized and owned by that server. The F icon indicates shared storage that is being used by another server.

FALCON:IPSTOR DISK

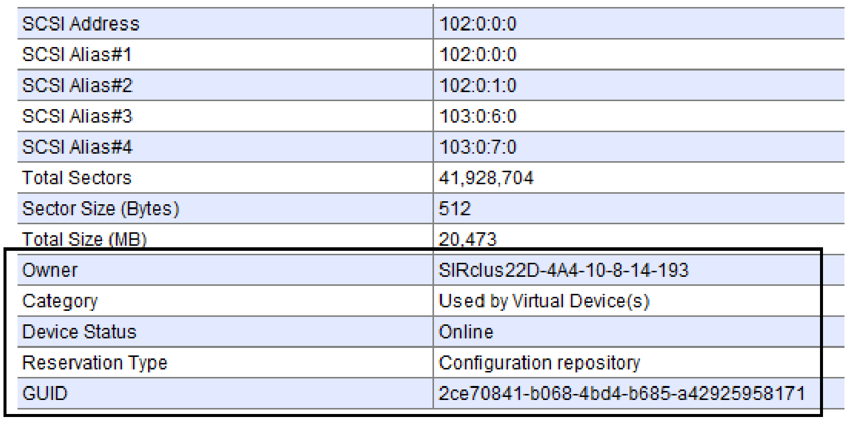

SCSI Address	102:0:0:0
SCSI Alias#1	102:0:0:0
SCSI Alias#2	102:0:1:0
SCSI Alias#3	103:0:6:0
SCSI Alias#4	103:0:7:0
Total Sectors	41,928,704
Sector Size (Bytes)	512
Total Size (MB)	20,473
Owner	SIRclus22D-4A4-10-8-14-193
Category	Used by Virtual Device(s)
Device Status	Online
Reservation Type	Configuration repository
GUID	2ce70841-b068-4bd4-b685-a42925958171

FALCON:IPSTOR DISK

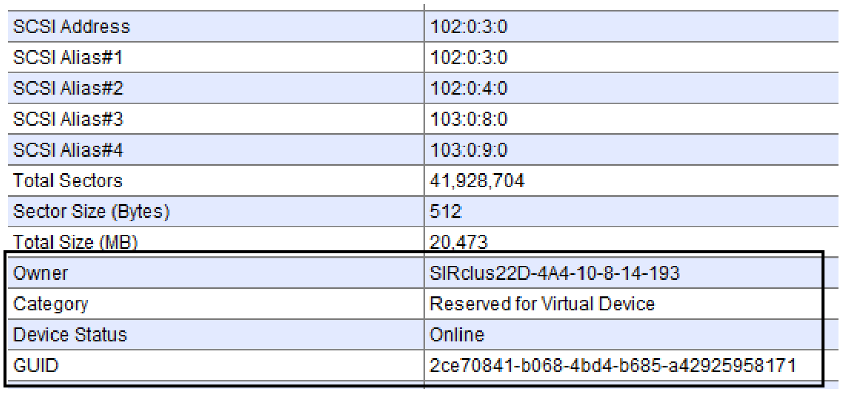

SCSI Address	102:0:3:0
SCSI Alias#1	102:0:3:0
SCSI Alias#2	102:0:4:0
SCSI Alias#3	103:0:8:0
SCSI Alias#4	103:0:9:0
Total Sectors	41,928,704
Sector Size (Bytes)	512
Total Size (MB)	20,473
Owner	SIRclus22D-4A4-10-8-14-193
Category	Reserved for Virtual Device
Device Status	Online
GUID	2ce70841-b068-4bd4-b685-a42925958171

- Both servers must have exactly the same VTL options licensed.
- Physical drives cannot be shared among VTL servers. Physical drives should be visible to the other server but not assigned to VTL.
- Both servers must be able to access the VTL database resource.
- Failover servers must be configured with a power control management interface on the motherboard. IPMI (Intelligent Platform Management Interface) and HP iLO are supported interfaces. Refer to 'Power control management options' for more information.

- Both servers must reside on the same network segment, because in the event of a failover, the secondary server must be reachable by the clients of the primary server. This network segment must have at least one other device that generates a network ping (such as a router, switch, or server). This allows the secondary server to detect the network in the event of a failure.
- If you are protecting multiple IP addresses, each protected IP address must be in its own subnet.
- You need to reserve an IP address for each network adapter in your failover servers. The IP address must be on the same subnet as the server. These IP addresses are used by the servers to monitor each other's health. The health monitoring IP address remains with the server in the event of failure so that the server’s health can be continually monitored. After failover, the health monitoring IP address still exists until the network services are restarted. Note that VTL clients and the console should not use the health monitoring IP address to connect to a server.
- We recommend using a direct network connection between servers to monitor heartbeats in a failover setup to make sure issues in network infrastructure, routers, and switches do not impact failover.
- Some switches do not automatically swap ports during failover. If your switch does not, you may need to configure port swapping. Refer to ‘Manage FC switch port swapping’ for more information.
- In addition to standard IP addresses, each server must have a power control IP address (used by power control software).
- The primary and secondary servers should use the exact same Target Port ID scheme on the matching WWPNs. We recommend using the same initiator adapter numbers on both sides to connect to the same storage, so the ACSL of all the devices will look identical on both sides.
- You must use static IP addresses for your failover configuration. We also recommend that the IP addresses of your servers be defined in a DNS server so they can be resolved.
- The first time you set up a failover configuration, the secondary server must not have any logical resources (including NAS resources, virtual tapes, drives, or libraries) or clients.
- Because the primary and secondary servers will become part of a failover group, the servers cannot belong to a multi-node group.
- If you have NDMP enabled on one server, you must enable it on the partner server before configuring failover.

- If you will be using Email Alerts, it is very important to enable them on each server so that an alert will be sent if failure occurs on any server. This should be done before configuring failover.

Notes:

- Once failover is configured, you cannot change the host name or IP addresses of the failover nodes.
- Once failover is configured, you cannot enable or disable iSCSI. If you want to change the state of iSCSI, you should do it before configuring failover.
- Once failover is configured, you cannot enable mirroring of the primary configuration repository.

VTL tape server failover requirements

The following are the additional requirements for VTL tape server failover:

- Both servers must have Fibre Channel Target Mode enabled.
- A standby initiator port needs to be available. This port needs to be connected to the same FC fabric as the target port and should not be zoned to anything else. In a multi-ID configuration, you can configure the target port to act as a standby initiator.
- The first time you set up a failover configuration, the secondary server must not have any deduplication server association. If failover has been configured previously, both VTL failover servers must be associated to the same deduplication server.
- If you will be using FalconStor OpenStorage Option (FSOST), you must enable it on both servers before configuring failover. If you are using the FSOST Direct to Tape Option, you must also enable it on both servers before configuring failover. In addition, OST must be enabled with the same role (VTL or NAS) on both servers. If you need to enable/disable FSOST (with or without Direct to Tape) in an active failover configuration, you must suspend failover first.

NAS failover requirements

The following are the additional requirements for NAS failover:

- NAS must be enabled on the secondary server. Enabling NAS on the secondary server can be done before or during failover configuration.
- Both servers must use the same authentication mode.
- If the primary server is using domain mode for authentication, the secondary server must have the ability to resolve the same Active Directory server as the primary. The /etc/hosts file should have a line for the Active Directory

server. During failover configuration, the wizard will confirm that the secondary server can resolve the same Active Directory server as the primary.

> **Note:** Once failover is configured for NAS-only servers, you cannot enable the VTL role on either server.

Power control management options

Power control management is a hardware-level interface that monitors various hardware functions on a server. Power control management is required in order for failover to occur. IPMI and HP iLO are supported interfaces.

At times, a server may become unresponsive, but, because of network or internal reasons, it may not release its resources or its IP address, thereby preventing failover from occurring. To allow for graceful failover, IPMI or HP iLO will reset the power of the primary server, forcing the release of the server's resources and IP address.

IPMI is included and configured on all DSI appliances.

If you are using a third-party appliance, you must determine if power control management is provided by your hardware vendor. To check for IPMI, use the `dmidecode` command and look for the *IPMI Device Information* section. You must also make sure that the *Interface Type* is defined.

If you determine that power control management is provided, you must follow the vendor's instructions to configure it and you must create an administrative user via your configuration tool. The IP address cannot be the virtual IP address that is set for failover.

Backup server failover configuration

While failover and failback are transparent for the VTL server, some configuration may be necessary for your backup server in order for physical devices to be accessed after failover. The configuration varies by backup application and operating system. To ensure that physical devices can be accessed after failover, we recommend you follow the instructions below.

AIX- All backup software

AIX uses the FC port ID to generate its local ID. VTL failover will change the FC switch port ID.

If you are using AIX 5.3 or above with a Brocade switch model 3900 or above, you need to enable the dynamic tracking option.

Configure failover

Notes:

- You will need to know the IP addresses of the failover servers, including the ones needed to handle the health monitoring process. It is a good idea to gather this information and find available IP addresses before you begin configuration.
- If you enable failover for a server configured for both Virtual Tape Library and NAS, only VTL resources will be protected; NAS share will not be available from the secondary server. NAS resources will be available after failback or server restart when failover is suspended.

1. Right-click the server that will become the primary server in your failover configuration and select *Failover --> Failover Setup Wizard*.

   If this is a VTL tape server, you will see a screen similar to the following that shows you a status of options on your server.

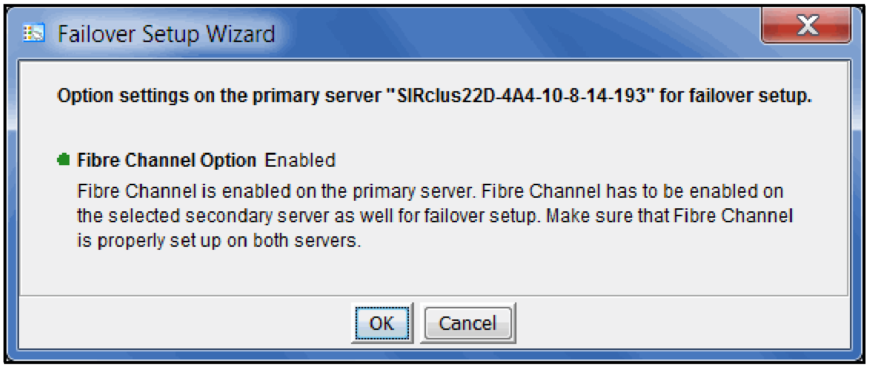

Option settings on the primary server "SIRclus22D-4A4-10-8-14-193" for failover setup.

● **Fibre Channel Option** Enabled

Fibre Channel is enabled on the primary server. Fibre Channel has to be enabled on the selected secondary server as well for failover setup. Make sure that Fibre Channel is properly set up on both servers.

2. Enter a group name for the failover pair.

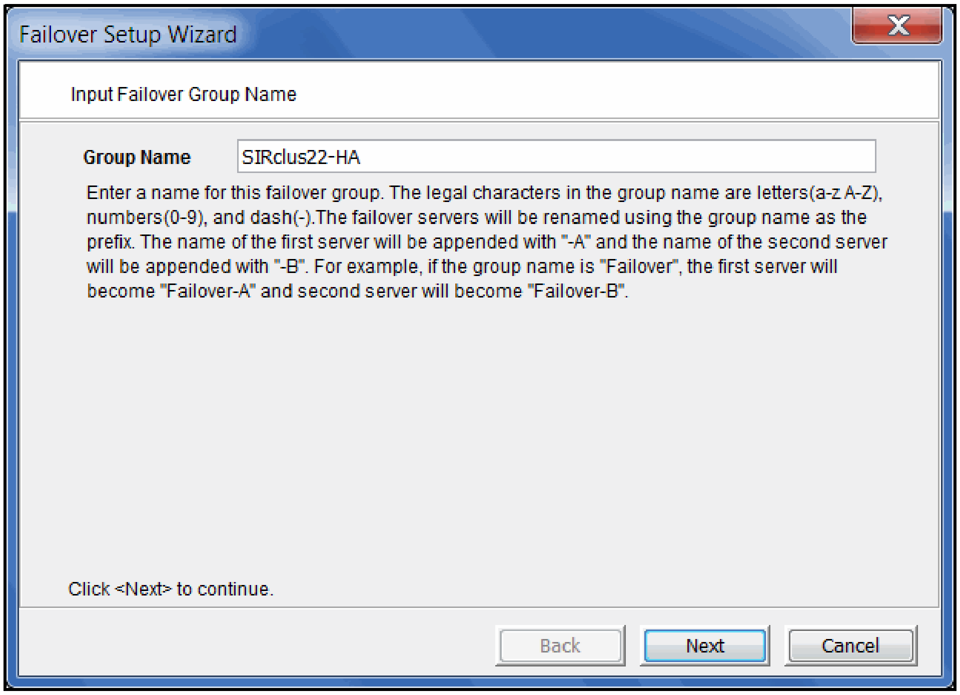

Input Failover Group Name

Group Name: SIRclus22-HA

Enter a name for this failover group. The legal characters in the group name are letters(a-z A-Z), numbers(0-9), and dash(-).The failover servers will be renamed using the group name as the prefix. The name of the first server will be appended with "-A" and the name of the second server will be appended with "-B". For example, if the group name is "Failover", the first server will become "Failover-A" and second server will become "Failover-B".

Click <Next> to continue.

For ease of identification, the primary and secondary failover servers in a set are grouped together in the console. If the servers are already part of a failover group, that group name will continue to be used and you will not see this dialog.

The failover group name can contain letters, numbers, a dash, or underscore. Spaces and other characters are not allowed.

3. Select the secondary server.

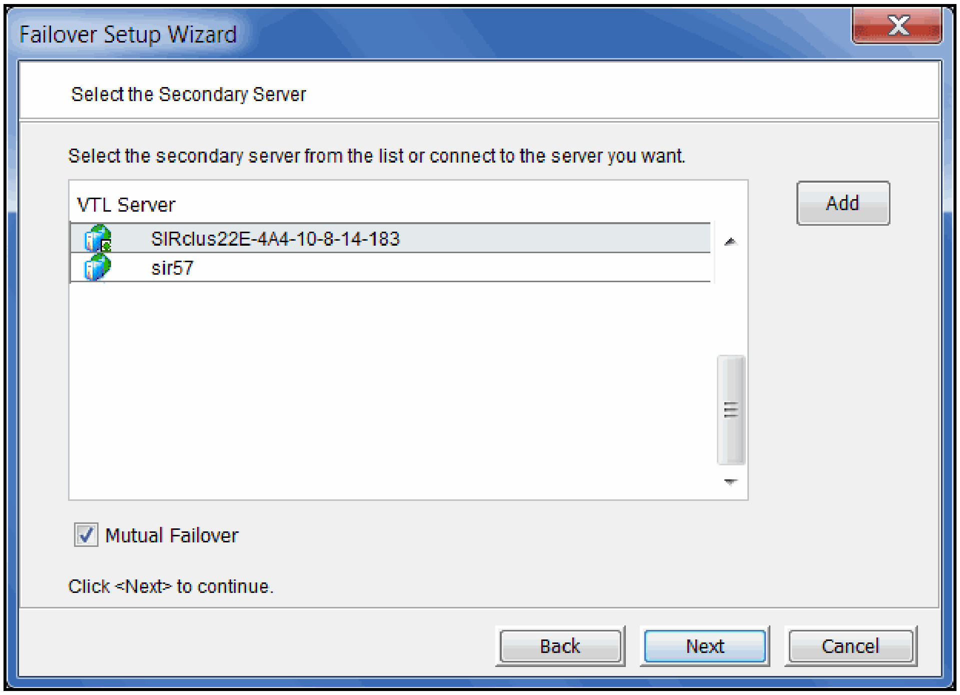

Select the Secondary Server

Select the secondary server from the list or connect to the server you want.

VTL Server
SIRclus22E-4A4-10-8-14-183
sir57

☑ Mutual Failover

Click <Next> to continue.

Select *Mutual Failover* if you want both servers to monitor each other. This is only valid for VTL resources; it is not applicable to NAS resources.

4. If prompted, determine if you want to change the host name of each failover server.

   If you choose to change the host name, the failover group name will be used as the prefix for the server name. The name of the first server is appended with "-A" and the name of the second server with "-B". For example, if the group name is "NewYork", the first server will be "NewYork-A" and the second server will be "NewYork-B". In order to rename a server, the server will have to be restarted.

5. Confirm or select the type of power control management interface your hardware is using.

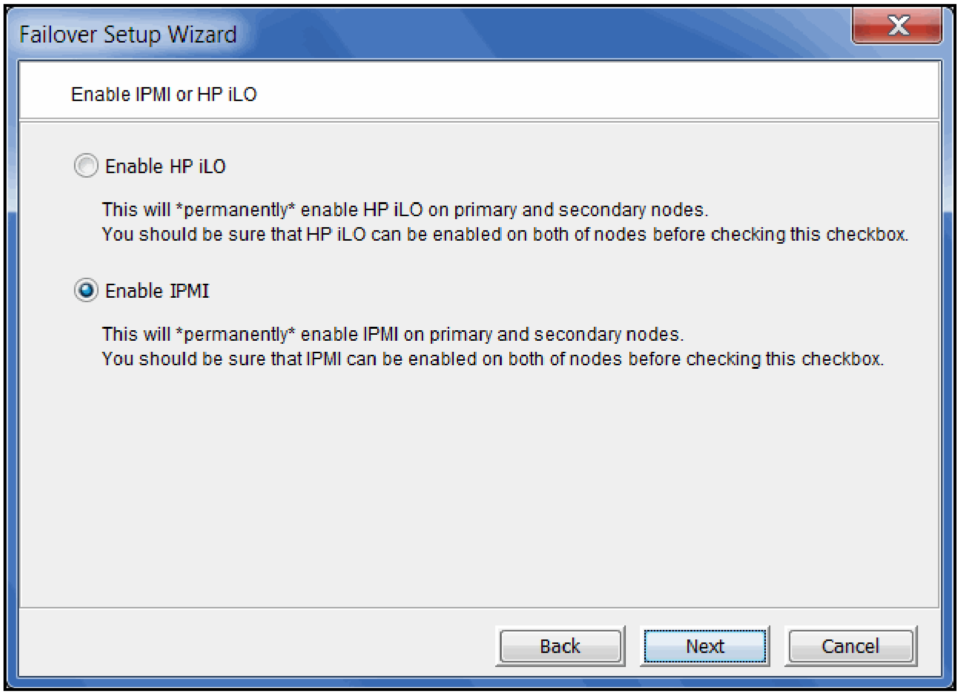

Enable IPMI or HP iLO

○ Enable HP iLO

This will *permanently* enable HP iLO on primary and secondary nodes.
You should be sure that HP iLO can be enabled on both of nodes before checking this checkbox.

● Enable IPMI

This will *permanently* enable IPMI on primary and secondary nodes.
You should be sure that IPMI can be enabled on both of nodes before checking this checkbox.

When you click *Next*, the driver will be loaded on each node.

6. Enter power control information.

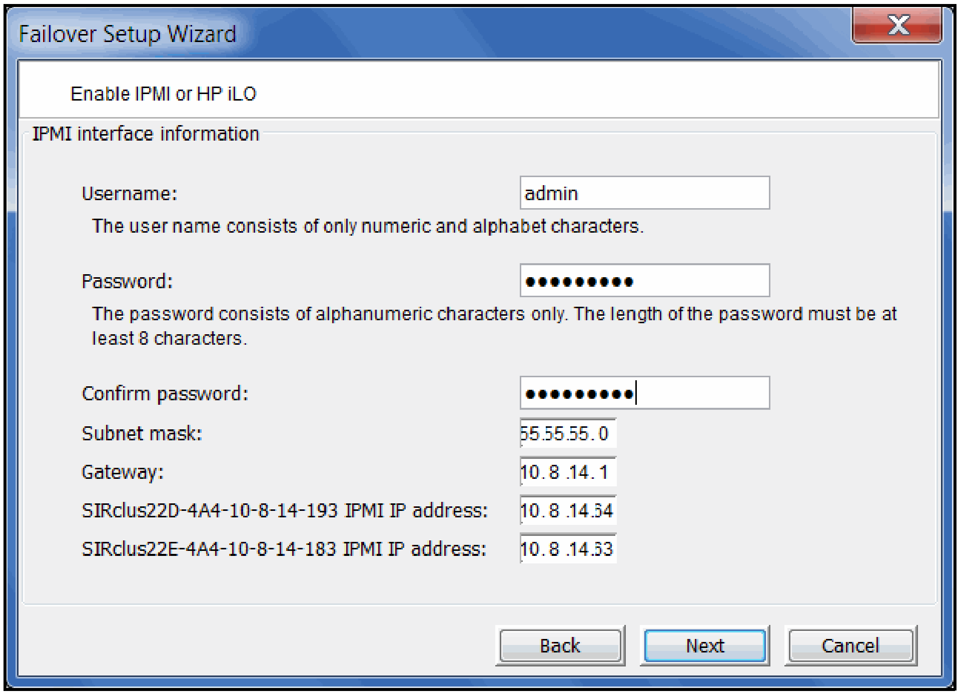

Enable IPMI or HP iLO

IPMI interface information

Username: admin

The user name consists of only numeric and alphabet characters.

Password: ●●●●●●●●●

The password consists of alphanumeric characters only. The length of the password must be at least 8 characters.

Confirm password: ●●●●●●●●●

Subnet mask: 55.55.55. 0

Gateway: 10. 8 .14. 1

SIRclus22D-4A4-10-8-14-193 IPMI IP address: 10. 8 .14.54

SIRclus22E-4A4-10-8-14-183 IPMI IP address: 10. 8 .14.53

The subnet, gateway, and IP address information is automatically pre-filled.

On DSI appliances, the IPMI user name and password are preset as *admin* and *falcon101*.

The user name and password you specify here will overwrite the existing information.

Note: After specifying IPMI interface information here, do not change them again using the IPMI-equivalent appliance management console.

7. Check the *Include this Network Adapter for failover* box if you want this network adapter monitored for failover.

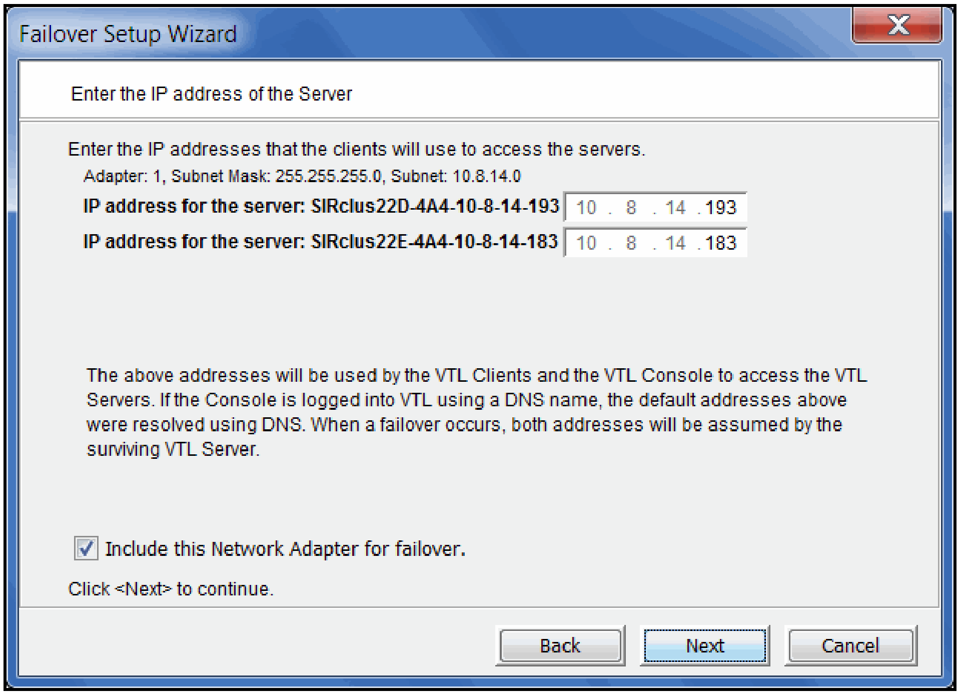

Enter the IP address of the Server

Enter the IP addresses that the clients will use to access the servers.

Adapter: 1, Subnet Mask: 255.255.255.0, Subnet: 10.8.14.0

IP address for the server: SIRclus22D-4A4-10-8-14-193 10 . 8 . 14 . 193

IP address for the server: SIRclus22E-4A4-10-8-14-183 10 . 8 . 14 . 183

The above addresses will be used by the VTL Clients and the VTL Console to access the VTL Servers. If the Console is logged into VTL using a DNS name, the default addresses above were resolved using DNS. When a failover occurs, both addresses will be assumed by the surviving VTL Server.

☑ Include this Network Adapter for failover.

Click <Next> to continue.

Select the IP addresses that clients will use to access the storage servers when using iSCSI, replication, and for console communication. Each IP address is known as a *Virtual IP address* and results in creation of a virtual Ethernet interface, (i.e., eth0:0, if the network interface eth0 is used). If this server fails over to its partner, the virtual IP address is transferred to the partner server so that console and client connections using this IP address remain active without any interruption.

If you uncheck the *Include this Network Adapter for failover* box, the wizard will display the next card it finds. You must choose at least one.

Notes:

- If you change the server IP addresses while the console is connected using those IP addresses, the Failover wizard will not be able to successfully create the configuration.
- Because failover can occur at any time, you should use only those IP addresses that are configured as part of the failover configuration to connect to the server.

8. Enter the health monitoring IP address you reserved for the selected network adapter.

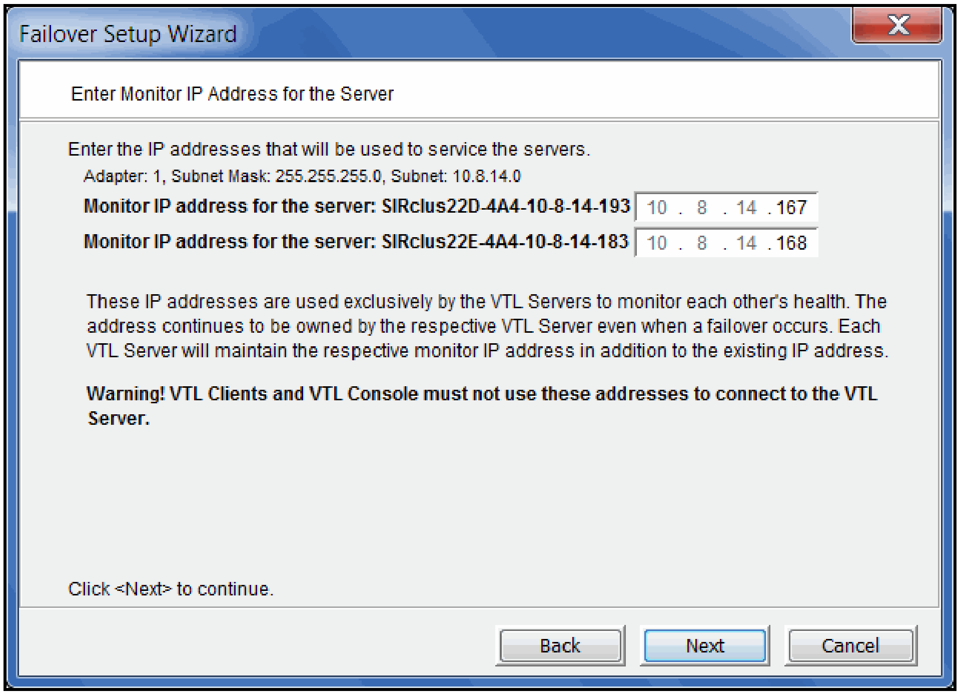

Enter Monitor IP Address for the Server

Enter the IP addresses that will be used to service the servers.

Adapter: 1, Subnet Mask: 255.255.255.0, Subnet: 10.8.14.0

Monitor IP address for the server: SIRclus22D-4A4-10-8-14-193 10 . 8 . 14 . 167

Monitor IP address for the server: SIRclus22E-4A4-10-8-14-183 10 . 8 . 14 . 168

These IP addresses are used exclusively by the VTL Servers to monitor each other's health. The address continues to be owned by the respective VTL Server even when a failover occurs. Each VTL Server will maintain the respective monitor IP address in addition to the existing IP address.

Warning! VTL Clients and VTL Console must not use these addresses to connect to the VTL Server.

Click <Next> to continue.

The Heartbeat IP address is used exclusively by the storage servers to monitor each other’s health. (These addresses must not be used for any other purpose.)

The health monitoring IP address, known as the *Heartbeat IP address*, remains with the server in the event of failure so that the server’s health can be continually monitored. Therefore it is recommended that you use static IP addresses. After failover, the health monitoring IP address still exists until the network services are restarted.

9. If you want to use additional network adapter cards, repeat steps 7 and 8.

10. If this is a VTL tape server, select the initiator on the secondary server that will function as a standby in case the target port on your primary server fails.

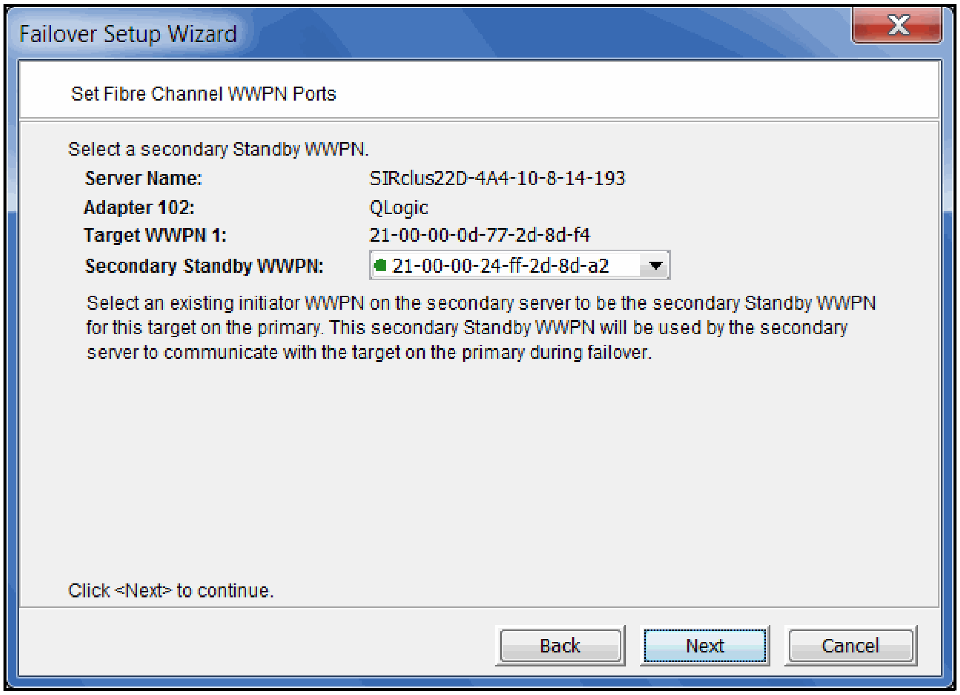

Set Fibre Channel WWPN Ports

Select a secondary Standby WWPN.

Server Name: SIRclus22D-4A4-10-8-14-193
Adapter 102: QLogic
Target WWPN 1: 21-00-00-0d-77-2d-8d-f4
Secondary Standby WWPN: 21-00-00-24-ff-2d-8d-a2

Select an existing initiator WWPN on the secondary server to be the secondary Standby WWPN for this target on the primary. This secondary Standby WWPN will be used by the secondary server to communicate with the target on the primary during failover.

Click <Next> to continue.

The proper adapter is usually selected for you, but you should confirm that the adapter shown is not the initiator on your secondary server that is connected to the storage array, and also that it is not the target adapter on your secondary server.

There should not be any devices attached to this initiator.

11. If this is a VTL tape server, set up the standby adapter for the secondary server.

12. Confirm all of the information and then click *Finish* to create the failover configuration.

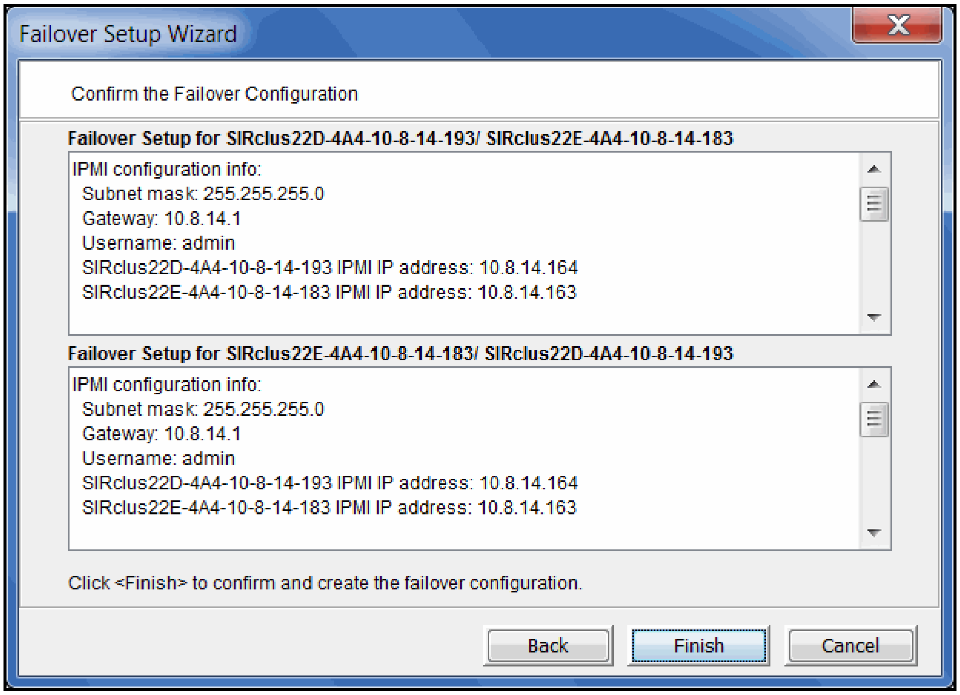

Confirm the Failover Configuration

Failover Setup for SIRclus22D-4A4-10-8-14-193/ SIRclus22E-4A4-10-8-14-183

```
IPMI configuration info:
  Subnet mask: 255.255.255.0
  Gateway: 10.8.14.1
  Username: admin
  SIRclus22D-4A4-10-8-14-193 IPMI IP address: 10.8.14.164
  SIRclus22E-4A4-10-8-14-183 IPMI IP address: 10.8.14.163
```

Failover Setup for SIRclus22E-4A4-10-8-14-183/ SIRclus22D-4A4-10-8-14-193

```
IPMI configuration info:
  Subnet mask: 255.255.255.0
  Gateway: 10.8.14.1
  Username: admin
  SIRclus22D-4A4-10-8-14-193 IPMI IP address: 10.8.14.164
  SIRclus22E-4A4-10-8-14-183 IPMI IP address: 10.8.14.163
```

Click <Finish> to confirm and create the failover configuration.

Once your configuration is complete, each time you connect to either server in the console, you will automatically be connected to the other as well.

Note: If the setup fails during the setup configuration stage (for example, the configuration is written to one server but then the second server is unplugged while the configuration is being written to it), use the *Remove Failover Configuration* option to delete the partially saved configuration. You can then create a new failover configuration.

Manage FC switch port swapping

In a failover setup, the client driver recognizes the drives by port ID instead of WWPN. When failover occurs, as the standby adapters on the secondary are connecting to different port IDs, the client will no longer be able to see the drives.

Some switches, such as Cisco, automatically swap ports. However, for some switches, such as Brocade, you need to create the following port swapping scripts:

- A pre-takeover script to switch port IDs when failover occurs.
- A pre-recovery script to change the port IDs back before failback occurs.

Notes:

- Tape drive multi-pathing is not supported with port swapping. Even though you can configure VTL to assign one drive to two paths, failover will not be transparent, meaning that when the backup job fails, the client will need to be reconfigured to use the device from the second path.
- Port swapping scripts on HP and AIX platforms are not supported with Multi-ID HBAs.
- The portswap command will only work within one switch. You can have multiple pairs of target and standby ports located on different switches but the target-standby ports of each pair need to be located on the same switch.

To configure port swapping after failover is configured:

1. Build SSH host-based authentication between VTL servers and the switch.
   - SSH to the switch using the "root" user account.
   - On both VTL servers, run: `ssh root@<`*`switch IP address`*`>`. After logging in to the switch, exit from SSH. This step will populate VTL servers as "known hosts" to the switch.
2. Copy `preRecovery.portswap` on the primary server as the pre-recovery script.

   ```
   #cd $ISHOME/bin
   ```

   ```
   #cp preRecovery preRecovery.bak
   ```

   ```
   #cp preRecovery.portswap preRecovery
   ```
3. Copy `preTakeOver.portswap` on the secondary server as the pre-takeover script.

   ```
   #cd $ISHOME/bin
   ```

   ```
   #cp preTakeOver.portswap preTakeOver.bak
   ```

   ```
   #cp preTakeOver.portswap preTakeOver
   ```
4. Use the `switchshow` command to collect the information that needs to be replaced in the `portswap.sh` lines of both scripts.
   - `switch_IP -` IP address of the switch.
   - `switch_password -` Password of root user for switch.

- `area id -` Switch port area ID. If no area ID, put any one of the ports to be swapped.
- `slot/port -` All ports to be swapped, one pair after another. Slot is optional.

> **Note:** You will need one portswap line for each switch.

5. Run `$ISHOME/bin/preTakeOver` to make sure port IDs are swapped correctly.
6. Run `$ISHOME/bin/preRecovery` to make sure port IDs are swapped back.

Check failover status

You can see the current status of your failover configuration, including all settings, by checking the *Failover Information* tab for the server.

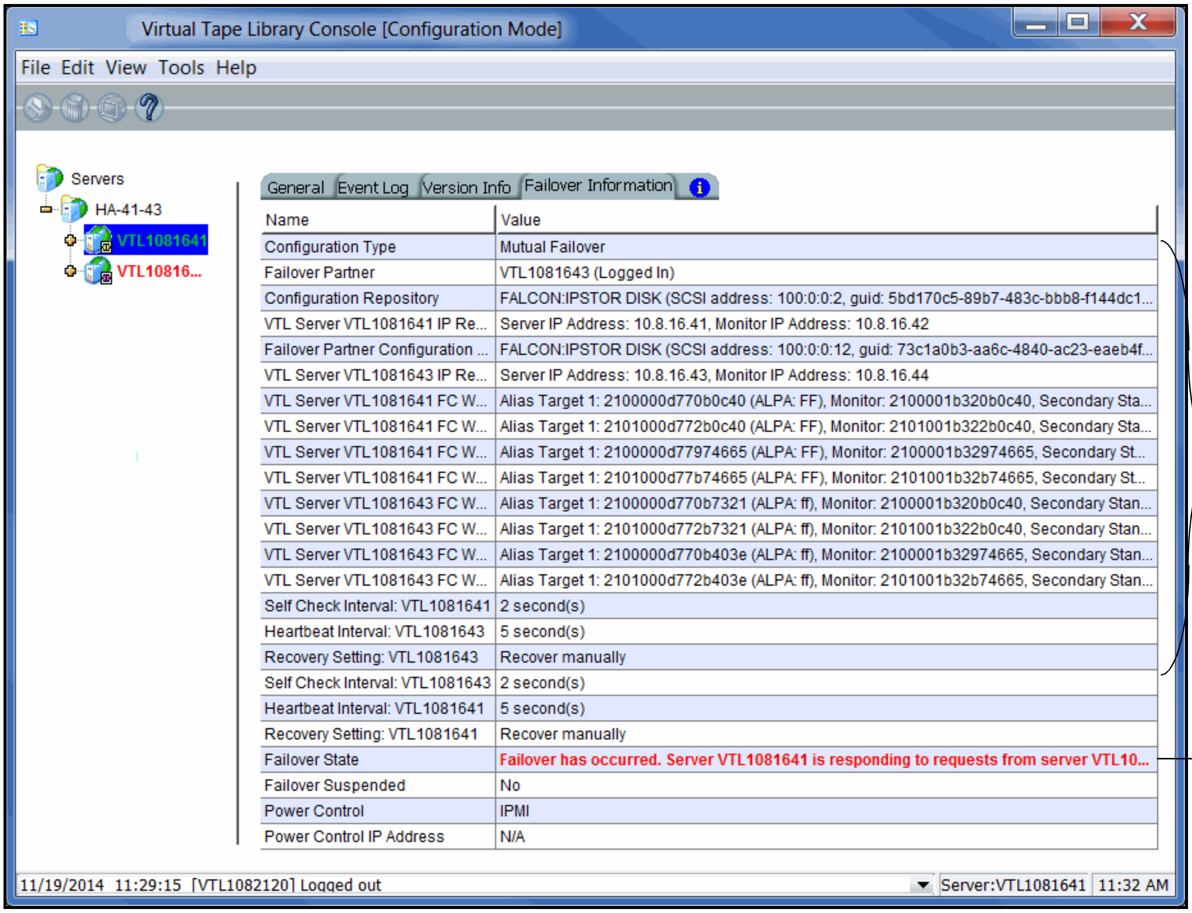

Name	Value
Configuration Type	Mutual Failover
Failover Partner	VTL1081643 (Logged In)
Configuration Repository	FALCON:IPSTOR DISK (SCSI address: 100:0:0:2, guid: 5bd170c5-89b7-483c-bbb8-f144dc1...
VTL Server VTL1081641 IP Re...	Server IP Address: 10.8.16.41, Monitor IP Address: 10.8.16.42
Failover Partner Configuration ...	FALCON:IPSTOR DISK (SCSI address: 100:0:0:12, guid: 73c1a0b3-aa6c-4840-ac23-eaeb4f...
VTL Server VTL1081643 IP Re...	Server IP Address: 10.8.16.43, Monitor IP Address: 10.8.16.44
VTL Server VTL1081641 FC W...	Alias Target 1: 2100000d770b0c40 (ALPA: FF), Monitor: 2100001b320b0c40, Secondary Sta...
VTL Server VTL1081641 FC W...	Alias Target 1: 2101000d772b0c40 (ALPA: FF), Monitor: 2101001b322b0c40, Secondary Sta...
VTL Server VTL1081641 FC W...	Alias Target 1: 2100000d77974665 (ALPA: FF), Monitor: 2100001b32974665, Secondary St...
VTL Server VTL1081641 FC W...	Alias Target 1: 2101000d77b74665 (ALPA: FF), Monitor: 2101001b32b74665, Secondary St...
VTL Server VTL1081643 FC W...	Alias Target 1: 2100000d770b7321 (ALPA: ff), Monitor: 2100001b320b0c40, Secondary Stan...
VTL Server VTL1081643 FC W...	Alias Target 1: 2101000d772b7321 (ALPA: ff), Monitor: 2101001b322b0c40, Secondary Stan...
VTL Server VTL1081643 FC W...	Alias Target 1: 2100000d770b403e (ALPA: ff), Monitor: 2100001b32974665, Secondary Stan...
VTL Server VTL1081643 FC W...	Alias Target 1: 2101000d772b403e (ALPA: ff), Monitor: 2101001b32b74665, Secondary Stan...
Self Check Interval: VTL1081641	2 second(s)
Heartbeat Interval: VTL1081643	5 second(s)
Recovery Setting: VTL1081643	Recover manually
Self Check Interval: VTL1081643	2 second(s)
Heartbeat Interval: VTL1081641	5 second(s)
Recovery Setting: VTL1081641	Recover manually
Failover State	Failover has occurred. Server VTL1081641 is responding to requests from server VTL10...
Failover Suspended	No
Power Control	IPMI
Power Control IP Address	N/A

```
11/19/2014  11:29:15  [VTL1082120] Logged out          Server:VTL1081641  11:32 AM
```

Failover settings, including which IP addresses are being monitored for failover.

Current status of failover configuration
.

In addition, VTL displays the server name in different colors to indicate the failover status:

- Black - Normal operations.
- Red - The server is currently in failover mode and has been taken over by the secondary server.
- Green - The server is currently in failover mode and has taken over the primary server's resources.
- Yellow - The user has suspended failover on this server. The current server will NOT take over the primary server's resources even if it detects an abnormal condition from the primary server.

Failover events are also written to the primary server's Event Log, so you can check there for status and operational information, as well as any errors. You should be aware that when a failover occurs, the console will show the failover partner’s Event Log for the server that failed.

Recover from failover

To recover a server after failover, perform the following steps:

1. Identify and fix the condition that caused the failure.
2. Restart the server.
3. Run the `sms` command to confirm the server is in a *Ready* state.
4. Check storage and network connectivity.

   During failback, the server must receive confirmation from its partner in order to recover its role as the primary server.

   If there is a communication problem between servers, the server does not receive confirmation and remains in a *Ready* state but does not recover its role as the primary server.
5. Select the *Stop Takeover* option from the console on the partner server to trigger recovery.

Change your failover configuration

The first time you set up your failover configuration, the secondary server cannot have any logical resources (including NAS resources, virtual tapes, drives, or libraries) or clients assigned to it. Afterwards, you may want to add resources, create virtual devices, and assign clients to the server.

In order to make any changes to a mutual failover configuration, you must be running the console with write access to both servers. VTL will automatically "log on" to the failover pair when you attempt any configuration on the failover set. While it is not required that both servers have the same user name and password, the system will try to connect to both servers using the same user name and password. If the servers have different usernames/passwords, it will prompt you to enter them before you can continue.

Notes:

- Once failover is configured, you cannot change the host name or IP addresses of the failover nodes.
- Once failover is configured for NAS-only servers, you cannot enable the VTL role on either server.
- When performing failover operations, you should only use a single console.

Change server hardware

If you make a change to a physical device (such as if you add a network card that will be used for failover), you will need to disable and then reconfigure failover.

If you make hardware changes on a server in a failover setup, the partner server will not be aware of these changes until you select the *Rescan* option. This will synchronize the hardware configuration.

Change the power control password

If you need to change the password that is used for power control, right-click the failover group and select *Failover* --> *Power Control* --> *Change Power Control Password.*

Start/stop failover or recovery

Start takeover

Right-click the server and select *Failover* --> *Start Takeover* to take over the other server. You may want to do this if you are taking a server offline, such as when you will be performing maintenance on it.

Stop takeover

Right-click the server and select *Failover* --> *Stop Takeover* if you want to manually initiate recovery to your primary server.

Suspend/resume failover

Right-click the failover group and select *Failover* --> *Suspend* to stop monitoring your servers. The server name will be displayed in yellow while failover is suspended. The server will NOT take over the primary server's resources even it detects an abnormal condition from the primary server.

Right-click the failover group and select *Failover* --> *Resume* to restart the monitoring.

Note: While failover is suspended, changes to the configuration of a library, drive, or client on either failover server will not be synchronized with the partner. Therefore, it is important to not change your library, drive, or client configuration while failover is suspended.

Disable failover

Right-click the failover group and select *Failover* --> *Disable* to disable failover.

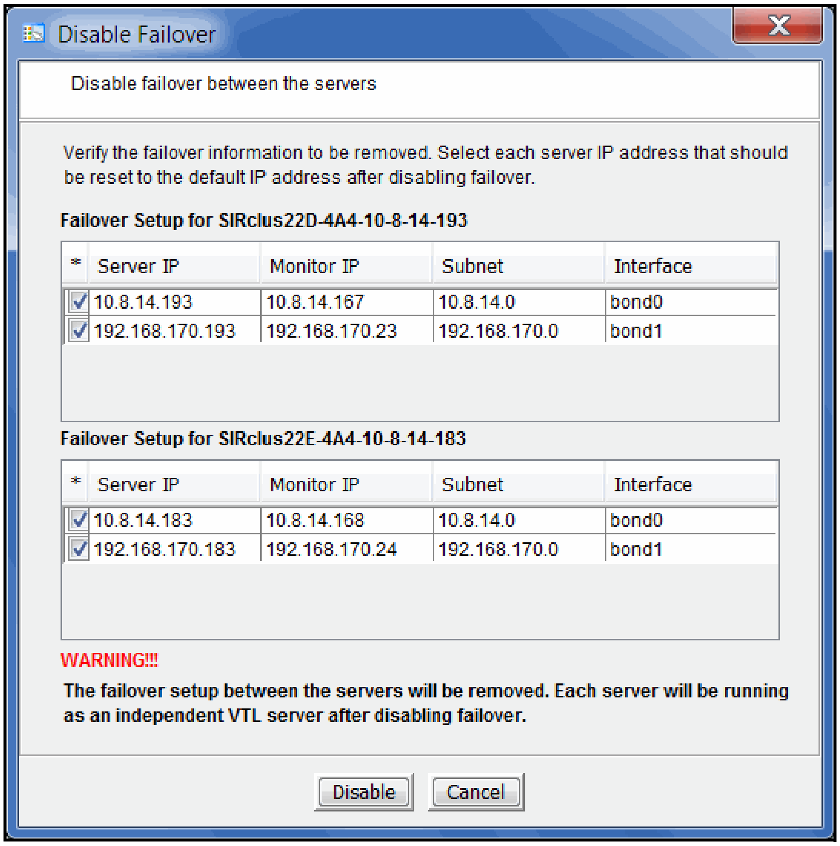

Disable failover between the servers

Verify the failover information to be removed. Select each server IP address that should be reset to the default IP address after disabling failover.

Failover Setup for SIRclus22D-4A4-10-8-14-193

*	Server IP	Monitor IP	Subnet	Interface
☑	10.8.14.193	10.8.14.167	10.8.14.0	bond0
☑	192.168.170.193	192.168.170.23	192.168.170.0	bond1

Failover Setup for SIRclus22E-4A4-10-8-14-183

*	Server IP	Monitor IP	Subnet	Interface
☑	10.8.14.183	10.8.14.168	10.8.14.0	bond0
☑	192.168.170.183	192.168.170.24	192.168.170.0	bond1

WARNING!!!

The failover setup between the servers will be removed. Each server will be running as an independent VTL server after disabling failover.

If everything is checked, this eliminates the failover relationship and removes the health monitoring IP addresses from the servers and restores the server IP addresses. If you uncheck the IP address(es) for a server, the health monitoring address becomes the server IP address.

If you want to remove the heartbeat IP addresses, restart the network to clear them.

Redundant Node Failover

Overview

Redundant node failover provides high availability for your deduplication servers.

In a redundant node configuration, a single standby server monitors the health of all servers in a deduplication cluster. This is referred to as “N+1” failover.

If an active server fails, the standby server sends a command to shut down the failed server. The standby server then assumes the identity of the failed server and the cluster continues to function normally.

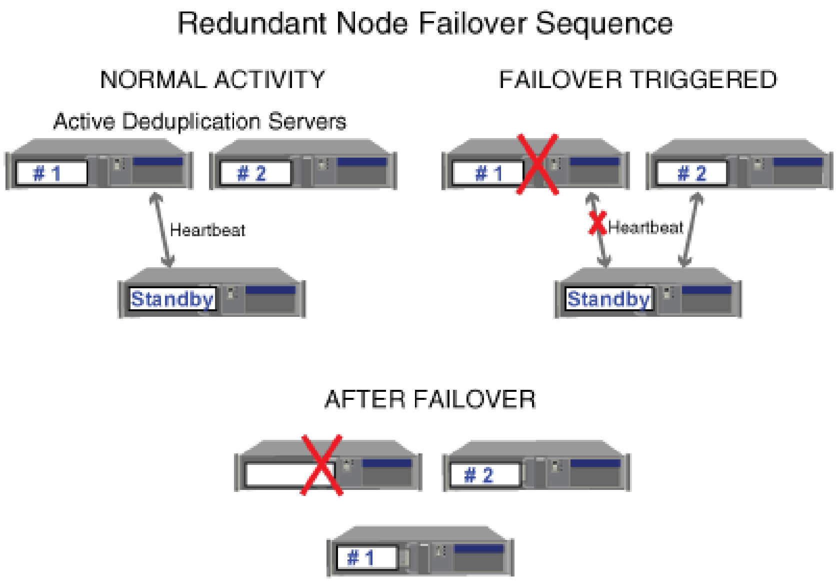

When a failed server is ultimately repaired and is turned back on, it assumes the identity of the standby server and begins checking the heartbeats of active servers.

Failover terminology

Active server

An active server is a functioning deduplication server. The health of each active server in a cluster is monitored by the standby server.

Standby server

(Also known as a *redundant server*) The job of the standby server is to monitor the health of the active servers in a cluster. If an active server fails to respond to the standby server, failover occurs.

Heartbeat

The IP address that identifies each server and is used to verify the health of each server.

Failover

(Also known as *Takeover*) Failover is the process that occurs when the standby server takes over the identity of an active server.

Failover triggers

Failover will occur under the following conditions when there is a heartbeat failure:

- Server lockup of active node
- Network failure of active node
- Deduplication node manager is stopped

Note: Automatic failover will not occur if IPMI or HP iLO is not accessible.

If the standby server cannot access an active node and it cannot determine the node's power status, it will not automatically take over the failed node. If failover is desired, it can be manually initiated with the `sirtakeover.sh` command. However, in order to prevent potential data integrity problems, you MUST confirm that the failed node has been shut down before forcing a takeover.

If you want to avoid triggering failover when you need to perform storage maintenance, such as updating the firmware, you should suspend failover. You can resume failover when you are done.

Failover behavior

During failover, any currently running deduplication jobs will fail. The jobs will restart at their next scheduled time.

If data is being restored from a deduplicated tape at the time redundant node failover occurs, the restore job will fail and will need to be restarted afterward.

We recommend running a file system consistency check on any server that has experienced a failure.

Server behavior during and after redundant node failover

When a failover occurs, the failed server is powered off by the standby server. The standby server starts the takeover process and assumes the identity (hostname, primary IP addresses, FC targets, and FC initiators) of the failed server. All server IP addresses remain the same as viewed from the console (excluding heartbeat IP addresses).

If you are in the console and are connected to the servers at the time they become unavailable, you will see the tree collapse and each server will be marked as "Disconnected" in the right pane.

Once the standby server completes the takeover process, it makes one attempt to turn the failed server back on. If the failed server is successfully powered back on, services will start automatically and the failed server will assume the name and identity of the standby server. If the failed server is not successfully powered on, you must repair the server and turn it back on.

When the failed server is ultimately repaired and is turned back on, it assumes the name and identity of the standby server. All server IP addresses remain the same (excluding heartbeat IP addresses). The new standby assumes the responsibility of checking heartbeats of active servers. The *General* tab for the deduplication server shows the current hostname and the original hostname after failover occurred.

If it is necessary to replace the failed node with new hardware, contact Technical Support.

Redundant node failover requirements

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

- You must have a standby deduplication server. The server must have an SIR role.
- All of the deduplication servers in the cluster, including the standby server, must be identical (same make, model, and BIOS version) with the same licenses and FC target mode enabled.
- The QLogic BIOS option that allows the system to boot from a FC LUN must be disabled. This is the default setting for QLogic cards shipped by DSI.
- In addition to a standard IP address, each server must have a heartbeat IP address (Identifies each server and is used to verify the health of each server) and a power control IP address (used by power control software).
- All IP addresses (heartbeat and power control) on the active server and the standby server need to be on the same subnet on the same NIC interface.
- If you are protecting multiple IP addresses, each protected IP address must be in its own subnet.
- The role of each port must be the same on each server. For example, if port 1 connects to storage on the active server, port 1 must connect to storage on the standby server.
- Each server must use the exact same Target Port ID scheme on the matching WWPNs. We recommend using the same initiator adapter numbers on both sides to connect to the same storage, so the ACSL of all the devices will look identical on both sides.
- All servers must be configured with a power control management interface on the motherboard. IPMI (Intelligent Platform Management Interface) and HP iLO are supported interfaces.
- If your storage binds to the World Wide Node Name (WWNN) **and** the World Wide Port Name (WWPN), it may be necessary to set the WWNN so that when failover occurs, the takeover server can maintain its storage connection. This is done in the *fshba.conf* file by adding a “wwnn” line for each Fibre Channel port on each node in your N+1 cluster. All “wwnn” lines can have identical node names for all ports and all nodes. For example, if you have four ports, you might add the following:

```
wwnn0=2100010001000100
wwnn1=2100010001000100
wwnn2=2100010001000100
wwnn3=2100010001000100
```

- If you want Email Alerts enabled on the standby server in a redundant node failover configuration, you must enable it on the standby server before redundant node failover is configured.

Notes:

- If you have upgraded from a version of VTL prior to v7.0, you will need to increase the size of your configuration repository from 500 MB to 4 GB for each active node. To reconfigure, right-click the *Configuration Repository* object, select *Reconfigure*, and select the new physical resource.
- Once failover is configured, you cannot change the host name or IP addresses of the cluster nodes.
- Once failover is configured, you cannot change the roles of your HBA ports.

Configure redundant node failover

To configure failover for a deduplication cluster:

1. Right-click the cluster and select *Deduplication* --> *Redundant Node* --> *Add*.
2. Select the standby server.

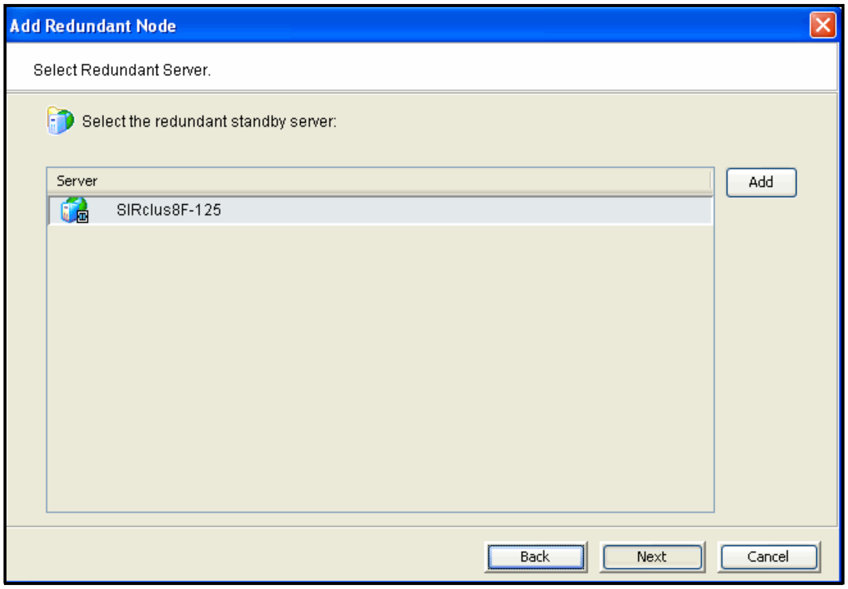

Add Redundant Node

Select Redundant Server.

Select the redundant standby server:

Server
SIRclus8F-125

3. Verify the type of power control management interface your hardware is using.

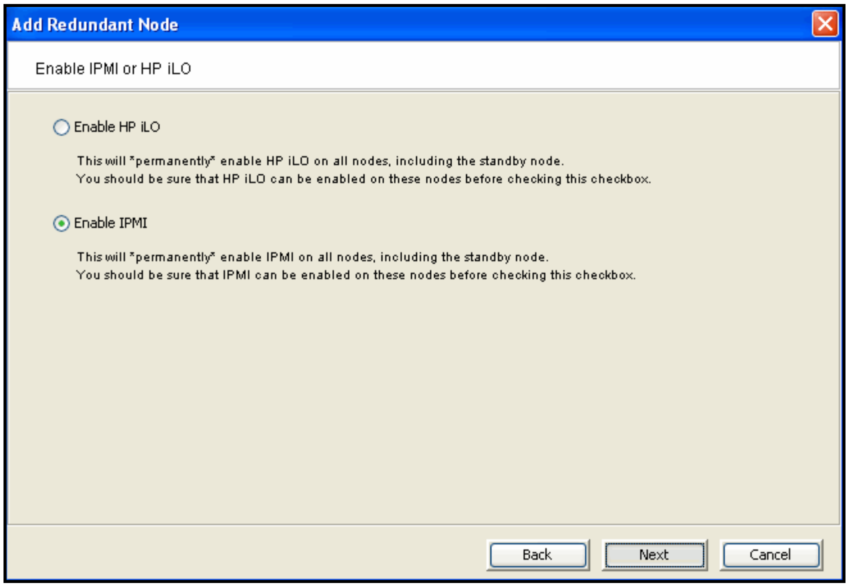

Add Redundant Node

Enable IPMI or HP iLO

○ Enable HP iLO

This will *permanently* enable HP iLO on all nodes, including the standby node.
You should be sure that HP iLO can be enabled on these nodes before checking this checkbox.

● Enable IPMI

This will *permanently* enable IPMI on all nodes, including the standby node.
You should be sure that IPMI can be enabled on these nodes before checking this checkbox.

The system automatically selects the correct one for you.

When you click *Next*, the driver will be loaded on each node.

4. Select which network interfaces or bond groups you want to protect.

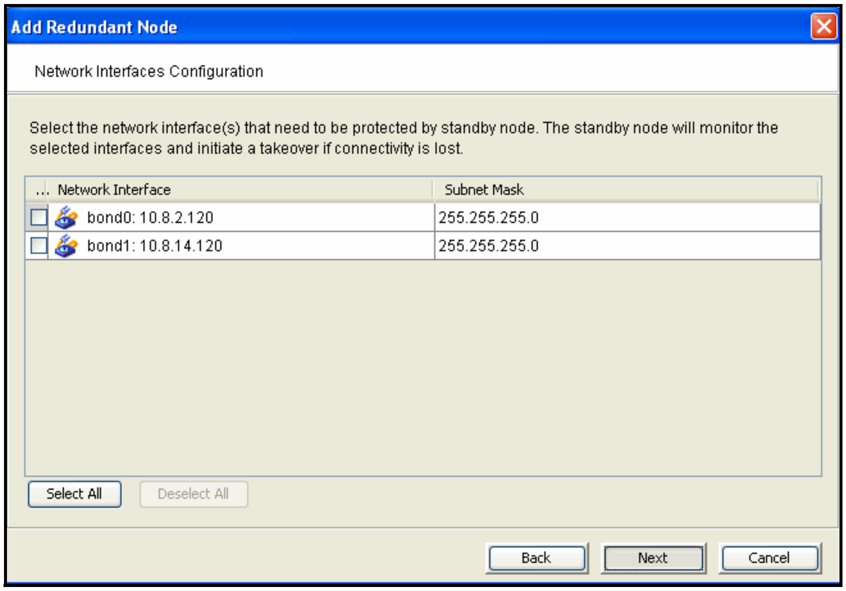

Network Interfaces Configuration

Select the network interface(s) that need to be protected by standby node. The standby node will monitor the selected interfaces and initiate a takeover if connectivity is lost.

...	Network Interface	Subnet Mask
□	bond0: 10.8.2.120	255.255.255.0
□	bond1: 10.8.14.120	255.255.255.0

5. Select the network interface or bond group that should be monitored as the heartbeat channel.

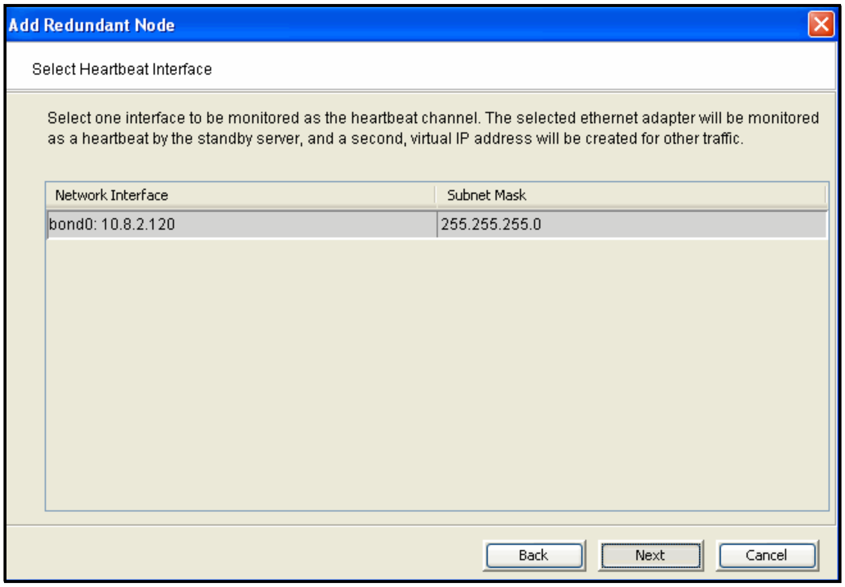

Select Heartbeat Interface

Select one interface to be monitored as the heartbeat channel. The selected ethernet adapter will be monitored as a heartbeat by the standby server, and a second, virtual IP address will be created for other traffic.

Network Interface	Subnet Mask
bond0: 10.8.2.120	255.255.255.0

6. Enter power control information.

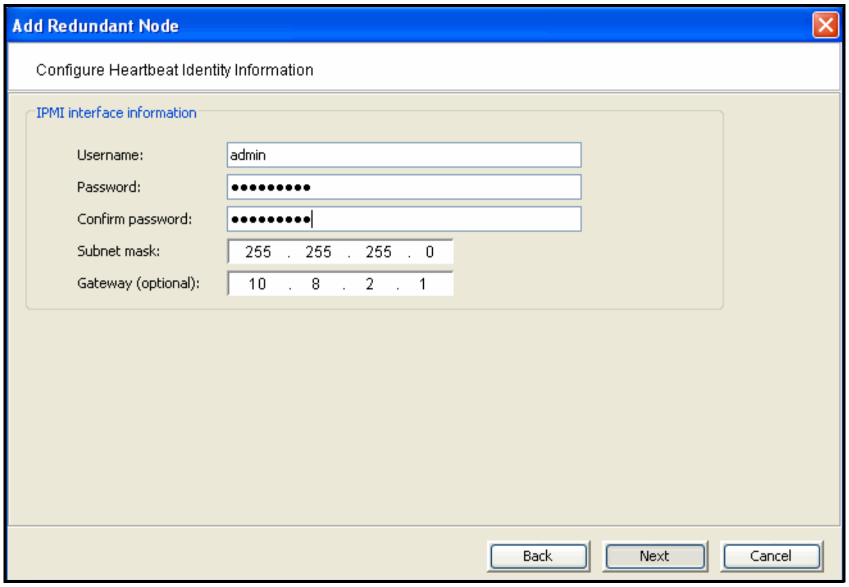

Configure Heartbeat Identity Information

IPMI interface information

Username:	admin
Password:	•••••••••
Confirm password:	•••••••••
Subnet mask:	255 . 255 . 255 . 0
Gateway (optional):	10 . 8 . 2 . 1

On DSI appliances, the IPMI user name and password are preset as *admin* and *falcon101*.

The user name and password you specify here will overwrite the existing information.

7. Enter the power control IP address and the heartbeat address.

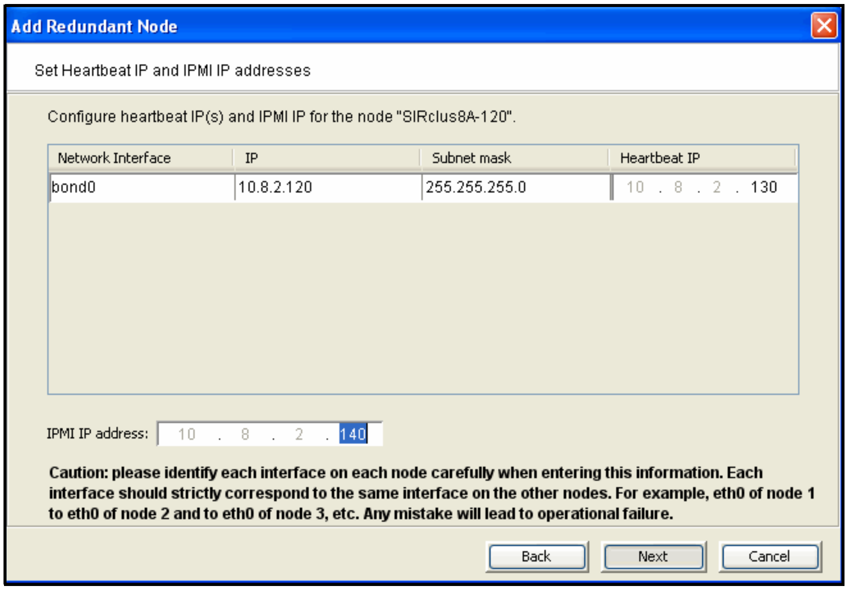

Set Heartbeat IP and IPMI IP addresses

Configure heartbeat IP(s) and IPMI IP for the node "SIRclus8A-120".

Network Interface	IP	Subnet mask	Heartbeat IP
bond0	10.8.2.120	255.255.255.0	10 . 8 . 2 . 130

IPMI IP address: 10 . 8 . 2 . 140

Caution: please identify each interface on each node carefully when entering this information. Each interface should strictly correspond to the same interface on the other nodes. For example, eth0 of node 1 to eth0 of node 2 and to eth0 of node 3, etc. Any mistake will lead to operational failure.

The subnet information is automatically pre-filled.

You will have to enter IP addresses for each server in the cluster.

8. When you have finished configuring all IP addresses, confirm all of the information and then click *Finish* to create the failover configuration.

Start nodes in deduplication cluster servers

Once redundant node failover has been configured, you must make sure the standby server meets one of the following conditions before starting any active node in a cluster:

- The standby server is powered on with services started.
- The standby server is powered off but power control is accessible. If power control cannot be accessed, contact DSI Technical Support.

Check redundant node failover status

After failover has been configured, the active servers and the standby server will be shown in the same cluster.

Servers marked with an A are active members while the server marked with an S is the standby server.

The server names will be displayed in yellow if failover has been suspended.

To view detailed information, including server, heartbeat, and IPMI addresses for each of the servers that are part of failover configuration, select the *Redundant Node* tab. The alias IP is listed here. During failover, this can be used to see which servers have changed their identity.

Change the power control password

If you need to change the password that is used for power control, right-click the cluster and select *Deduplication --> Redundant Node --> Change Power Control Password.*

Suspend/resume redundant node failover

When failover is suspended, the standby server will NOT assume the identity of a failed server. You need to do this when you patch the server or if you have to replace storage.

To suspend failover, right-click the cluster and select *Deduplication --> Redundant Node --> Suspend.*

The server names will be displayed in yellow while failover is suspended.

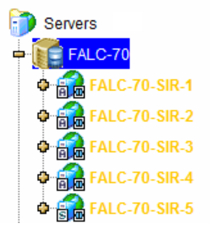

To resume failover, right-click the cluster and select *Deduplication --> Redundant Node --> Resume.*

Remove a redundant node failover configuration

To remove a failover configuration, right-click the cluster and select *Deduplication --> Redundant Node --> Remove.*

Data Replication

Replication protects data by maintaining a copy of the data on the same VTL server or on another VTL server.

There are several methods for replicating data in VTL; four provide automatic replication and one is a manual process:

Method	Description
Virtual tape replication	Replicates *changed* data from a primary virtual tape to the same server or another server at prescribed intervals, based on user defined policies.
Deduplicated tape replication	Replicates *changed* data from a primary VIT to another server at prescribed intervals, based on user-defined policies (defined within a deduplication policy). If advanced replication (cascaded or parallel) is configured, data can be replicated to an additional target server.
NAS resource replication	Replicates stub files of deduplicated files based on user-defined policies.
Auto replication	Replicates the contents of a single tape whenever a virtual tape is ejected from a virtual library and moved to the virtual vault (manually or by backup software).
Remote copy	Replicates the contents of a single tape *on demand.*

Note: In order to configure replication to a target server with a VTL version that is higher than the source server, you must use a console installed with the higher version.

Replication of virtual tapes

Replication is a process that protects the information on a virtual tape by maintaining a copy of a virtual tape on the same VTL server or on another VTL server.

At prescribed intervals, when the tape is not in use, changed data from the *primary* virtual tape on the source server is transmitted to the *replica resource* on the target server so that they are synchronized. The target VTL server is usually located at a remote location. The backup server does not have access to the replica resource on the target server.

If a disaster occurs and the replica is needed, the administrator can *promote* the replica to become the primary virtual tape so that backup servers can access it.

VTL offers two types of replication: *Remote Replication* and *Local Replication*.

Remote Replication

Remote Replication allows fast data synchronization of storage volumes from one VTL server to another over the IP network.

With Remote Replication, the replica disk is located on a separate target VTL server.

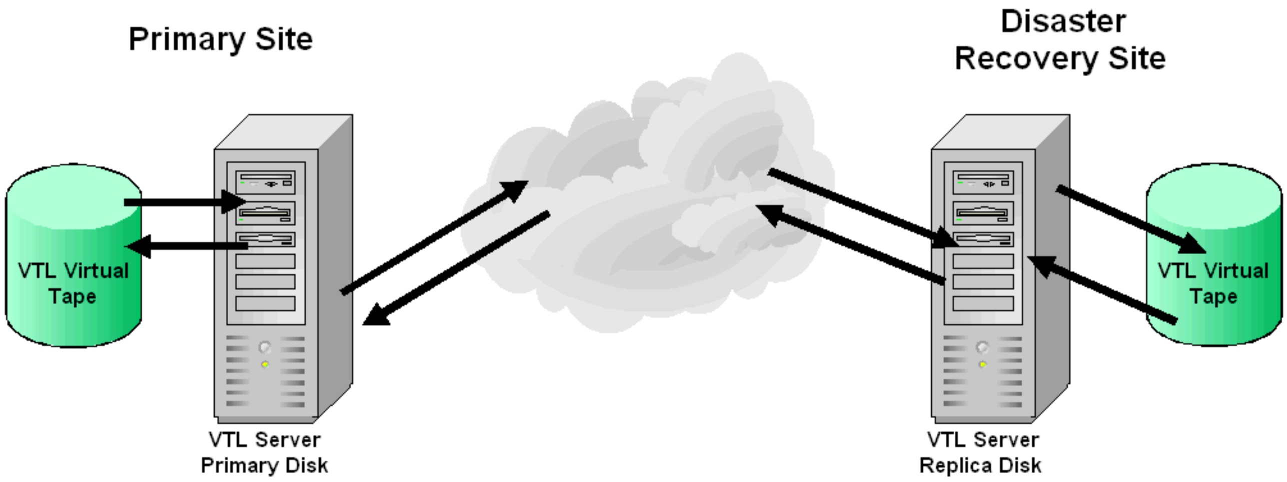

Local Replication

Local Replication allows fast, data synchronization of storage volumes within one VTL server. Because there is only one VTL server, the primary and target servers are the same server.

Local Replication can be used to maintain a local copy of virtual tape data or it can be used to maintain a remote copy within metropolitan area Fibre Channel SANs.

With Local Replication, the replica disk can be connected to the VTL server via a gateway using edge routers or protocol converters.

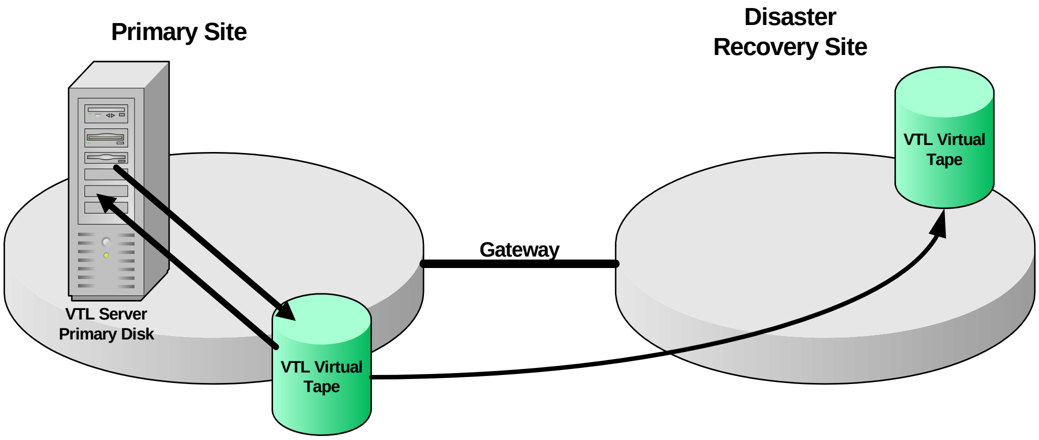

Replication requirements for virtual tapes

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

General requirements

- You must have enough space on the target server for the replica resource.

Remote replication requirements

- You must have two VTL servers.
- You must have administrative rights on both servers.
- If a virtual tape is encrypted, encryption must be enabled on the target server; the key used by source tape must exist on both servers and be identical. This means that the keys have the same name and were created using the same secret phrase. If the secret phrase is not the same, you can export a key from the source server and import it to the target.

Configure replication for virtual tapes

You can configure replication at the tape level or for a whole virtual tape library.

Note: If you need to change the IP address of your VTL appliance, you must do so before configuring replication.

1. Right-click one or more virtual tapes in a virtual tape library or in the virtual vault and select *Replication* --> *Add*.

   You can also right-click the virtual tape library and select *Replication* --> *Add.*

   Each virtual tape can only have one replica resource.

   Note: If you get a message that Replication cannot be enabled because *Auto Archive/Auto Replication* is enabled, you must first disable *Auto Archive/Auto Replication* for the tape. To do this, right-click the tape (or virtual tape library for all tapes), select *Properties,* and go to the *Auto Archive/Replication* tab.

2. Indicate whether you want to use remote replication or local replication.

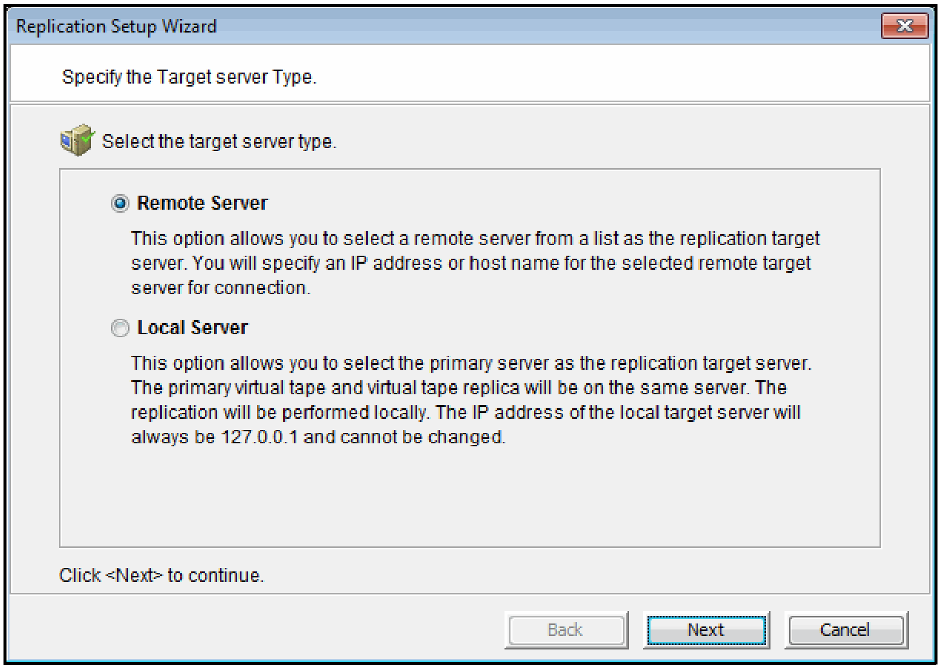

Replication Setup Wizard

Specify the Target server Type.

Select the target server type.

● **Remote Server**

This option allows you to select a remote server from a list as the replication target server. You will specify an IP address or host name for the selected remote target server for connection.

○ **Local Server**

This option allows you to select the primary server as the replication target server. The primary virtual tape and virtual tape replica will be on the same server. The replication will be performed locally. The IP address of the local target server will always be 127.0.0.1 and cannot be changed.

Click <Next> to continue.

3. Select the server that will contain the replica.

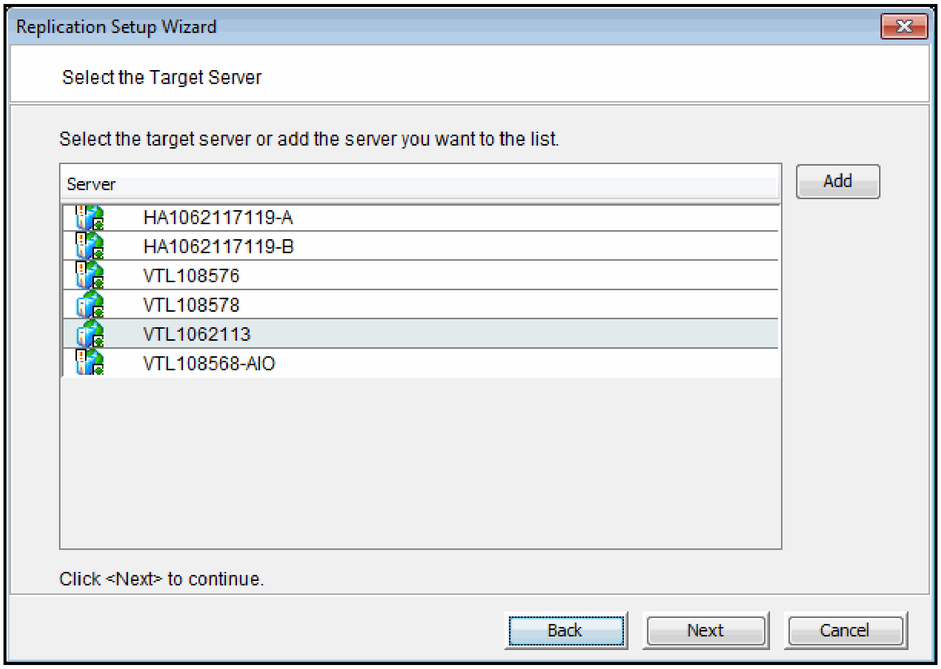

Replication Setup Wizard

Select the Target Server

Select the target server or add the server you want to the list.

Server
HA1062117119-A
HA1062117119-B
VTL108576
VTL108578
VTL1062113
VTL108568-AIO

Click <Next> to continue.

If the server you want does not appear on the list, click the *Add* button.

4. Confirm/enter the target server's IP address.

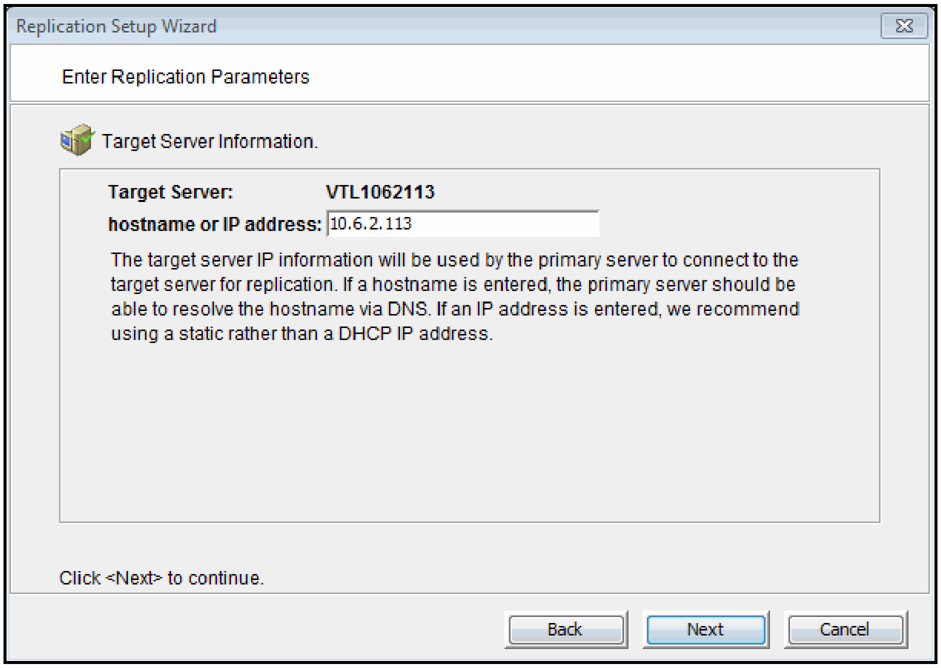

Replication Setup Wizard

Enter Replication Parameters

Target Server Information.

Target Server: **VTL1062113**

hostname or IP address: 10.6.2.113

The target server IP information will be used by the primary server to connect to the target server for replication. If a hostname is entered, the primary server should be able to resolve the hostname via DNS. If an IP address is entered, we recommend using a static rather than a DHCP IP address.

Click <Next> to continue.

5. (Tape caching is enabled for these tapes) Configure when replication should occur.

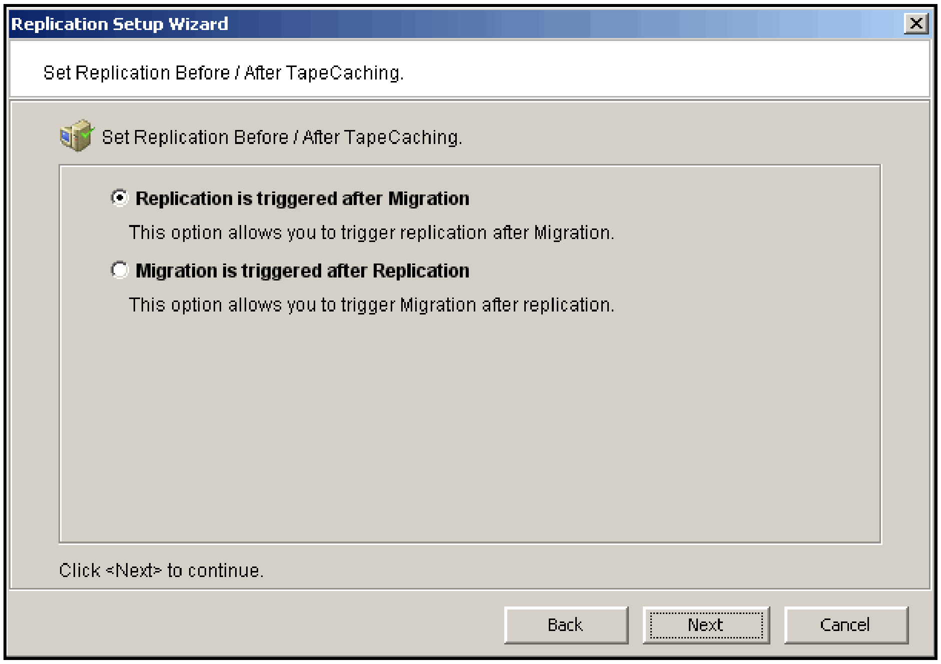

Replication Setup Wizard

Set Replication Before / After TapeCaching.

Set Replication Before / After TapeCaching.

● **Replication is triggered after Migration**

This option allows you to trigger replication after Migration.

○ **Migration is triggered after Replication**

This option allows you to trigger Migration after replication.

Click <Next> to continue.

6. (Tape caching is not enabled for these tapes) Configure how often, and under what circumstances, replication should occur.

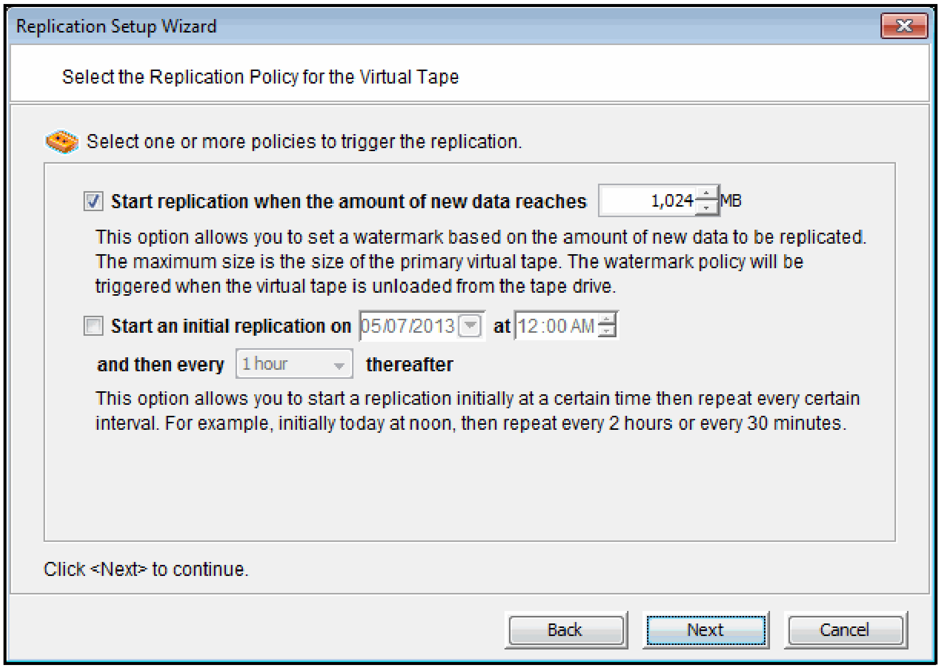

Select the Replication Policy for the Virtual Tape

Select one or more policies to trigger the replication.

☑ **Start replication when the amount of new data reaches** 1,024 MB

This option allows you to set a watermark based on the amount of new data to be replicated. The maximum size is the size of the primary virtual tape. The watermark policy will be triggered when the virtual tape is unloaded from the tape drive.

☐ **Start an initial replication on** 05/07/2013 **at** 12:00 AM

and then every 1 hour **thereafter**

This option allows you to start a replication initially at a certain time then repeat every certain interval. For example, initially today at noon, then repeat every 2 hours or every 30 minutes.

Click <Next> to continue.

You must select at least one policy, but you can have multiple.

Start replication when the amount of new data reaches - If you enter a watermark value, when the value is reached, replication of the changed data will begin as soon as the virtual tape is ejected from the tape drive after backup.

Start an initial replication on mm/dd/yyyy at hh:mm and then every n hours/ minutes thereafter - Indicate when replication should begin and how often it should be repeated.

If replication is already occurring when the next time interval is reached, the new replication request will be ignored.

7. Indicate what to do if a replication attempt fails.

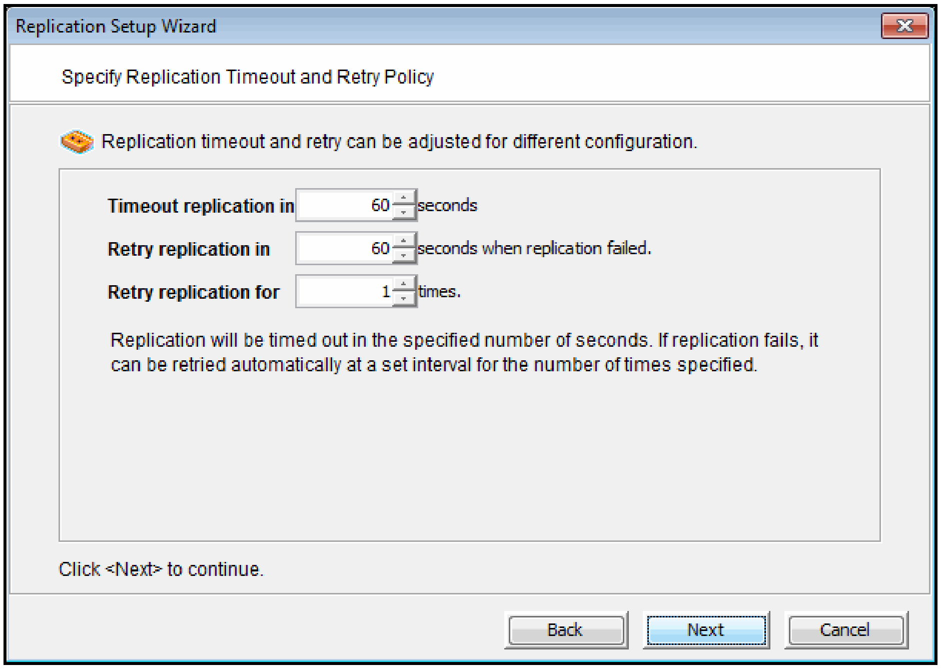

Replication Setup Wizard

Specify Replication Timeout and Retry Policy

Replication timeout and retry can be adjusted for different configuration.

Timeout replication in 60 seconds

Retry replication in 60 seconds when replication failed.

Retry replication for 1 times.

Replication will be timed out in the specified number of seconds. If replication fails, it can be retried automatically at a set interval for the number of times specified.

Click <Next> to continue.

Replication can only occur when the virtual tape is not in a tape drive. Indicate how long the system should attempt to replicate data before timing out and how often it should attempt to retry before skipping a scheduled replication.

8. (Remote Replication only) Indicate if you want to use *Compression* and/or *Encryption*.

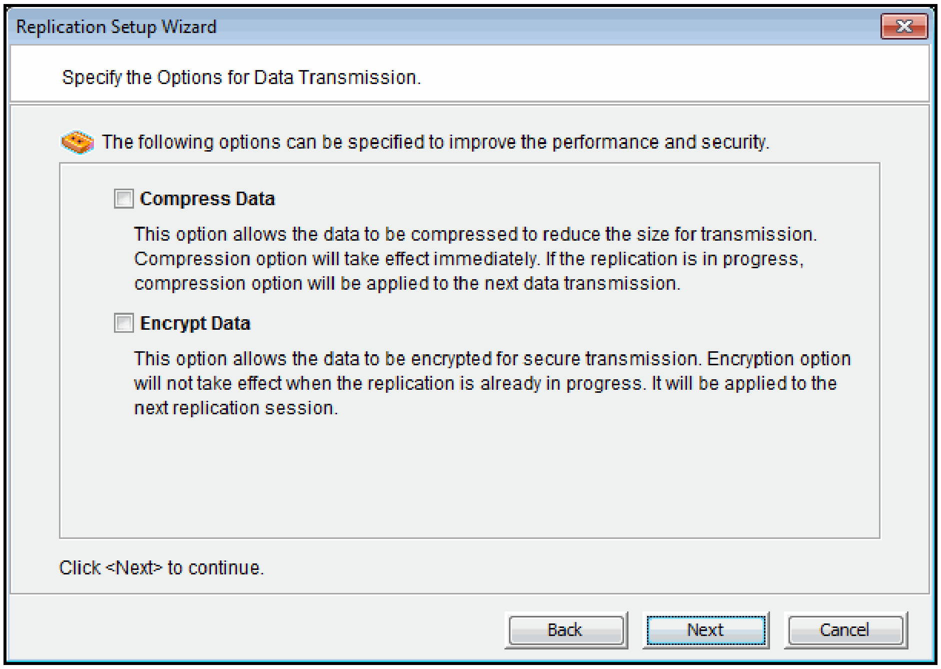

Replication Setup Wizard

Specify the Options for Data Transmission.

The following options can be specified to improve the performance and security.

□ **Compress Data**

This option allows the data to be compressed to reduce the size for transmission. Compression option will take effect immediately. If the replication is in progress, compression option will be applied to the next data transmission.

□ **Encrypt Data**

This option allows the data to be encrypted for secure transmission. Encryption option will not take effect when the replication is already in progress. It will be applied to the next replication session.

Click <Next> to continue.

Compression and encryption are only available for virtual tapes that are not using virtual tape encryption. If you are enabling replication for multiple tapes (some encrypted, some not), compression and encryption during replication will reduce overall performance.

The *Compression* option provides enhanced throughput during replication by compressing the data stream.

The *Encryption* option 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.

9. (Local Replication only) Enter a name for the replica resource.

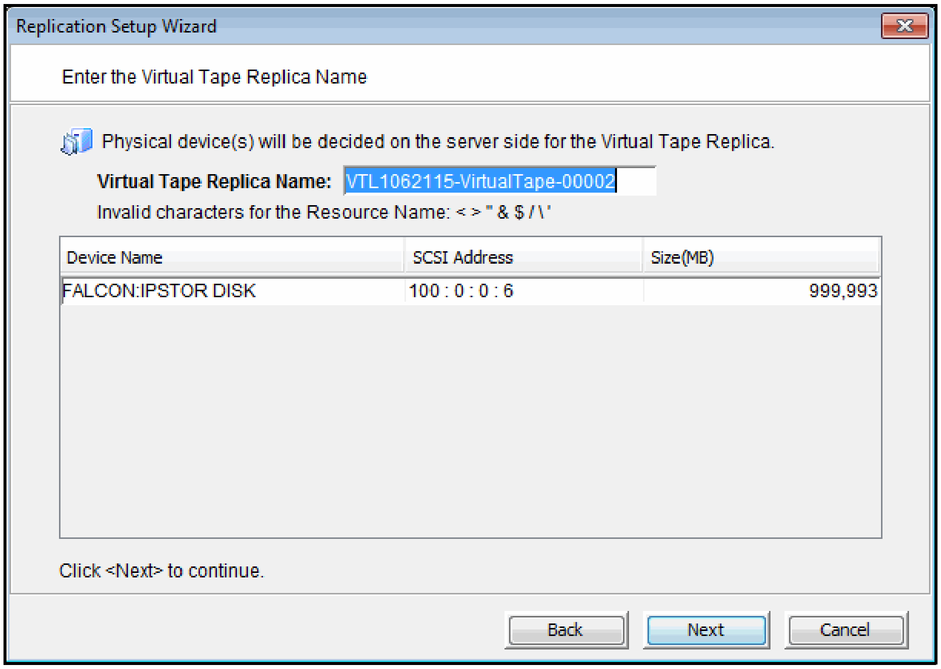

Enter the Virtual Tape Replica Name

Physical device(s) will be decided on the server side for the Virtual Tape Replica.

Virtual Tape Replica Name: VTL1062115-VirtualTape-00002

Invalid characters for the Resource Name: < > " & $ / \ '

Device Name	SCSI Address	Size(MB)
FALCON:IPSTOR DISK	100 : 0 : 0 : 6	999,993

Click <Next> to continue.

The name is not case sensitive.

10. Confirm that all information is correct and then click *Finish* to create the replication configuration.

Set replication throttling for virtual tapes

You can set global replication options that affect available network bandwidth during replication on VTL servers. If throttling is not used, replication will use the maximum bandwidth that is available.

1. Right-click a VTL server and select *Properties*.
2. On the *Performance* tab, enable replication throttling and then enter the maximum number of KBs per second that should be used for bandwidth.

   Transmission will not exceed the set value. This is a global server parameter and affects all virtual tapes.

   Once enabled, the default is 10 KBs per second. Besides 0, valid input is 10-1,000,000 KB/s (1G).

Check replication status for virtual tapes

There are several ways to check replication status.

Replication tab (source server)

The *Replication* tab of the primary virtual tape displays information about the target replica server, the policies set for replication, and the replication status.

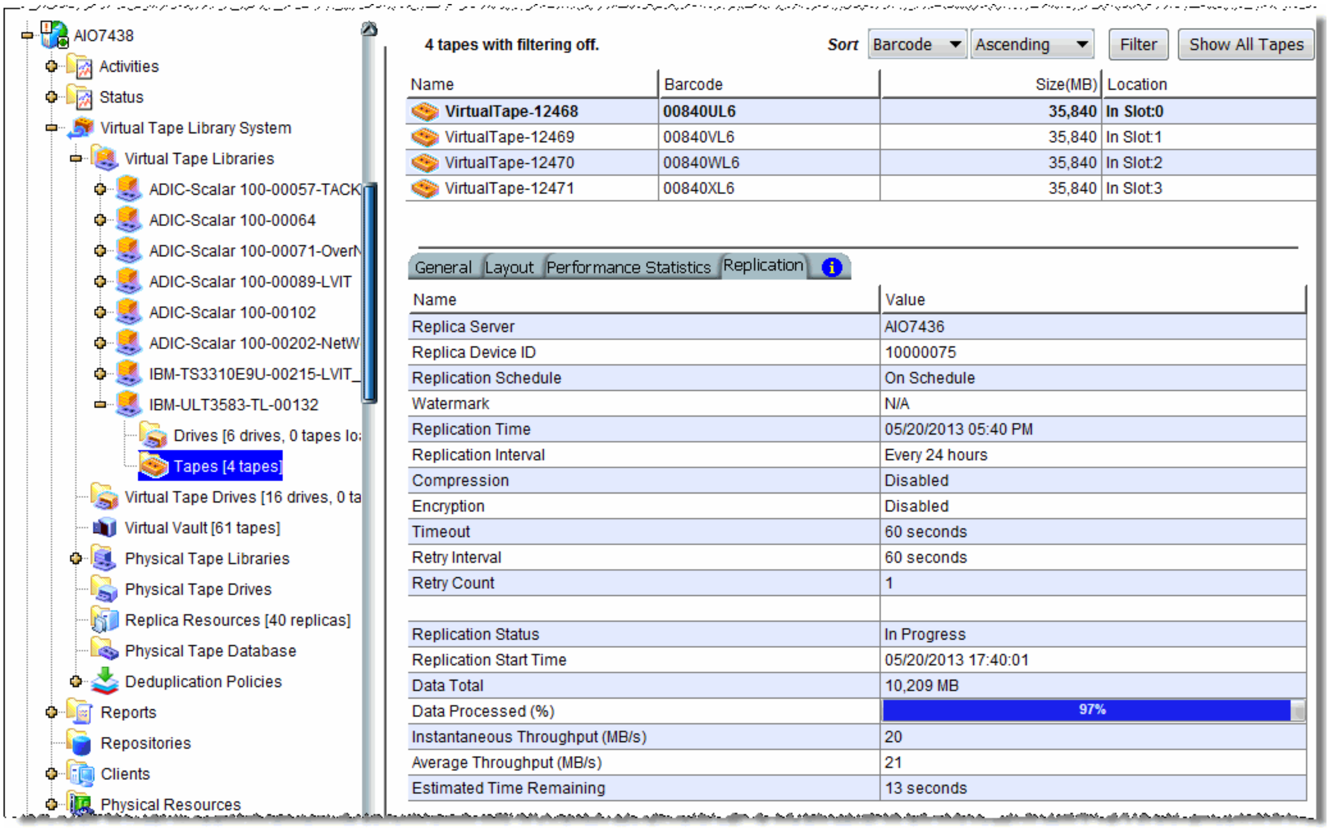

4 tapes with filtering off.

Name	Barcode	Size(MB)	Location
VirtualTape-12468	**00840UL6**	**35,840**	**In Slot:0**
VirtualTape-12469	00840VL6	35,840	In Slot:1
VirtualTape-12470	00840WL6	35,840	In Slot:2
VirtualTape-12471	00840XL6	35,840	In Slot:3

Name	Value
Replica Server	AIO7436
Replica Device ID	10000075
Replication Schedule	On Schedule
Watermark	N/A
Replication Time	05/20/2013 05:40 PM
Replication Interval	Every 24 hours
Compression	Disabled
Encryption	Disabled
Timeout	60 seconds
Retry Interval	60 seconds
Retry Count	1
Replication Status	In Progress
Replication Start Time	05/20/2013 17:40:01
Data Total	10,209 MB
Data Processed (%)	97%
Instantaneous Throughput (MB/s)	20
Average Throughput (MB/s)	21
Estimated Time Remaining	13 seconds

Replication Status Report (source server)

The Replication Status Report (run from the *Reports* object) provides a centralized view for displaying real-time replication status for all virtual tapes enabled for replication. It can be generated for an individual tape, multiple tapes, source server or target server, for any range of dates. This report is useful for administrators managing multiple servers that either replicate data or are the recipients of replicated data. Refer to ‘Replication Status’ for more information.

***Replica Resources* object (target server)**

The *Replica Resources* object on the target server displays the status of replication jobs. Note that in order to check if a replica is encrypted you must check the soure tape.

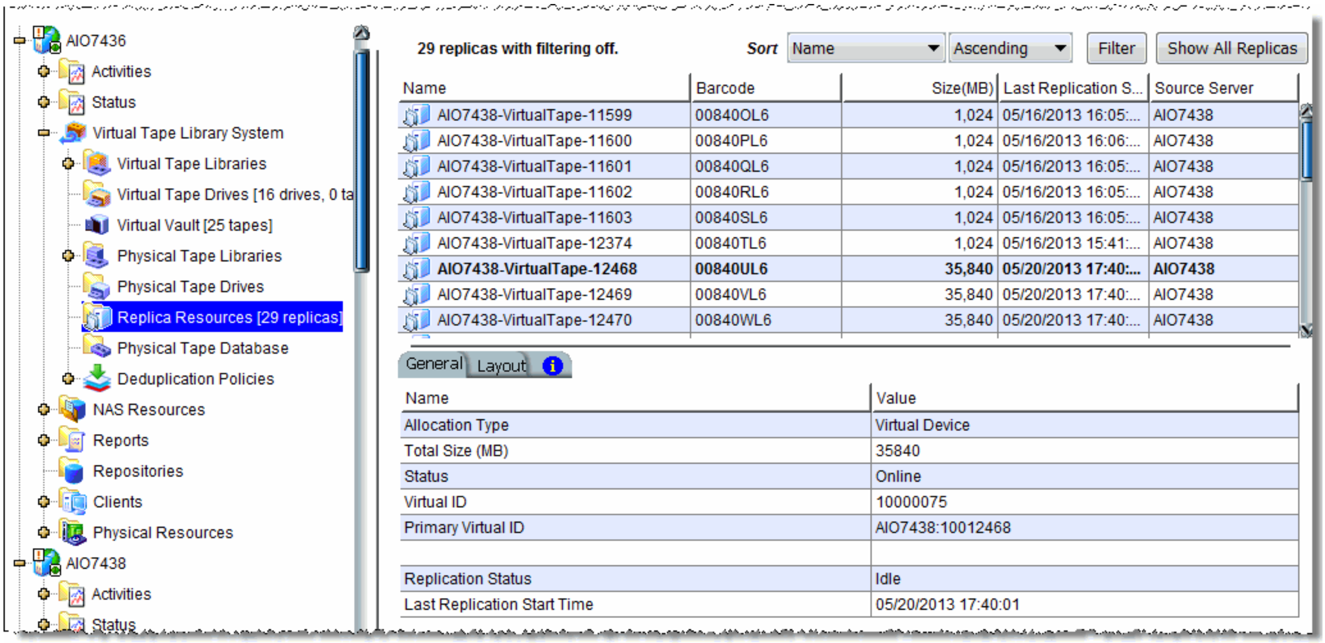

29 replicas with filtering off.

Sort Name Ascending

Name	Barcode	Size(MB)	Last Replication S...	Source Server
AIO7438-VirtualTape-11599	00840OL6	1,024	05/16/2013 16:05:...	AIO7438
AIO7438-VirtualTape-11600	00840PL6	1,024	05/16/2013 16:06:...	AIO7438
AIO7438-VirtualTape-11601	00840QL6	1,024	05/16/2013 16:05:...	AIO7438
AIO7438-VirtualTape-11602	00840RL6	1,024	05/16/2013 16:05:...	AIO7438
AIO7438-VirtualTape-11603	00840SL6	1,024	05/16/2013 16:05:...	AIO7438
AIO7438-VirtualTape-12374	00840TL6	1,024	05/16/2013 15:41:...	AIO7438
AIO7438-VirtualTape-12468	**00840UL6**	**35,840**	**05/20/2013 17:40:...**	**AIO7438**
AIO7438-VirtualTape-12469	00840VL6	35,840	05/20/2013 17:40:...	AIO7438
AIO7438-VirtualTape-12470	00840WL6	35,840	05/20/2013 17:40:...	AIO7438

Name	Value
Allocation Type	Virtual Device
Total Size (MB)	35840
Status	Online
Virtual ID	10000075
Primary Virtual ID	AIO7438:10012468
Replication Status	Idle
Last Replication Start Time	05/20/2013 17:40:01

Promote a virtual tape replica resource

If a replica resource is needed, the administrator can *promote* the replica to become a usable virtual tape. After promotion, the virtual tape is put into the virtual vault so that you can move it to any virtual library on *that* server (formerly the target server). If you need to get the virtual tape back to the former primary server, you must replicate it back to that server.

Promoting a replica resource breaks the replication configuration. Once a replica resource is promoted, it cannot revert back to a replica resource.

You must have a valid replica resource in order to promote it. For example, if a problem occurred (such as a transmission problem or the replica resource failing) during the first and only replication, the replicated data would be compromised and therefore could not be promoted to a primary virtual tape.

You cannot promote a replica resource while replication is in progress.

1. Locate the target server, right-click the appropriate replica resource, and select *Replication* --> *Promote*.
2. Confirm the promotion and click *OK*.
3. From the backup server, rescan devices or restart the backup server to see the promoted virtual tape.

Promote a virtual tape replica resource without breaking a replication configuration

Under normal circumstances, when replica storage is needed, the administrator promotes the replica to become a usable virtual tape, thereby breaking the replication configuration.

However, there may be times, such as for disaster recovery testing, when you want to promote replica storage *without* breaking the replication configuration.

When you promote a replica without breaking the replication configuration, you will have a *read-only* version of the tape on the replica server. This tape can then be used for testing or for file recovery.

You must have a valid replica storage in order to promote it. For example, if a problem occurred (such as a transmission problem or the replica storage failing) during the first and only replication, the replicated data would be compromised and therefore could not be promoted to a primary virtual tape.

You cannot promote replica storage while replication is in progress.

1. Locate the target server, right-click the appropriate replica resource and select *Replication* --> *Test Mode Promote*.
2. Confirm the promotion and click *OK*.

Change your virtual tape replication configuration options

You can change the following for your replication configuration:

- Static IP address of your target server
- Policies that trigger replication (watermark, interval, time)
- Timeout and retry policies
- Data transmission options (encryption, compression)

To change the configuration:

1. Right-click the primary virtual tape and select *Replication* --> *Properties*.
2. Make the appropriate changes and click *OK*.

Suspend/resume virtual tape replication schedule

You can suspend 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. To suspend/resume replication, right-click the primary virtual tape and select *Replication* --> *Suspend* (or *Resume*).

You can see the current settings by checking the *Replication Schedule* field on *Replication* tab of the primary virtual tape.

Start/stop replication of a virtual tape

To force replication that is not scheduled, select *Replication* --> *Synchronize*.

To stop replication of a virtual tape that is currently in progress, right-click the primary virtual tape and select *Replication* --> *Stop.*

Note that you do not need to stop an active replication job so that a backup can occur. When a virtual tape is mounted in a virtual tape drive, the active replication job will automatically be cancelled so that the backup application can write to the tape. Replication will continue when the next replication trigger occurs.

Replication of virtual tapes and failover

If replication is in progress and a failover occurs at the same time, the replication will stop. After failover, replication will start at the next normally scheduled interval. This is also true in reverse, if replication is in progress and a recovery occurs at the same time.

If you use the console to manually initiate take over of a primary server and then stop the take over while a replication job is running, you will need to “stop” the replication job before restarting it.

Remove a virtual tape replication configuration

This allows you to remove the replication configuration on the primary and either delete or promote the replica resource on the target server at the same time.

1. Right-click the primary virtual tape and select *Replication* --> *Remove*.
2. Determine if you want to promote or delete the replica.
3. If deleting, confirm that you want to remove the replica.

Consolidate virtual tapes from multiple locations to a single data center

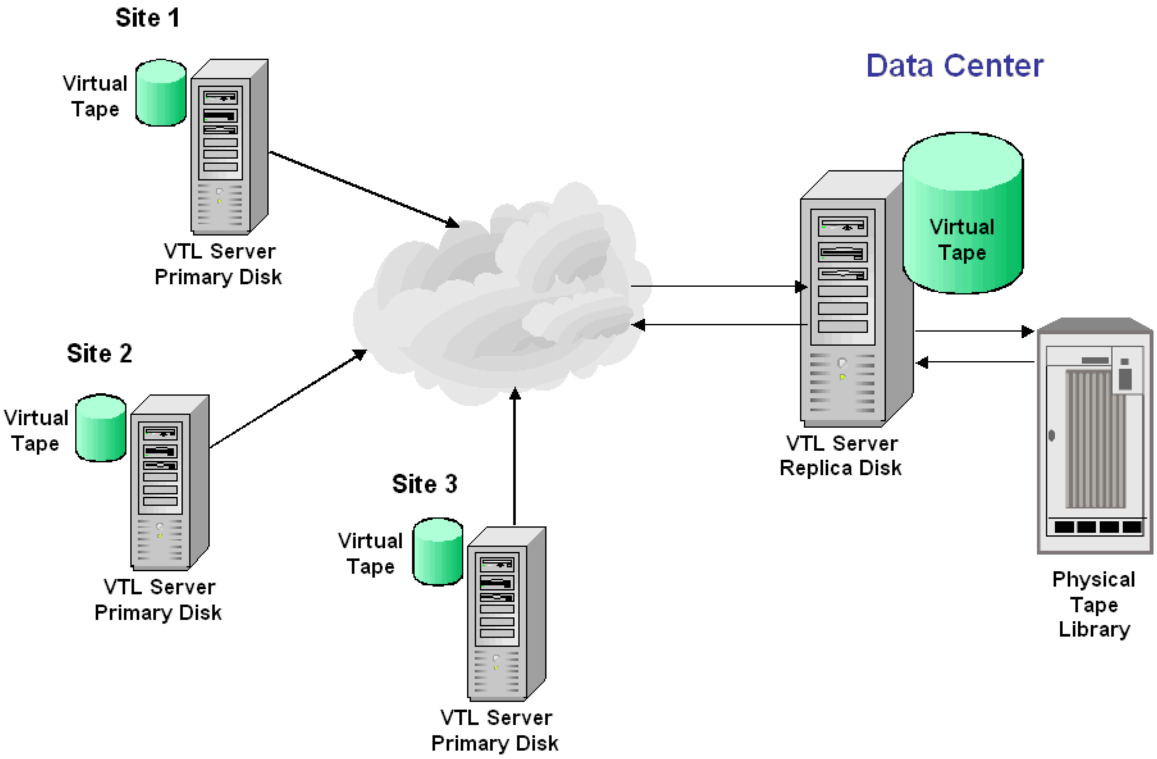

The following information is for environments with multiple VTL locations *without* physical tape libraries that replicate tape data to a remote VTL server that *has* a physical tape library that supports barcodes.

In this environment, if you will be exporting tapes from the remote VTL server to the physical tape library, you want to make sure that when you create tapes on the primary servers (at the multiple VTL locations *without* physical tape libraries), you match the barcodes of the tapes on the physical library attached to the target server.

Replication of deduplicated tapes

When you create a deduplication policy, you can configure replication for the tapes in the policy. If you do this for all tapes in all deduplication policies, you effectively replicate the entire deduplication repository.

Replication from the source server to the target server occurs via TCP (and optionally via Fibre Channel). The target server is located at a remote location. If advanced replication (cascaded or parallel) is configured, data can be replicated to an additional remote location.

If there is more than one FC connection between sites, VTL will automatically use all of the paths to transmit the data. If one path goes down, VTL will detect it automatically and use the remaining paths to continue transmitting data.

Each cluster can contain one, two, or four deduplication nodes. The clusters can each have a different number of deduplication nodes. For example, a two-node cluster can replicate to a one-node cluster, and vice versa.

Replication can occur in several configurations:

- One-to-one configuration - one cluster replicates to one cluster.
- One-to-many configuration - the source server replicates data to multiple targets.
    - *Cascaded* replication - the source server replicates data to Replication Target 1, which in turn replicates data to Replication Target 2.
    - *Parallel* replication - the source server replicates data to two target servers, 1 and 2. This can occur concurrently (replication to both target servers at the same time), or serially (replication first to Replication Target 1 and then to Replication Target 2).
- Many-to-one configuration - multiple clusters replicate to a single cluster, which would be the case where remote offices replicate data to a corporate data center.

Unlike virtual tape replication, where the actual tape data is replicated directly to another virtual tape, replicating a deduplicated tape involves copying the virtual index tape and the missing data from the repository to the target server.

In this way, data duplicated across remote sites is deduplicated at the central site, enabling only globally unique data to be stored.

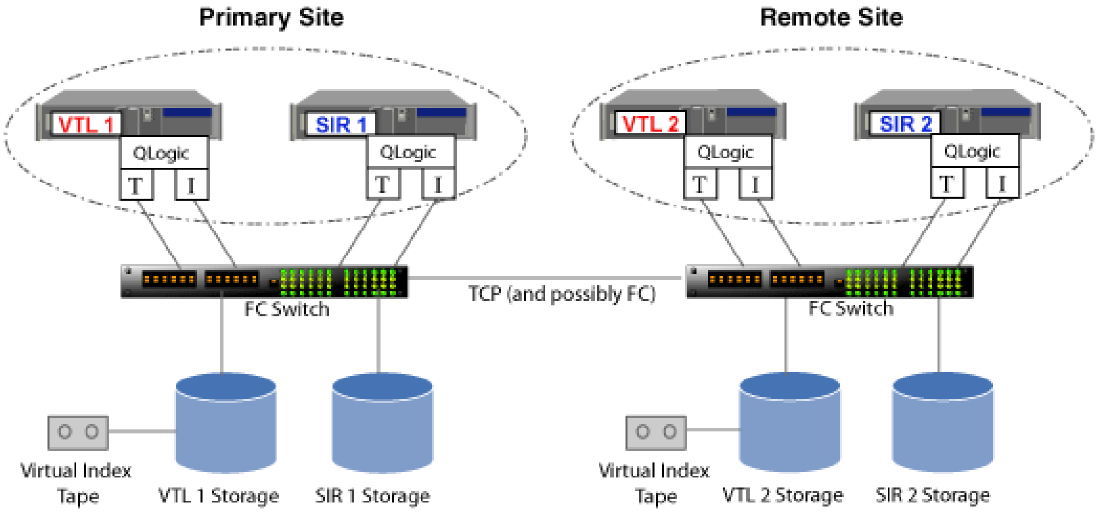

Replication of deduplicated data occurs in several phases, which you can see identified in replication status displays in the VTL console:

> **Note:** Replication will occur only during the period of time you specified when you created the deduplication 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.

- 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), which you can see when you select the *Replica Resources* object of the target server.
- 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.
- 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 virtual vault or in a virtual tape library, depending upon how the deduplication 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 virtual vault or in the virtual tape library. The name of the LVIT corresponds to the name of the FVIT. The image below shows the VTL target server, with the new VITs for replicated data under the *Virtual Vault* object. A green "R" icon indicates that the tape has been successfully resolved. A red "R" icon indicates that the last attempt at resolving the tape failed or the tape is currently being resolved or has

not been resolved.

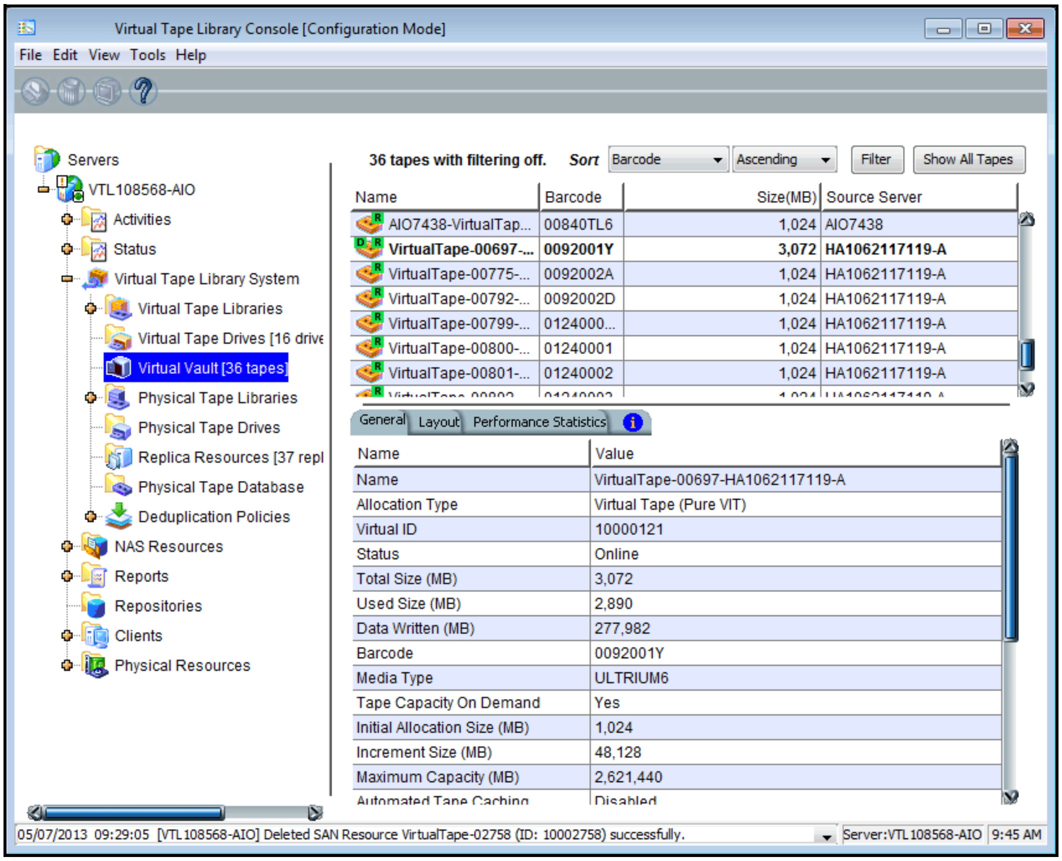

Name	Barcode	Size(MB)	Source Server
AIO7438-VirtualTap...	00840TL6	1,024	AIO7438
VirtualTape-00697-...	**0092001Y**	**3,072**	**HA1062117119-A**
VirtualTape-00775-...	0092002A	1,024	HA1062117119-A
VirtualTape-00792-...	0092002D	1,024	HA1062117119-A
VirtualTape-00799-...	0124000...	1,024	HA1062117119-A
VirtualTape-00800-...	01240001	1,024	HA1062117119-A
VirtualTape-00801-...	01240002	1,024	HA1062117119-A

Name	Value
Name	VirtualTape-00697-HA1062117119-A
Allocation Type	Virtual Tape (Pure VIT)
Virtual ID	10000121
Status	Online
Total Size (MB)	3,072
Used Size (MB)	2,890
Data Written (MB)	277,982
Barcode	0092001Y
Media Type	ULTRIUM6
Tape Capacity On Demand	Yes
Initial Allocation Size (MB)	1,024
Increment Size (MB)	48,128
Maximum Capacity (MB)	2,621,440
Automated Tape Caching	Disabled

05/07/2013 09:29:05 [VTL108568-AIO] Deleted SAN Resource VirtualTape-02758 (ID: 10002758) successfully.

Server:VTL108568-AIO 9:45 AM

- The FVITs are listed when you select the *Replica Resources* object. You can sort the FVITs by tape name, barcode, last replication start time, and source server. Replica resources for deduplicated tapes can also be filtered to only display tapes from a specific source server from which replication has occurred.
- On the source server, you can see that replication is complete by checking the *Replication* tab for a virtual tape in a tape library. The *Resolved* field will display *true*.
- 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 deduplication tape drives on the target server and the amount of data on the FVIT.

Replication requirements for deduplicated tapes

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

- You must have at least two VTL servers.
- You must have at least two functional deduplication clusters.
- You must enable replication between the deduplication clusters by adding the target deduplication cluster to the primary cluster.
- You must have administrative rights on all servers.
- You must have enough space on the target server for the replica data.
- Each virtual tape you want to replicate must be included in a deduplication policy.
- Just as you have to associate your deduplication server with your VTL server before you can create deduplication policies, the deduplication server on the target side must be associated with your local deduplication server before you can configure replication.
- *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.
- While you can configure replication for a virtual tape that has not been deduplicated, replication will not run until at least one deduplication has taken place.
- An IP connection is required (even if you are using Fibre Channel for replication).
- On the target server, you must prepare physical resources for deduplication use and enable deduplication.
- Network firewalls should allow access through TCP ports 11782 and 11781 for replication encryption and unencryption.

Configure replication for deduplicated tapes

Note: If you need to change the IP address or host name of your VTL appliance, you must do so before configuring replication for tapes in a deduplication policy.

Overview of steps to configure replication for deduplicated tapes

You must do the following to configure replication with deduplication:

1. Associate each VTL server with a deduplication cluster.

   The source VTL server should be associated with the source deduplication cluster and the target VTL server should be associated with the target deduplication cluster. This is done when you enable deduplication.

2. Add a target deduplication replication server to enable replication between the source and target deduplication servers.

   If you are using Cascaded replication, you must add a target replication server to your primary server and another target replication server to the second server.

   If you are using Parallel replication, you must add both target replication servers to your primary server.

3. Create a deduplication policy and enable replication.

   When you enable replication, you will have to add a VTL target server. This is either the VTL target server that is associated with the *target* deduplication cluster or it is a VTL-S server. Refer to ‘Create tape deduplication policies’ for more information.

Add a target deduplication replication server

If you are using deduplication, before you can configure replication for tapes in a deduplication policy, you must enable replication between the source and target deduplication servers. If the replica is a VTL-S appliance, that appliance is your target.

To do this:

1. Right-click the primary cluster and select *Deduplication* --> *Replication* --> *Add Target*.

2. Select the target cluster or click *Add Target* if the cluster is not listed.

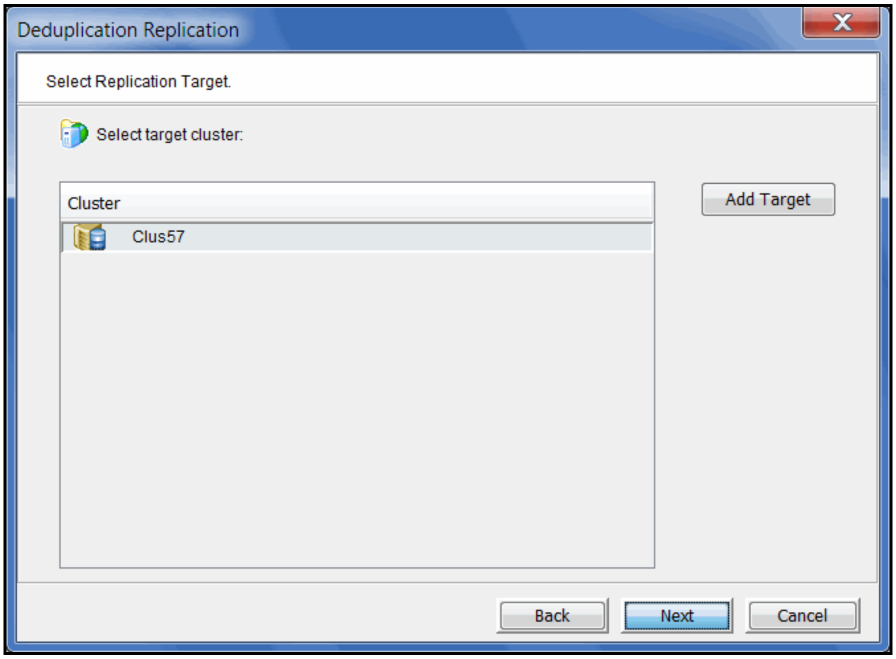

3. Select the protocol to use for replication.

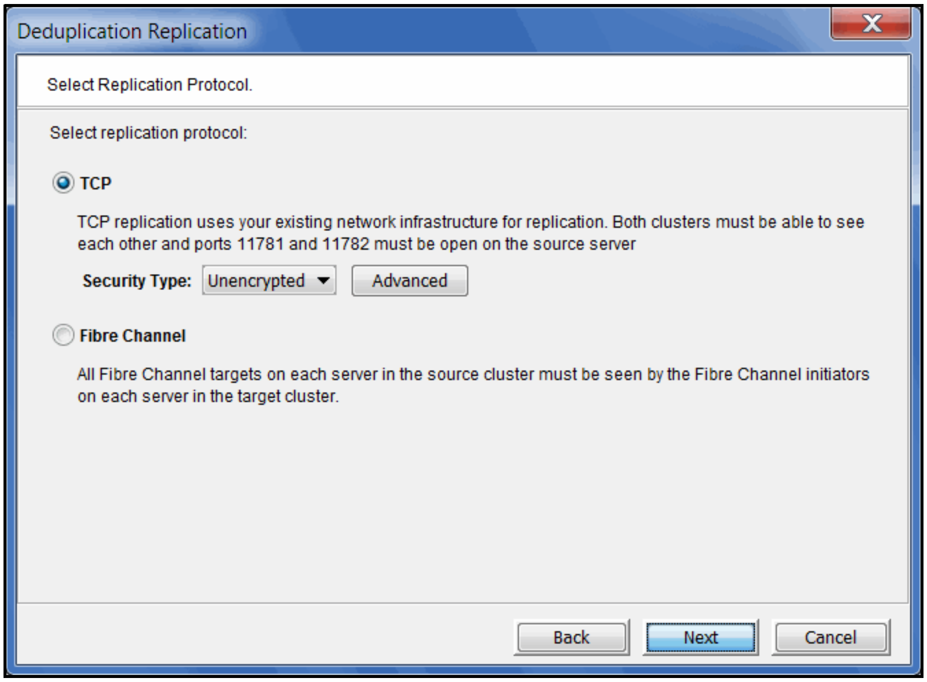

Deduplication Replication
Select Replication Protocol.
Select replication protocol:
Back
Next
Cancel

● TCP

TCP replication uses your existing network infrastructure for replication. Both clusters must be able to see each other and ports 11781 and 11782 must be open on the source server

Security Type: Unencrypted

○ Fibre Channel

All Fibre Channel targets on each server in the source cluster must be seen by the Fibre Channel initiators on each server in the target cluster.

For backward compatibility purposes, if you are replicating to a deduplication server that uses an SIR version prior to 7.50, you will have the older TCP options instead of the ones described below. Refer to the user guide that came with your software version for information about those options.

In order to use TCP replication, both servers must be able to communicate with each other and the following port must be open on the source server:

- Port 11781 - for unencrypted data replication

- Port 11782 - for encrypted data replication
- Port 11780 - if the replica uses SIR software prior to version 7.5

If you select TCP, you can choose to use encryption from the *Security Type* drop-down box. The system uses 256-bit AES encryption.

If you select to use encryption, you can click the *Advanced* button to set additional options that control the maximum number of TCP connections to use and the timeout.

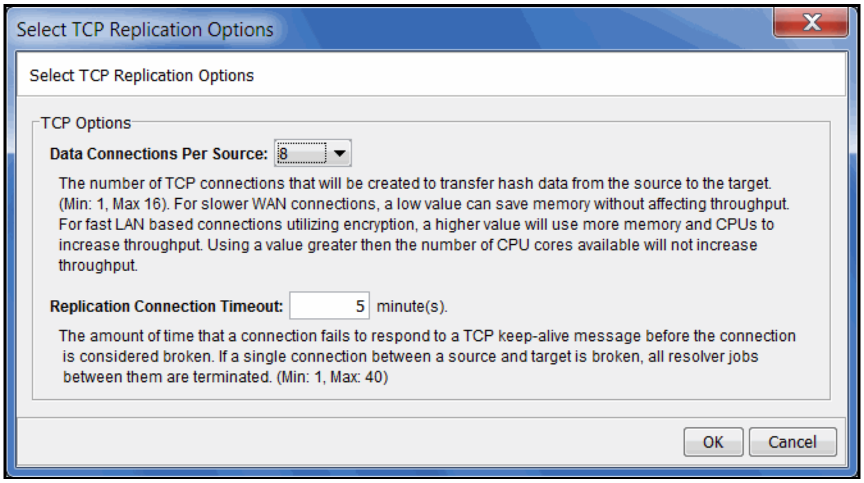
Select TCP Replication Options
X
Select TCP Replication Options
TCP Options
OK
Cancel

Data Connections Per Source: 8

The number of TCP connections that will be created to transfer hash data from the source to the target. (Min: 1, Max 16). For slower WAN connections, a low value can save memory without affecting throughput. For fast LAN based connections utilizing encryption, a higher value will use more memory and CPUs to increase throughput. Using a value greater then the number of CPU cores available will not increase throughput.

Replication Connection Timeout: 5 minute(s).

The amount of time that a connection fails to respond to a TCP keep-alive message before the connection is considered broken. If a single connection between a source and target is broken, all resolver jobs between them are terminated. (Min: 1, Max: 40)

The suggested defaults should be sufficient for most configurations.

Data Connections Per Source - The number of TCP connections that will be created to transfer hash data from the source to the target. (Min: 1, Max: 16). For slower WAN connections, a low value can save memory without affecting throughput. For fast LAN based connections utilizing encryption, a higher value will use more memory and CPUs to increase throughput. Using a value greater than the number of CPU cores available will not increase throughput. The default value is 8.

Replication Connection Timeout - The amount of time that a connection fails to respond before the connection is considered broken. If a single connection between a source and target is broken, all resolver jobs between them are terminated. (Min: 1, Max: 40). The default value is 5.

4. If you selected TCP protocol, select the IP address to use for replication on the source server.

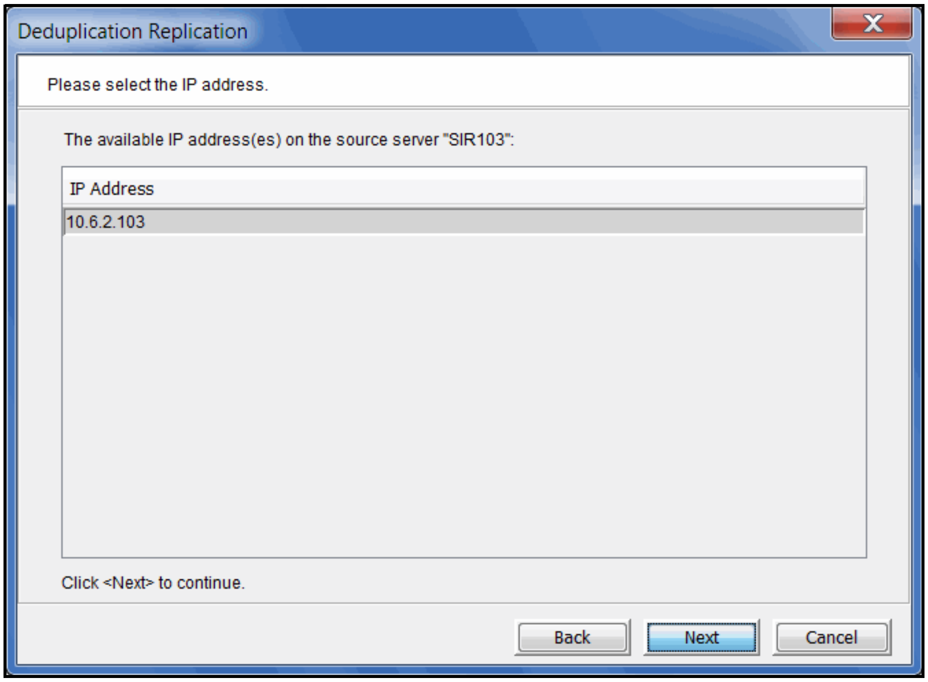

Please select the IP address.

The available IP address(es) on the source server "SIR103":

IP Address
10.6.2.103

Click <Next> to continue.

You will see this dialog for each node (other than a standby redundant node) in the source cluster.

Note: If you are using network address translation (NAT), contact Technical Support before continuing.

5. If you selected TCP protocol, select the IP address to use for replication on the target server.

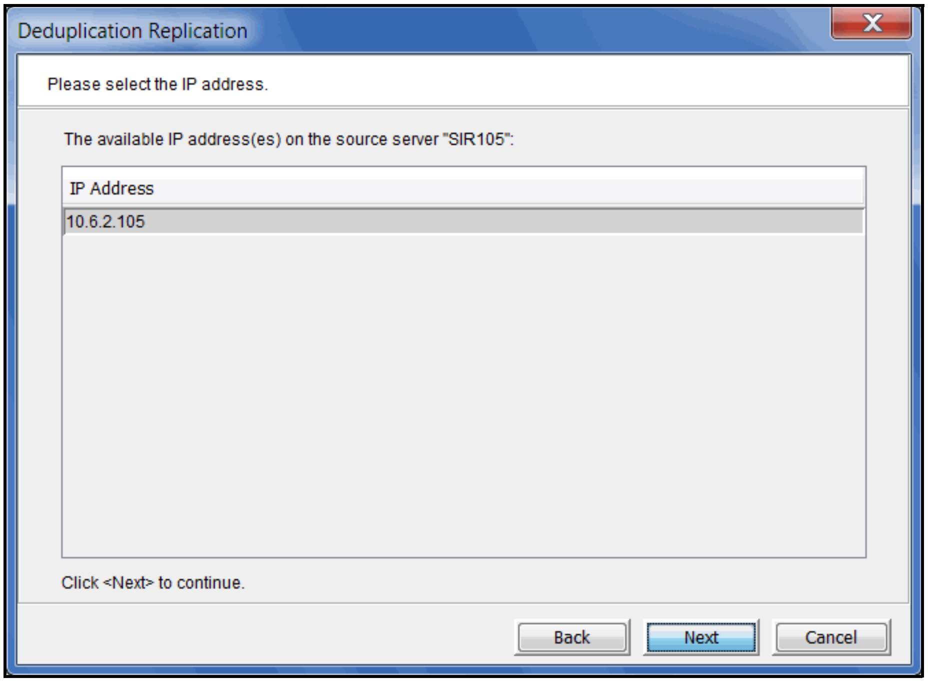

Please select the IP address.

The available IP address(es) on the source server "SIR105":

IP Address
10.6.2.105

Click <Next> to continue.

You will see this dialog for each node in the target cluster.

6. Confirm the information and click *Finish*.

   The target deduplication server is configured to be a replication target.

   You can now configure replication for tapes when you add a deduplication policy. Refer to ‘Create tape deduplication policies’ for more information.

Edit a replication target

If your replication configuration uses TCP protocol, you can change the security level and/or advanced options. To do this:

1. Right-click the cluster object and select *Deduplication* --> *Replication* --> *Edit target*.

2. Select the target.

   If you are not connected to that server, you will have to enter login information for the target server.

3. Change the security level or click *Advanced* to modify advanced options recommended by your network administrator.

4. Confirm the information and click *Finish*.

   The target server is reconfigured. The revised TCP options will apply to new replication jobs.

Set replication throttling for deduplicated tapes

You can set global replication options that affect available network bandwidth used by the VIT resolver on deduplication servers. If throttling is not used, replication will use the maximum bandwidth that is available.

1. Right-click a deduplication cluster and select *Properties*.

2. On the *Performance* tab, enable VIT resolver throttling and then enter the maximum number of KBs per second that should be used for bandwidth.

   You can limit the amount of available network bandwidth that is used for replication of VITs on the source server side or for the VIT resolver on the target deduplication server. Transmission will not exceed the set value. This is a global server parameter and affects all resources.

   Once enabled, the default is 10 KBs per second. Besides 0, valid input is 10-1,000,000 KB/s (1G).

Check replication status for deduplicated tapes

There are several ways to check replication status.

Active Policies tab (source server)

The *Active Policies* tab of a deduplication policy displays information about currently running replication jobs. While replication is occurring, you will see status displays related to the Index and unique replication phases. Refer to ‘Active Policies tab’ for more information.

Deduplication Replication Status Report (source server)

The Deduplication Replication Status Report (run from the *Reports* object) provides a centralized view for displaying replication status for all deduplication policies. Refer to ‘Deduplication Replication Status’ for more information.

Unique Replication Queue tab (target server)

The *Unique Replication Queue* tab (under *Activities* on the target server) displays replication information for deduplicated tapes. It lists the tapes currently replicating (after the index has been replicated) and those awaiting replication.

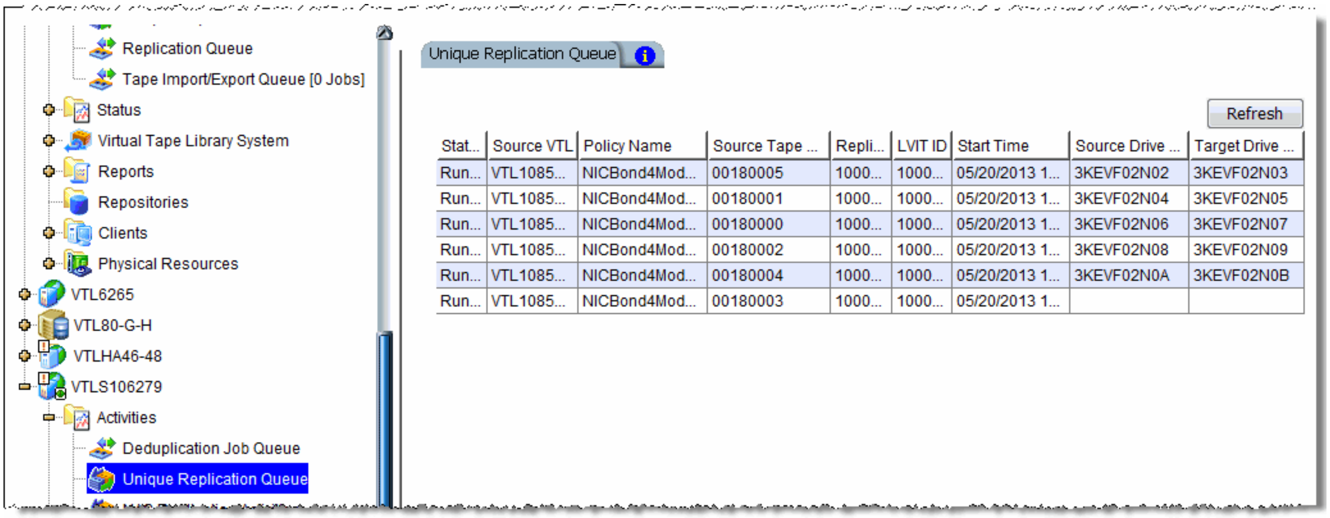

Stat...	Source VTL	Policy Name	Source Tape ...	Repli...	LVIT ID	Start Time	Source Drive ...	Target Drive ...
Run...	VTL1085...	NICBond4Mod...	00180005	1000...	1000...	05/20/2013 1...	3KEVF02N02	3KEVF02N03
Run...	VTL1085...	NICBond4Mod...	00180001	1000...	1000...	05/20/2013 1...	3KEVF02N04	3KEVF02N05
Run...	VTL1085...	NICBond4Mod...	00180000	1000...	1000...	05/20/2013 1...	3KEVF02N06	3KEVF02N07
Run...	VTL1085...	NICBond4Mod...	00180002	1000...	1000...	05/20/2013 1...	3KEVF02N08	3KEVF02N09
Run...	VTL1085...	NICBond4Mod...	00180004	1000...	1000...	05/20/2013 1...	3KEVF02N0A	3KEVF02N0B
Run...	VTL1085...	NICBond4Mod...	00180003	1000...	1000...	05/20/2013 1...		

Virtual vault (target server)

When replication is complete, the replica VIT will be visible in the virtual vault (or the virtual library) on the target server.

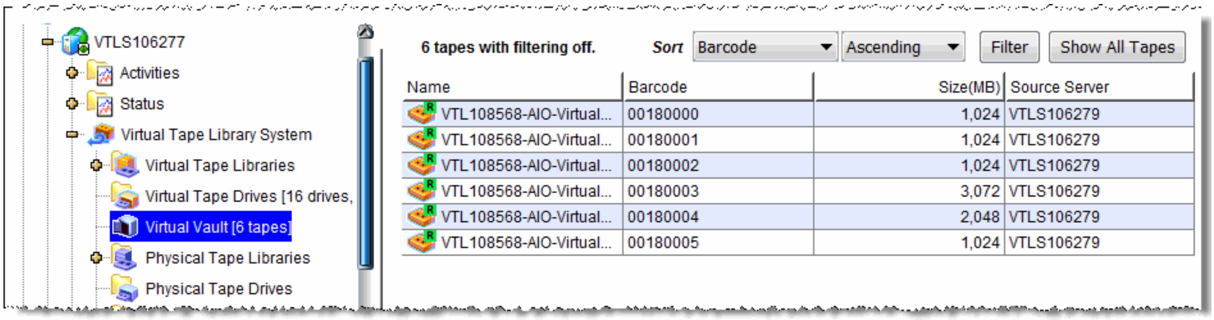

6 tapes with filtering off.

Name	Barcode	Size(MB)	Source Server
VTL108568-AIO-Virtual...	00180000	1,024	VTLS106279
VTL108568-AIO-Virtual...	00180001	1,024	VTLS106279
VTL108568-AIO-Virtual...	00180002	1,024	VTLS106279
VTL108568-AIO-Virtual...	00180003	3,072	VTLS106279
VTL108568-AIO-Virtual...	00180004	2,048	VTLS106279
VTL108568-AIO-Virtual...	00180005	1,024	VTLS106279

A green “R” icon indicates that the tape has been successfully resolved. Red indicates that the last attempt at resolving the tape failed or that the tape is currently being resolved or has not been resolved.

Access data on a replicated VIT

If a replicated virtual tape is needed (due to a failure at the primary site), the administrator can do one of the following so that the data can be accessed by backup software:

- Move the virtual tape from the virtual vault to a virtual library on the target server.
- If the primary cluster has been repaired, you can replicate the entire dedupicated virtual tape back to that server. Depending upon how much data exists in the cluster’s repository, this may be time consuming. However, if all of the data is there, only the VIT will need to be replicated back.
- Export the data on the virtual tape to a physical tape. Refer to the ‘Export data to physical tape’ section for more information.
- If you move a local VIT out of the vault, replication of this VIT will be discontinued until the tape is moved back to the vault. **It is important to note** that any new data added to the tape while it is not in the vault will be overwritten when the tape is returned to the vault and replication proceeds.

Stop replication of a VIT

To stop replication of a VIT, right-click the policy and select *Stop*.

Remove replication for deduplicated tapes

To remove replication for tapes in a deduplication policy, edit the policy and uncheck the *Enable Replication* option.

Replication of NAS resources

Replication protects data on NAS resources by maintaining a copy of data on another VTL 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 VTL 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.

Replication requirements for NAS resources

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

- You must have two VTL servers with NAS enabled.
- You must have administrative rights on both servers.
- You must have enough space on the target server for the replica data.

Configure replication for NAS resources

In this section, replication applies to data on *regular* NAS resources. Replication of OST resources is handled by the FalconStor OpenStorage Option plug-in.

Replication of NAS shares and OST resources can co-exist. However, when you configure replication through FSOST, it will overwrite any existing outgoing replication configuration so that both types of replication will share the same parameters.

Incoming NAS replication

Note: You should have created at least one NAS resource before you can configure incoming replication.

Do the following on the server that will hold replicated files from a source server:

1. Right-click the *NAS Resources* object and select *Replication* --> *Incoming* --> *Select Volume*.

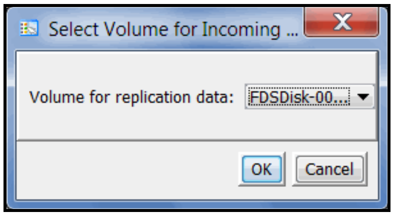

Volume for replication data: FDSDisk-00...

2. Select the NAS resource to use for the replicated data.

> **Note:** If you are using the FalconStor OpenStorage Option, you must select *ANY VOLUME* in order for NetBackup replication to complete successfully.

Once replicated data has been received on this server, you will see a new *Replica* object in the tree beneath the NAS resource being used to hold incoming replication.

> **Note:** If you designate a specific NAS resource for incoming replication, you will not be able to delete that NAS resource. If you ever need to delete it, you will have to change your incoming replication volume to *ANY VOLUME*.

Outgoing NAS replication

Do the following on the source server (the server from which you will replicate files):

1. Right-click the *NAS Resources* object and select *Replication --> Outgoing --> Enable*.

2. Select the replica server (the server that will hold the replicated data).

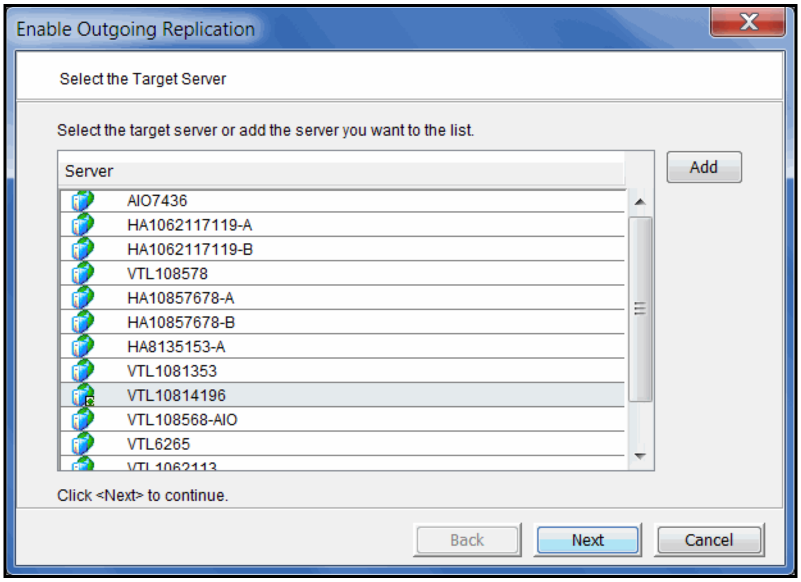

Enable Outgoing Replication

Select the Target Server

Select the target server or add the server you want to the list.

Server

AIO7436
HA1062117119-A
HA1062117119-B
VTL108578
HA10857678-A
HA10857678-B
HA8135153-A
VTL1081353
VTL10814196
VTL108568-AIO
VTL6265
VTL1062113

Click <Next> to continue.

3. Enter login information for the replica server and determine if you want to use encryption during replication.

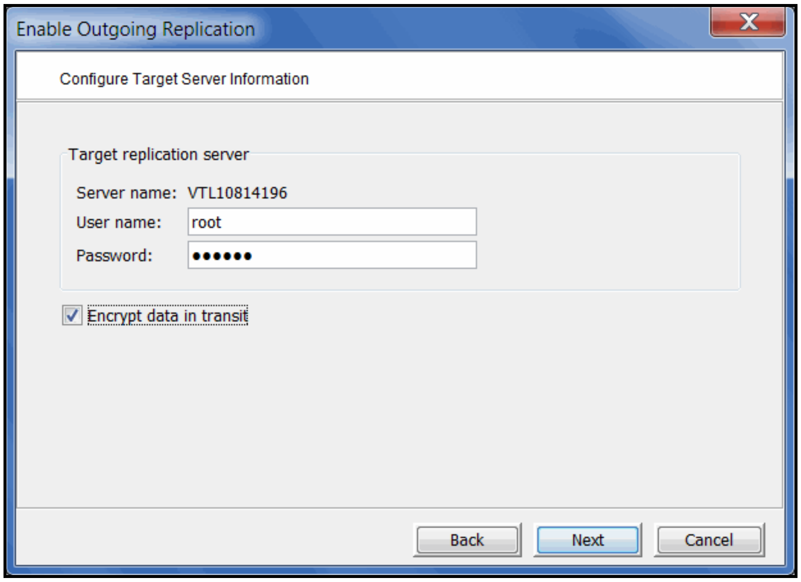

Configure Target Server Information

Target replication server

Server name: VTL10814196

User name: root

Password: ••••••

☑ Encrypt data in transit

You can select which IP address to use for replication. The IP address you enter here should be one that the management console can connect to.

4. Select a specific IP address on the replica server to use for replication.

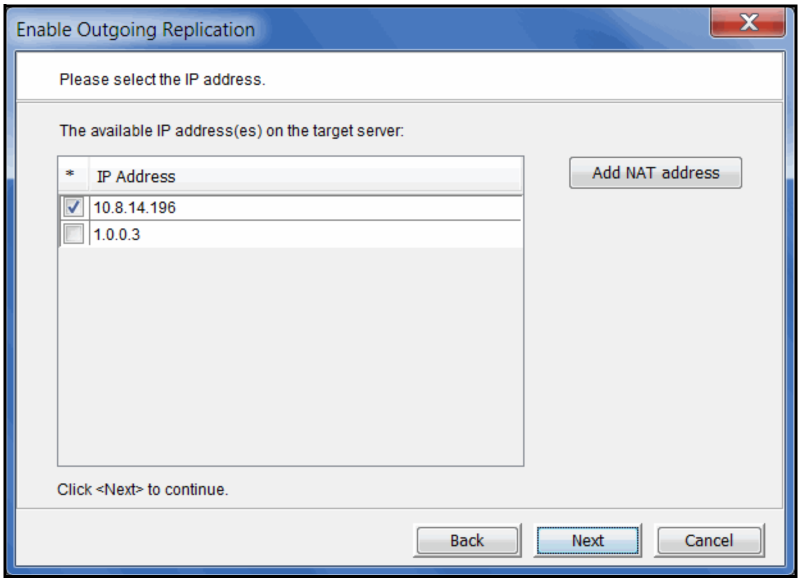

Please select the IP address.

The available IP address(es) on the target server:

*	IP Address
☑	10.8.14.196
☐	1.0.0.3

Click <Next> to continue.

5. Specify if you will be creating a single global policy or multiple policies.

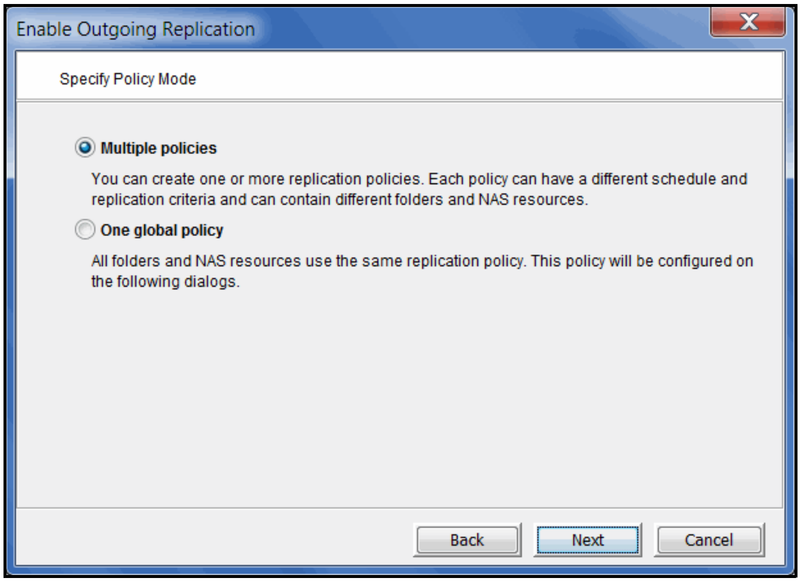

Enable Outgoing Replication

Specify Policy Mode

● **Multiple policies**

You can create one or more replication policies. Each policy can have a different schedule and replication criteria and can contain different folders and NAS resources.

○ **One global policy**

All folders and NAS resources use the same replication policy. This policy will be configured on the following dialogs.

Multiple policies - Replication can be performed at the folder level and file system level. Each policy can include one or more folders or file system resources and you will be able to customize the schedule of each replication policy. If you select this mode, you will create your policies after this wizard completes.

One global policy - All folders and resources will be included for replication. You will need to manually exclude folders that you do not want to replicate. If you select this mode, you will configure the global policy on the subsequent dialogs in this wizard.

6. (Global policy only) Determine when replication should occur.

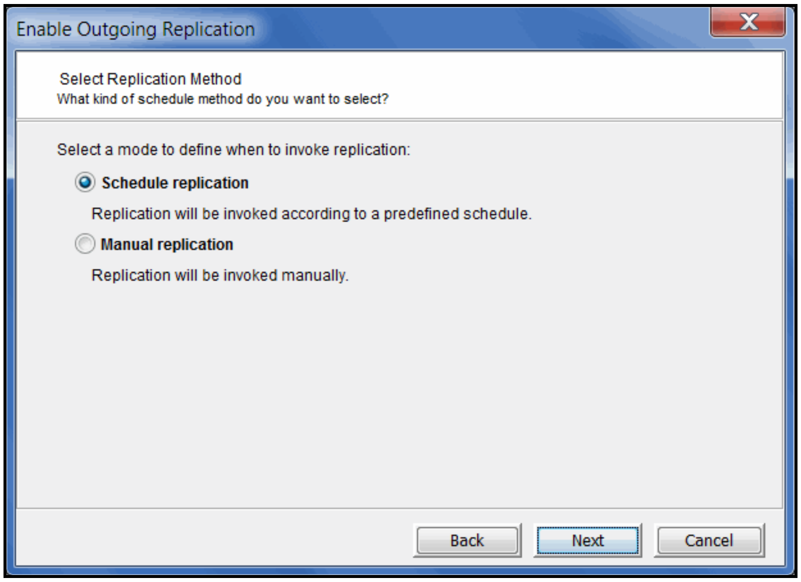

Enable Outgoing Replication

Select Replication Method
What kind of schedule method do you want to select?

Select a mode to define when to invoke replication:

● **Schedule replication**

Replication will be invoked according to a predefined schedule.

○ **Manual replication**

Replication will be invoked manually.

7. (Global policy only) If you selected to schedule replication, set the schedule.

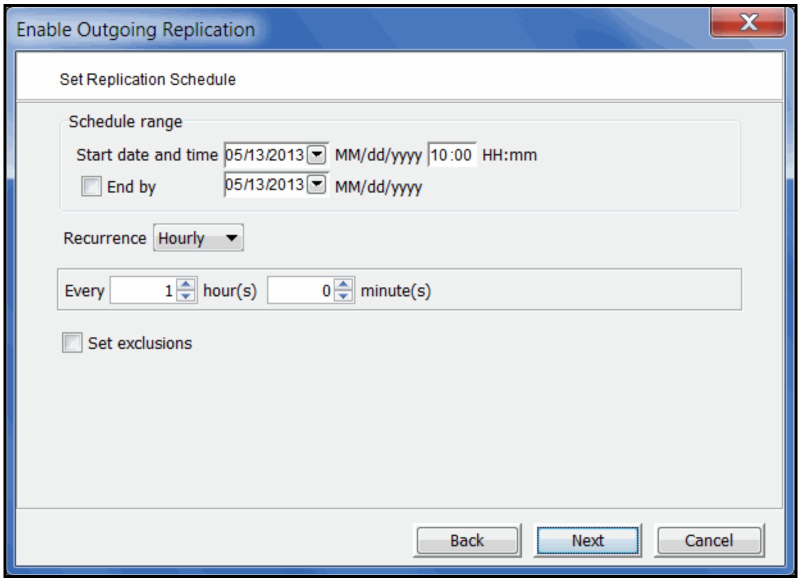

Enable Outgoing Replication

Set Replication Schedule

Schedule range

Start date and time 05/13/2013 MM/dd/yyyy 10:00 HH:mm

□ End by 05/13/2013 MM/dd/yyyy

Recurrence Hourly

Every 1 hour(s) 0 minute(s)

□ Set exclusions

Specify when replication should begin and, optionally, when it should end. Also specify the frequency. If you want to exclude specific days/hours/months, select *Set exclusions*.

Replication will run based on the schedule you set. If you need to start it manually, you can right-click the server object and select *Replication --> Outgoing --> Start.*

8. (Global policy only) If you selected *Set exclusions*, select the days, hours, or months during which replication should not occur.

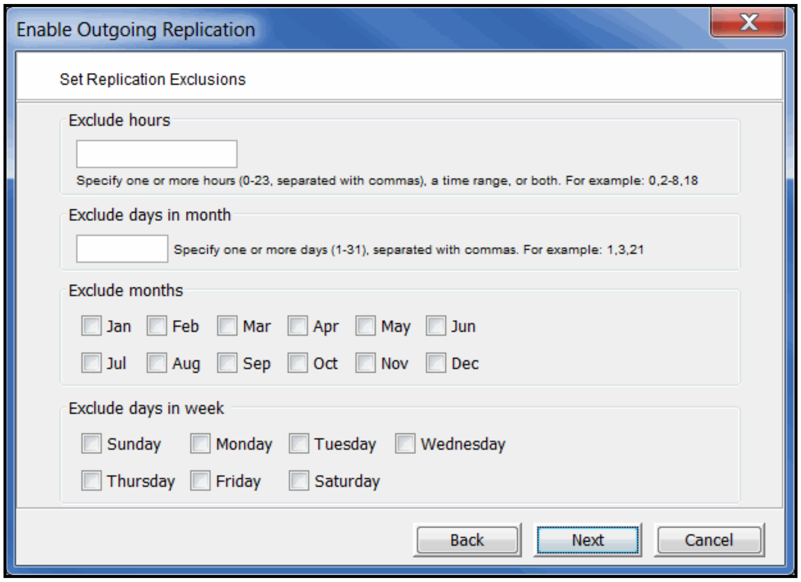

Enable Outgoing Replication

Set Replication Exclusions

Exclude hours

Specify one or more hours (0-23, separated with commas), a time range, or both. For example: 0,2-8,18

Exclude days in month

Specify one or more days (1-31), separated with commas. For example: 1,3,21

Exclude months

□ Jan □ Feb □ Mar □ Apr □ May □ Jun

□ Jul □ Aug □ Sep □ Oct □ Nov □ Dec

Exclude days in week

□ Sunday □ Monday □ Tuesday □ Wednesday

□ Thursday □ Friday □ Saturday

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

9. (Global policy only) Specify replication criteria.

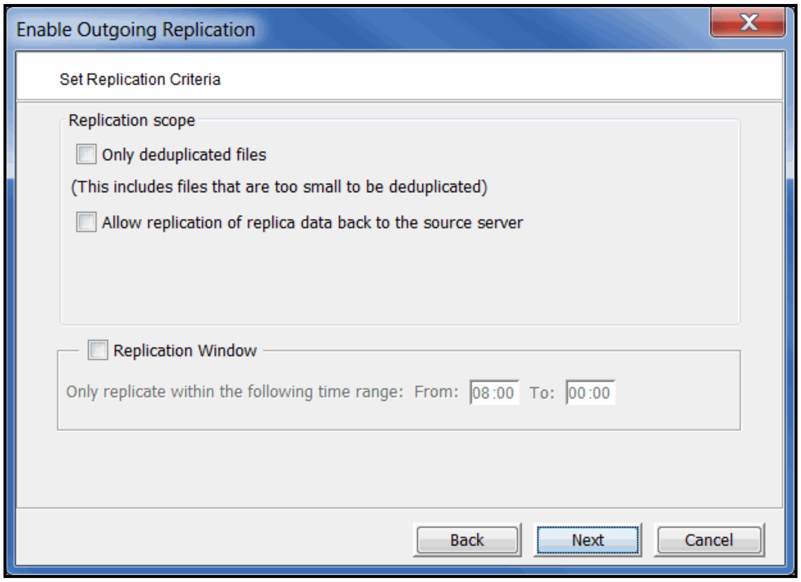

Enable Outgoing Replication

Set Replication Criteria

Replication scope

□ Only deduplicated files

(This includes files that are too small to be deduplicated)

□ Allow replication of replica data back to the source server

□ Replication Window

Only replicate within the following time range: From: 08:00 To: 00:00

Determine if you want to only include deduplicated files or if you want to replicate all files regardless of whether or not they have been deduplicated. Note that small files (<8K) will be replicated regardless of whether or not they were deduplicated.

The *Allow replication of replica data back to the source server* option is only applicable for disaster recovery purposes, such as if the source server has been damaged and you need to replicate all of its data back to it. For example, if server A normally replicates to server B and then server A is destroyed, you can set this option on server B and all of the data will be replicated back to server A.

There is no need to select this option for an initial replication configuration.

If you scheduled replication, you can specify a window to limit when replication can run.

10. Confirm all information and click *Finish* to configure replication.

If you selected to create multiple policies, you will be prompted to create them.

Create NAS replication policies

You can only create replication policies if you selected the *Multiple Policies* option during configuration.

1. Right-click the *NAS Resources* object and select *Replication --> Outgoing --> Policies --> New.*
2. Specify a name for the policy.
3. Determine when replication should occur.

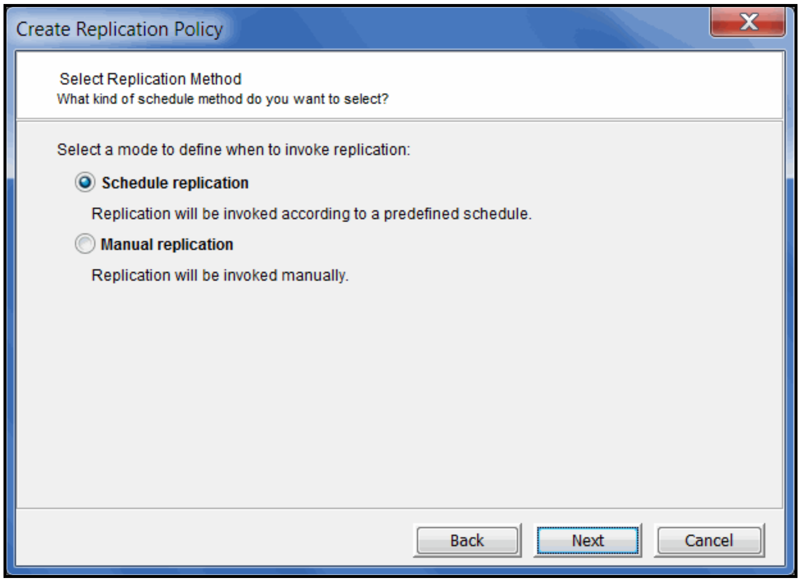

Create Replication Policy

Select Replication Method
What kind of schedule method do you want to select?

Select a mode to define when to invoke replication:

● **Schedule replication**

Replication will be invoked according to a predefined schedule.

○ **Manual replication**

Replication will be invoked manually.

4. If you selected to schedule replication, set the schedule.

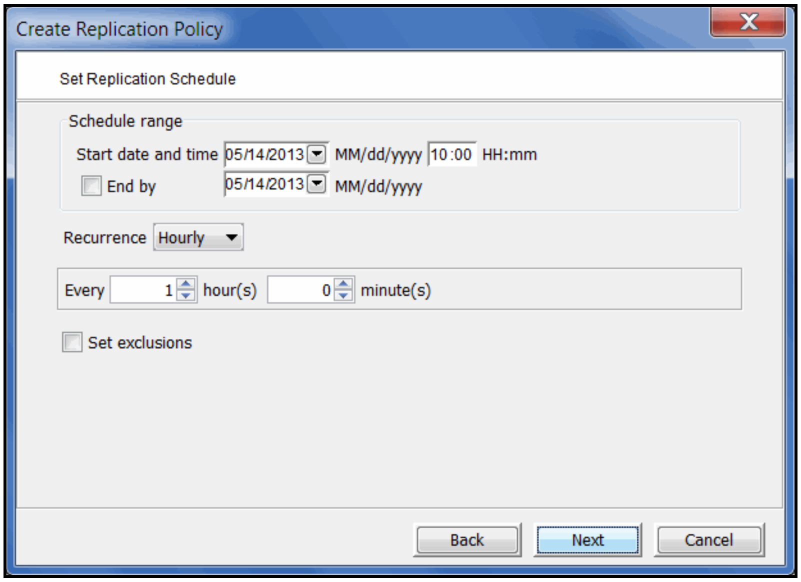

Create Replication Policy

Set Replication Schedule

Schedule range

Start date and time 05/14/2013 MM/dd/yyyy 10:00 HH:mm

□ End by 05/14/2013 MM/dd/yyyy

Recurrence Hourly

Every 1 hour(s) 0 minute(s)

□ Set exclusions

Specify when replication should begin and, optionally, when it should end. Also specify the frequency. If you want to exclude specific days/hours/months, select *Set exclusions*.

Replication will run based on the schedule you set. If you need to start it manually, you can right-click the server object and select *Replication* --> *Outgoing* --> *Start.*

5. If you selected *Set exclusions*, select the days, hours, or months during which replication should not occur.

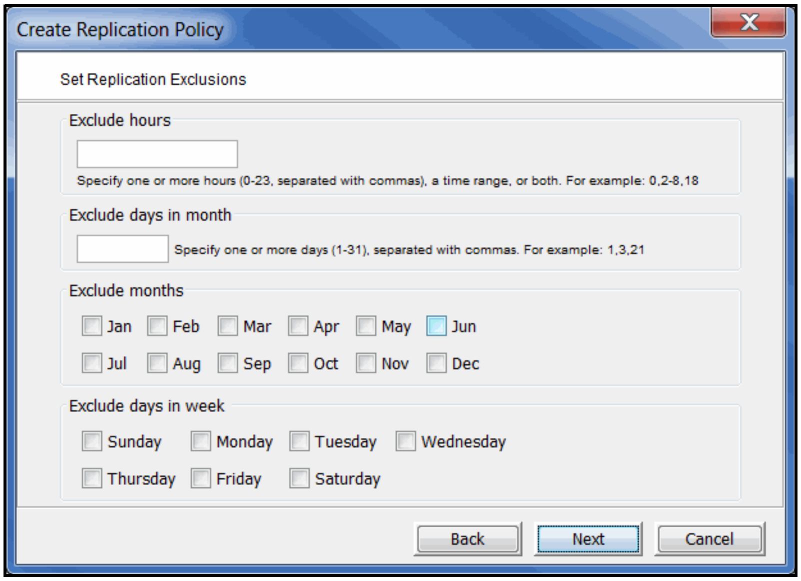

Create Replication Policy

Set Replication Exclusions

Exclude hours

Specify one or more hours (0-23, separated with commas), a time range, or both. For example: 0,2-8,18

Exclude days in month

Specify one or more days (1-31), separated with commas. For example: 1,3,21

Exclude months

□ Jan □ Feb □ Mar □ Apr □ May □ Jun
□ Jul □ Aug □ Sep □ Oct □ Nov □ Dec

Exclude days in week

□ Sunday □ Monday □ Tuesday □ Wednesday
□ Thursday □ Friday □ Saturday

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

6. Specify replication criteria.

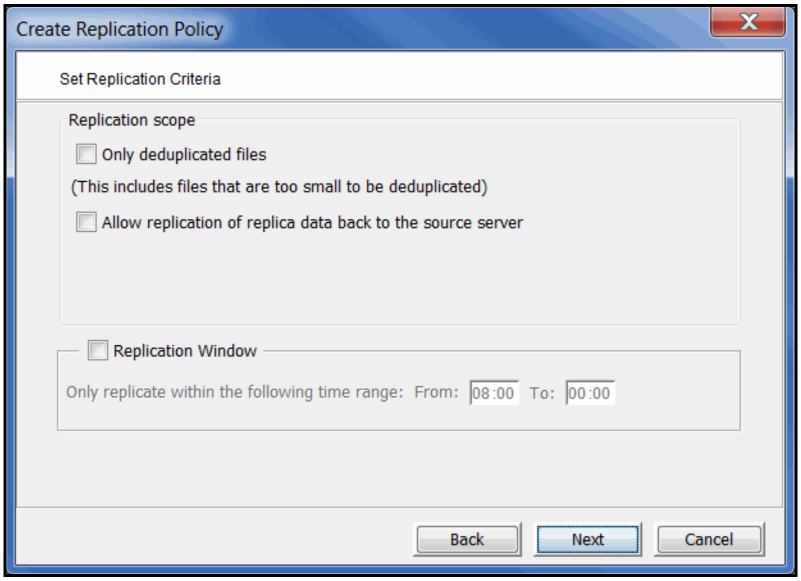

Create Replication Policy

Set Replication Criteria

Replication scope

□ Only deduplicated files

(This includes files that are too small to be deduplicated)

□ Allow replication of replica data back to the source server

□ Replication Window

Only replicate within the following time range: From: 08:00 To: 00:00

Determine if you want to only include deduplicated files or if you want to replicate all files regardless of whether or not they have been deduplicated. Note that small files (<8K) will be replicated regardless of whether or not they were deduplicated.

The *Allow replication of replica data back to the source server* option is only applicable for disaster recovery purposes, such as if the source server has been damaged and you need to replicate all of its data back to it. For example, if server A normally replicates to server B and then server A is destroyed, you can set this option on server B and all of the data will be replicated back to server A.

If you scheduled replication, you can specify a window to limit when replication can run.

7. Select the paths to be included in this replication policy.

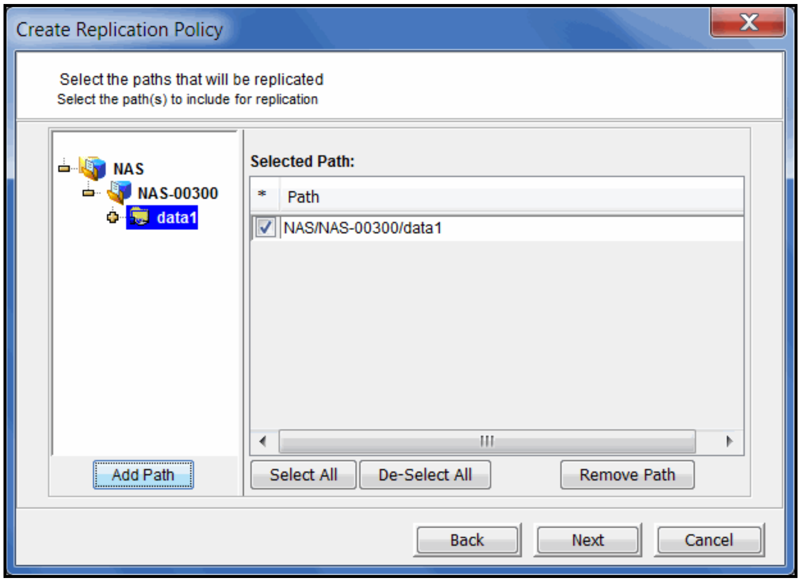

Create Replication Policy

Select the paths that will be replicated

Select the path(s) to include for replication

Selected Path:

*	Path
☑	NAS/NAS-00300/data1

The left pane lists all NAS resources. Select a file system or folder that you want to include and click the *Add Path* button.

8. Confirm all information and click *Finish* to create this policy.

Exclude NAS folders from replication

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

Exclude

To exclude a specific folder (and its sub-folders), right-click the folder and select *Replication --> Exclude*.

You will see a gray X icon on every folder that is excluded.

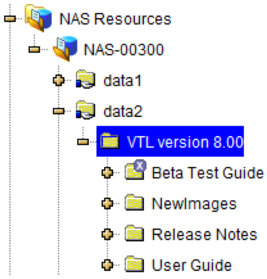

You can see a list of all excluded folders on the *Excluded Paths* tab of the *NAS Resources* object.

Include

To include a folder that was previously excluded, right-click the folder and select *Replication --> Include*.

Set replication throttling for NAS resources

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.

1. Right-click the *NAS Resources* object and select *Replication --> Outgoing --> Throttling*.

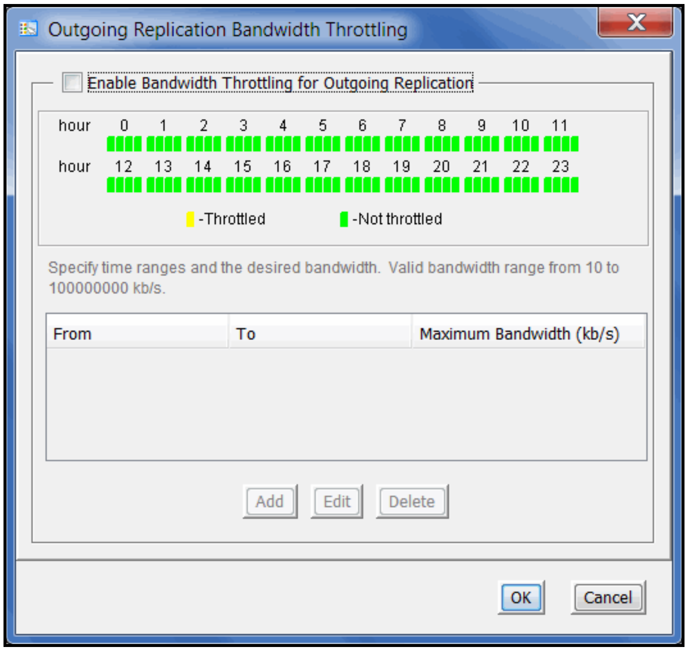

Outgoing Replication Bandwidth Throttling

□ Enable Bandwidth Throttling for Outgoing Replication

hour	0	1	2	3	4	5	6	7	8	9	10	11
	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
hour	12	13	14	15	16	17	18	19	20	21	22	23
	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩

🟨 -Throttled 🟩 -Not throttled

Specify time ranges and the desired bandwidth. Valid bandwidth range from 10 to 100000000 kb/s.

From	To	Maximum Bandwidth (kb/s)

2. Select the *Enable Bandwidth Throttling for Outgoing Replication* checkbox.

3. Click the *Add* button.

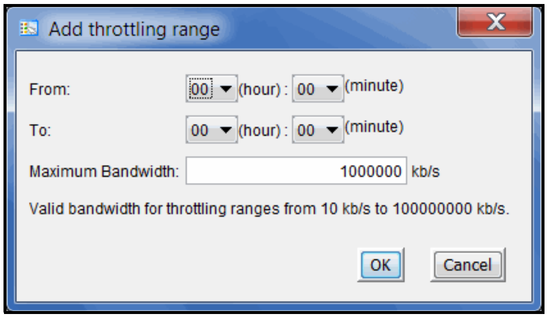

Add throttling range

From: 00 (hour) : 00 (minute)

To: 00 (hour) : 00 (minute)

Maximum Bandwidth: 1000000 kb/s

Valid bandwidth for throttling ranges from 10 kb/s to 100000000 kb/s.

4. 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 is 10-100,000,000 kb/s.

Check replication status for NAS resources

There are several ways to check replication status from the source server.

Replication tab or *NAS Replication Activities* tab

Highlight the *Replication* tab to see replication statistics for all NAS resources. You can also see your replication configuration settings on this tab.

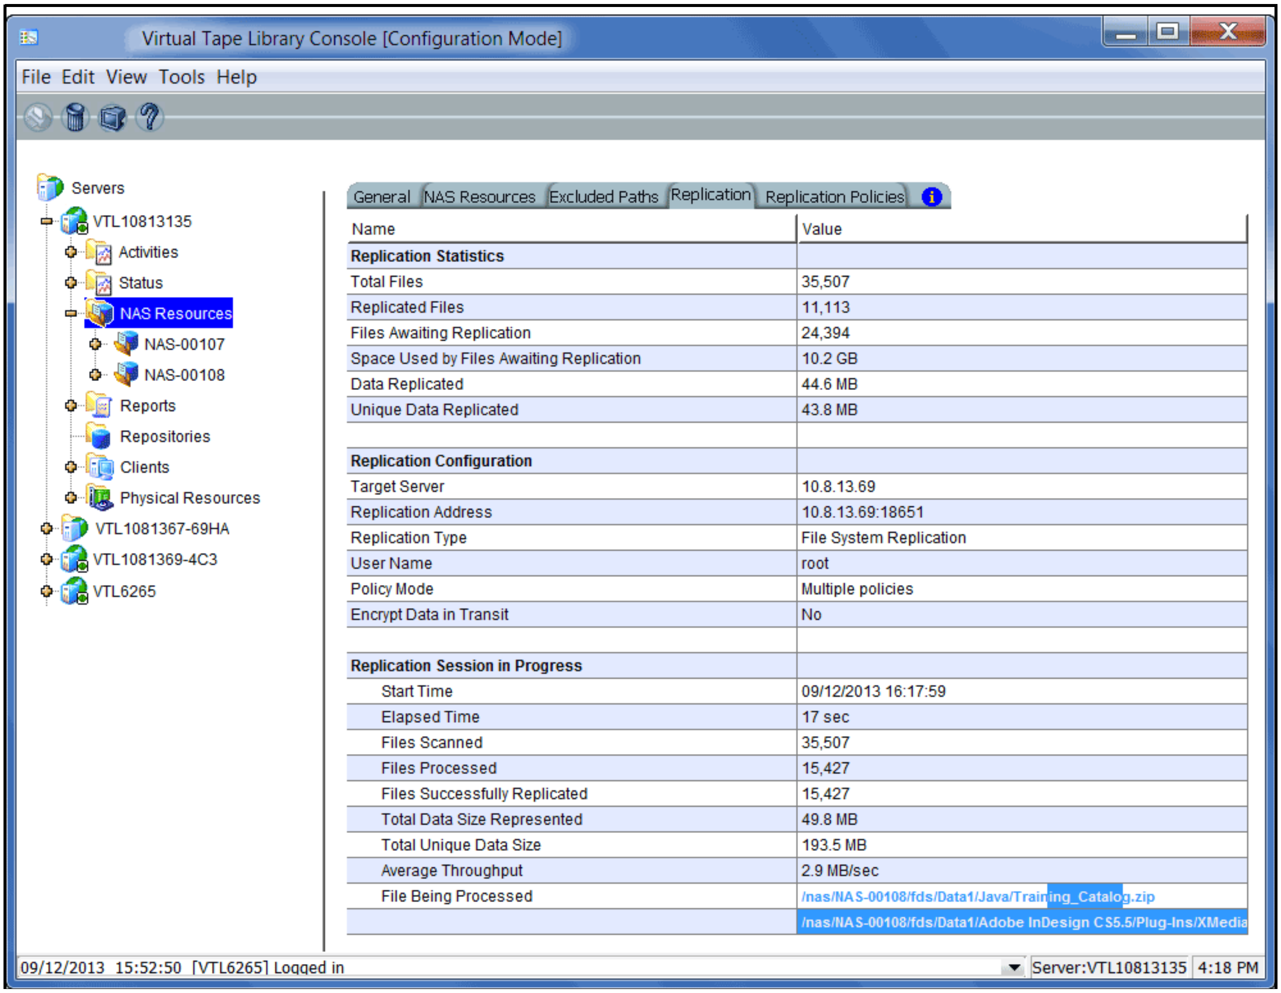

Name	Value
Replication Statistics	
Total Files	35,507
Replicated Files	11,113
Files Awaiting Replication	24,394
Space Used by Files Awaiting Replication	10.2 GB
Data Replicated	44.6 MB
Unique Data Replicated	43.8 MB
Replication Configuration	
Target Server	10.8.13.69
Replication Address	10.8.13.69:18651
Replication Type	File System Replication
User Name	root
Policy Mode	Multiple policies
Encrypt Data in Transit	No
Replication Session in Progress	
Start Time	09/12/2013 16:17:59
Elapsed Time	17 sec
Files Scanned	35,507
Files Processed	15,427
Files Successfully Replicated	15,427
Total Data Size Represented	49.8 MB
Total Unique Data Size	193.5 MB
Average Throughput	2.9 MB/sec
File Being Processed	/nas/NAS-00108/fds/Data1/Java/Training_Catalog.zip
	/nas/NAS-00108/fds/Data1/Adobe InDesign CS5.5/Plug-Ins/XMedia

09/12/2013 15:52:50 [VTL6265] Logged in

Server:VTL10813135 4:18 PM

Total Files - Total number of files. This includes files that have been replicated, files that were excluded, and files awaiting replication.

Replicated Files - Total number of files replicated. Depending upon how replication was configured, this may not include excluded files.

Files Awaiting Replication - Total number of files that have not yet been replicated.

Space Used by Files Awaiting Replication - Total size of files that have not yet been replicated.

Data Replicated - Total amount of data replicated.

Unique Data Replicated - Total amount of unique data replicated.

If replication is currently taking place, you can see the status in the bottom section of the screen (below your replication settings). Status includes when replication started, how long it has been running, how many files were processed, how many were actually replicated, total data size represented, amount of unique data, current throughput, and the file currently being processed.

Note that you can also view NAS replication statistics on the *NAS Replication Activities* tab on the source server.

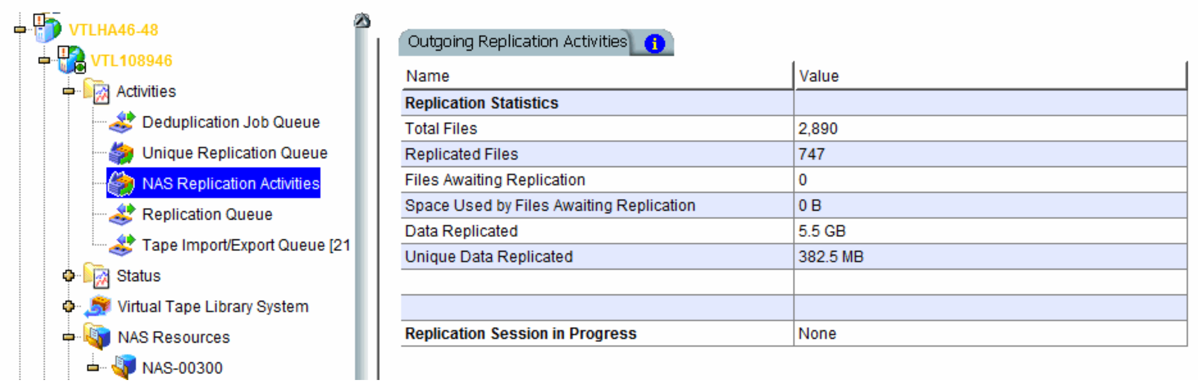

Name	Value
Replication Statistics	
Total Files	2,890
Replicated Files	747
Files Awaiting Replication	0
Space Used by Files Awaiting Replication	0 B
Data Replicated	5.5 GB
Unique Data Replicated	382.5 MB
Replication Session in Progress	None

Deduplication & Replication tab

The *Deduplication & Replication* tab for a NAS folder displays a list of all files and folders in that folder. The *Replication Status* field indicates whether each file has been replicated yet.

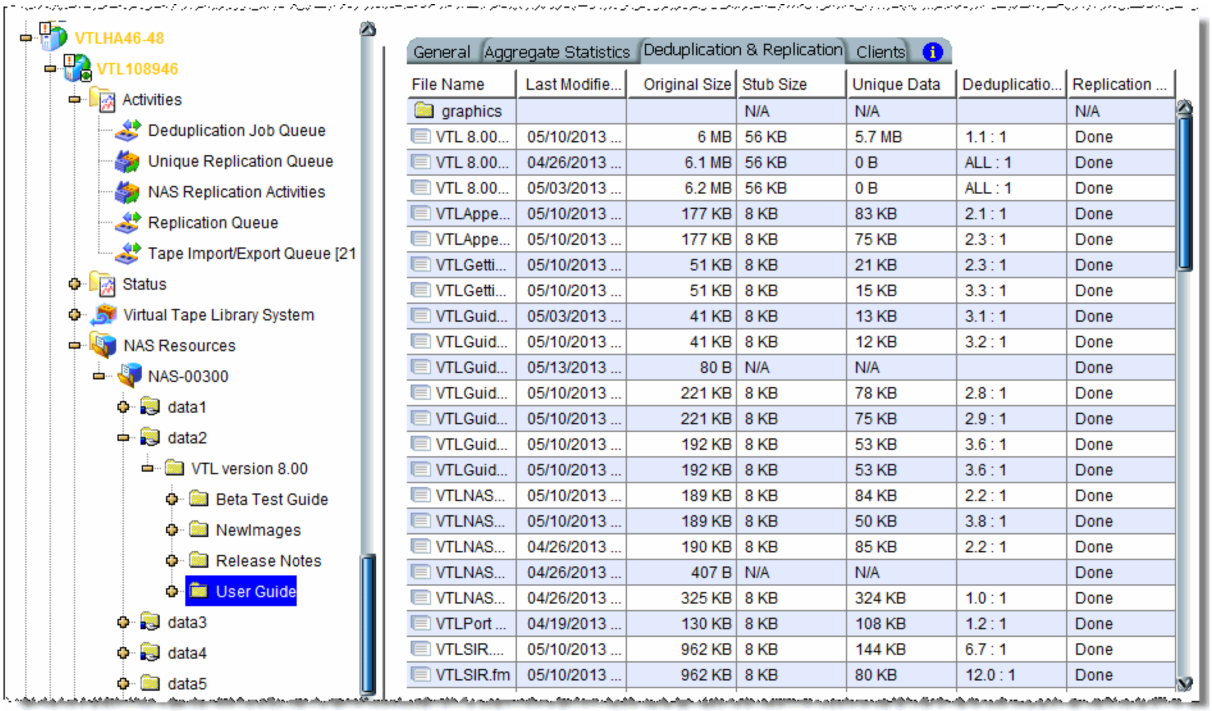

File Name	Last Modifie...	Original Size	Stub Size	Unique Data	Deduplicatio...	Replication ...
graphics			N/A	N/A		N/A
VTL 8.00...	05/10/2013 ...	6 MB	56 KB	5.7 MB	1.1 : 1	Done
VTL 8.00...	04/26/2013 ...	6.1 MB	56 KB	0 B	ALL : 1	Done
VTL 8.00...	05/03/2013 ...	6.2 MB	56 KB	0 B	ALL : 1	Done
VTLAppe...	05/10/2013 ...	177 KB	8 KB	83 KB	2.1 : 1	Done
VTLAppe...	05/10/2013 ...	177 KB	8 KB	75 KB	2.3 : 1	Done
VTLGetti...	05/10/2013 ...	51 KB	8 KB	21 KB	2.3 : 1	Done
VTLGetti...	05/10/2013 ...	51 KB	8 KB	15 KB	3.3 : 1	Done
VTLGuid...	05/03/2013 ...	41 KB	8 KB	13 KB	3.1 : 1	Done
VTLGuid...	05/10/2013 ...	41 KB	8 KB	12 KB	3.2 : 1	Done
VTLGuid...	05/13/2013 ...	80 B	N/A	N/A		Done
VTLGuid...	05/10/2013 ...	221 KB	8 KB	78 KB	2.8 : 1	Done
VTLGuid...	05/10/2013 ...	221 KB	8 KB	75 KB	2.9 : 1	Done
VTLGuid...	05/10/2013 ...	192 KB	8 KB	53 KB	3.6 : 1	Done
VTLGuid...	05/10/2013 ...	192 KB	8 KB	53 KB	3.6 : 1	Done
VTLNAS...	05/10/2013 ...	189 KB	8 KB	84 KB	2.2 : 1	Done
VTLNAS...	05/10/2013 ...	189 KB	8 KB	50 KB	3.8 : 1	Done
VTLNAS...	04/26/2013 ...	190 KB	8 KB	85 KB	2.2 : 1	Done
VTLNAS...	04/26/2013 ...	407 B	N/A	N/A		Done
VTLNAS...	04/26/2013 ...	325 KB	8 KB	324 KB	1.0 : 1	Done
VTLPort ...	04/19/2013 ...	130 KB	8 KB	108 KB	1.2 : 1	Done
VTLSIR....	05/10/2013 ...	962 KB	8 KB	144 KB	6.7 : 1	Done
VTLSIR.fm	05/10/2013 ...	962 KB	8 KB	80 KB	12.0 : 1	Done

NAS Statistics Report

The NAS Statistics Report (run from the *Reports* object) displays deduplication and/or replication session summary information and settings for each NAS deduplication/replication session. Refer to 'NAS Statistics Summary' for more information.

Recover data from a replicated NAS resource

Recover at the replica site

To recover files at the replica site, you simply need to share the replica folder that is on the target server and then map/mount the appropriate share.

Do the following on the target server:

1. Right-click the desired replicated resource or share and select *Sharing*.

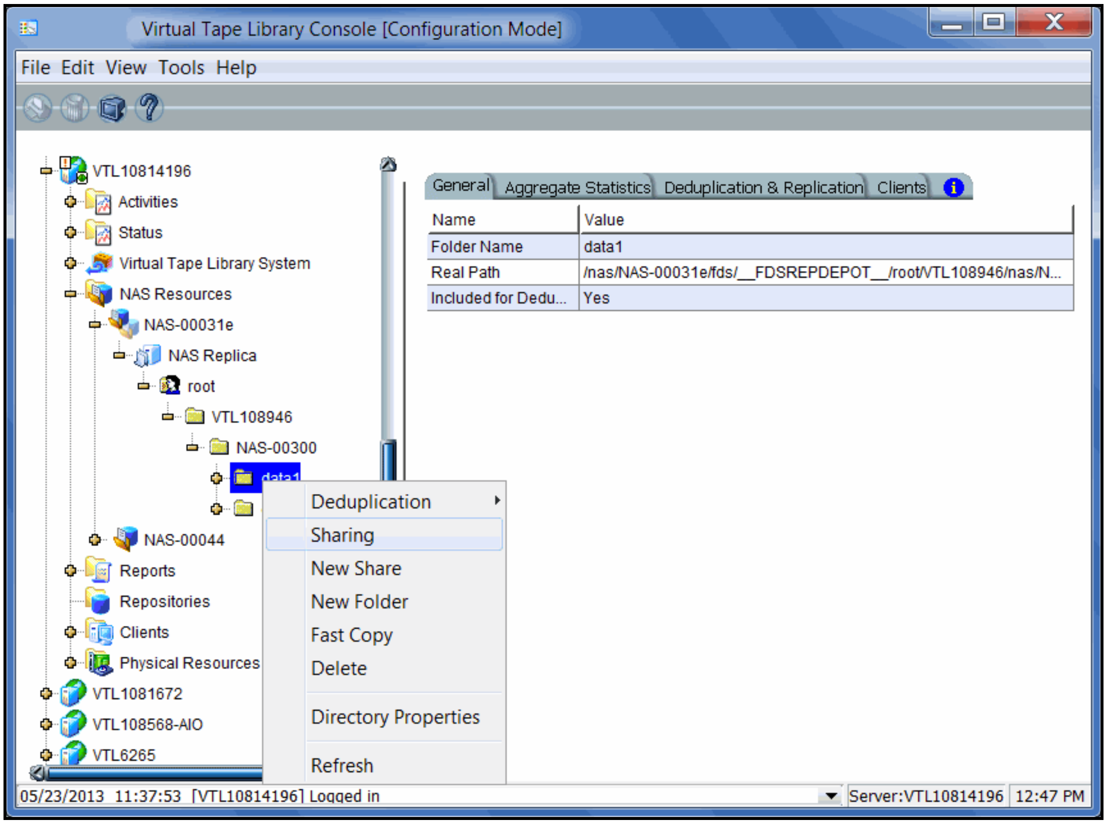

Name	Value
Folder Name	data1
Real Path	/nas/NAS-00031e/fds/__FDSREPDEPOT__/root/VTL108946/nas/N...
Included for Dedu...	Yes

```
05/23/2013  11:37:53  [VTL10814196] Logged in
```

Replicated data can be found under the *NAS Replica* object.

2. On the Windows share, select *Enable Windows Share*, enter a new share name if desired, and set login information.

   You can set this folder as a Windows or NFS share, or both.

3. Map/mount the share.

   You can now recover files from the share.

Recover back to the source site

If your source server (server A) has been damaged, you may need to replicate all of its data back to it from the target server (server B).

To do this, configure *outgoing* replication from the target server (server B) to the source server (server A) and set the *Allow replication of replica data back to the source server* option. When this option is set, the replica data will be replicated back to the source server (server A) when replication is run.

Once the data is on the source server, you can share the replica folder and then map/mount the appropriate share. Follow the same instructions as those listed in the 'Recover at the replica site' section.

Start replication of NAS resources

To force replication that is not scheduled when you have multiple policies, highlight the *NAS Resources* object. On the *Replication Policies* tab, right-click the policy and select *Start* or *Synchronize*.

To force replication that is not scheduled when you have one global policy, right-click the *NAS Resources* object or any NAS resource or share and select *Start* or *Synchronize*.

When you perform synchronization, files/directories between the source and replica server will be synchronized. If a replicated file has been deleted from the source, the file will be deleted from the replica.

In general, only one outgoing replication job can be active at any time. If multiple replication jobs are pending, the next job will become active after the current job finishes, until all replication jobs have finished.

Change a NAS replication policy

To modify the properties of a policy, highlight the *NAS Resources* object. On the *Replication Policies* tab, right-click the policy and select *Edit*. You can change the policy name, schedule, replication criteria, and paths. You cannot change the replication target for a policy.

Suspend/resume a NAS replication policy

To suspend/resume replication for a policy, highlight the *NAS Resources* object. On the *Replication Policies* tab, right-click the policy and select *Suspend* (or *Resume*).

Delete a NAS replication policy

To completely remove a policy, highlight the *NAS Resources* object. On the *Replication Policies* tab, right-click the policy and select *Delete*.

Change your NAS replication configuration

You can change the following for your replication configuration:

- IP address(es) and login information for your target server
- Encryption
- Policy type (global or multiple) - Note that you cannot switch from multiple policies to a global policy if you have existing policies.
- (Global policy only) Replication schedule
- (Global policy only) Replication criteria

To change the configuration:

1. Right-click the *NAS Resources* object and select *Replication* --> *Outgoing* --> *Configure*.
2. Make the appropriate changes and click *Finish*.

Disable NAS replication

Disabling replication removes the current replication configuration. To disable replication, highlight the *NAS Resources* object and select *Replication* --> *Outgoing* --> *Disable*.

Auto Replication

Auto Replication replicates the contents of a single tape whenever a virtual tape is ejected from a virtual library and moved to the virtual vault (manually or by backup software).

Auto Replication can be enabled when you create a virtual tape library. If it is enabled for a library, you can enable/disable *Auto Replication* for individual tapes when you create tapes for the library.

Notes:

- Do not enable auto replication for libraries or tapes for which you will be defining a deduplication policy. This feature is not supported for virtual index tapes (VITs).
- If virtual tape encryption is used, encryption must be enabled on the target server; all keys used by the source tapes must exist on both servers and be identical. This means that the keys have the same name and were created using the same secret phrase. If the secret phrase is not the same, you can export a key from the source server and import it to the target.

If you want to enable *Auto Replication* for an existing library:

1. Right-click a virtual tape library and select *Properties*.
2. Select *Auto Replication*.
3. Select whether you want the virtual tape copied (retained) or moved (removed) after the data is replicated.

   If you select to move it, indicate how long to wait before deleting it.
4. Select the target server.

Remote Copy

You can copy the contents of a single tape whenever you need to. Because the *Remote Copy* feature replicates the full tape rather than appending to an existing virtual tape, you can only copy a tape if there is no virtual tape on the target server with the same barcode. Therefore, if you have copied this tape before, you must delete the copy from the target server before continuing.

Notes:

- You cannot copy a deduplicated tape or a virtual tape that is configured for scheduled replication or *Auto Replication/Auto Archive*.
- If virtual tape encryption is used, encryption must be enabled on the target server; the key used by source tape must exist on both servers and be identical. This means that the keys have the same name and were created using the same secret phrase. If the secret phrase is not the same, you can export a key from the source server and import it to the target.

1. Right-click a tape and select *Remote Copy*.
2. Select if you want to copy to a local or remote server.

   If you select to copy to a remote server, you will have to select the server. If the server you want does not appear on the list, click the *Add* button.
3. Confirm/enter the target server's IP address.
4. Select a location for the copied tape.

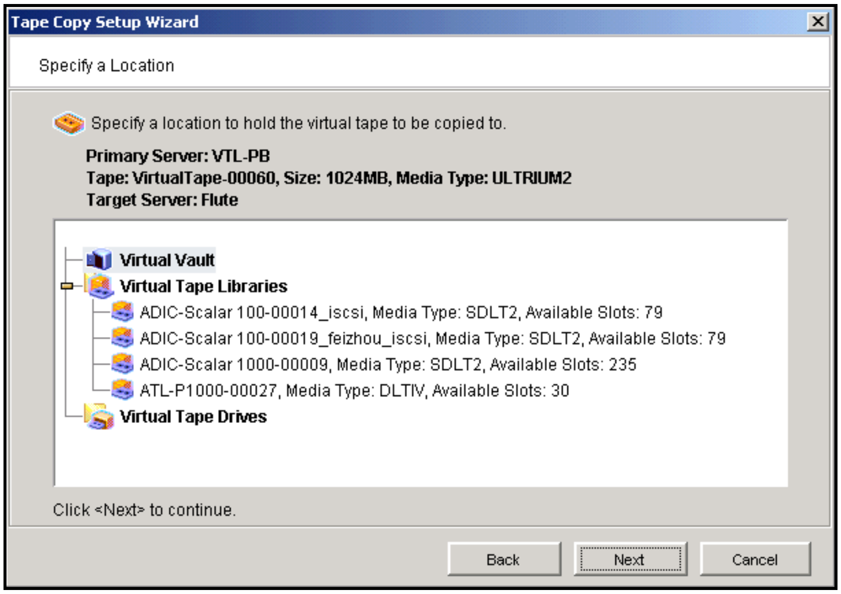

Tape Copy Setup Wizard

Specify a Location

Specify a location to hold the virtual tape to be copied to.

Primary Server: VTL-PB
Tape: VirtualTape-00060, Size: 1024MB, Media Type: ULTRIUM2
Target Server: Flute

- Virtual Vault
- Virtual Tape Libraries
  - ADIC-Scalar 100-00014_iscsi, Media Type: SDLT2, Available Slots: 79
  - ADIC-Scalar 100-00019_feizhou_iscsi, Media Type: SDLT2, Available Slots: 79
  - ADIC-Scalar 1000-00009, Media Type: SDLT2, Available Slots: 235
  - ATL-P1000-00027, Media Type: DLTIV, Available Slots: 30
- Virtual Tape Drives

Click <Next> to continue.

You can select a tape library or the virtual vault.

However, if the target server is at version 7.50, the target tape will be located in the virtual vault. To use the tape, you must manually move it from the virtual vault to a tape library.

If you select a tape library, the media must be compatible with the original media.

5. Confirm that all information is correct and then click *Finish* to create the copy.

   Once replication is completed, the replica is promoted.

Automated Tape Caching

Overview

Automated Tape Caching enhances the functionality of VTL by acting as a cache to your physical tape library, providing transparent access to data regardless of its location.

With Automated Tape caching, tapes will always appear to be inside virtual libraries and will be visible to the backup application regardless of whether the data is actually on disk or on physical tape.

Automated Tape Caching also provides advanced flexibility that allows you to set up policies that automatically trigger data migration to physical tapes based on criteria, such as the number of days that data has been on disk or the amount of used disk space.

With Automated Tape Caching, you can not only determine which events will activate the action, but also when it will occur. For example, you can set the policy to migrate the data immediately or at a specific time or day. This enables data to be written to physical tapes as a background process without impacting production servers.

You can also set up a reclamation policy that allows you to specify how and when the data that has been migrated to physical tape can be deleted from the disk to make space for new backups.

In order to use Automated Tape Caching, you must enable the feature for your virtual library, set your migration and reclamation policies, and create a cache for each of your physical tapes. You may have done this during the initial setup wizard when you first launched VTL or when you first created a virtual tape library. You can also create an automated tape caching policy for an existing virtual tape library for which Automated Tape Caching was not previously enabled.

Note: You can use Automated Tape Caching **or** Auto Archive/Replication on a virtual tape library, but not both. Enabling Tape Caching on a library for which you previously selected Auto Archive/Replication will disable those features on all tapes in the library.

Tape caching policies

A tape caching policy contains the data migration triggers and reclamation triggers for a virtual tape library. The tape caching policy affects how data will be read/written from/to tapes.

Scenario 1: Data on virtual tape. Data not written to physical tape.

If the data has not been written to physical tape, reads will be from the virtual tape. Writes will either append or rewrite the virtual tape.

Scenario 2: Data written to physical tape. Virtual tape not reclaimed.

If the data has been written to physical tape but is still retained on the virtual tape, reads will be from the virtual tape. Writes to the tape will either append or rewrite the virtual tape (and restarts the clock on the migration policy).

Scenario 3: Data written to physical tape. Virtual tape reclaimed.

If the data has been written to physical tape and the virtual tape has been reclaimed, reads will be directly from the *direct link* tape. A *direct link* tape is not an actual tape but a link to a physical tape. If you overwrite the beginning of the tape, VTL will create a new virtual tape, which breaks the *direct link* tape and restarts the clock on the migration policy. If you try to append to the tape, VTL will append data on the physical tape.

Create/change a tape caching policy

To create or change a tape caching policy:

1. Right-click a virtual tape library and select *Automated Tape Caching*.
2. If necessary, select the *Enable Automated Tape Caching* check box.
3. On the *Data Migration Triggers* tab, select the type of data migration triggers that you want to set.

   Data migration triggers control when data in the cache will be copied to physical tape.

   Note: Regardless of which triggers you set, there must be at least 1 MB of data on the tape in order to trigger data migration.

For *Time Based* triggers, specify when data migration should actually occur.

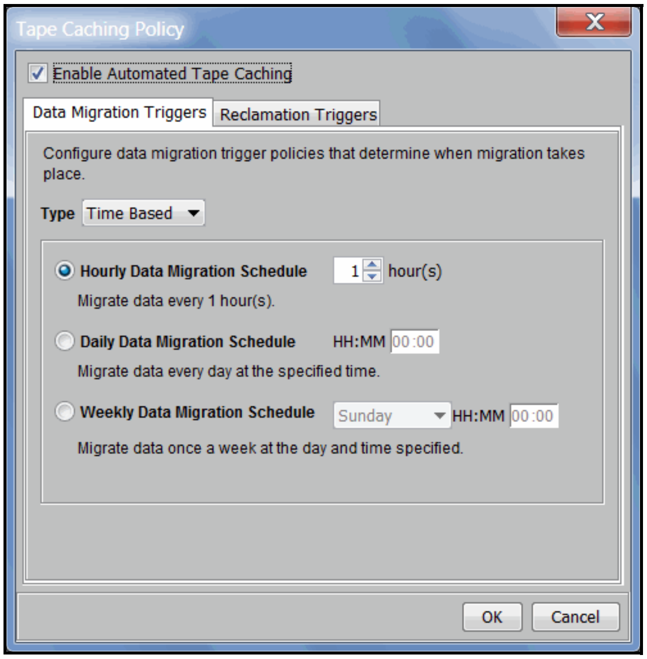

Tape Caching Policy

☑ Enable Automated Tape Caching

Configure data migration trigger policies that determine when migration takes place.

Type: Time Based

● Hourly Data Migration Schedule 1 hour(s)
Migrate data every 1 hour(s).

○ Daily Data Migration Schedule HH:MM 00:00
Migrate data every day at the specified time.

○ Weekly Data Migration Schedule Sunday HH:MM 00:00
Migrate data once a week at the day and time specified.

Hourly Data Migration Schedule - Migration occurs every *n* hours.

Daily Data Migration Schedule - Migration occurs at a specific time of day. Type the hour and minute (in 24-hour format) in the box. Note that if the trigger occurs past the specified time, the migration will occur at that time on the next day.

Weekly Data Migration Schedule - Migration occurs on a specific day of the week. Specify the day of the week from the list and type the hour and minute (in 24-hour format) in the text box. Note that if the trigger occurs past the specified time, the migration will occur at the next scheduled day and time.

For *Policy Based* triggers, determine what criteria will trigger migration.

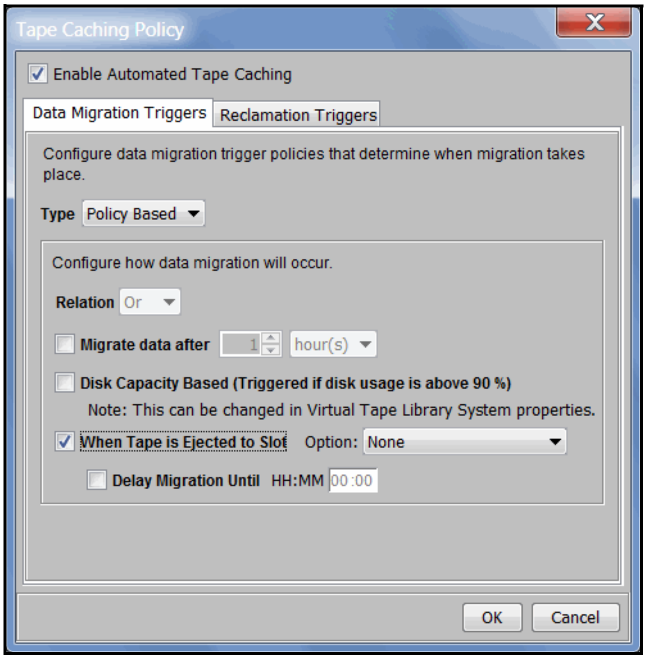

Tape Caching Policy

☑ Enable Automated Tape Caching

Configure data migration trigger policies that determine when migration takes place.

Type: Policy Based

Configure how data migration will occur.

Relation: Or

□ Migrate data after 1 hour(s)

□ Disk Capacity Based (Triggered if disk usage is above 90 %)
Note: This can be changed in Virtual Tape Library System properties.

☑ When Tape is Ejected to Slot Option: None

□ Delay Migration Until HH:MM 00:00

If multiple triggers are set, select *And* if all the triggers must be met to initiate data migration or select *Or* if meeting any one of them will initiate data migration.

For example, if you select both *Migrate data after* and *Disk Capacity Based*, and you select *And*, data migration will occur only when both the specified number of days/hours has elapsed and the specified disk capacity has been reached. If you select *Or*, the occurrence of either one of those events will trigger the data migration.

Migrate data after - Migration will occur when the data has been on the virtual disk for a specified number of hours or days. Specify the desired number of hours/days in the list box.

Disk Capacity Based - Migration will occur when the used disk space exceeds the specified disk capacity. The actual percentage is a global variable which is set for all virtual tape libraries. To change the number, right-click *Virtual Tape Library System* in the tree, click *Properties*, and type the desired percentage in the *Tape Caching Policy Disk Capacity Migration Threshold* box.

> **Note:** The *Tape Caching Policy Disk Capacity Threshold* setting affects other capacity-based actions as well

When Tape is Ejected to Slot - Migration will occur when a backup has completed and the virtual tape is ejected to a slot. If you select the option *Only When Tape is Full*, migration will only occur if the tape is full.

Delay Migration Until - Migration will be delayed until the time you specify after one of the above policies has been triggered. You may want to select a time when system usage is very light. Type the hour and minute (in 24-hour format) in the box.

4. Click the *Reclamation Triggers* tab and specify when the data that has been migrated to physical tape can be deleted to free up cache disk space.

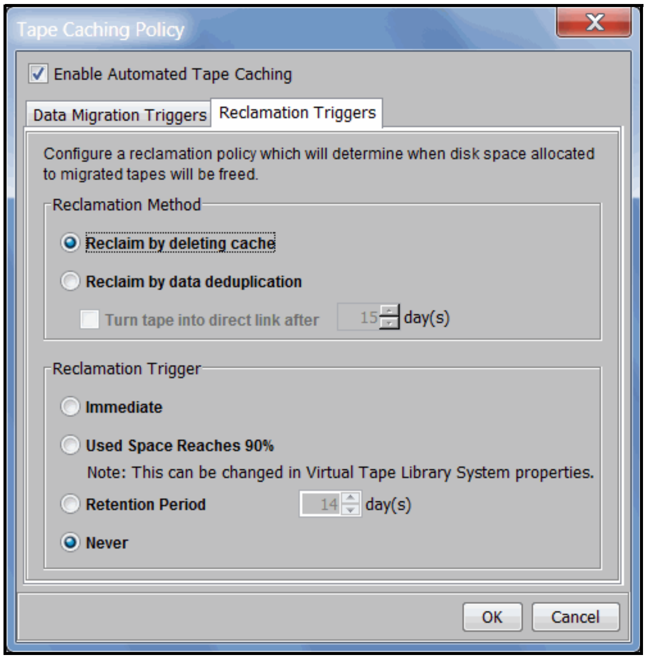

Tape Caching Policy

☑ Enable Automated Tape Caching

Data Migration Triggers | Reclamation Triggers

Configure a reclamation policy which will determine when disk space allocated to migrated tapes will be freed.

Reclamation Method

● Reclaim by deleting cache

○ Reclaim by data deduplication

□ Turn tape into direct link after 15 day(s)

Reclamation Trigger

○ Immediate

○ Used Space Reaches 90%

Note: This can be changed in Virtual Tape Library System properties.

○ Retention Period 14 day(s)

● Never

Reclamation methods:

Reclaim by deleting cache - After the cache is deleted, tapes become *direct link* tapes. A direct link tape is not an actual tape but a link to a physical tape. If your backup application ever overwrites a direct link tape, VTL will automatically start caching the physical tape. Direct link tapes are only deleted if you disable tape caching on the library or if you delete the tape. If you unassign a physical library, the direct link tape will not be deleted.

Reclaim by data deduplication - Deduplication is triggered and single instances of unique data are copied to the deduplication repository. After deduplication, the virtual tape is replaced with a virtual index tape (VIT) pointing to deduplication storage. In order to use this feature, you must have a deduplication cluster associated with this VTL server. After selecting this option, you will be prompted to create a deduplication policy. By default, this policy will have the same name as the tape library and all tape caching tapes created under this library will automatically be added to the policy. Regular virtual tapes will not be added to the policy and cannot be added manually to this policy. Also note that deduplication replication and migration with replication are mutually exclusive. Therefore, migration with replication will not be allowed.

Turn tape into direct link after n days - When used with the *Reclaim by data deduplication* option, VIT tapes will become direct link tapes after the specified number of days since the last successful migration. If the tape is in the middle of being deduplicated, the system will wait until deduplication completes.

Note: A physical tape can be ejected to a mail slot by a backup application only if it is associated with a direct link tape that was created using the *Reclaim by deleting cache* method. If the physical tape is associated with a tape reclaimed using the *Reclaim by data deduplication* method, it cannot be ejected by the backup application

Reclamation triggers:

Immediate - Cache disk space is freed up as soon as the data migration is complete.

Used Space Reaches n% - Cache disk space is freed up when the used space reaches this threshold. The actual percentage is a global variable which is set for all virtual tape libraries. To change the number, right-click *Virtual Tape Library System* in the tree, click *Properties*, and type the desired percentage in the *Tape Caching Policy Disk Capacity Reclamation Threshold* box.

Retention Period - Cache disk space is freed up after a specified number of days has elapsed. Specify the number of days that the data should be retained.

Never - Cache disk space is never freed up.

5. Click *OK*.

   The policy takes effect immediately.

Note: When you move a tape from the virtual tape library to a vault, it retains the Tape Caching policy associated with the original virtual tape library.

Set global tape caching options

You can set global tape caching options for all virtual tape libraries. To do this:

1. Right-click *Virtual Tape Library System* and select *Properties*.

   If the server is a member of a group, right-click the group and select *VTL Properties*.

2. Set the global migration and reclamation thresholds.

   Migration Threshold - Migration will occur when the used disk space exceeds the specified disk capacity.

   Reclamation Threshold - Cache disk space is freed up when the used space reaches this threshold.

Disable a policy

When Automated Tape Caching is disabled for a library, all of the tapes in the library with tape caching policies will be disabled.

> **Note:** You cannot disable tape caching for a library if:
>
> - The library has any direct link tapes (in a slot, drive, or the virtual vault). A list of all direct link tapes will be displayed and you will need to manually delete them before disabling tape caching.
> - There is a caching export job running on a cached tape from the library.
> - A cached tape is loaded in a drive.

To disable a tape caching policy:

1. Right-click a virtual tape library and click *Automated Tape Caching*.
2. Clear the *Enable Automated Tape Caching* check box.

   All the options that you previously set are retained, but data migration will not occur automatically until you select this check box again.

3. Click *OK*.

Create a cache for your physical tapes

Automated Tape Caching stores data on disk before it is migrated to physical tape. In order for this to happen, you must create a cache for each of your physical tapes. This is typically done after you create a virtual tape library. If this has not been done, you must sync the library to create a cache for your physical tapes. You also need to sync the library if a direct link tape was deleted and the associated physical tape is needed to recover data.

1. Right-click your virtual tape library and select *Sync Library*.
2. If you have multiple libraries, select the appropriate physical library.

   A physical tape library can only be synchronized to one virtual tape library at a time, unless the ACSLS or IBM 3494 option is being used.
3. Select the physical tape(s) for which you want to create a cache.

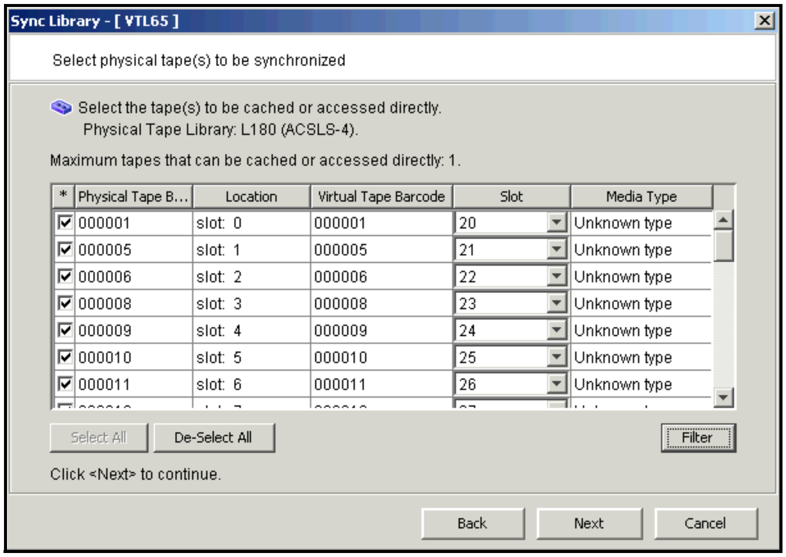

Select physical tape(s) to be synchronized

Select the tape(s) to be cached or accessed directly.
Physical Tape Library: L180 (ACSLS-4).

Maximum tapes that can be cached or accessed directly: 1.

*	Physical Tape B...	Location	Virtual Tape Barcode	Slot	Media Type
☑	000001	slot: 0	000001	20	Unknown type
☑	000005	slot: 1	000005	21	Unknown type
☑	000006	slot: 2	000006	22	Unknown type
☑	000008	slot: 3	000008	23	Unknown type
☑	000009	slot: 4	000009	24	Unknown type
☑	000010	slot: 5	000010	25	Unknown type
☑	000011	slot: 6	000011	26	Unknown type

Click <Next> to continue.

To display specific tapes, click the *Filter* button to select a single barcode or a range of barcodes. If you want to specify a particular starting/ending range number, select *Start With* or *End With* in the *From/To* fields. You can then type the number in the box to the right.

> **Note:** Make sure that you select physical tapes that use the same media type as your virtual tapes.

4. Select *Create Cache* and indicate if you want to use encryption.

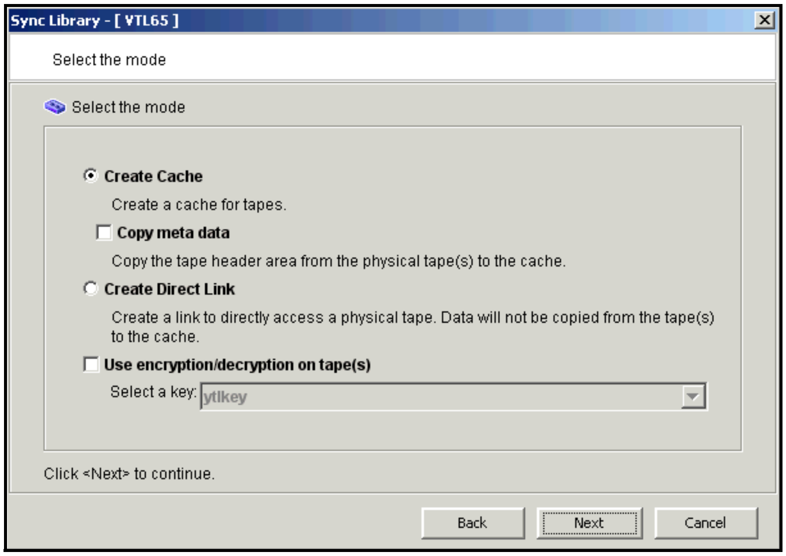

Sync Library - [VTL65]
Select the mode
Select the mode
● Create Cache
Create a cache for tapes.
□ Copy meta data
Copy the tape header area from the physical tape(s) to the cache.
○ Create Direct Link
Create a link to directly access a physical tape. Data will not be copied from the tape(s) to the cache.
□ Use encryption/decryption on tape(s)
Select a key: ytlkey
Click <Next> to continue.
Back
Next
Cancel

Copy meta data - Copies the tape header from the physical tape to the cache. Select this option if your backup application requires a tape header to identify a tape.

Use encryption/decryption on tape(s) - Select if you want to encrypt the data on the tape. You can select this option only if at least one key has been created. If you select this option, you must select the key to use. All the data on the tape will be indecipherable until is imported back to a virtual tape library and decrypted using the same key. For more information about encryption, refer to 'Manage encryption keys'.

5. If applicable, specify how to create the cache.

6. If applicable, specify a prefix and size.

7. Confirm all information and click *Finish*.

Create uncached virtual tapes

Even though you are using Automated Tape Caching for your tape library, you can still create *uncached* virtual tapes that will not be migrated to physical tapes. This can be useful for a single backup that is not part of your normal backup routine. You can create one or more virtual tapes by right-clicking a virtual tape library or on the *Tapes* object and selecting *New Tape(s)*.

Note that if you create virtual tapes, they cannot match the barcodes of your physical tapes.

Enable tape caching for existing virtual tapes

Tape caching can be enabled for existing uncached virtual tapes. This provides an easy way to migrate data on existing uncached virtual tapes to physical tapes without having to back up the data to new tapes.

Note that in order to migrate data to physical tapes, physical tape(s) with corresponding barcodes must exist in the physical library.

1. Right-click your virtual tape library and select *Enable Tape Caching on Tape(s)*.
2. Select the tapes you want and click *OK*.

   A job will be kicked off to migrate the data to physical tape.

Manually migrate cached data to physical tape

You can manually cause data in a cache to be migrated to physical tape even if the tape has been previously migrated. To do this, right-click a virtual tape cache and select *Migrate to Physical Tape*.

Force migration of an entire tape to physical tape

You can migrate an entire tape to physical tape. To do this, right-click a virtual tape cache and select *Force Migrate to Tape*.

Note: This will overwrite all data on the physical tape.

Reclaim disk space manually

You can manually cause the data that has been migrated to physical tape to be deleted to free up cache disk space. To do this for a single cache, right-click a virtual tape cache and select *Reclaim Disk Space*. Note that this will delete data to free up cache disk space.

To do this for multiple tape caches, right-click the *Virtual Tape Library System* object and select *Reclaim Disk Space*.

Renew cache for a direct link tape

If your backup application ever overwrites the direct link tape, VTL will automatically start caching the physical tape.

You can also manually renew the cache for a direct link tape. To do this, right-click a direct link tape and select *Renew Cache*.

Recover data using Automated Tape Caching

In a cached environment, tapes are always visible to the backup application regardless of whether the data is actually on disk or on a physical tape in the physical tape library. When it comes time to restore data, your backup application will seamlessly read the data from disk (if it is still there) or from the physical tape.

If the data is no longer on disk, when the direct link tape is to be mounted, the corresponding physical tape also needs to be mounted. When loading the physical tape, VTL will look for an available drive. If there are no free drives available, VTL will cancel any import or export jobs in an attempt to free a tape drive.

Note: If a direct link tape was deleted and the associated physical tape is needed to recover data, you will have to sync the library to create a direct link to the physical tape.

Fibre Channel Configuration

Overview

Just as the VTL server supports different types of storage devices (such as SCSI, Fibre Channel, and iSCSI), the VTL server is protocol-independent and supports multiple outbound target protocols, including Fibre Channel Target Mode.

This chapter provides configuration information for Fibre Channel Target Mode as well as the associated Fibre Channel SAN equipment.

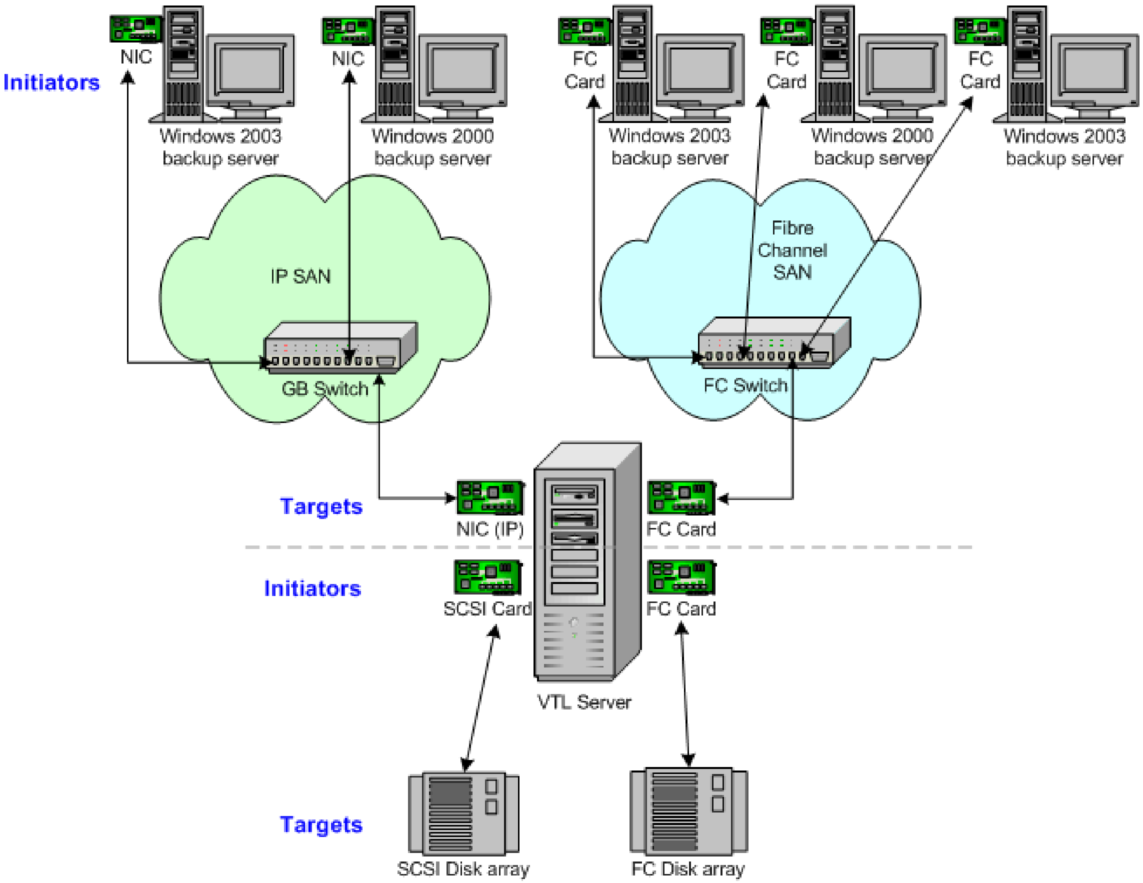

As you can see from the illustration above, an application server can be either an iSCSI client or a Fibre Channel client, but not both. Using separate cards and switches, you can have all types of VTL Clients (FC and iSCSI) on your network.

Configure Fibre Channel hardware on server

VTL supports the use of QLogic HBAs for the VTL server. Refer to the certification matrix on the DSI website for a list of HBAs that are currently certified.

Ports

Your VTL appliance will be equipped with several Fibre Channel ports. Some of these ports will interface with storage arrays. Others will interface with physical tape libraries, while the remaining ports will interface with backup (media) servers.

The ports that connect to storage arrays are commonly known as *Initiator Ports*.

The ports that will interface with the backup application servers' FC initiator ports will run in a different mode known as *Target Mode*.

The ports that are connected to physical tape libraries are known as *Library Connection Ports*.

HBA driver

QLogic NPIV driver

NPIV (N_Port ID Virtualization) is the default driver for VTL servers. NPIV allows a port to have the role of both initiator and target in full-duplex mode.

The following is required in order to use NPIV:

- You must have a supported HBA. VTL supports both 4Gb and 8Gb HBAs. Check the DSI certification matrix for a list of supported HBAs.
- The fabric switch must support NPIV.
- If a QLogic FC switch is being used, you must disable “IOStreamGuard” for any switch port that connects to an NPIV target port.

When using the NPIV driver, there are two WWPNs, the *base* port and the *alias*.

Notes:

- With dual mode, clients will need to be zoned to the alias port (called *Target WWPN*). If they are zoned to the base port, clients will not see any devices.
- You will only see the alias port when that port is in target mode.
- You will only see the alias once all of the VTL services are started.

QLogic driver

The QLogic driver is the single-mode, point-to-point driver where targets and initiators reside on separate ports.

Zoning

Note: If a port is connected to a switch, we highly recommend the port be in at least one zone so it will display in your SNS table.

There are two types of zoning that can be configured on each switch, soft zoning (based on WWPNs), and hard zoning (based on port #).

Soft zoning

Soft zoning is required for the QLogic NPIV driver and uses the WWPN in the configuration. The WWPN remains the same in the zoning configuration regardless of the port location. If a port fails, you can simply move the cable from the failed port to another valid port without having to reconfigure the zoning.

The rules for soft zoning between VTL and deduplication servers using the QLogic NPIV driver are:

- No VTL or deduplication server target can be zoned with another of its own targets.
- No VTL or deduplication server target can be zoned with one of its own initiators.
- No VTL or deduplication server initiators can be zoned with one of its own initiators.

Hard zoning

Hard zoning is only supported for QLogic drivers without NPIV. It uses the port number of the switches for zoning. With hard zoning, if a zone has two ports (0 and 1) and port 0 goes down for some reason, you will need to remove the current zoning configuration, move the plug to another valid port, re-zone, and then enable the new zoning configuration.

If hard zoning is used, it is necessary to create zones for each standby target, doubling the number of upstream zones. This extra set of zones is not necessary in the case of soft zoning because zones are defined by WWPN combinations. In a failover event, the standby ports assume the WWPNs of the target ports of the failed VTL server. Therefore, the single set of soft zones is still valid.

General zoning requirements

VTL recommends isolated zoning, where one initiator is zoned to one target in order to minimize I/O interruptions by non-related FC activities, such as port login/out and resets. This does not apply in the case of FC connectivity between VTL and deduplication appliances.

Additionally, make sure that storage devices to be used by VTL are not zoned to clients (backup application servers). Ports on storage devices to be used by VTL should be zoned to VTL’s initiator ports while the clients are zoned to VTL’s target ports. Make sure that from the storage unit’s management GUI (such as SANtricity and NaviSphere), the LUNs are re-assigned to VTL as the “host”. VTL will virtualize these LUNS. VTL can then define virtual tapes out of these LUNS and further provision them to the clients.

Persistent binding

Persistent binding is automatically enabled for all QLogic HBAs connected to storage device targets upon the discovery of the device (via a Console physical device rescan with the *Discover New Devices* option enabled). However, persistent binding will not be SET until the HBA is reloaded. You can reload HBAs by restarting VTL with the command:

```
vtl restart all
```

Without persistent binding, there is a risk that the wrong storage controller port will be accessed when the VTL appliance is rebooted (or VTL HBA driver is reloaded).

FSHBA.CONF file

The *fshba.conf* file is found in $ISHOME/etc and is used to adjust settings for FC adapters installed on the VTL appliance.

1. Determine the HBA settings to change.
2. Back up the *fshba.conf* file:

   ```
   cp fshba.conf fshba.conf.bak
   ```

3. Modify *fshba.conf* using the *vi* editor.
4. Save the *fshba.conf* file.
5. Start or restart VTL and its HBA module with the following command:

   ```
   vtl start all
   ```

   You must restart the HBA drivers and VTL modules on the VTL appliance for the changes in the *fshba.conf* file to take effect and to recognize the new settings.

Link speed

In the *fshba.conf* file, the link speed is set to auto-negotiate by default for every FC port. You must manually update this and match the link speed with the switch speed. It may be necessary to manually set the port switch speed on the FC switch as well.

If you are attaching a tape library or storage array directly to the VTL appliance, adjust the link speed for all FC ports (VTL and/or tape library). Check with your vendor to obtain any recommended FC HBA settings.

Device identification

Typically, Linux will assign its own device numbers, such as SCSI adapter0 and SCSI adapter1, etc. Therefore, if you have a single port QLogic HBA loading up AFTER two internal SCSI devices, it will become SCSI adapter2.

However, the VTL appliance may not identify the same devices in the same way. VTL will identify SCSI devices as hba0, hba1, hba2, and so on in the *fshba.conf* file. Settings for each individual FC port (for example, hba0 or hba1) can be modified in fshba.conf.

To identify which adapter belongs to which HBA in fshba.conf:

1. Run the following commands:

   ```
   ls /proc/scsi/qla2xxx
   ```

   This will output all adapter numbers (i.e. 100, 101, 102). Then, match up the adapter numbers: 100->hba0, 101->hba1, etc.

2. Run the following command to display the WWPN for that adapter:

   ```
   grep BIOSWWPN /proc/scsi/qla2xxx/###
   ```

3. Run the following command to determine which physical port belongs to each adapter number in fshba.conf:

   ```
   tail -f /var/log/messages
   ```

   and then unplug the FC port. You will see a loop down message like the one below. 100 is the adapter number in this example:

   ```
   Jan 23 13:40:03 <hostname> kernel: scsi(100): LOOP DOWN detected
   ```

Data rate

1. Scroll down to the appropriate adapter section.
2. Search for *data_rate-hbaX*.

   It should look like this:

   ```
   #data_rate-hbaX=2
   #comment=this option allow driver software to select a fixed rate or
   #        request that the firmware negotiate the
   #        data rate (1-2G,2-auto,3-4G,4-8G)
   #range=0 or 1 or 2
   #===============
   data_rate-hba0=2
   data_rate-hba1=2
   ```

3. For the adapter to be configured (i.e., hba0), change the value:

   ```
   data_rate-hba0=0 or 1 NOT 2 (auto)
   ```

4. Repeat for each adapter to be configured.

WWNN

If you are using redundant node failover and your storage binds to the World Wide Node Name (WWNN) **and** the World Wide Port Name (WWPN), it may be necessary to set the WWNN so that when failover occurs, the takeover server can maintain its storage connection. To set the WWNN, add a “wwnn” line for each Fibre Channel port on each node in your N+1 cluster. All “wwnn” lines can have identical node names for all ports and all nodes. For example, if you have four ports, you might add the following:

```
wwnn0=2100010001000100
wwnn1=2100010001000100
wwnn2=2100010001000100
wwnn3=2100010001000100
```

Configure Fibre Channel hardware on clients

Fabric topology

(For all clients *except* Solaris SPARC clients) When setting up clients on a Fibre Channel network using a Fabric topology, we recommend that you set the topology that each HBA will use to log into your switch to *Point-to-Point Only*.

> **Note:** We recommend hard coding the link speed of the HBA to be in line with the switch speed.

Load balance the path for each downstream storage LUN

For optimal performance, if you have more than one path available for your storage LUNs, you can set VTL to evenly distribute I/O between all storage LUNs. To do this:

1. Right-click on a Fibre Channel device under *Physical Resources --> Storage Devices --> Fibre Channel Devices* and select *Properties*.

2. On the *I/O Path* tab, highlight a path and use the arrow keys to move it up or down in the list.

Verify your hardware configuration

After all of your Fibre Channel hardware has been configured, you should verify that everything is set correctly. You can do this in the VTL console by highlighting *Storage HBAs* under *Physical Resources*.

General tab

The *General* tab displays information about the port, including mode (dual, target, or initiator), status, and WWPN.

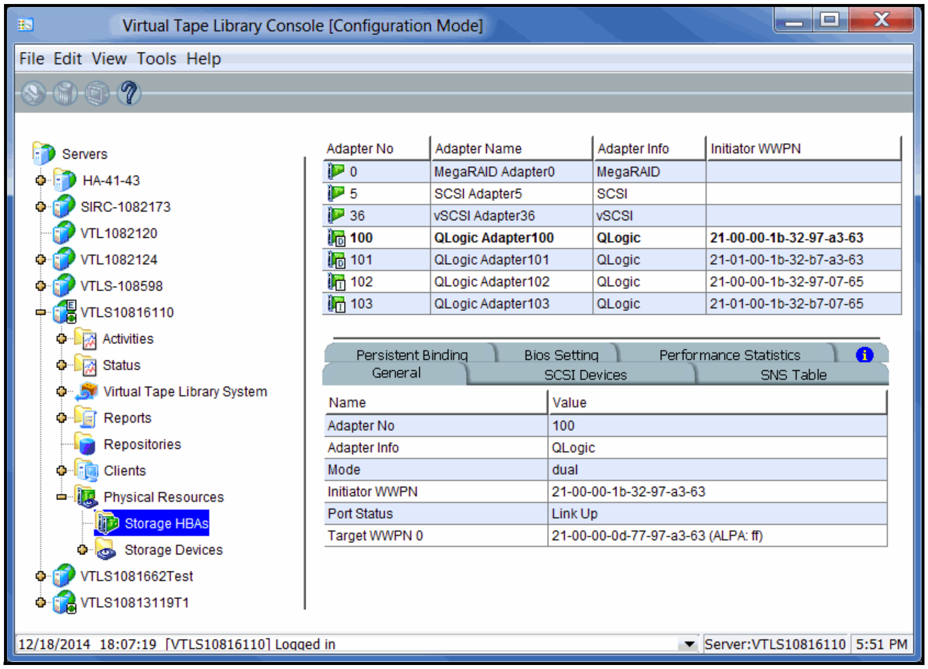

Adapter No	Adapter Name	Adapter Info	Initiator WWPN
0	MegaRAID Adapter0	MegaRAID	
5	SCSI Adapter5	SCSI	
36	vSCSI Adapter36	vSCSI	
100	**QLogic Adapter100**	**QLogic**	**21-00-00-1b-32-97-a3-63**
101	QLogic Adapter101	QLogic	21-01-00-1b-32-b7-a3-63
102	QLogic Adapter102	QLogic	21-00-00-1b-32-97-07-65
103	QLogic Adapter103	QLogic	21-01-00-1b-32-b7-07-65

Name	Value
Adapter No	100
Adapter Info	QLogic
Mode	dual
Initiator WWPN	21-00-00-1b-32-97-a3-63
Port Status	Link Up
Target WWPN 0	21-00-00-0d-77-97-a3-63 (ALPA: ff)

12/18/2014 18:07:19 [VTLS10816110] Logged in — Server:VTLS10816110 — 5:51 PM

SCSI Devices tab

The SCSI Devices tab lists the SCSI storage devices attached to this adapter. If you expect to see a device that is not listed, right-click the adapter and select *Rescan*.

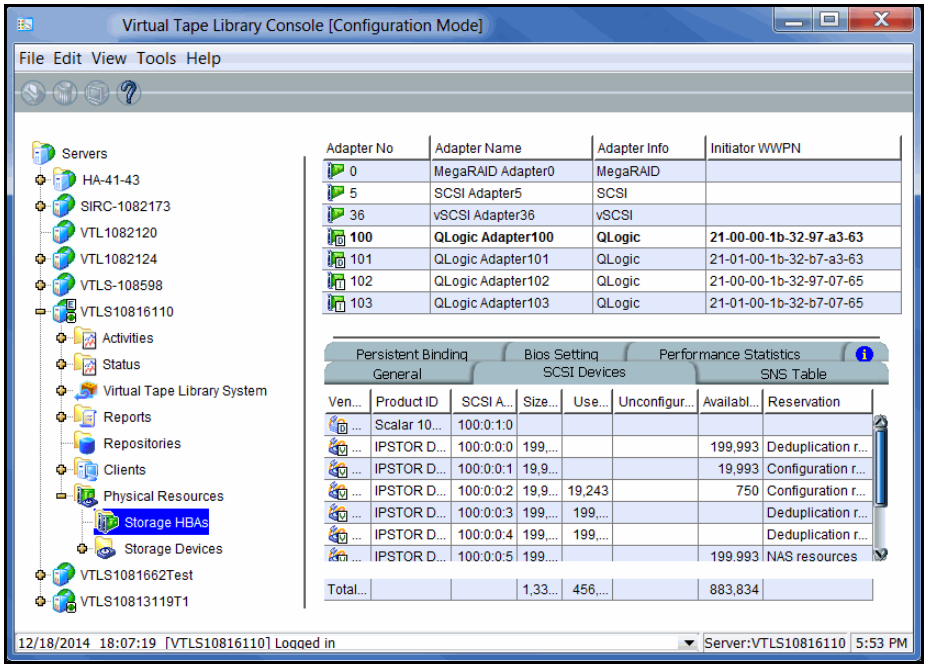

Adapter No	Adapter Name	Adapter Info	Initiator WWPN
0	MegaRAID Adapter0	MegaRAID	
5	SCSI Adapter5	SCSI	
36	vSCSI Adapter36	vSCSI	
100	**QLogic Adapter100**	**QLogic**	**21-00-00-1b-32-97-a3-63**
101	QLogic Adapter101	QLogic	21-01-00-1b-32-b7-a3-63
102	QLogic Adapter102	QLogic	21-00-00-1b-32-97-07-65
103	QLogic Adapter103	QLogic	21-01-00-1b-32-b7-07-65

Ven...	Product ID	SCSI A...	Size...	Use...	Unconfigur...	Availabl...	Reservation
...	Scalar 10...	100:0:1:0					
...	IPSTOR D...	100:0:0:0	199,...			199,993	Deduplication r...
...	IPSTOR D...	100:0:0:1	19,9...			19,993	Configuration r...
...	IPSTOR D...	100:0:0:2	19,9...	19,243		750	Configuration r...
...	IPSTOR D...	100:0:0:3	199,...	199,...			Deduplication r...
...	IPSTOR D...	100:0:0:4	199,...	199,...			Deduplication r...
...	IPSTOR D...	100:0:0:5	199,...			199,993	NAS resources
Total...			1,33...	456,...		883,834	

12/18/2014 18:07:19 [VTLS10816110] Logged in — Server:VTLS10816110 — 5:53 PM

SNS Table tab

The SNS Table tab lists the ports to which this adapter is zoned. VTL queries the switch for its Simple Name Server (SNS) database and displays this information. If you expect to see a WWPN that is not listed, right-click the adapter and select *Refresh SNS.*

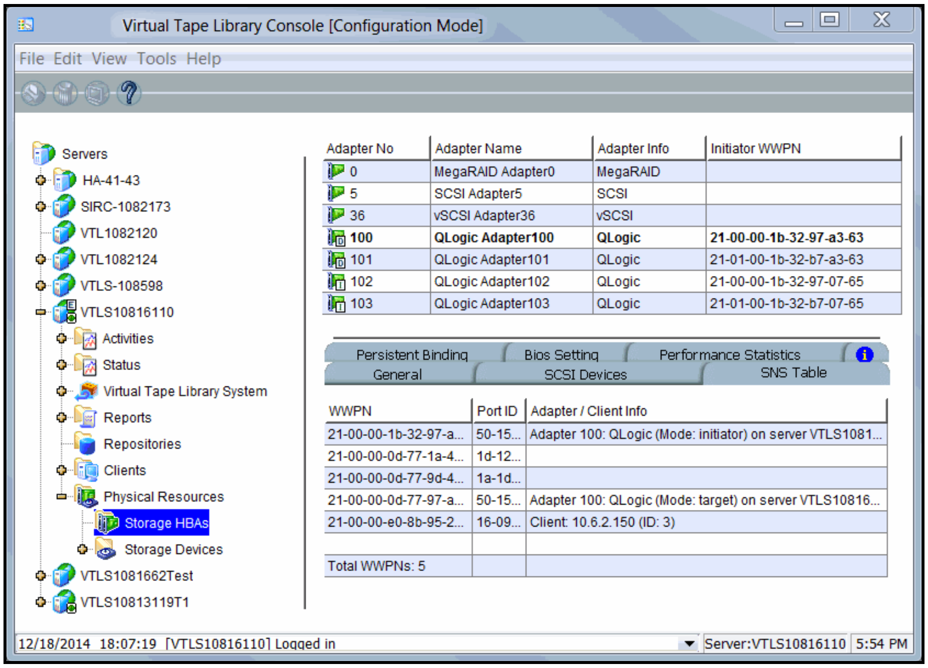

Adapter No	Adapter Name	Adapter Info	Initiator WWPN
0	MegaRAID Adapter0	MegaRAID	
5	SCSI Adapter5	SCSI	
36	vSCSI Adapter36	vSCSI	
100	**QLogic Adapter100**	**QLogic**	**21-00-00-1b-32-97-a3-63**
101	QLogic Adapter101	QLogic	21-01-00-1b-32-b7-a3-63
102	QLogic Adapter102	QLogic	21-00-00-1b-32-97-07-65
103	QLogic Adapter103	QLogic	21-01-00-1b-32-b7-07-65

WWPN	Port ID	Adapter / Client Info
21-00-00-1b-32-97-a...	50-15...	Adapter 100: QLogic (Mode: initiator) on server VTLS1081...
21-00-00-0d-77-1a-4...	1d-12...	
21-00-00-0d-77-9d-4...	1a-1d...	
21-00-00-0d-77-97-a...	50-15...	Adapter 100: QLogic (Mode: target) on server VTLS10816...
21-00-00-e0-8b-95-2...	16-09...	Client: 10.6.2.150 (ID: 3)
Total WWPNs: 5		

12/18/2014 18:07:19 [VTLS10816110] Logged in | Server:VTLS10816110 | 5:54 PM

Persistent Binding tab

(Initiator ports only) The Persistent Binding tab lists all of the target ports to which this adapter is bound.

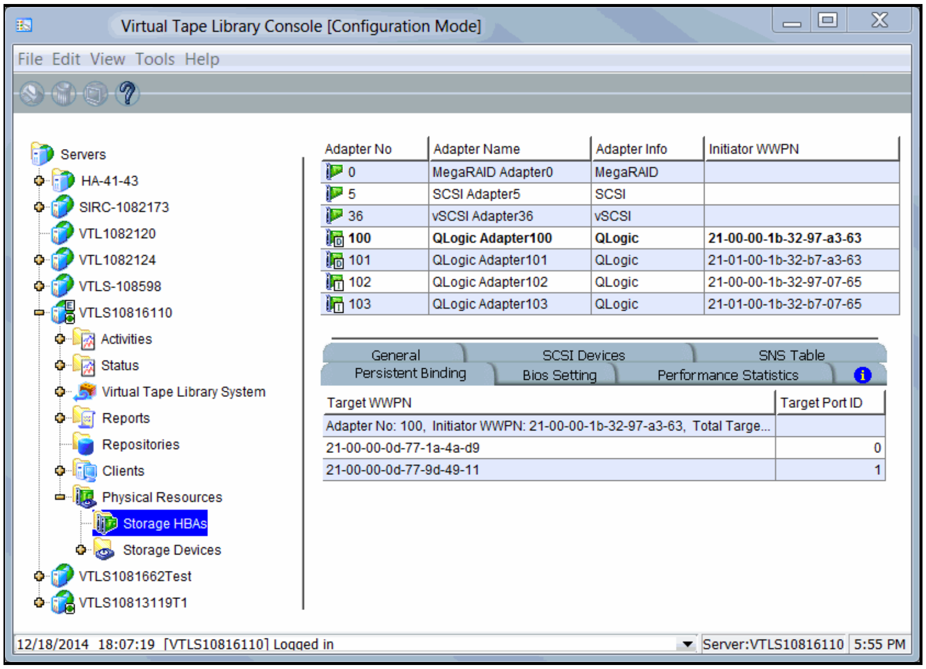

Adapter No	Adapter Name	Adapter Info	Initiator WWPN
0	MegaRAID Adapter0	MegaRAID	
5	SCSI Adapter5	SCSI	
36	vSCSI Adapter36	vSCSI	
100	**QLogic Adapter100**	**QLogic**	**21-00-00-1b-32-97-a3-63**
101	QLogic Adapter101	QLogic	21-01-00-1b-32-b7-a3-63
102	QLogic Adapter102	QLogic	21-00-00-1b-32-97-07-65
103	QLogic Adapter103	QLogic	21-01-00-1b-32-b7-07-65

Target WWPN	Target Port ID
Adapter No: 100, Initiator WWPN: 21-00-00-1b-32-97-a3-63, Total Targe...	
21-00-00-0d-77-1a-4a-d9	0
21-00-00-0d-77-9d-49-11	1

12/18/2014 18:07:19 [VTLS10816110] Logged in | Server:VTLS10816110 | 5:55 PM

Bios Setting tab

The Bios Setting tab lists all of the HBA settings for this adapter so that you can confirm what is set.

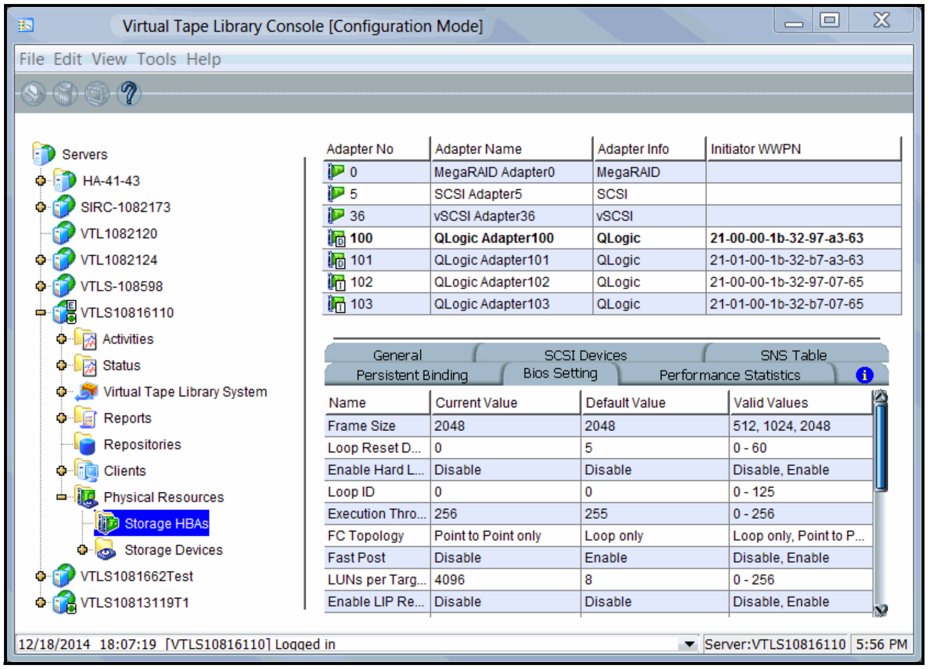

Adapter No	Adapter Name	Adapter Info	Initiator WWPN
0	MegaRAID Adapter0	MegaRAID	
5	SCSI Adapter5	SCSI	
36	vSCSI Adapter36	vSCSI	
100	**QLogic Adapter100**	**QLogic**	**21-00-00-1b-32-97-a3-63**
101	QLogic Adapter101	QLogic	21-01-00-1b-32-b7-a3-63
102	QLogic Adapter102	QLogic	21-00-00-1b-32-97-07-65
103	QLogic Adapter103	QLogic	21-01-00-1b-32-b7-07-65

Name	Current Value	Default Value	Valid Values
Frame Size	2048	2048	512, 1024, 2048
Loop Reset D...	0	5	0 - 60
Enable Hard L...	Disable	Disable	Disable, Enable
Loop ID	0	0	0 - 125
Execution Thro...	256	255	0 - 256
FC Topology	Point to Point only	Loop only	Loop only, Point to P...
Fast Post	Disable	Enable	Disable, Enable
LUNs per Targ...	4096	8	0 - 256
Enable LIP Re...	Disable	Disable	Disable, Enable

12/18/2014 18:07:19 [VTLS10816110] Logged in | Server:VTLS10816110 | 5:56 PM

Performance Statistics tab

The *Performance Statistics* tab displays a chart showing read and write throughput for the last 60 minutes. Current performance is also displayed. All information is displayed in MB per second.

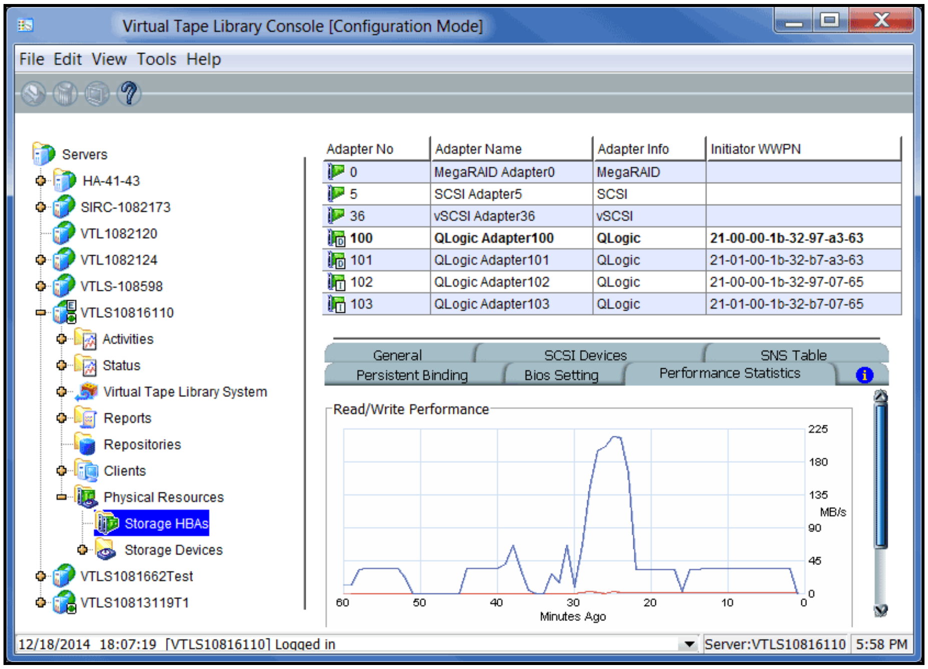

Adapter No	Adapter Name	Adapter Info	Initiator WWPN
0	MegaRAID Adapter0	MegaRAID	
5	SCSI Adapter5	SCSI	
36	vSCSI Adapter36	vSCSI	
100	**QLogic Adapter100**	**QLogic**	**21-00-00-1b-32-97-a3-63**
101	QLogic Adapter101	QLogic	21-01-00-1b-32-b7-a3-63
102	QLogic Adapter102	QLogic	21-00-00-1b-32-97-07-65
103	QLogic Adapter103	QLogic	21-01-00-1b-32-b7-07-65

12/18/2014 18:07:19 [VTLS10816110] Logged in | Server:VTLS10816110 | 5:58 PM

Set QLogic ports to target mode

Multi port QLogic HBAs

With a multi-ID HBA, each port can be both a target and an initiator. To use target mode, you must enable target mode on a port.

For VTL Failover, if you are using multi-ID HBAs, you need a minimum of one port for client connectivity and one for storage.

To set target mode:

1. In the Console, expand *Physical Resources*.
2. Right-click a multi-ID HBA and select *Options* --> *Enable Target Mode*.
3. Click *OK* to enable.

   Afterwards, you will see two WWPNs listed for the port. The first is the base WWPN and the second is the Target WWPN (also known as the alias port). Clients need to be zoned to this port in order to see devices.

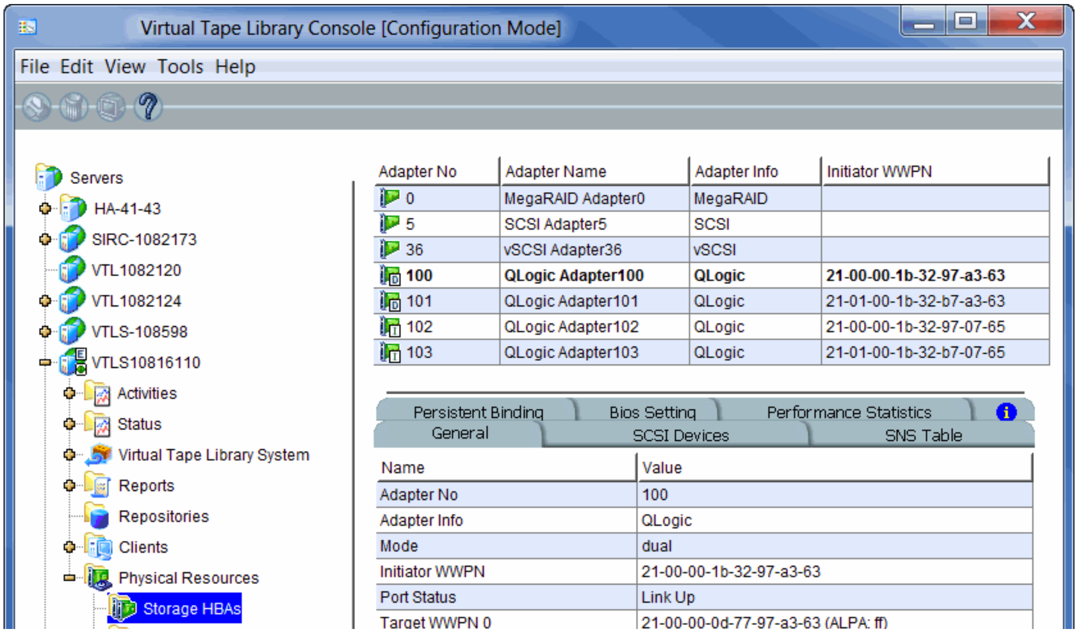

Adapter No	Adapter Name	Adapter Info	Initiator WWPN
0	MegaRAID Adapter0	MegaRAID	
5	SCSI Adapter5	SCSI	
36	vSCSI Adapter36	vSCSI	
100	**QLogic Adapter100**	**QLogic**	**21-00-00-1b-32-97-a3-63**
101	QLogic Adapter101	QLogic	21-01-00-1b-32-b7-a3-63
102	QLogic Adapter102	QLogic	21-00-00-1b-32-97-07-65
103	QLogic Adapter103	QLogic	21-01-00-1b-32-b7-07-65

Name	Value
Adapter No	100
Adapter Info	QLogic
Mode	dual
Initiator WWPN	21-00-00-1b-32-97-a3-63
Port Status	Link Up
Target WWPN 0	21-00-00-0d-77-97-a3-63 (ALPA: ff)

Single port QLogic HBAs

By default, all QLogic point-to-point ports are set to initiator mode, which means they will initiate requests rather than receive them. Determine which ports you want to use in target mode and set them to become target ports so that they can receive requests from your Fibre Channel Clients.

For VTL Failover, if you are using single port HBAs, you minimally need two ports for client connectivity (one for normal operation, one for standby) and one initiator port to connect to storage.

You need to switch one of those initiators into target mode so your clients will be able to see the VTL server. You will then need to select the equivalent adapter on the secondary server and switch it to target mode.

> **Note:** If a port is in initiator mode and has devices attached to it, that port cannot be set for target mode.

To set a port:

1. In the console, expand *Physical Resources*.
2. Right-click an HBA and select *Options* --> *Enable Target Mode*.

   You will get a *Loop Up* message on your VTL server if the port has successfully been placed in target mode.
3. When done, make a note of all of your WWPNs.

   It may be convenient for you to highlight your server and take a screenshot of the console.

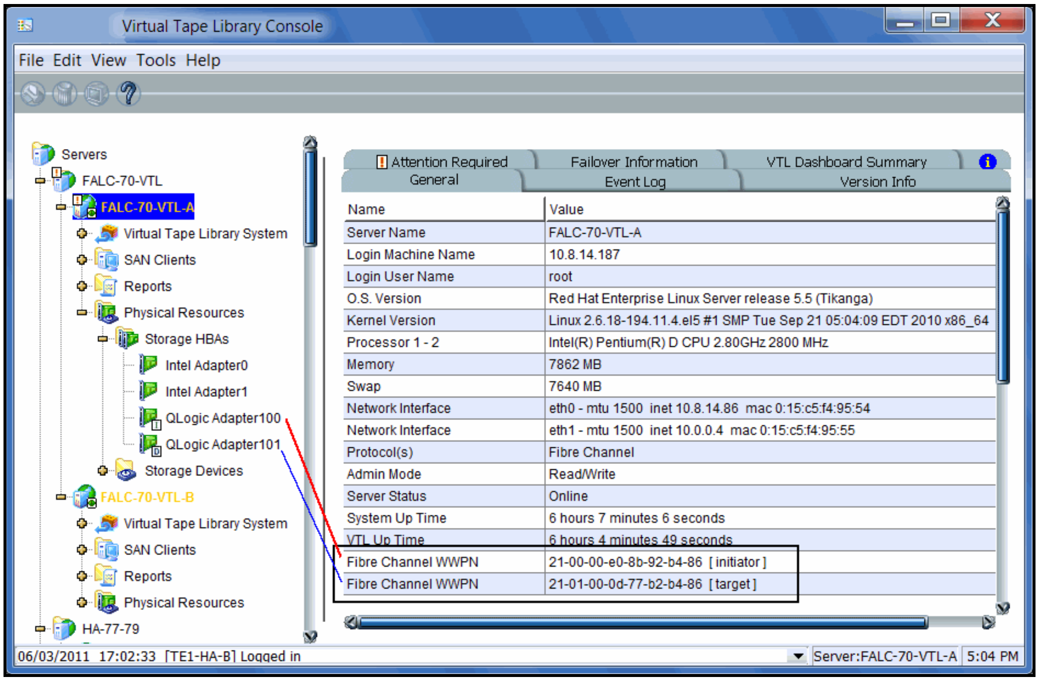

Name	Value
Server Name	FALC-70-VTL-A
Login Machine Name	10.8.14.187
Login User Name	root
O.S. Version	Red Hat Enterprise Linux Server release 5.5 (Tikanga)
Kernel Version	Linux 2.6.18-194.11.4.el5 #1 SMP Tue Sep 21 05:04:09 EDT 2010 x86_64
Processor 1 - 2	Intel(R) Pentium(R) D CPU 2.80GHz 2800 MHz
Memory	7862 MB
Swap	7640 MB
Network Interface	eth0 - mtu 1500 inet 10.8.14.86 mac 0:15:c5:f4:95:54
Network Interface	eth1 - mtu 1500 inet 10.0.0.4 mac 0:15:c5:f4:95:55
Protocol(s)	Fibre Channel
Admin Mode	Read/Write
Server Status	Online
System Up Time	6 hours 7 minutes 6 seconds
VTL Up Time	6 hours 4 minutes 49 seconds
Fibre Channel WWPN	21-00-00-e0-8b-92-b4-86 [initiator]
Fibre Channel WWPN	21-01-00-0d-77-b2-b4-86 [target]

06/03/2011 17:02:33 [TE1-HA-B] Logged in

Server:FALC-70-VTL-A 5:04 PM

Associate World Wide Port Names with clients

Similar to an IP address, the WWPN uniquely identifies a port in a Fibre Channel environment. Unlike an IP address, the WWPN is vendor assigned and is hardcoded and embedded.

Depending upon whether or not you are using a switched Fibre Channel environment, determining the WWPN for each port *may* be difficult.

- If you are using a switched Fibre Channel environment, VTL will query the switch for its Simple Name Server (SNS) database and will display a list of all available WWPNs. You will still have to identify which WWPN is associated with each machine.
- If you are not using a switched Fibre Channel environment, you can manually determine the WWPN for each of your ports. There are different ways to determine it, depending upon the hardware vendor. You may be able to get the WWPN from the BIOS during boot up or you may have to read it from the physical card. Check with your hardware vendor for their preferred method.

Do the following for each client for which you want to assign specific virtual devices:

1. Highlight the Fibre Channel Client in the console.
2. Right-click the protocol under the client and select *Properties*.

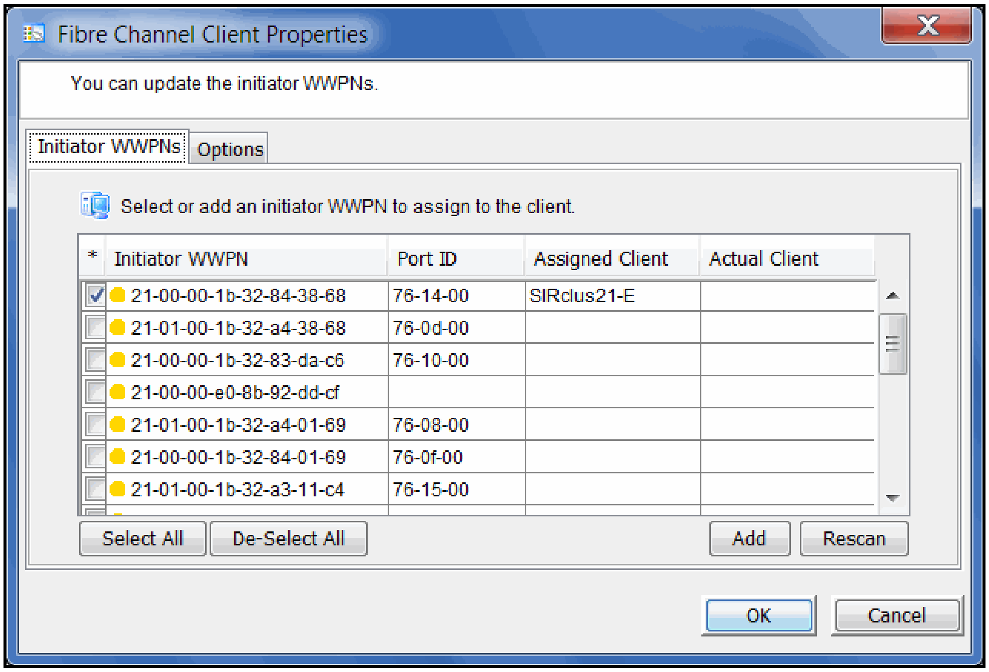

Fibre Channel Client Properties

You can update the initiator WWPNs.

Select or add an initiator WWPN to assign to the client.

*	Initiator WWPN	Port ID	Assigned Client	Actual Client
☑	● 21-00-00-1b-32-84-38-68	76-14-00	SIRclus21-E	
☐	● 21-01-00-1b-32-a4-38-68	76-0d-00		
☐	● 21-00-00-1b-32-83-da-c6	76-10-00		
☐	● 21-00-00-e0-8b-92-dd-cf			
☐	● 21-01-00-1b-32-a4-01-69	76-08-00		
☐	● 21-00-00-1b-32-84-01-69	76-0f-00		
☐	● 21-01-00-1b-32-a3-11-c4	76-15-00		

3. Select the Initiator WWPN(s) belonging to your client.

   Here are some methods to determine the WWPN of your clients:

   - Most Fibre Channel switches allow administration of the switch through an Ethernet port. These administration applications have utilities to reveal or allow you to change the following: Configuration of each port on the switch, zoning configurations, the WWPNs of connected Fibre Channel cards, and the current status of each connection. You can use this utility to view the WWPN of each Client connected to the switch.

- When starting up your Client, there is usually a point at which you can access the BIOS of your Fibre Channel card. The WWPN can be found there.
- The first time a new Client connects to the VTL server, the following message appears on the server screen:
  FSQLtgt: New Client WWPN Found: 21 00 00 e0 8b 43 23 52

4. If necessary, click *Add* to add WWPNs for the client.

   You will see the following dialog if there are no WWPNs in the server's list. This could occur because the client machines were not turned on or because all WWPNs were previously associated with clients.

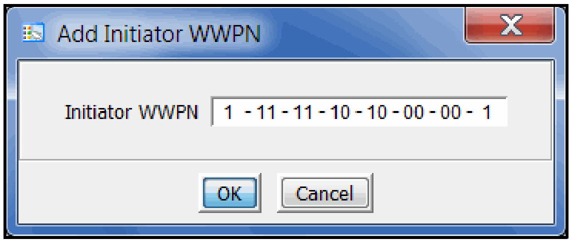

Add Initiator WWPN

Initiator WWPN: 1 - 11 - 11 - 10 - 10 - 00 - 00 - 1

iSCSI Configuration

Overview

The VTL server is protocol-independent and supports multiple outbound target protocols, including iSCSI Target Mode.

iSCSI builds on top of the regular SCSI standard by using the IP network as the connection link between various entities involved in a configuration. iSCSI inherits many of the basic concepts of SCSI. For example, just like SCSI, the entity that makes requests is called an *initiator*, while the entity that responds to requests is called a *target*. Only an initiator can make requests to a target; not the other way around. Each entity involved, initiator or target, is uniquely identified.

By default, when a client machine is added as an iSCSI client of a VTL server, it becomes an iSCSI initiator.

The initiator name is important because it is the main identity of an iSCSI initiator.

Supported platforms

iSCSI target mode is supported for the following platforms:

- Windows
- Linux

iSCSI users

VTL iSCSI Users are used for iSCSI protocol login authentication from iSCSI backup application servers. When you configure access for backup application servers, you designate users who can authenticate for the client.

There are several ways to create iSCSI users:

- Use the *Account Management* function in the VTL console and select *VTL iSCSI User* from the *Group* list. Create at least one unique user for each client.
- Add users when the *Add Client* function requires you to add/select users who can authenticate for the client.
- Add users to an existing client in *iSCSI Client Properties*.

Windows configuration

Requirements

- A VTL server with an Ethernet adapter installed.
- A Windows client machine.
- iSCSI software initiator installed on each backup application server. iSCSI initiator software/hardware is available from many sources. You can download the Microsoft iSCSI initiator from Microsoft's website: http://www.microsoft.com/windowsserversystem/storage/iscsi.mspx

Prepare client initiators to access your VTL server

Before a backup application server (the client initiator) can communicate with a VTL server, the two entities need to mutually recognize each other. Use an iSCSI initiator on every backup application server that will access the VTL server using iSCSI. This will let you add the VTL server as a target portal and log the client onto the iSCSI target you create on the VTL server.

The following steps are for the Microsoft iSCSI Initiator. If you are using a different iSCSI initiator, refer to the documentation provided by the vendor.

1. Run *Microsoft iSCSI Initiator* on the backup application server.

   You can find the program in the Control Panel or on your desktop (if you are the user that installed it).

2. Click the *Discovery* tab, then click *Add* under the *Target Portals* group box.

3. Enter the VTL server's IP address or name (if resolvable).

   To determine the IP address, go to the VTL console. Select the VTL server object. The IP address is on the *Login Machine Name* line in the right-hand pane of the Console.

   Use the default port (3260) and then click OK to add the client.

Enable iSCSI

Note: You cannot enable or disable iSCSI if you have already configured VTL failover. If you want to change the state of iSCSI, you need to remove your VTL failover configuration first.

In order to add a client using the iSCSI protocol, you must enable iSCSI for your VTL server.

If you haven't already done so, right-click your VTL server in the VTL console and select *Options* --> *Enable iSCSI*.

The following sections take you through the process of configuring iSCSI clients to work with the VTL server.

Add an iSCSI client

1. In the VTL console, right-click *Clients* and select *Add*.
2. Enter the client name.
3. Select *iSCSI*.
4. Select the initiator that this client uses.

   iSCSI clients correspond to specific iSCSI client initiators, and consequently, the client machines that own the specific initiator names. When a client connects to the VTL server, it can access only the resources assigned to a specific initiator name.

   By default, when a backup application server is added as an iSCSI client of a VTL server, it becomes an iSCSI initiator. The initiator name is important because it is the main identity of an iSCSI initiator. If you already added the VTL server as a Target Portal using the iSCSI initiator on your backup application server, the initiator name and backup application server IP address appear in the dialog.

   Otherwise, click *Add* and add the initiator name manually. (The IP address will not display.

   An available initiator shows a green dot; select the initiator name that is associated with the backup application server's IP address.

5. Add/select users who can authenticate for this client.

   To define authenticated access (using CHAP), select *Select or add users who can authenticate for the client.* iSCSI users you have already created in the VTL console are displayed. You can select one of these users or select *Add* to create a new user.

   More than one username/password pair can be assigned to the client, but they will be useful only when coming from the machine with an authorized initiator name.

   For unauthenticated access, select *Allow unauthenticated access*. The VTL server will recognize the client as long as it has an authorized initiator name.

6. Confirm all information and click *Finish*.

Create targets for the iSCSI client to log onto

1. In the VTL console, create at least one virtual iSCSI device (i.e. a virtual tape library) that can be used for iSCSI clients but do not assign it/them to the iSCSI clients until a target is created.

2. Expand the *Clients* object until you see the *iSCSI* object.

3. Right-click the *iSCSI* object and select *Create Target*.

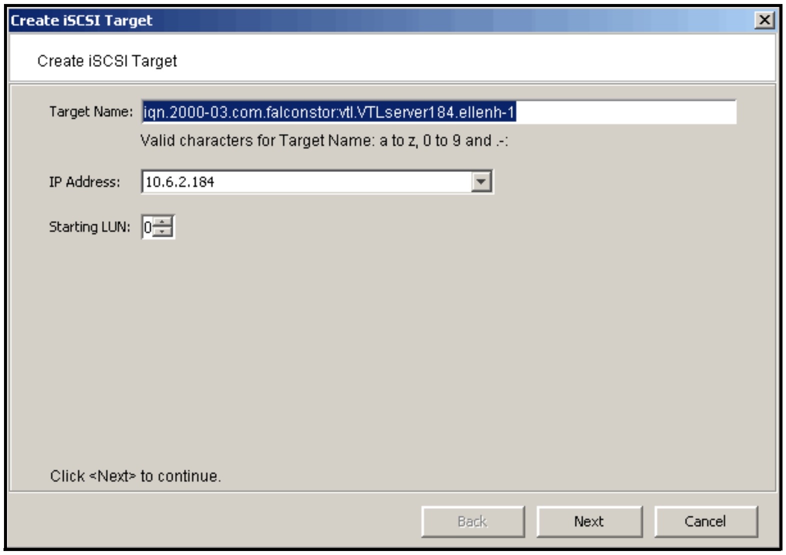

Create iSCSI Target

Create iSCSI Target

Target Name: iqn.2000-03.com.falconstor:vtl.VTLserver184.ellenh-1

Valid characters for Target Name: a to z, 0 to 9 and .-:

IP Address: 10.6.2.184

Starting LUN: 0

Click <Next> to continue.

4. Enter a name for the target or accept the default and select the IP address of the adapter on the VTL server.

   The list includes all Ethernet adapters you have configured on the server.

   > **Note:** Network adapter(s) on the backup application server need to be on the same subnet(s) as the selected adapter(s) on the VTL server.

   If you are using failover, be sure to select the server’s IP address, not the heartbeat IP address, so that clients can see devices while in failover mode.

5. Use the default starting LUN.

   LUN IDs must start with zero.

   Once the iSCSI target is created for a client, LUNs can be assigned under the target using available virtual iSCSI devices.

6. Confirm all information and click *Finish*.

7. Select *Yes* to assign a resource (virtual tape library) to the new target.

Assign a virtual tape library to the iSCSI target

1. Select the virtual library to be assigned to the client.

   You can also select *Allow tape drives in the tape library to be assigned individually* to display the virtual drives in the library.

   You can only assign a device to a client once even if the client has multiple targets.

2. On the next screen, change the LUN for the resource if you need to resolve a conflict.

3. Confirm all information and click *Finish*.

Log the client onto the target

The following steps are for the Microsoft iSCSI Initiator. If you are using a different iSCSI initiator, refer to the documentation provided by the vendor.

1. To see the iSCSI targets from the client machine, run *Microsoft iSCSI Initiator* again.

2. 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.*

3. Click *Advanced* and select *CHAP logon information* in the *Advanced Settings* dialog. Replace the initiator name with any of the usernames you selected as an iSCSI user for this client.

   In *Target Secret*, enter the password associated with that username.

   Click *OK*.

4. Click *OK* to log on to the target.

   The status for the target will change from *Inactive to Connected.*

   The *Targets* tab lists all iSCSI targets, whether or not they are connected. To log off a backup application server from its connection, select the target, click *Details*, select the *Target Identifier*, and then click *Log Off*.

   If you selected the option to *Automatically restore this connection*, the iSCSI target is listed in the *Persistent Targets* tab.

Disable iSCSI

To disable iSCSI for a VTL server, right-click the server node in the VTL console, and select *Options* --> *Disable iSCSI*.

Note that before disabling iSCSI, all iSCSI initiators and targets for this VTL server must be removed.

Linux configuration

Prepare the iSCSI initiator

You must install and configure an iSCSI software initiator on each of your Linux client machines.

1. Download the latest version of the iSCSI initiator package.
    - If you are running Red Hat Enterprise Linux 5.x on a server connected to the internet, you can install the iSCSI package by running the following as root:

      ```
      yum install iscsi-initiator-utils
      ```
    - If you are using a 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
    - If you are using another 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. If your Linux vendor does not provide a binary package of the iSCSI initiator, the source code is freely available from http://sourceforge.net/projects/linux-iscsi/. Note that you will have to manually compile and install the iSCSI initiator if you chose this method.

2. Edit the `/etc/iscsi.conf` file.

   If you are **not using CHAP**, add the following line to the end of the file:

   ```
   DiscoveryAddress=IP address of VTL server
   ```

   For example: `DiscoveryAddress=192.10.10.1`

   If you are **using CHAP**, add the following lines to the end of the file:

   ```
   DiscoveryAddress=IP address of VTL server
   OutgoingUsername=CHAP username
   OutgoingPassword=CHAP password
   ```

   You must make a note of the CHAP username and password because you will have to enter it in the VTL console.

3. Start the initiator by typing:

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

Enable iSCSI

Refer to ‘Enable iSCSI’ for more informations.

Add an iSCSI client

Refer to ‘Add an iSCSI client’.

Create targets for the iSCSI client to log onto

Refer to ‘Create targets for the iSCSI client to log onto’.

Assign a virtual tape library to the iSCSI target

1. Select the virtual library to be assigned to the client.

   You can also select *Allow tape drives in the tape library to be assigned individually* to display the virtual drives in the library.

   You can only assign a device to a client once even if the client has multiple targets.

2. On the next screen, change the LUN for the resource if you need to resolve a conflict.

3. Confirm all information and click *Finish*.

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
```

Afterwards, you can display a list of all the disks that this client can access (including the target) by typing:

```
cat /proc/scsi/scsi
```

IBM 3494 Configuration

The VTL IBM 3494 option is only available for physical VTL systems; it is not available for virtual appliances

Overview

The IBM Enterprise Tape Library 3494 is an automated tape library that can be shared by multiple backup servers.

While the VTL IBM 3494 option does not emulate the IBM 3494 tape library, it does emulate the tape drives inside (such as IBM 3590 drives). This makes it possible to import data from physical tapes and export data on virtual tapes to physical tapes.

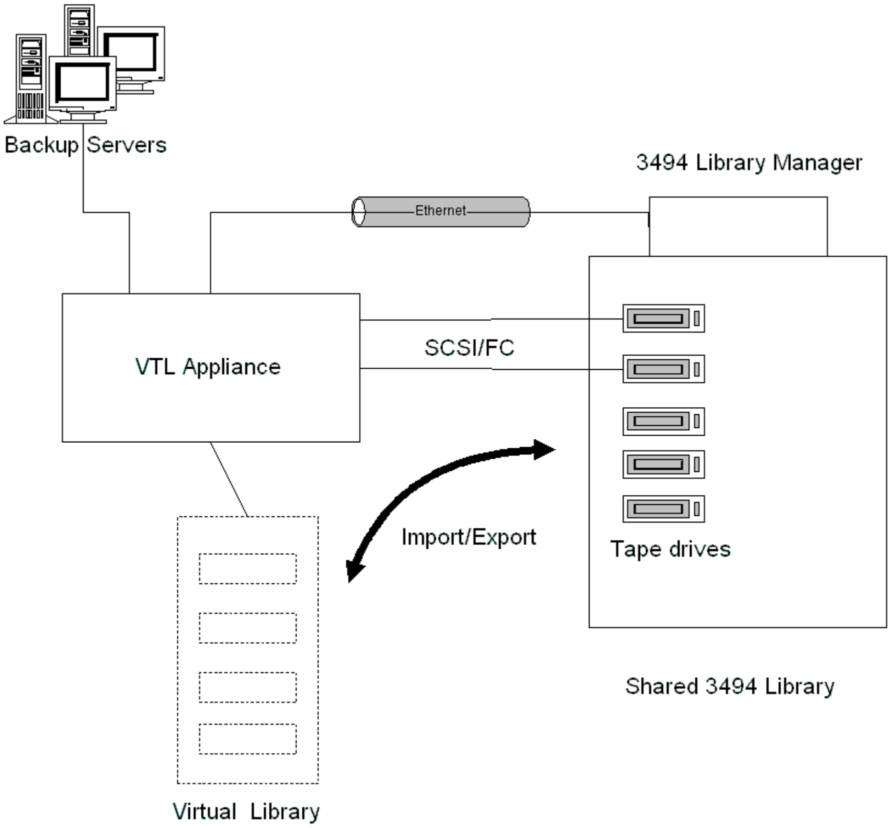

Configuration

Because it has a library manager that manages the tape cartridge inventory and interfaces with attached hosts, VTL requires some special configuration in order to work with the 3494.

1. Install the *IBM Linux Tape Library Driver* (ibmatl.#.#.#.#.i386.rpm.bin) on your VTL server.

2. In the VTL Console, right-click *Physical Tape Libraries* and select *Assign IBM 3494*.

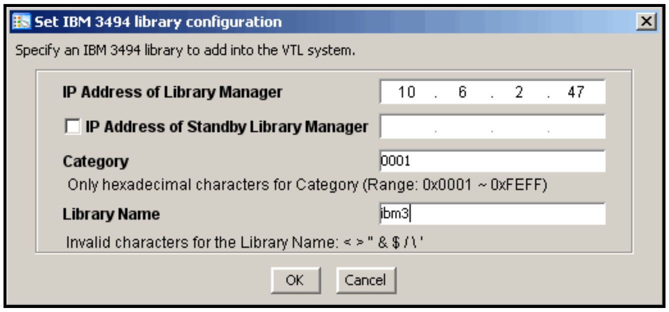

Specify an IBM 3494 library to add into the VTL system.

IP Address of Library Manager 10 . 6 . 2 . 47

□ **IP Address of Standby Library Manager** . . .

Category 0001

Only hexadecimal characters for Category (Range: 0x0001 ~ 0xFEFF)

Library Name ibm3

Invalid characters for the Library Name: < > " & $ / \ '

3. Enter the IP address of the IBM 3494 library manager.

4. Enter the IP address of the standby Library Manager, if available.

5. Use the default category or set to an existing category.

6. Assign the physical tape drives to your IBM 3494 tape library.

7. Inventory the IBM 3494 by right-clicking on the physical library and selecting *Inventory*.

Adding/removing tapes

Whenever you add or remove tapes in the VTL category from your 3494, you must inventory the tapes through the VTL Console (right-click the physical library/drive and select *Inventory)*.

IBM System i Configuration

Overview

IBM System i servers support several types of tape libraries, ranging from relatively simple solutions that can automatically load tapes during operation and maintain a limited cartridge inventory to tape automation systems capable of supporting many systems and managing vast cartridge inventories.

DSI’s Virtual Tape Library for IBM System i allows IBM System i systems to connect to a VTL appliance which emulates IBM 03590E11, IBM TS3500L32(03584L32), and the IBM ULT3583-TL tape libraries with IBM 3580-TD1, IBM 3580-TD2, and IBM 3580-TD3 tape drives.

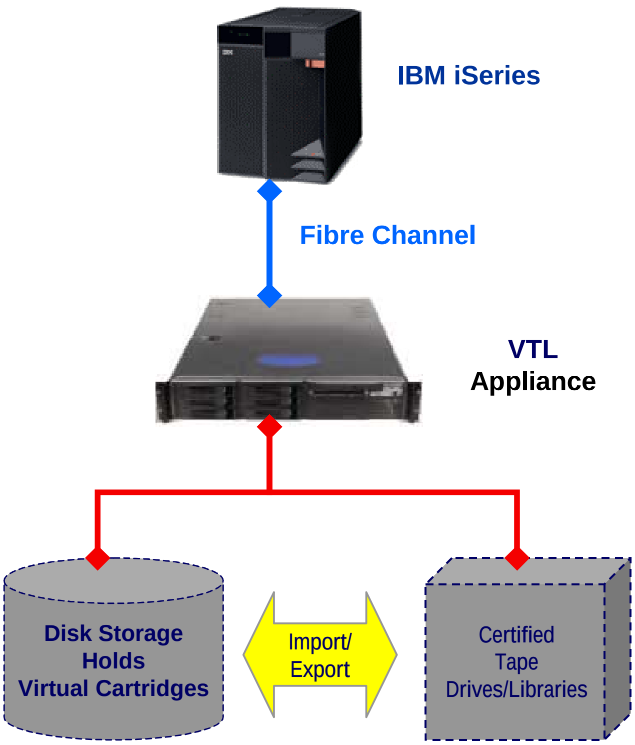

Before you begin

Before you can use Virtual Tape Library for IBM System i, your environment must meet the following criteria:

- You must be using a supported version of IBM Backup Recovery & Media Services. Refer to the DSI Certification Matrix for a list of supported versions.
- There must be a Fibre Channel connection between the System i host and the VTL appliance.
- IBM System i and the VTL appliance must use certified FC HBAs. Refer to the certification matrix on the DSI website for a list of HBAs that are currently certified.

Set up the tape library

With Virtual Tape Library for IBM System i, you use the procedures described earlier in this guide to create a virtual tape library and assign it to an System i host.

> **Note:** When you create a virtual tape library for use with an System i host, you must select IBM 03590E11, IBM TS3500L32(03584L32), or IBM ULT3583-TL virtual tape library emulation.

In addition, for VTL import/export, you must use one of the supported tape drives in the physical tape library. This ensures that you have 1:1 data transfer between the virtual volume and the physical tape media.

Once you have created a virtual tape library and assigned it to the host, you should perform the following tasks to ensure that the System i system can see and properly work with the library. (For more information about working with the System i, refer to your System i documentation.)

1. At the System i system, display the library status functions.

   To do this, access the command line and type the following command:

   ```
   WRKMLBSTS
   ```

2. Make resources available to the tape drive.

   In the option field next to each resource that you want to make available to the tape drive, type 4 (ALLOCATE) and press *Enter*.

3. Inventory the tape library.

   In the option field next to the tape library, type 9 (INVENTORY) and press Enter.

4. Add a tape to the inventory by typing the following command at the command line:

   ```
   ADDTAPCTG DEV(library device name) CTG(cartridge identifier)
   CGY(*NOSHARE) CHKVOL(*NO)
   ```

   Alternatively, you can use *SHARE400 for the CGY parameter.

   After you issue this command, the tape status changes from INSERT to AVAILABLE.

5. Mount a tape onto a drive by typing the following command:

   ```
   CHKTAP DEV(device name) VOL(volume identifier)
   ```

   After you issue this command, the tape status changes from AVAILABLE to MOUNTED.

6. Back up a library object by typing the following command:

   ```
   SAVLIB LIB(library name) DEV(tape media library device name) VOL(volume
   identifier)
   ```

7. Create a library object by typing the following command:

   ```
   CRTLIB LIB(library name)
   ```

8. Restore a library object by typing the following command:

   ```
   RSTLIB SAVLIB(original library name) DEV(tape media library device name)
   VOL(volume identifier) RSTLIB(destination library name)
   ```

9. To confirm that the restore worked, display the library object content by typing the following command:

   ```
   DSPLIB LIB(library name)
   ```

10. Delete a library object by typing the following command:

    ```
    DLTLIB LIB(library name)
    ```

11. Unmount a tape by typing the following command:

    ```
    CHKTAP DEV(device name) VOL(volume identifier) ENDOPT(*UNLOAD)
    ```

    After you issue this command, the tape status changes from MOUNTED to AVAILABLE.

Import cartridges

The process of adding cartridges to the tape library inventory is called *importing*. Most tape libraries provide an I/O station for adding cartridges without interrupting any automated operation.

To import cartridges:

1. From the VTL console, move the tape from vault to library.
2. At the AS/400, re-inventory the library as described in step 3 in 'Set up the tape library'.
3. Add the tape into inventory as described in step 4 in 'Set up the tape library'.

Export cartridges (move to vault)

Cartridges that have been removed from the tape library inventory are referred to as *exported*.

To export a cartridge, type the following command at the command line:

RMVTAPCTG DEV(*library device name*) CTG(*cartridge identifier*)

After you issue this command, if you re-inventory the library from the AS/400, the tape is no longer there. From the VTL console, you can see the tape in the virtual vault.

Hosted Backup

Overview

The Hosted Backup option eliminates the need for a dedicated backup server by making virtual tape libraries and drives available to the local system, allowing certified backup applications to be installed directly onto the VTL appliance.

Note: Your backup application should be installed on the /apps partition on your VTL appliance.

While VTL natively accelerates backup from the backup server to virtual tape, data transfer between application servers and the backup application is accelerated because the backup application is hosted on the VTL appliance itself. This shortens the data path between the application server and the backup application/server and therefore enhances backup performance.

Configure Hosted Backup

To configure Hosted Backup:

1. Right-click the VTL server in the console and select *Options* --> *Enable Hosted Backup*.

   After it has been enabled, a new client called *HostedBackupClient* appears under *Clients*.

   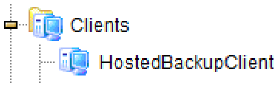

2. Right-click *HostedBackupClient* and select *Assign* to assign virtual libraries to this client.

3. Select the virtual libraries or drives that this client will use.

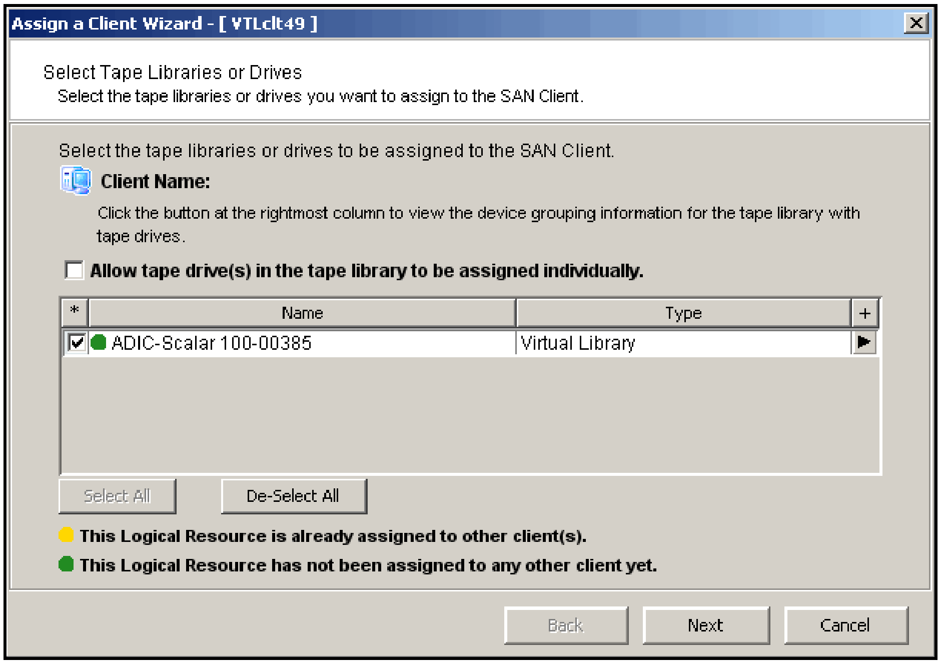

Assign a Client Wizard - [VTLclt49]

Select Tape Libraries or Drives

Select the tape libraries or drives you want to assign to the SAN Client.

Select the tape libraries or drives to be assigned to the SAN Client.

Client Name:

Click the button at the rightmost column to view the device grouping information for the tape library with tape drives.

☐ **Allow tape drive(s) in the tape library to be assigned individually.**

*	Name	Type	+
☑	● ADIC-Scalar 100-00385	Virtual Library	▶

● **This Logical Resource is already assigned to other client(s).**

● **This Logical Resource has not been assigned to any other client yet.**

Libraries assigned to the *HostedBackupClient* can also be assigned to other clients.

4. Confirm all information and click *Finish*.

If installed, the backup application will now see the devices as local devices. You can use Linux's *cat /proc/scsi/scsi* command on your VTL appliance to see the library and all of the drives in the library listed as local devices.

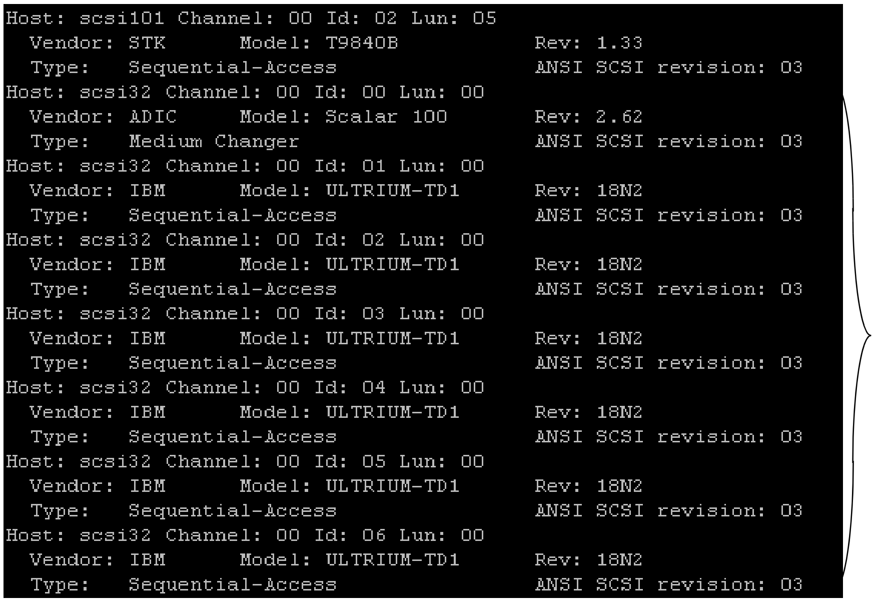

```
Host: scsi101 Channel: 00 Id: 02 Lun: 05
  Vendor: STK      Model: T9840B          Rev: 1.33
  Type:   Sequential-Access               ANSI SCSI revision: 03
Host: scsi32 Channel: 00 Id: 00 Lun: 00
  Vendor: ADIC     Model: Scalar 100      Rev: 2.62
  Type:   Medium Changer                  ANSI SCSI revision: 03
Host: scsi32 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: ULTRIUM-TD1     Rev: 18N2
  Type:   Sequential-Access               ANSI SCSI revision: 03
Host: scsi32 Channel: 00 Id: 02 Lun: 00
  Vendor: IBM      Model: ULTRIUM-TD1     Rev: 18N2
  Type:   Sequential-Access               ANSI SCSI revision: 03
Host: scsi32 Channel: 00 Id: 03 Lun: 00
  Vendor: IBM      Model: ULTRIUM-TD1     Rev: 18N2
  Type:   Sequential-Access               ANSI SCSI revision: 03
Host: scsi32 Channel: 00 Id: 04 Lun: 00
  Vendor: IBM      Model: ULTRIUM-TD1     Rev: 18N2
  Type:   Sequential-Access               ANSI SCSI revision: 03
Host: scsi32 Channel: 00 Id: 05 Lun: 00
  Vendor: IBM      Model: ULTRIUM-TD1     Rev: 18N2
  Type:   Sequential-Access               ANSI SCSI revision: 03
Host: scsi32 Channel: 00 Id: 06 Lun: 00
  Vendor: IBM      Model: ULTRIUM-TD1     Rev: 18N2
  Type:   Sequential-Access               ANSI SCSI revision: 03
```

The virtual library and its drives are listed here as local devices.

If the library is not listed in the console, right-click the *Physical Resources* object and select *Discover New Devices* in the dialog.

5. Verify that the following rpm packages are installed on the VTL server:

   All backup software:

   - audit-libs-1.7.7-6.el5.i386.rpm
   - cracklib-2.8.9-3.3.i386.rpm
   - glibc-2.5-34.i686.rpm
   - glibc-2.5-34.x86_64.rpm
   - glibc-common-2.5-34.x86_64.rpm
   - libacl-2.2.39-3.el5.i386.rpm
   - libattr-2.4.32-1.1.i386.rpm
   - libICE-1.0.1-2.1.i386.rpm
   - libselinux-1.33.4-5.1.el5.i386.rpm
   - libsepol-1.15.2-1.el5.i386.rpm
   - libSM-1.0.1-3.1.i386.rpm
   - libXp-1.0.0-8.1.el5.i386.rpm
   - libXt-1.0.2-3.1.fc6.i386.rpm
   - libXtst-1.0.1-3.1.i386.rpm
   - ncurses-5.5-24.20060715.i386.rpm
   - pam-0.99.6.2-4.el5.i386.rpm
   - pam-0.99.6.2-4.el5.x86_64.rpm
   - xinetd-2.3.14-10.el5.x86_64.rpm

   If they are not installed, Oracle Linux RPMs can be downloaded from http://public-yum.oracle.com/repo/EnterpriseLinux/EL5/.

   Refer to your backup application’s own documentation to determine if there are any additional packages that may be needed.

6. If not yet installed, install your backup application and configure it.

   Your backup application should be installed on the /apps partition on your VTL appliance.

   If the operating system sees the hosted backup devices but the backup application does not, you may need to rescan devices from the backup application or restart the backup application services in order to see the devices.

Note: Any time you add or remove physical or virtual devices on the VTL server and then reboot the server, you will need to check that the paths of the assigned tape and library devices are in the correct sequence in your backup software. If they are not in the correct sequence, you will need to reconfigure the paths.

Stop VTL processes with Hosted Backup

If you are using the Hosted Backup option, you must make sure to stop the backup application before stopping VTL.

NDMP Backup Support

Overview

The *NDMP Backup Support* option allows certified backup applications and industry standard NAS devices (i.e. NetApp filers) to perform backup and restore using the NDMP protocol over an IP network.

With the *NDMP Backup Support* option, the VTL appliance acts as an NDMP server, centralizing management by eliminating locally attached tape devices from each NAS device. When a backup occurs, data is moved from the NAS device directly to the virtual library.

The *NDMP Backup Support* option supports the following:

- NDMP v2, 3, 4
- The following data transfer models:
    - Filer to direct-attach tape drive (local filer)
    - Filer to another filer attached tape drive (filer-to-filer)
    - Filer to *NetVault* Client/Server attached tape drive (filer-to-server)
    - *NetVault* Client/Server to filer attached tape drive (server/client-to-filer)
- Library and tape sharing
- Direct Access Restores (DAR)

Notes:

- The NDMP Backup Support option cannot be used with the FalconStor OpenStorage Option (FSOST). If Symantec NetBackup has configured VTL as an OST server, NetBackup prevents it from being added as an NDMP host.
- The NDMP Backup Support option is not needed when presenting a virtual tape library over FC to a NAS filer as a replacement for a physical library.
- Before you begin configuration, you must define the hostname of the VTL server in the /etc/hosts file in the format "IPAddress Hostname".
  For example: 10.7.2.41   Server41
- Because some backup applications use NDMP, if you are running backup software on the VTL server, it should be started after VTL has started and should be stopped before VTL is stopped. Otherwise, the NDMP service that is loaded by the backup software may interfere with VTL's NDMP service.

Configure NDMP support

To configure NDMP support:

1. Right-click your VTL server and select *Options* --> *NDMP* --> *Enable NDMP*.

2. Enter the user name and password the backup server will use to talk to NDMP.

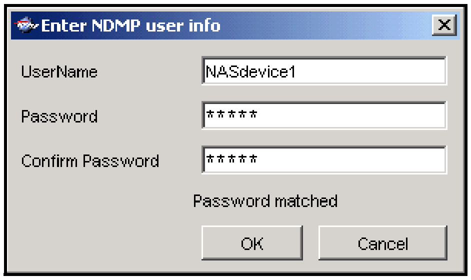

Enter NDMP user info

UserName: NASdevice1

Password: *****

Confirm Password: *****

Password matched

You must enter the same user name/password into the NDMP module in your backup application.

3. Right-click *HostedBackupClient* and select *Assign* to assign virtual libraries to this client.

4. Select the virtual libraries or drives that this client will use.

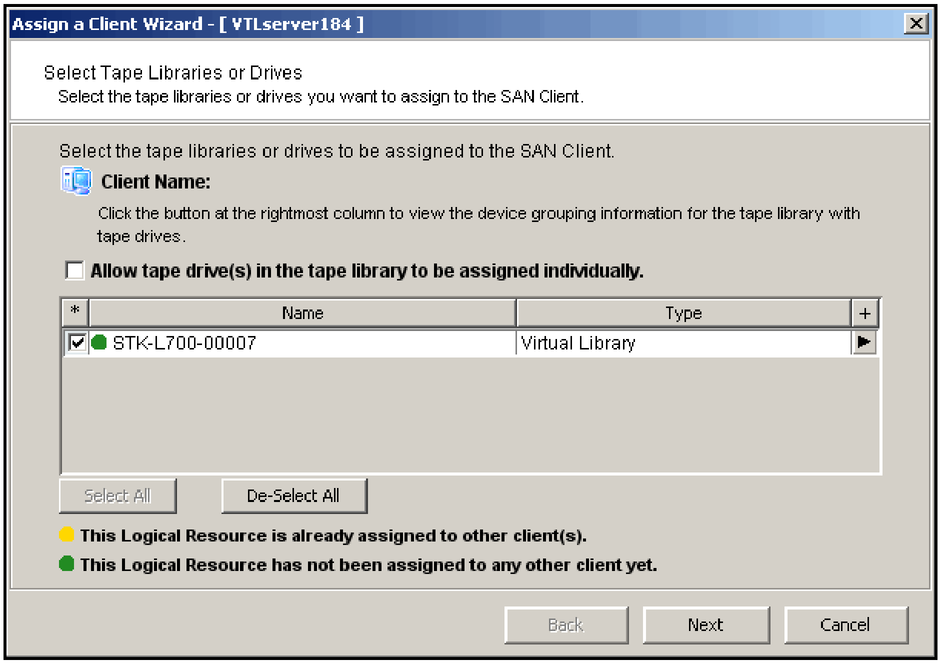

Assign a Client Wizard - [VTLserver184]

Select Tape Libraries or Drives

Select the tape libraries or drives you want to assign to the SAN Client.

Select the tape libraries or drives to be assigned to the SAN Client.

Client Name:

Click the button at the rightmost column to view the device grouping information for the tape library with tape drives.

☐ **Allow tape drive(s) in the tape library to be assigned individually.**

*	Name	Type	+
☑	● STK-L700-00007	Virtual Library	▶

● This Logical Resource is already assigned to other client(s).

● This Logical Resource has not been assigned to any other client yet.

HostedBackupClient can have any number of virtual libraries assigned to it. Conversely, libraries assigned to the *HostedBackupClient* can also be assigned to other clients.

5. Confirm all information and click *Finish*.

The backup application will now see the devices as local devices. You can use Linux’s *cat /proc/scsi/scsi* command on your VTL appliance to see the library and all of the drives in the library listed as local devices.

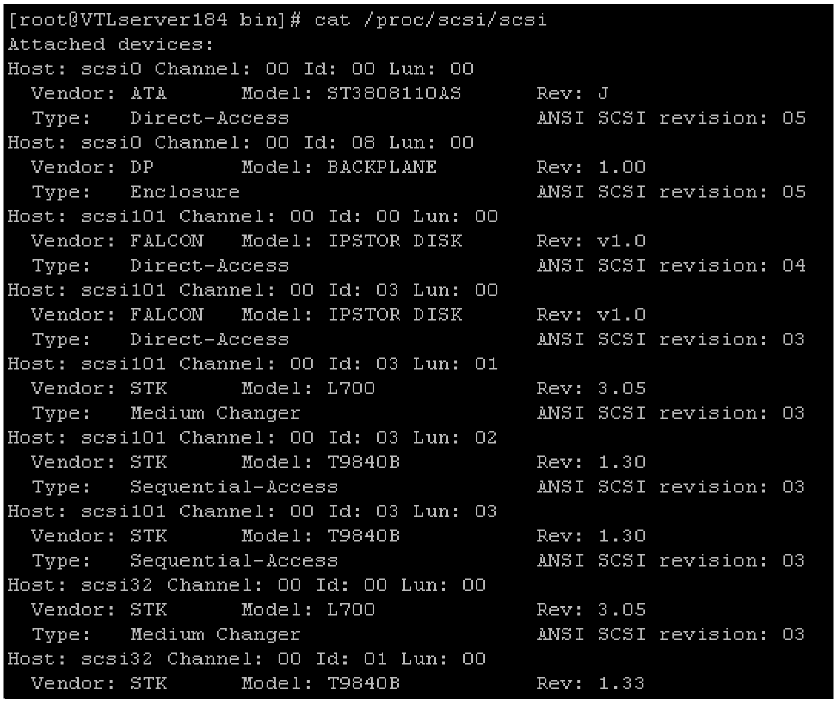

```
[root@VTLserver184 bin]# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: ATA      Model: ST380811AS       Rev: J
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 08 Lun: 00
  Vendor: DP       Model: BACKPLANE        Rev: 1.00
  Type:   Enclosure                        ANSI SCSI revision: 05
Host: scsi101 Channel: 00 Id: 00 Lun: 00
  Vendor: FALCON   Model: IPSTOR DISK      Rev: v1.0
  Type:   Direct-Access                    ANSI SCSI revision: 04
Host: scsi101 Channel: 00 Id: 03 Lun: 00
  Vendor: FALCON   Model: IPSTOR DISK      Rev: v1.0
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi101 Channel: 00 Id: 03 Lun: 01
  Vendor: STK      Model: L700             Rev: 3.05
  Type:   Medium Changer                   ANSI SCSI revision: 03
Host: scsi101 Channel: 00 Id: 03 Lun: 02
  Vendor: STK      Model: T9840B           Rev: 1.30
  Type:   Sequential-Access                ANSI SCSI revision: 03
Host: scsi101 Channel: 00 Id: 03 Lun: 03
  Vendor: STK      Model: T9840B           Rev: 1.30
  Type:   Sequential-Access                ANSI SCSI revision: 03
Host: scsi32 Channel: 00 Id: 00 Lun: 00
  Vendor: STK      Model: L700             Rev: 3.05
  Type:   Medium Changer                   ANSI SCSI revision: 03
Host: scsi32 Channel: 00 Id: 01 Lun: 00
  Vendor: STK      Model: T9840B           Rev: 1.33
```

The virtual library and its drives are listed here as local devices.

Important notes about NDMP devices:

- Any time you add or remove physical or virtual devices on the VTL server and then reboot the server, you will need to check that the paths of the assigned NDMP devices are in the correct sequence in Veritas NetBackup. If they are not in the correct sequence, you will need to reconfigure the paths.
- If you are using a physical library assigned as an NDMP device to your client, make sure that this library is not assigned for any other VTL functions.

ACSLS and Library Station Configuration

The ACSLS/Library Station option is only available for physical VTL systems; it is not available for virtual appliances.

Overview

ACSLS Manager™ and Library Station software manage heterogeneous StorageTek tape libraries.

The ACSLS/Library Station option works with ACSLS/Library Station-managed tape libraries, allowing the system to share ACSLS/Library Station-managed libraries among the VTL server and your backup servers. This makes it possible to import data from physical tapes and export data on virtual tapes to physical tapes.

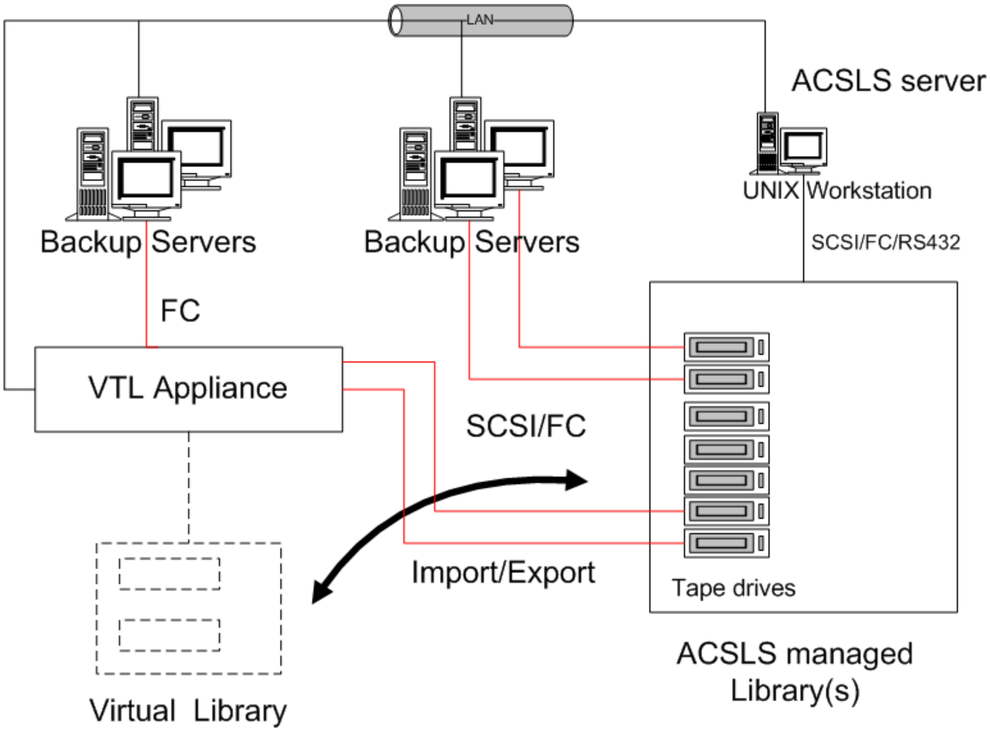

Hardware configuration

1. Physically connect the tape drives that will be assigned to the VTL appliance.

   > **Note:** Physical tape drives cannot be shared with VTL appliances or other applications.

2. (ACSLS only) Create at least one storage pool on the ACSLS server for VTL and assign tapes to it.

   If you have already created a pool, you can use that one.

3. Make a note of the following:
   - ACS IDs, LSM IDs, Panel IDs, and Device IDs of the libraries that hold the tape drives connected to VTL. You can run the `cmd_proc` utility on your ACSLS server to determine the IDs. For Library Station users, check with your Library Station administrator to determine the IDs.
   - IP address of the ACSLS/Library Station server.
   - (ACSLS only) IDs of the storage pools to be assigned to the VTL appliance.

4. Make sure that the VTL server and the ACSLS/Library Station server can communicate with each other.

Configure VTL to work with ACSLS

The following instructions give you an overview of the steps you must follow to configure your ACSLS/Library Station option. Refer to the appropriate sections in this User Guide for more detailed information.

1. Launch the VTL console and connect to the VTL appliance.

2. Right-click your VTL server and select *System Maintenance* --> *Network Configuration* to make sure DNS is configured properly.

   Enter the *Domain Name*, select *Append suffix to DNS lookup* and enter the DNS server IP address.

3. Right-click the *Physical Tape Libraries* object and select *Add ACSLS/Library Station.*

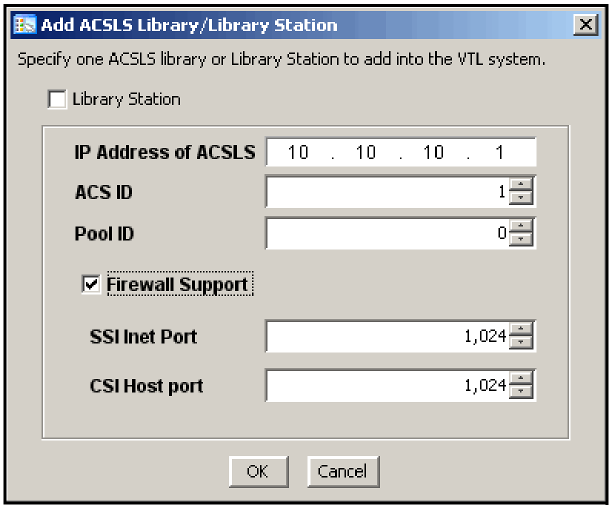

Add ACSLS Library/Library Station

Specify one ACSLS library or Library Station to add into the VTL system.

□ Library Station

IP Address of ACSLS: 10 . 10 . 10 . 1

ACS ID: 1

Pool ID: 0

☑ Firewall Support

SSI Inet Port: 1,024

CSI Host port: 1,024

4. Enter the IP address of the ACSLS/Library Station server; the ACS ID and the Pool ID of the ACSLS/Library Station library.

   Once completed, the server automatically does an inventory to obtain a list of physical tapes.

5. If applicable, enter your firewall information.

   If you select the *Firewall Support* option, you need to edit the `$ISHOME/etc/acsls_ls_cdk.conf` file and add the following lines:

   ```
   CSI_HOSTPORT=[ACSLS assigned port number]
   ```

   ```
   SSI_INET_PORT=[ACSLS assigned port number]
   ```

   If you do not select the Firewall Support option, the portmap service needs to be running. Otherwise, the system will fail to assign or retrieve the library's status after restarting VTL services or rebooting.

   To enable the portmap service on Linux, run the following command:

   ```
   chkconfig --add portmap
   ```

   By design, when a firewall is set for an ACSLS server, only one ACS physical library can be assigned using the same SSI_INET_PORT defined in the `acsls_ls_cdk.conf` file. However, if the CSI and SSI INET ports are not set in the configuration file, you can use portmapper to assign several physical libraries. For example, use the following configuration to assign two ACS physical libraries from an ACSLS behind a firewall. Open all of the ports in your firewall configuration but do not define them in the `acsls_ls_cdk.conf` file.

   For the first ACS use:

   - SSI inet port: 1024
   - CSI host port: 30031

   For the second ACS use:

   - SSI inet port: 1025

- CSI host port: 30031

6. Assign your physical tape drives to your ACSLS/Library Station tape library.

   You will have to enter the *Drive ID* for each. The *Drive ID* is comprised of the drive’s ACS ID, LSM ID, Panel ID, and Device ID in the format `n,n,n,n`

   You can run the `cmd_proc` utility on your ACSL server to determine the IDs. For Library Station users, check with your Library Station administrator to determine the IDs. You may want to supply the administrator with the drive’s SCSI address to help him determine the IDs.

Set eject policy

You need to edit the $ISHOME`/etc/acsls_ls_cdk.conf` file to set how many tapes will be ejected after each backup.

Enter your policy information in the following format:

```
RepositoryID:ACSLS/LS IP:acsid:poolid:eject
```

In the following example:

```
4:10.6.2.62:1:2:5
```

- 4 is the Repository/Database ID. (This can be found in the VTL console by clicking on the *Database* object.)
- 10.6.2.62 is the IP address of the ACSLS/Library Station.
- 1 is the ACS ID.
- 2 is the Pool ID.
- 5 is the eject value. In this example, after each backup, the system will eject five tapes at a time. This is usually set to the same value as the *CAP size*.

Filter tapes displayed in the VTL console

If you pool physical tapes in Library Station and you *only* want to display a range of Library Station barcodes in the VTL console, you will need to edit the `$ISHOME/etc/acsls_ls_cdk.conf` file.

Locate the following sections:

```
#Library Station barcode filters...
#end Library Station barcode filters
```

Enter your range information **between** these sections (before the `#end` line) in the following format:

```
RepositoryID:ACSLS/LS IP:acsid:startrange-endrange:startrange-
endrange
```

The barcode length should be six characters and spaces should not be used. Any number of barcode ranges can be specified. In the following example:

```
4:10.6.2.62:6:000000-004444:AA0000-AA8888
```

- 4 is the Repository ID. (This can be found in the VTL console by clicking on the *Database* object.)
- 10.6.2.62 is the IP address of the ACSLS/Library Station.
- 6 is the ACS ID.
- 000000-004444 is the first range to display.
- AA0000-AA8888 is the next range to display.

Configure ACSLS to work with VTL Failover

When VTL Failover is enabled and ACSLS is configured, you must make the following modification to the `/etc/hosts` file so that ACSLS will continue to function after a takeover (when the secondary server takes over for the primary).

1. Change the virtual IP to the heartbeat IP of both servers in the set.

2. Run the following command on both servers:

   ```
   vtl restart tleupcall
   ```

Add/remove tapes

Whenever you add or remove tapes from an ACSLS/Library Station pool, you must inventory the tapes through the VTL console (right-click the physical library and select *Inventory)*.

Replace ACSLS physical drive

This section explains how to replace ACSLS physical drives that are assigned to VTL for tape import/export purposes.

Notes:

- Take an X-ray of the VTL server before making any changes.
- Import/export jobs should not access any physical drive that is being replaced. If possible, stop all import/export jobs until the procedure is complete.

1. From the DSI management console, expand *Virtual Tape Library System* --> *Physical Tape Libraries*, and highlight the physical drive.

   Make a note of the following information:

   - SCSI ID
   - ACSLS device ID
   - Serial number of the physical drive

2. Unassign the physical drive by right-clicking the drive to be removed and select *Unassign*.

3. Remove and replace the physical drive in the physical library.

4. Use the ACSLS command interface to configure the new drive.

   This can be done using: `config drive ’panel_id’`

   For example, if the device ID is 0,2,1,12, you would use the following command: `config drive 0,2,1,12`

   Refer to the ACSLS library administrator’s guide for more information.

5. From the DSI management console, rescan for the new device.

   To do this, right-click the HBA with the new physical drive and select *Rescan.* Use the *Discover New Devices* option.

6. From the DSI management console, assign the new drive to VTL.

   To do this, right-click the physical library and select *Assign*. Add the device ID from step 1.

   Afterward, highlight the physical tape drive to confirm that the serial number of this new drive matches.

Server Maintenance

VTL servers are designed to require little or no maintenance.

All day-to-day administrative functions can be performed through the console. However, there may be situations when direct access to the server is required, particularly during initial setup and configuration of physical storage devices attached to the server or for troubleshooting purposes.

If access to the server’s operating system is required, it can be done either directly or remotely from computers on the network. You can log in from a terminal connected directly to a VTL server. There is no graphical user interface (GUI) shell required. By default, only the root user has login privileges to the operating system. Other VTL administrators do not. To log in, enter the password for the root user.

Start/stop server modules

The following commands are available:

- `vtl start` - Starts the VTL server modules.
- `vtl restart -` Stops and then starts the VTL server modules.
- `vtl status` - Checks the status of the VTL server modules. You will see a list of modules that are currently running.
- `vtl stop` - Stops the VTL server modules.

Important notes about stopping a VTL server

Warning: Stopping the VTL modules will detach all virtual devices. To prevent data loss, we recommend stopping all VTL client services prior to shutdown.

Note: If you are using the Hosted Backup option, you must make sure to stop the backup application before stopping VTL.

Server modules

When you run a server command, you will see a list of modules. The list will vary, depending upon whether it is a VTL or SIR server and which options you are using. A description of each module is listed below.

Module	Server	Description
Authentication Module	All	Manages authentication requests between replica servers.
Base Module	All	Provides basic memory management and SCSI device management to IO modules.
Cipher Module	All	Performs data encryption from virtual tape to physical tape and data decryption from physical tape to virtual tape.
CLI Proxy Module	All	Facilitates communication between the CLI utility and the VTL server.
Communication Module	All	Handles console-to-server communication and manages overall system configuration information.
Compression Module	All	Provides software compression for replication and deduplication.
Deduplication Module	SIR, VTL-S	Provides the repository server.
Email Alerts Module	VTL	Sends alerts via email to indicate alarming situations.
Event Module	All	Provides message logging interface to the system log.
Failover Module	VTL	Checks the partner server’s health in a failover setup.
FC Initiator Module	All	Represents the QLogic Fibre Channel initiator module, which provides interaction between the VTL server and the FC storage.
FC Target Module	All	Provides Fibre Channel target functionality.
FUSE Module	VTL	Represents ‘File system in Userspace (FUSE)’ module to allow creation of file systems in the user space without editing the Linux kernel code.
Hifn HW Compression Module	All	Provides hardware compression via Hifn adapters for tape data.
Hosted Backup Module	VTL	Exposes virtual libraries and drives to the local Linux system for local backup.
Ingest Throttling Module	VTL	Provides ingest throttling for NAS.
IO Core Module	All	Provides core IO services.
iSCSI Target Module	VTL	Provides iSCSI target functionality via a network adapter.
iSCSI Daemon	VTL	Represents a user daemon, which handles the login process to VTL iSCSI targets from iSCSI initiators.
Logger Module	All	Provides the information logging function for VTL reports.

Module	Server	Description
Memory Map Module	All	Allows mapping of physical memory from kernel space to user space.
NAS MGTD Module	VTL	(NAS management module) Synchronizes local user/group access rights with Active Directory service.
NAS NMBD Module	VTL	Replies to netbios name service requests.
NAS SMBD Module	VTL	Samba (CIFS) server implementation.
NAS UNFSD Module	VTL	NFS server daemon process.
Node Manager	SIR, VTL-S	Facilitates communication among VTL and SIR servers.
Outgoing Replication Throttling Module	VTL	Provides outgoing replication throttling for NAS.
Path Manager Module	All	Manages I/O traffic over multiple paths to multi-path storage devices.
Reclamation Triggering Module	SIR, VTL-S	Monitors deduplication usage to trigger reclamation at pre-set thresholds or when scheduled.
Replication Server Module	VTL	Provides file replication for NAS.
ROFS Rescan Module	VTL	Monitors the file system status on NAS resources; a file system might become read-only due to a physical device failure; when the failed device becomes online, this module will mount the file system back in read-write mode.
Self-Monitoring Module	All	Checks the server's own health.
SNMPD Module	All	Interacts with SNMP management software to reply to MIB browsing queries and to send SNMP traps for abnormal situations.
Statistics Database Daemon	SIR, VTL-S	Provides deduplication statistics.
TCP Stream Replication Source Module	SIR, VTL-S	Provides outgoing deduplication unique data replication support.
TCP Stream Replication Target Module	SIR, VTL-S	Provides incoming deduplication unique data replication support.
TLE Core Module	VTL	Provides the tape drive/library emulation and interaction to physical tape libraries.
TLE Upcall Kernel Module	VTL	Represents the kernel module, handling interactions between kernel mode and user mode components, which manage tape operations.
TLE Upcall User Module	VTL	Represents the user module, handling interactions between kernel mode and user mode components, which manage tape operations.
Transport Module	All	Handles I/O transport for replication.

Module	Server	Description
Upcall Module	All	Handles interactions between kernel mode and user mode components.
Virtual Device FSVSHOST Module	SIR, VTL-S	Provides device access to the repository.
Virtual SCSI Device Driver	SIR, VTL-S	Provides a block device interface to a deduplication virtual device.

Reports

VTL provides a wide variety of pre-defined reports that help you manage your VTL/deduplication clusters and NAS resources. Some reports focus on server conditions and individual hardware component configuration and behavior, such as disk space usage, physical resource allocation, and Fibre Channel configuration. Others are related to VTL features and gather comprehensive status information about virtual tapes and libraries, deduplication, and replication.

You can generate all reports from the *Reports* object in the DSI Management Console navigation tree. You can also create a schedule to repeatedly run a report automatically.

If you have configured a multi-node group, the *Group Reports* object is available under the group object. Reports generated from this object reflect only the servers in the group.

The report wizard displayed from either object lets you choose a report to create and then specify a variety of parameters and options to filter data.

You can also set general report properties such as whether all or selected reports should be emailed and how long to retain them on the server.

Reports about NAS resources and activity are available only if NAS is enabled on the VTL server.

The amount of data you can include in a report depends on settings in server properties, in the *Activity Database Maintenance* tab. Before you create reports, make sure that the size of the activity data file and the number of days of activities kept are appropriate for the amount of data you intend to include (refer to 'Server properties').

Report types

Reports are available in these categories:

Information

- Deduplication - Policy Status
- Deduplication - Tape Activity
- Deduplication Replication Status
- Deduplication Repository - Reclamation
- Import/Export Jobs
- Replication Status
- Virtual Library and Drive Assignment
- Virtual Library Information
- Virtual Tape Activity
- Virtual Tape Information
- NAS Statistics Summary

Usage	• Deduplication - Tape Usage
	• Deduplication Repository - Memory and Space Usage
	• Physical Tape Usage
	• Disk Space Usage History
	• LUNs
	• NAS Resource Usage
	• NAS CIFS Share Usage
Allocation	• Disk Space Allocation for Virtual Tapes in Libraries
	• Physical Resource Allocation
Configuration	• Fibre Channel Adapters Configuration
	• Physical Resources Configuration
Performance	• Deduplication Repository - Performance
	• VTL Performance

Create a report

Create a one-time report

Each report can be created for a specific server and will run only once.

> **Note:** If you plan to email reports, email must be configured in *Report Properties* for the server (refer to 'Set report properties'*).*

1. To create a report, right-click the *Reports* object and select *New*. The Reports Wizard is displayed.

2. Select a report type.

3. If applicable, choose the period of time to include in the report and provide other selection criteria, based on server dates:
   - *Today* - activity for the current server date
   - *Yesterday* - activity for the day prior to the current server date
   - *Past 7 days* - activity for the 7 days prior to the current server date
   - *Past 30 days* - activity for the 30 days prior to the current server date
   - *Past 365 days* - activity for the 365 days prior to the current server date
   - *Date range* - specify a beginning and end date within the 365 days prior to the current server date.

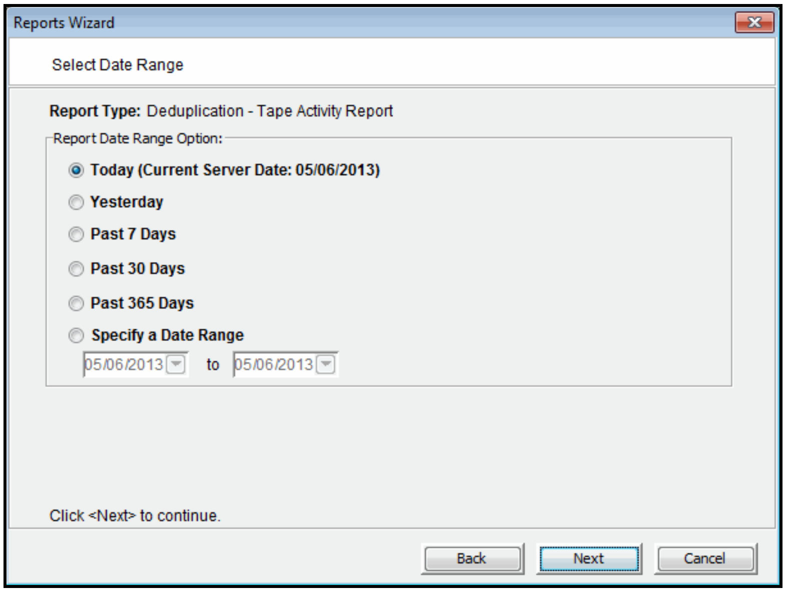

Reports Wizard

Select Date Range

Report Type: Deduplication - Tape Activity Report

Report Date Range Option:

- ● Today (Current Server Date: 05/06/2013)
- ○ Yesterday
- ○ Past 7 Days
- ○ Past 30 Days
- ○ Past 365 Days
- ○ Specify a Date Range

05/06/2013 to 05/06/2013

Click <Next> to continue.

4. If applicable, choose the interval between data points in the report or, depending on the type of report, the interval represented by a bar in a bar chart. This can be an hour, a day, a week, a month, or a quarter. Depending on the report, the data point/bar can represent one of the following:
   - an average of the data measured during the interval
   - the total data measured during the interval
   - the total measured as of a specific point in time

   Note: When selecting an interval value for a report, consider the length of time the server or cluster has been in operation. Do not select an interval that is larger than that period of time. For instance, if your system has not been in operation for at least three months prior to the date on which you are creating the report, do not choose the *quarter* interval.

5. If applicable, indicate what items to include in the report, such as specific tape libraries/drives, physical resources, barcodes, deduplication policies, tape locations, adapters, devices, and/or SCSI devices.

6. Enter a name for the report.

7. If you have configured email for reports, indicate if you want to email this report.

   You will have to enter the recipient(s) and a subject. You can also include text for the body of the email and specify a format (.txt, .csv, .pdf, .xls, .html) for the report attachment.

8. Confirm all information and click *Finish* to create the report.

Schedule a report

You can schedule most reports to run automatically at regular intervals.

1. Right-click the *Scheduled Reports* object under *Reports* and select *New*.
2. Select a report.
3. Set the schedule for how often this report should run.

   You can run the report on an hourly, daily, or weekly basis and you must indicate a starting time. If you select *weekly*, you must also select which day to run the report. If you select *hourly*, you must select the frequency (in hours).
4. Depending upon which report you select, additional windows appear to allow you to filter the information it will include.

   For instance, set the date or date range for the report and select display options, as described for one-time reports.
5. Enter a name for the report.
6. If email is configured for reports, provide email options.
7. Confirm all information and click *Finish* to save the report schedule.

View a report schedule

Select the *Scheduled Jobs* object to display a list of defined jobs in the upper section of the right-hand pane; select a job to display its schedule in the lower section of the display.

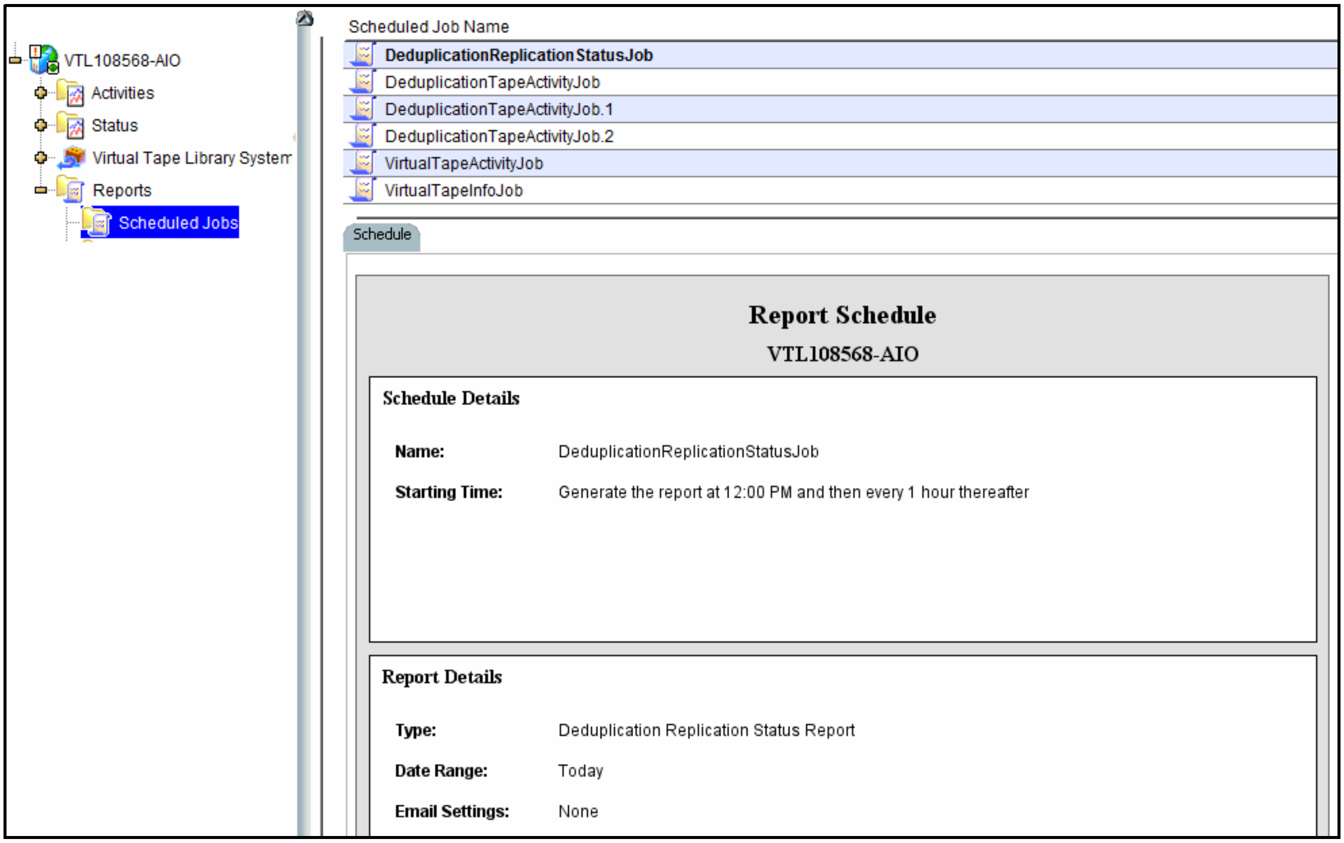

Scheduled Job Name
DeduplicationReplicationStatusJob
DeduplicationTapeActivityJob
DeduplicationTapeActivityJob.1
DeduplicationTapeActivityJob.2
VirtualTapeActivityJob
VirtualTapeInfoJob

Report Schedule

VTL108568-AIO

Schedule Details

Name:	DeduplicationReplicationStatusJob
Starting Time:	Generate the report at 12:00 PM and then every 1 hour thereafter

Report Details

Type:	Deduplication Replication Status Report
Date Range:	Today
Email Settings:	None

Manage job

You cannot modify a job schedule; if you no longer want a job, you must delete it. You can then create a new job.

To delete a job, right-click it in the list in the right-pane and select *Delete*.

Create a group report

After you configure a multi-node group, the *Group Reports* object is available below the group object.

1. Right-click the *Group Reports* object under the group object and select *New*.

2. Choose one of the two group report options:

   Regular Report - All standard reports are available. The report will be generated for specified servers in the group. In the console, the report will be listed below the *Reports* object for each individual server.

   Consolidated Report - Generates the *Group Disk Space Allocation for Virtual Tapes in Libraries* report (refer to 'Disk Space Allocation for Virtual Tapes in Libraries' for details), which will collect information from all servers in the group and present it in a single report that will be listed below the *Group Reports* object.

3. For a *Regular Report*, choose the report you want to create and the period of time and options for data you want the report to include.

   - Select the servers (in the group) that you want to include in the report.

   The *Consolidated Report* has two variations: *Current disk space allocation* or *Historical library space allocation*.

   - If you choose *Current disk space allocation*, there are no additional options.
   - If you choose *Historical library space allocation*, choose the time frame - in days or within a range of dates - the report should represent, as well as the interval (the amount of time) between datapoints.

4. Enter a name for the report.

5. If email is configured for reports, provide email options.

6. Confirm all options and click *Finish* to generate the report.

View a report

The first time you create any type of report, a *report category* is created in the navigation tree, below the *Reports* (or *Group Reports*) object. A report category associates all reports of the same type with a single heading. Additional reports of the same type are associated with the appropriate report category. Expand the *Reports* (or *Group Reports*) object to see a list of categories for generated reports.

When you select a report category, a list of available reports is displayed in the upper section of the right-hand pane. Select a report to display it below the list.

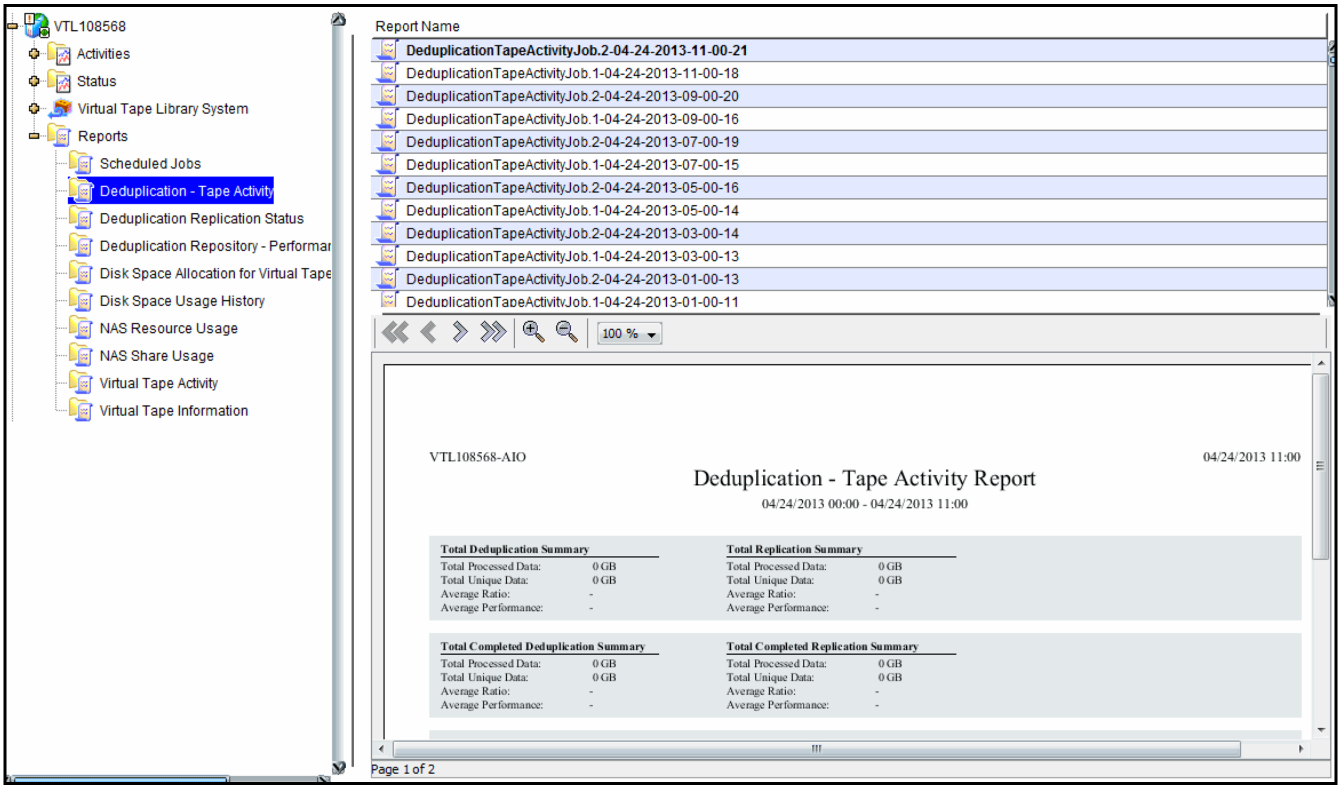
VTL108568
Activities
Status
Virtual Tape Library System
Reports
Scheduled Jobs
Deduplication - Tape Activity
Deduplication Replication Status
Deduplication Repository - Performar
Disk Space Allocation for Virtual Tape
Disk Space Usage History
NAS Resource Usage
NAS Share Usage
Virtual Tape Activity
Virtual Tape Information
100 %

Report Name
DeduplicationTapeActivityJob.2-04-24-2013-11-00-21
DeduplicationTapeActivityJob.1-04-24-2013-11-00-18
DeduplicationTapeActivityJob.2-04-24-2013-09-00-20
DeduplicationTapeActivityJob.1-04-24-2013-09-00-16
DeduplicationTapeActivityJob.2-04-24-2013-07-00-19
DeduplicationTapeActivityJob.1-04-24-2013-07-00-15
DeduplicationTapeActivityJob.2-04-24-2013-05-00-16
DeduplicationTapeActivityJob.1-04-24-2013-05-00-14
DeduplicationTapeActivityJob.2-04-24-2013-03-00-14
DeduplicationTapeActivityJob.1-04-24-2013-03-00-13
DeduplicationTapeActivityJob.2-04-24-2013-01-00-13
DeduplicationTapeActivityJob.1-04-24-2013-01-00-11

VTL108568-AIO 04/24/2013 11:00

Deduplication - Tape Activity Report

04/24/2013 00:00 - 04/24/2013 11:00

Total Deduplication Summary		**Total Replication Summary**	
Total Processed Data:	0 GB	Total Processed Data:	0 GB
Total Unique Data:	0 GB	Total Unique Data:	0 GB
Average Ratio:	-	Average Ratio:	-
Average Performance:	-	Average Performance:	-

Total Completed Deduplication Summary		**Total Completed Replication Summary**	
Total Processed Data:	0 GB	Total Processed Data:	0 GB
Total Unique Data:	0 GB	Total Unique Data:	0 GB
Average Ratio:	-	Average Ratio:	-
Average Performance:	-	Average Performance:	-

Page 1 of 2

The report includes a toolbar that lets you navigate through pages in the report sequentially, or skip to the last or first page in the report. Two viewing tools are available - zoom-in/zoom-out, and display magnification.

In every report, the server name appears in the upper left corner and the date on which the report was created appears in the upper right corner. The period of time represented in the report is displayed below the report title. Any display options appear below the report table.

Most reports organize data in a tabular format; some reports include a graphical view of data, as a pie-chart or bar graph.

Manage reports

Set report properties

You can set email and retention properties for all or individual reports. To do this:

1. Right-click the *Reports* object and select *Properties.*

   If this is a multi-node group, right-click *Group Reports* and select *Properties*.

2. If you will be emailing reports, enter information about your SMTP configuration.

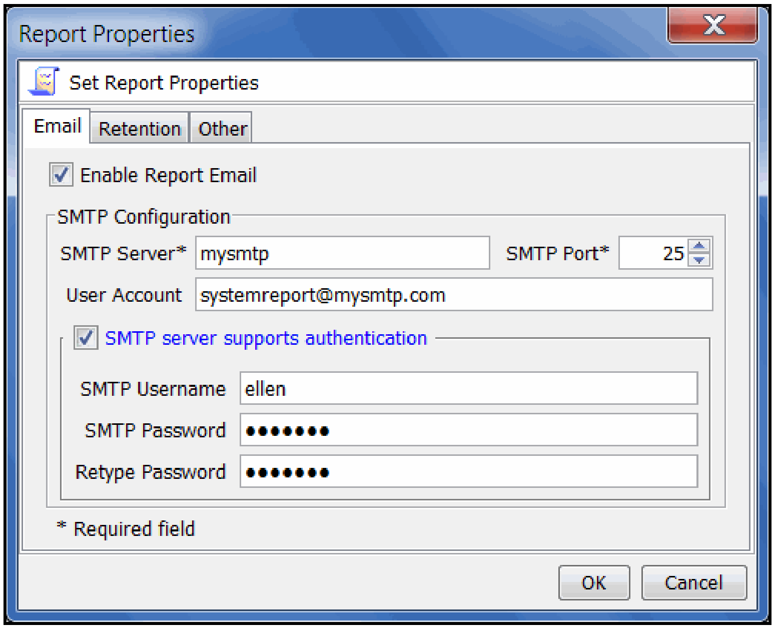

☑ Enable Report Email

SMTP Configuration

SMTP Server* mysmtp

SMTP Port* 25

User Account systemreport@mysmtp.com

☑ SMTP server supports authentication

SMTP Username ellen

SMTP Password •••••••

Retype Password •••••••

* Required field

SMTP Server - Specify the mail server that should be used. You can enter an IP address or hostname consisting of alphabet letters, numbers, "_", "-", or ".". The maximum length is 255 characters.

SMTP Port - Specify the mail server port that should be used.

User Account - Specify the email account that will be used in the “From” field of emails.

SMTP server supports authentication - Indicate if the SMTP server supports authentication.

SMTP Username/Password - Specify the user account that will be used to log into the mail server.

3. On the *Retention* tab, specify how long generated reports should be retained.

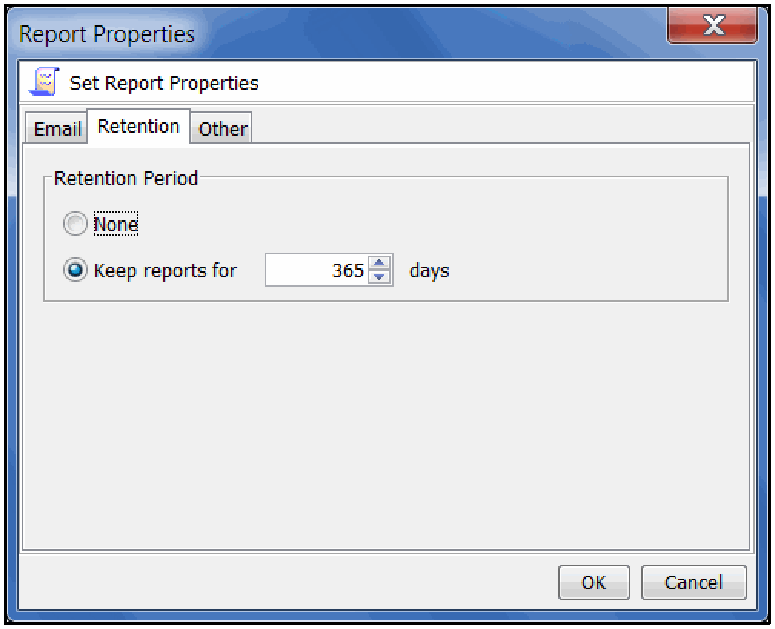

Retention Period

○ None

● Keep reports for 365 days

4. On the *Other* tab, define how a day is defined in your environment: from midnight to midnight, or from noon to noon.

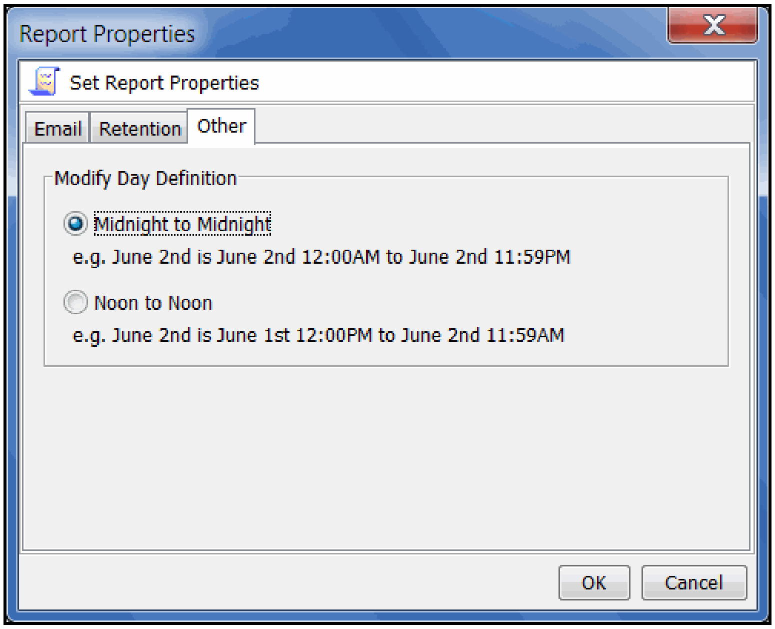

Modify Day Definition

● Midnight to Midnight

e.g. June 2nd is June 2nd 12:00AM to June 2nd 11:59PM

○ Noon to Noon

e.g. June 2nd is June 1st 12:00PM to June 2nd 11:59AM

This affects all reports, those generated from the console and at the command line.

Export data from a report

You can export a report from the server to another location. To do this, right-click a generated report object and select *Export,* then choose from one of the available formats: comma delimited (.csv), tab delimited (.txt) text, Excel spreadsheet (.xls), PDF (.pdf), web page (.html, a .zip file is created).

Email a report

In order to be able to email a report, email properties for reports must be configured before you create the report (refer to 'Set report properties'). If this has been done, you can set the report to be emailed in the Reports Wizard. To email a previously created report:

1. Right-click a report that is generated and select *Email.*
2. Specify a recipient and a subject and then click *Send*.

Refresh report display

You can refresh the list of displayed reports. This will update the list to include reports that have been generated automatically during the time you have been using the console. To do this, right-click *Reports* (or *Group Reports*) and select *Refresh.*

Print a report

You can print individual reports. To do this, right-click an existing report and select *Print.*

Delete a report

You can delete one or more reports. To delete a single report, right-click the report object in the right-hand pane and select *Delete.*

To delete multiple reports, right-click a report category or the *Reports* (or *Group Reports*) object and select *Delete*, then choose individual reports to delete.

Reports and VTL Failover

Reports cannot be generated from a server when it is in a failed state (the server name is displayed in red in the management console). After failback, reports generated from the original node can include data generated during the failover period, as long as the specified report dates include that period of time.

Information reports

Deduplication - Policy Status

This history report lists all deduplication policies and provides summary information for all replication and deduplication jobs run during the specified period of time. Summary information for each deduplication policy includes the number of tapes in the job, the name and IP address of the deduplication cluster, and the deduplication schedule. If replication is enabled for the policy, the summary shows schedule information and the name and IP address of the all replication target servers.

For each deduplication job, the report displays the run date and time, the number of tapes, total and unique data size, deduplication ratio, job duration, performance rate, and status (complete, incomplete, error). The *Type* column indicates the deduplication trigger (such as inline, manual, or end of backup).

If replication is enabled, all categories except for status are also displayed. If Advanced Replication is enabled, the names for target servers 1 and 2 are displayed.

This report is available only as a one-time report.

To see details for specific tapes in deduplication jobs, create the *Deduplication - Tape Activity Report*.

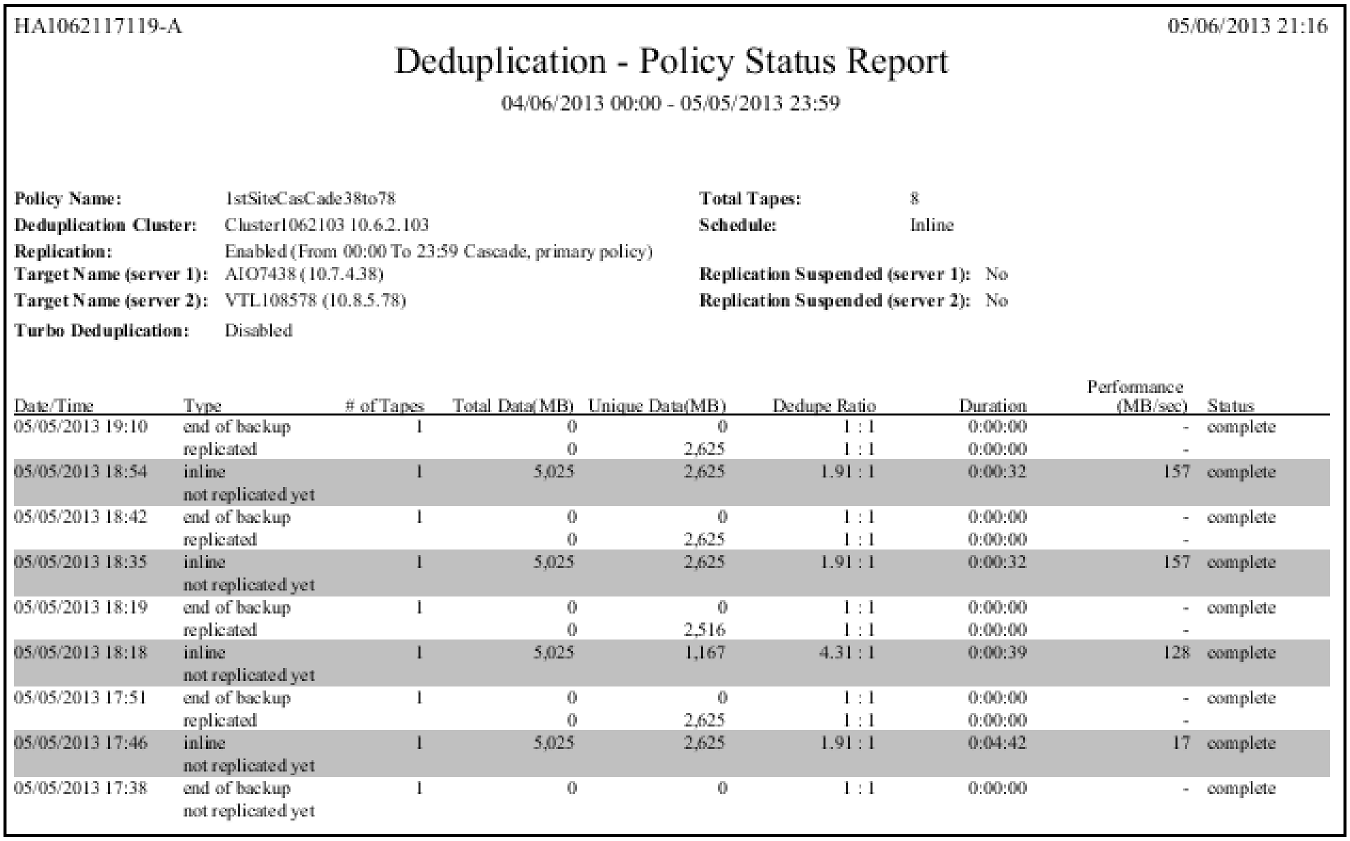

HA1062117119-A 05/06/2013 21:16

Deduplication - Policy Status Report

04/06/2013 00:00 - 05/05/2013 23:59

Policy Name:	1stSiteCasCade38to78	**Total Tapes:**	8
Deduplication Cluster:	Cluster1062103 10.6.2.103	**Schedule:**	Inline
Replication:	Enabled (From 00:00 To 23:59 Cascade, primary policy)		
Target Name (server 1):	AIO7438 (10.7.4.38)	**Replication Suspended (server 1):**	No
Target Name (server 2):	VTL108578 (10.8.5.78)	**Replication Suspended (server 2):**	No
Turbo Deduplication:	Disabled		

Date/Time	Type	# of Tapes	Total Data(MB)	Unique Data(MB)	Dedupe Ratio	Duration	Performance (MB/sec)	Status
05/05/2013 19:10	end of backup	1	0	0	1 : 1	0:00:00	-	complete
	replicated		0	2,625	1 : 1	0:00:00	-	
05/05/2013 18:54	inline	1	5,025	2,625	1.91 : 1	0:00:32	157	complete
	not replicated yet							
05/05/2013 18:42	end of backup	1	0	0	1 : 1	0:00:00	-	complete
	replicated		0	2,625	1 : 1	0:00:00	-	
05/05/2013 18:35	inline	1	5,025	2,625	1.91 : 1	0:00:32	157	complete
	not replicated yet							
05/05/2013 18:19	end of backup	1	0	0	1 : 1	0:00:00	-	complete
	replicated		0	2,516	1 : 1	0:00:00	-	
05/05/2013 18:18	inline	1	5,025	1,167	4.31 : 1	0:00:39	128	complete
	not replicated yet							
05/05/2013 17:51	end of backup	1	0	0	1 : 1	0:00:00	-	complete
	replicated		0	2,625	1 : 1	0:00:00	-	
05/05/2013 17:46	inline	1	5,025	2,625	1.91 : 1	0:04:42	17	complete
	not replicated yet							
05/05/2013 17:38	end of backup	1	0	0	1 : 1	0:00:00	-	complete
	not replicated yet							

Deduplication - Tape Activity

This history report provides detailed information about tape activity for specific deduplication jobs run during the specified period of time. By default, results include summary information for all deduplication and replication jobs run within the past 24 hours. *Active* jobs are not included by default.

When any type of *Advanced Replication* is enabled, results include all jobs on all target servers.

In addition to selecting a date range for the report, you can also choose which deduplication policies to include. You can then customize report results by choosing from an array of job options.

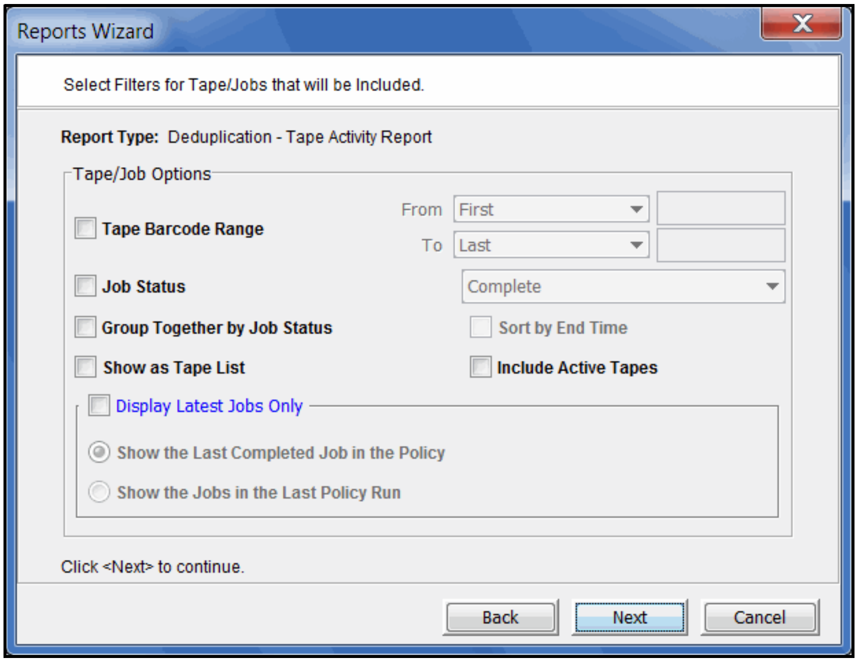

Select Filters for Tape/Jobs that will be Included.

Report Type: Deduplication - Tape Activity Report

Tape/Job Options

□ Tape Barcode Range — From: First; To: Last

□ Job Status — Complete

□ Group Together by Job Status — □ Sort by End Time

□ Show as Tape List — □ Include Active Tapes

□ Display Latest Jobs Only

● Show the Last Completed Job in the Policy

○ Show the Jobs in the Last Policy Run

Click <Next> to continue.

Tape Barcode Range - The report will include all completed jobs, regardless of their final job status, for those tapes with a barcode in the provided range. By default, the range includes all barcodes.

Job Status - Include only jobs with a specific job status in this report.

- *Complete* - Include only completed jobs.
- *Error* - Include only failed jobs.
- *Incomplete* - Include all jobs that are incomplete or were cancelled by the user or by the system.
- *New* - Include tapes that have never been deduplicated in their current policy. If a tape was previously deduplicated in a different policy, that information is not listed.

Group Together by Job Status - By default, the report will show the information grouped by policy (ID) and sorted by the job start time. If this check box is selected, the report will group the information by the policy (ID) and by the final job status (in this order: complete, error, incomplete) and sort it by descending start time.

Sort by End Time - When *Group Together by Job Status* is selected, this option replaces the start time with the end time and sorts the information in descending order.

Show as Tape List - This option displays all of the deduplication jobs regardless of the current policies. The policy information is not displayed and the report layout is a continuous list of jobs.

Include Active Tapes - This option will also display the deduplication jobs that are active at the time the report is generated. Those jobs are shown at the top of the policy or report. If you want the summary section of the report to include all currently active deduplication and replication jobs, as well as jobs that failed during the specified period of the report, you must include the current time in the date/time range.

Show the Last Completed Job in the Policy - This option lists only the last successfully completed job for each policy.

Show the Jobs in the Last Policy Run - This option lists only the last deduplication job executed in the specified interval for each policy, regardless of the job status.

Report results

The summary section includes totals for all deduplication and replication jobs.

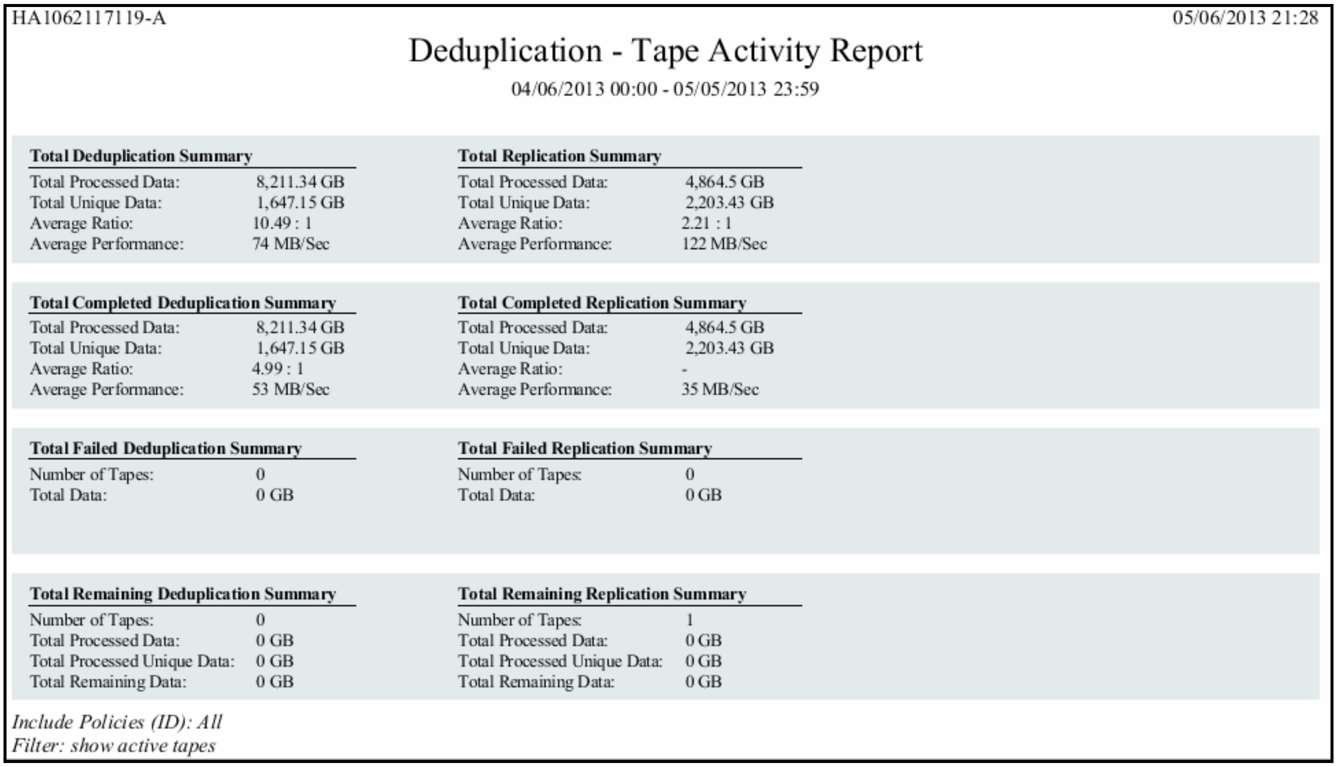

HA1062117119-A 05/06/2013 21:28

Deduplication - Tape Activity Report

04/06/2013 00:00 - 05/05/2013 23:59

Total Deduplication Summary		**Total Replication Summary**	
Total Processed Data:	8,211.34 GB	Total Processed Data:	4,864.5 GB
Total Unique Data:	1,647.15 GB	Total Unique Data:	2,203.43 GB
Average Ratio:	10.49 : 1	Average Ratio:	2.21 : 1
Average Performance:	74 MB/Sec	Average Performance:	122 MB/Sec

Total Completed Deduplication Summary		**Total Completed Replication Summary**	
Total Processed Data:	8,211.34 GB	Total Processed Data:	4,864.5 GB
Total Unique Data:	1,647.15 GB	Total Unique Data:	2,203.43 GB
Average Ratio:	4.99 : 1	Average Ratio:	-
Average Performance:	53 MB/Sec	Average Performance:	35 MB/Sec

Total Failed Deduplication Summary		**Total Failed Replication Summary**	
Number of Tapes:	0	Number of Tapes:	0
Total Data:	0 GB	Total Data:	0 GB

Total Remaining Deduplication Summary		**Total Remaining Replication Summary**	
Number of Tapes:	0	Number of Tapes:	1
Total Processed Data:	0 GB	Total Processed Data:	0 GB
Total Processed Unique Data:	0 GB	Total Processed Unique Data:	0 GB
Total Remaining Data:	0 GB	Total Remaining Data:	0 GB

Include Policies (ID): All
Filter: show active tapes

Results also include a separate section for each deduplication policy, in which the tapes in the policy are listed under their associated policy, sorted by start time (or end time if you selected that option), in descending order.

When replication is enabled in the policy, results for each tape occupy two lines because the deduplication job is executed in two phases - scanning and resolving - which are displayed in the *Tape* and *Replica* lines of the table, respectively. The *resolving* phase also includes index data replication. The *Job Status* column shows the completion status of each phase. A deduplication-with-replication job is considered complete when both phases are completed successfully.

When Advanced Replication is enabled in the policy, the *Target Server* column will show *1* or *2*, depending upon which replication target server was involved.

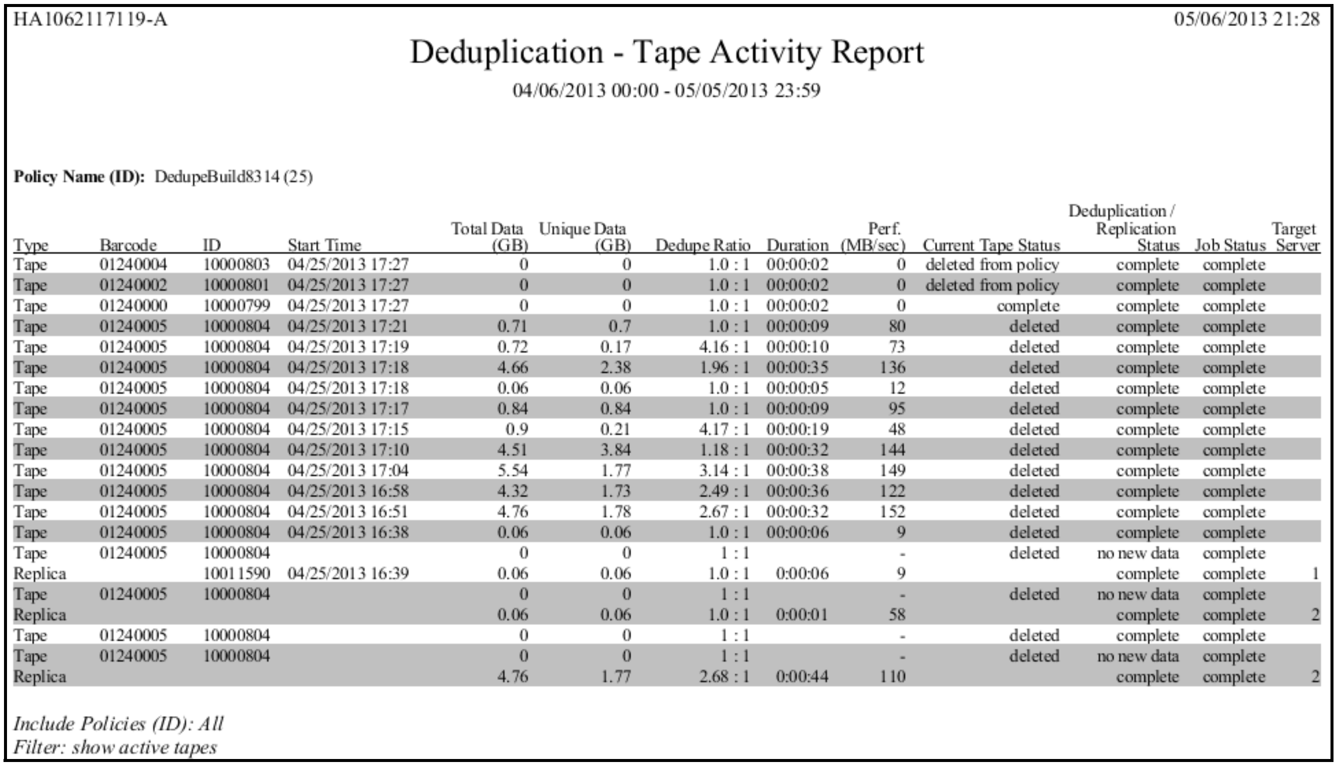

HA1062117119-A 05/06/2013 21:28

Deduplication - Tape Activity Report

04/06/2013 00:00 - 05/05/2013 23:59

Policy Name (ID): DedupeBuild8314 (25)

Type	Barcode	ID	Start Time	Total Data (GB)	Unique Data (GB)	Dedupe Ratio	Duration	Perf. (MB/sec)	Current Tape Status	Deduplication / Replication Status	Job Status	Target Server
Tape	01240004	10000803	04/25/2013 17:27	0	0	1.0 : 1	00:00:02	0	deleted from policy	complete	complete	
Tape	01240002	10000801	04/25/2013 17:27	0	0	1.0 : 1	00:00:02	0	deleted from policy	complete	complete	
Tape	01240000	10000799	04/25/2013 17:27	0	0	1.0 : 1	00:00:02	0	complete	complete	complete	
Tape	01240005	10000804	04/25/2013 17:21	0.71	0.7	1.0 : 1	00:00:09	80	deleted	complete	complete	
Tape	01240005	10000804	04/25/2013 17:19	0.72	0.17	4.16 : 1	00:00:10	73	deleted	complete	complete	
Tape	01240005	10000804	04/25/2013 17:18	4.66	2.38	1.96 : 1	00:00:35	136	deleted	complete	complete	
Tape	01240005	10000804	04/25/2013 17:18	0.06	0.06	1.0 : 1	00:00:05	12	deleted	complete	complete	
Tape	01240005	10000804	04/25/2013 17:17	0.84	0.84	1.0 : 1	00:00:09	95	deleted	complete	complete	
Tape	01240005	10000804	04/25/2013 17:15	0.9	0.21	4.17 : 1	00:00:19	48	deleted	complete	complete	
Tape	01240005	10000804	04/25/2013 17:10	4.51	3.84	1.18 : 1	00:00:32	144	deleted	complete	complete	
Tape	01240005	10000804	04/25/2013 17:04	5.54	1.77	3.14 : 1	00:00:38	149	deleted	complete	complete	
Tape	01240005	10000804	04/25/2013 16:58	4.32	1.73	2.49 : 1	00:00:36	122	deleted	complete	complete	
Tape	01240005	10000804	04/25/2013 16:51	4.76	1.78	2.67 : 1	00:00:32	152	deleted	complete	complete	
Tape	01240005	10000804	04/25/2013 16:38	0.06	0.06	1.0 : 1	00:00:06	9	deleted	complete	complete	
Tape	01240005	10000804		0	0	1 : 1		-	deleted	no new data	complete	
Replica		10011590	04/25/2013 16:39	0.06	0.06	1.0 : 1	0:00:06	9		complete	complete	1
Tape	01240005	10000804		0	0	1 : 1		-	deleted	no new data	complete	
Replica				0.06	0.06	1.0 : 1	0:00:01	58		complete	complete	2
Tape	01240005	10000804		0	0	1 : 1		-	deleted	complete	complete	
Tape	01240005	10000804		0	0	1 : 1		-	deleted	no new data	complete	
Replica				4.76	1.77	2.68 : 1	0:00:44	110		complete	complete	2

Include Policies (ID): All
Filter: show active tapes

A summary for the policy is displayed at the end of each policy section. In addition to information about data processed and performance, the summary identifies the replication target servers.

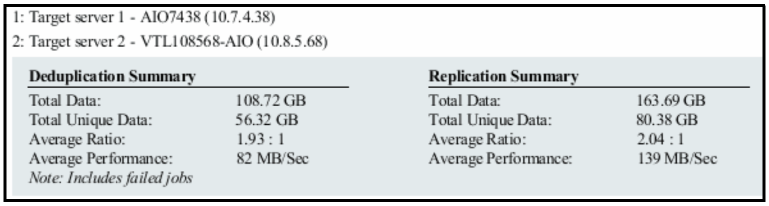

1: Target server 1 - AIO7438 (10.7.4.38)
2: Target server 2 - VTL108568-AIO (10.8.5.68)

Deduplication Summary		**Replication Summary**	
Total Data:	108.72 GB	Total Data:	163.69 GB
Total Unique Data:	56.32 GB	Total Unique Data:	80.38 GB
Average Ratio:	1.93 : 1	Average Ratio:	2.04 : 1
Average Performance:	82 MB/Sec	Average Performance:	139 MB/Sec

Note: Includes failed jobs

Deduplication Replication Status

This status report is run from a VTL target server and displays information about replication of tapes in deduplication policies (deduplication replication) on selected source servers, making it possible to determine how well data on those servers is protected.

This report will run only for the current day and can also be configured as a scheduled report.

Replication summary information for all tapes on the selected servers is always included in a pie chart at the beginning of the tabular report. Three detail categories are available for the tabular report: *Replicated Tapes, Not-Replicated Tapes,* and *Zero-Data Tapes*. When you select one or more categories, the table includes details for tapes in that category.

The pie chart summary shows the percentage of tapes that have been replicated and the percentage that have not been replicated. The *Not-Replicated* category includes tapes for which replication is in progress, are currently being resolved, or which have never been replicated. Tapes with zero data are counted in the total number of tapes but are not represented in the chart.

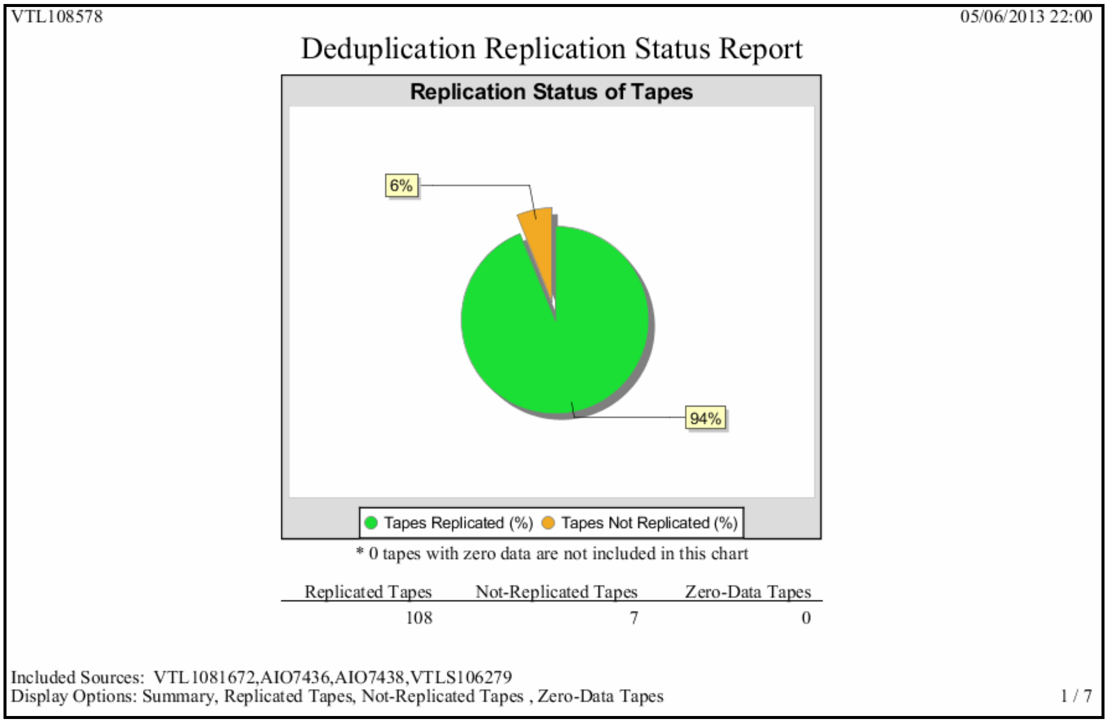

VTL108578 05/06/2013 22:00

Deduplication Replication Status Report

Replication Status of Tapes
6%
94%
● Tapes Replicated (%) ● Tapes Not Replicated (%)

* 0 tapes with zero data are not included in this chart

Replicated Tapes	Not-Replicated Tapes	Zero-Data Tapes
108	7	0

Included Sources: VTL1081672,AIO7436,AIO7438,VTLS106279
Display Options: Summary, Replicated Tapes, Not-Replicated Tapes , Zero-Data Tapes 1 / 7

A summary area at the top of the first page of the tabular report identifies the source VTL servers included in the report and for each source server, indicates the number of tapes that are configured for replication to the target, the number of replicated tapes, not-replicated tapes, and zero-data tapes, the percentage of data that has been replicated on those tapes, and whether this data is based on a report successfully retrieved from the source server.

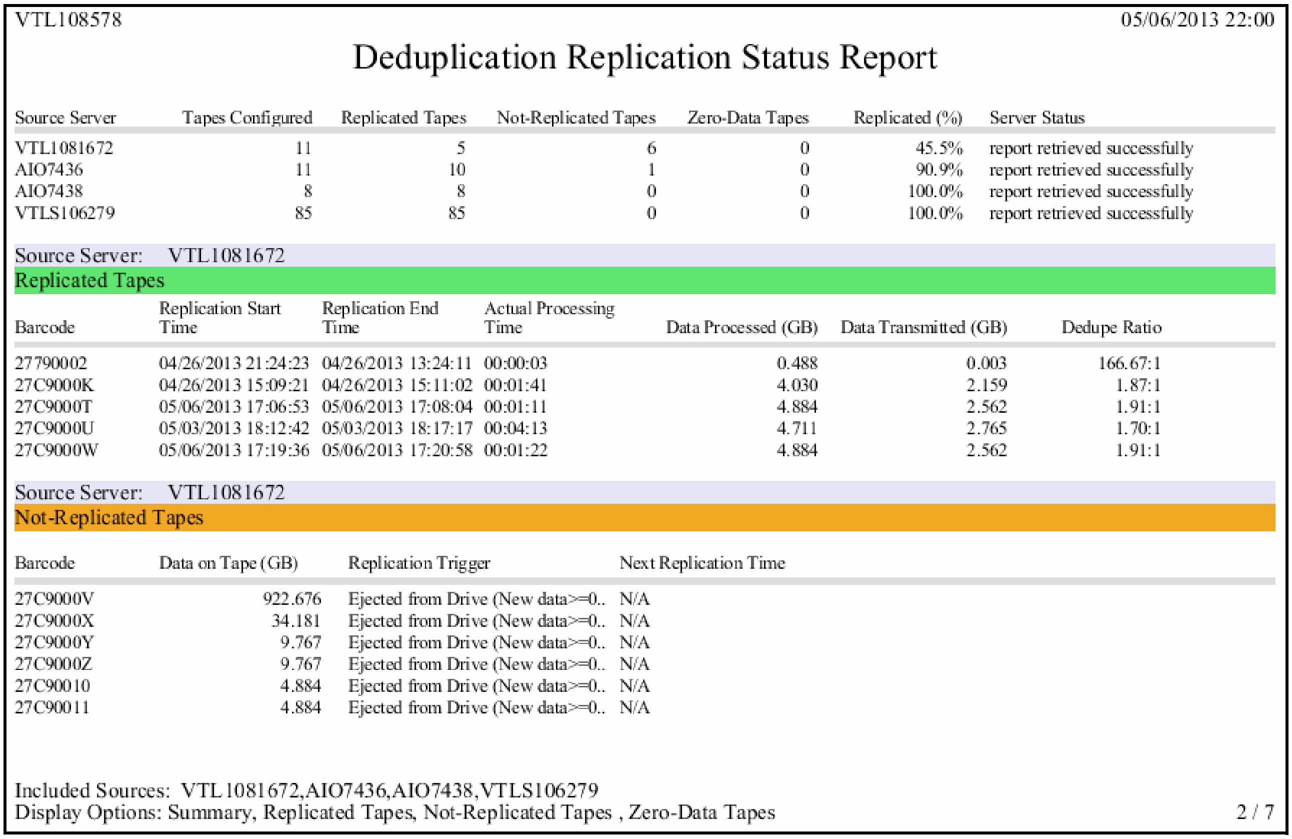

VTL108578 05/06/2013 22:00

Deduplication Replication Status Report

Source Server	Tapes Configured	Replicated Tapes	Not-Replicated Tapes	Zero-Data Tapes	Replicated (%)	Server Status
VTL1081672	11	5	6	0	45.5%	report retrieved successfully
AIO7436	11	10	1	0	90.9%	report retrieved successfully
AIO7438	8	8	0	0	100.0%	report retrieved successfully
VTLS106279	85	85	0	0	100.0%	report retrieved successfully

Source Server: VTL1081672

Replicated Tapes

Barcode	Replication Start Time	Replication End Time	Actual Processing Time	Data Processed (GB)	Data Transmitted (GB)	Dedupe Ratio
27790002	04/26/2013 21:24:23	04/26/2013 13:24:11	00:00:03	0.488	0.003	166.67:1
27C9000K	04/26/2013 15:09:21	04/26/2013 15:11:02	00:01:41	4.030	2.159	1.87:1
27C9000T	05/06/2013 17:06:53	05/06/2013 17:08:04	00:01:11	4.884	2.562	1.91:1
27C9000U	05/03/2013 18:12:42	05/03/2013 18:17:17	00:04:13	4.711	2.765	1.70:1
27C9000W	05/06/2013 17:19:36	05/06/2013 17:20:58	00:01:22	4.884	2.562	1.91:1

Source Server: VTL1081672

Not-Replicated Tapes

Barcode	Data on Tape (GB)	Replication Trigger	Next Replication Time
27C9000V	922.676	Ejected from Drive (New data>=0..	N/A
27C9000X	34.181	Ejected from Drive (New data>=0..	N/A
27C9000Y	9.767	Ejected from Drive (New data>=0..	N/A
27C9000Z	9.767	Ejected from Drive (New data>=0..	N/A
27C90010	4.884	Ejected from Drive (New data>=0..	N/A
27C90011	4.884	Ejected from Drive (New data>=0..	N/A

Included Sources: VTL1081672,AIO7436,AIO7438,VTLS106279

Display Options: Summary, Replicated Tapes, Not-Replicated Tapes , Zero-Data Tapes 2 / 7

Details for *Zero-Data Tapes* include tape barcodes, the replication trigger, and the next time replication is scheduled to occur.

For *Replicated Tapes*, details include tape barcodes, the time replication started and ended and how long it took to complete, plus the amount of data processed (the total size of data written to the tape) and transmitted to the target, and the deduplication ratio for the latest successful replication.

For *Not-Replicated Tapes*, details include tape barcodes, the amount of data on the tape, the replication trigger, and the next replication time.

Notes:

- The *Data Processed* value should be the same as the amount reported for the source tape by the backup application.
- When replication is enabled for a VTL deduplication policy, replication is automatically set for each tape as it is added to the policy. Replication of tapes in deduplication policies for the FalconStor OpenStorage option (FSOST) works differently. When FSOST deduplication policies are configured for replication, replication is not configured on individual source tapes in the policy until replication is actually triggered. However, since all tapes in the policy *will be* replicated, they are all counted as configured for replication even if replication has not yet occurred.

Deduplication Repository - Reclamation

The Dashboard Summary tab (displayed for a deduplication server) and the Repository Dashboard Summary tab (displayed for a VTL-S server) display information about repository reclamation. This status report is an alternate way to view the reclamation information on these tabs.

Repository reclamation is a process by which the VTL system reclaims space no longer needed on the deduplication data and metadata disks, making that space available for use. Reclamation also frees up repository index cache that is no longer being used. For details on the reclamation process, refer to ‘Reclaim disk space’.

This history report shows information about all reclamation operations within the specified period of time.

Because information for this report is found on the associated deduplication server, an association between the backup server and a deduplication server must exist in order for this report to be available.

This report is available only as a one-time report.

The report display consists of four bar graphs, each one representing an area in which reclamation can be performed: index cache, folder disk, index disk, and data disk. Each bar represents a reclamation job and displays the resulting reclaimed space (in GB for folder/index/data disk) and index cache consumed by the repository (as a percentage). The example below shows all reclamation jobs that occurred during the 30 days prior to the current day.

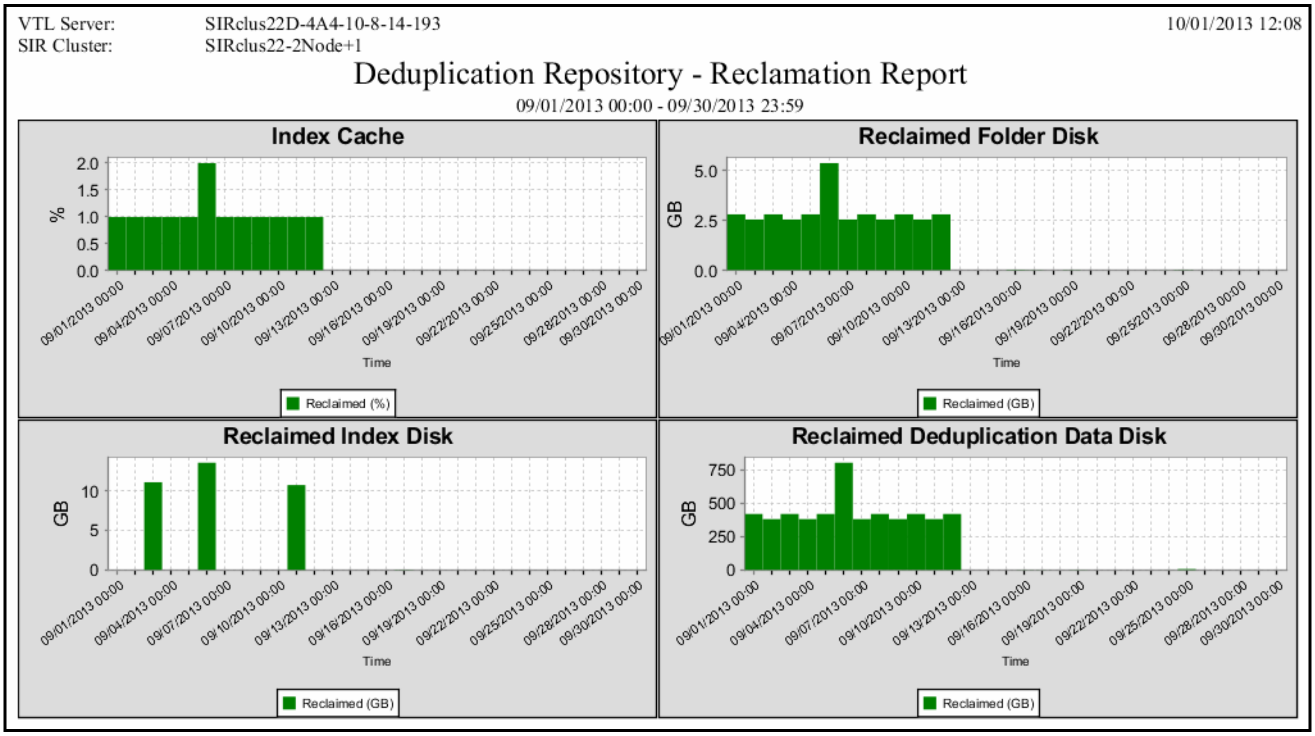

VTL Server: SIRclus22D-4A4-10-8-14-193 10/01/2013 12:08
SIR Cluster: SIRclus22-2Node+1

Deduplication Repository - Reclamation Report

09/01/2013 00:00 - 09/30/2013 23:59

Results are also displayed in tabular form. Each row in the table includes details about the reclamation operation performed on one of the deduplication resources.

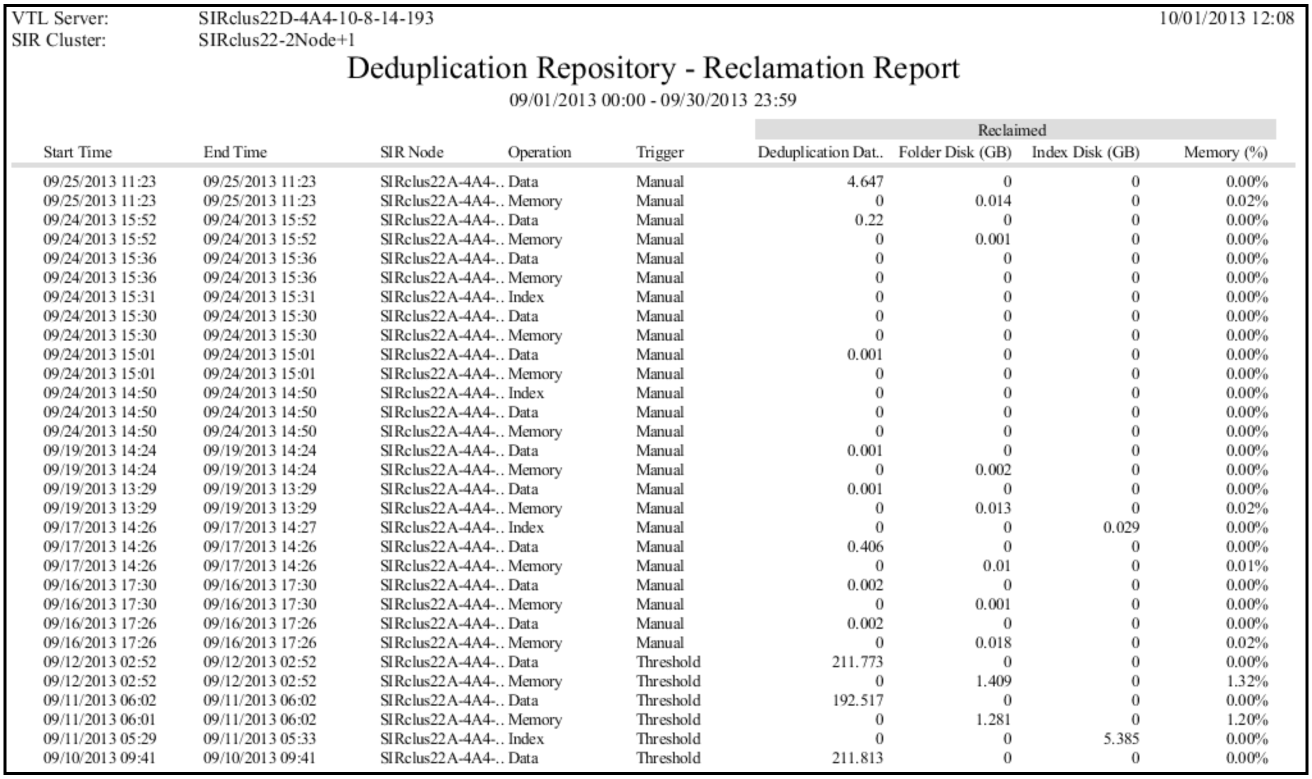

VTL Server: SIRclus22D-4A4-10-8-14-193 10/01/2013 12:08
SIR Cluster: SIRclus22-2Node+1

Deduplication Repository - Reclamation Report

09/01/2013 00:00 - 09/30/2013 23:59

					Reclaimed			
Start Time	End Time	SIR Node	Operation	Trigger	Deduplication Dat..	Folder Disk (GB)	Index Disk (GB)	Memory (%)
09/25/2013 11:23	09/25/2013 11:23	SIRclus22A-4A4-..	Data	Manual	4.647	0	0	0.00%
09/25/2013 11:23	09/25/2013 11:23	SIRclus22A-4A4-..	Memory	Manual	0	0.014	0	0.02%
09/24/2013 15:52	09/24/2013 15:52	SIRclus22A-4A4-..	Data	Manual	0.22	0	0	0.00%
09/24/2013 15:52	09/24/2013 15:52	SIRclus22A-4A4-..	Memory	Manual	0	0.001	0	0.00%
09/24/2013 15:36	09/24/2013 15:36	SIRclus22A-4A4-..	Data	Manual	0	0	0	0.00%
09/24/2013 15:36	09/24/2013 15:36	SIRclus22A-4A4-..	Memory	Manual	0	0	0	0.00%
09/24/2013 15:31	09/24/2013 15:31	SIRclus22A-4A4-..	Index	Manual	0	0	0	0.00%
09/24/2013 15:30	09/24/2013 15:30	SIRclus22A-4A4-..	Data	Manual	0	0	0	0.00%
09/24/2013 15:30	09/24/2013 15:30	SIRclus22A-4A4-..	Memory	Manual	0	0	0	0.00%
09/24/2013 15:01	09/24/2013 15:01	SIRclus22A-4A4-..	Data	Manual	0.001	0	0	0.00%
09/24/2013 15:01	09/24/2013 15:01	SIRclus22A-4A4-..	Memory	Manual	0	0	0	0.00%
09/24/2013 14:50	09/24/2013 14:50	SIRclus22A-4A4-..	Index	Manual	0	0	0	0.00%
09/24/2013 14:50	09/24/2013 14:50	SIRclus22A-4A4-..	Data	Manual	0	0	0	0.00%
09/24/2013 14:50	09/24/2013 14:50	SIRclus22A-4A4-..	Memory	Manual	0	0	0	0.00%
09/19/2013 14:24	09/19/2013 14:24	SIRclus22A-4A4-..	Data	Manual	0.001	0	0	0.00%
09/19/2013 14:24	09/19/2013 14:24	SIRclus22A-4A4-..	Memory	Manual	0	0.002	0	0.00%
09/19/2013 13:29	09/19/2013 13:29	SIRclus22A-4A4-..	Data	Manual	0.001	0	0	0.00%
09/19/2013 13:29	09/19/2013 13:29	SIRclus22A-4A4-..	Memory	Manual	0	0.013	0	0.02%
09/17/2013 14:26	09/17/2013 14:27	SIRclus22A-4A4-..	Index	Manual	0	0	0.029	0.00%
09/17/2013 14:26	09/17/2013 14:26	SIRclus22A-4A4-..	Data	Manual	0.406	0	0	0.00%
09/17/2013 14:26	09/17/2013 14:26	SIRclus22A-4A4-..	Memory	Manual	0	0.01	0	0.01%
09/16/2013 17:30	09/16/2013 17:30	SIRclus22A-4A4-..	Data	Manual	0.002	0	0	0.00%
09/16/2013 17:30	09/16/2013 17:30	SIRclus22A-4A4-..	Memory	Manual	0	0.001	0	0.00%
09/16/2013 17:26	09/16/2013 17:26	SIRclus22A-4A4-..	Data	Manual	0.002	0	0	0.00%
09/16/2013 17:26	09/16/2013 17:26	SIRclus22A-4A4-..	Memory	Manual	0	0.018	0	0.02%
09/12/2013 02:52	09/12/2013 02:52	SIRclus22A-4A4-..	Data	Threshold	211.773	0	0	0.00%
09/12/2013 02:52	09/12/2013 02:52	SIRclus22A-4A4-..	Memory	Threshold	0	1.409	0	1.32%
09/11/2013 06:02	09/11/2013 06:02	SIRclus22A-4A4-..	Data	Threshold	192.517	0	0	0.00%
09/11/2013 06:01	09/11/2013 06:02	SIRclus22A-4A4-..	Memory	Threshold	0	1.281	0	1.20%
09/11/2013 05:29	09/11/2013 05:33	SIRclus22A-4A4-..	Index	Threshold	0	0	5.385	0.00%
09/10/2013 09:41	09/10/2013 09:41	SIRclus22A-4A4-..	Data	Threshold	211.813	0	0	0.00%

Import/Export Jobs

This history report lists all import/export and tape caching jobs that were placed in the queue during the specified period of time, regardless of job status. Options in the wizard include job types and statuses.

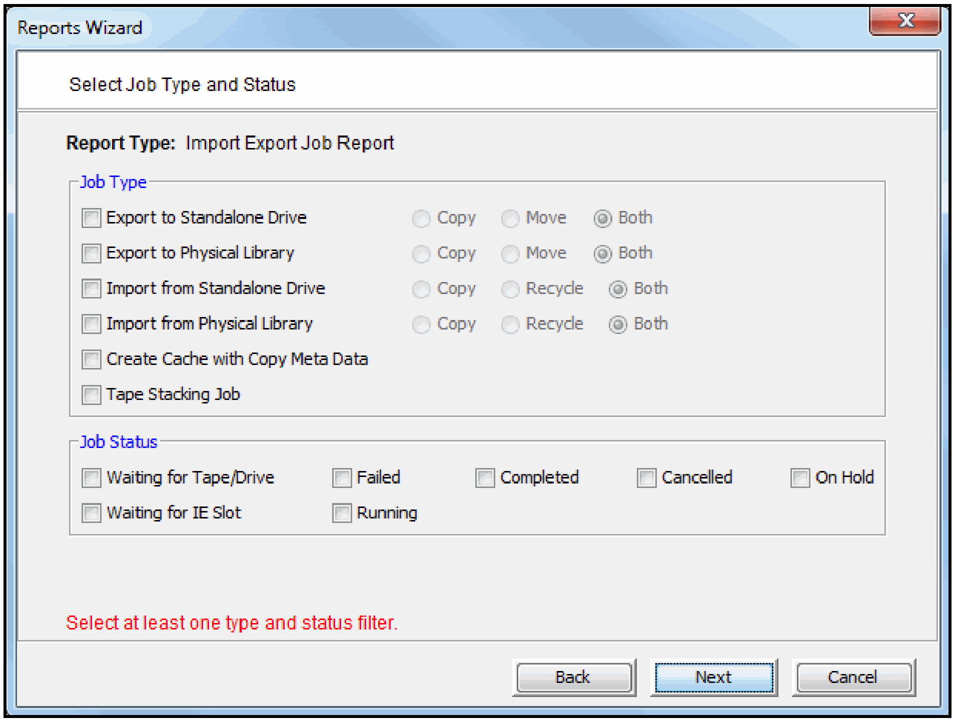

Select Job Type and Status

Report Type: Import Export Job Report

Job Type

□ Export to Standalone Drive ○ Copy ○ Move ● Both
□ Export to Physical Library ○ Copy ○ Move ● Both
□ Import from Standalone Drive ○ Copy ○ Recycle ● Both
□ Import from Physical Library ○ Copy ○ Recycle ● Both
□ Create Cache with Copy Meta Data
□ Tape Stacking Job

Job Status

□ Waiting for Tape/Drive □ Failed □ Completed □ Cancelled □ On Hold
□ Waiting for IE Slot □ Running

Select at least one type and status filter.

You must select at least one job type and one job status.

Job type

- Export to Standalone Drive/Export to Physical Library - For these jobs, you can include results for jobs that used *Copy Mode*, *Move Mode*, or both.
- Import from Standalone Drive/Import from Physical Library - For these jobs, you can include results for jobs that used *Copy Mode*, *Recycle Mode*, or both.

Job status

For all selected job types, the report will include information on jobs with the selected status(es).

The summary page displays the number of jobs found for all job types and all job statuses.

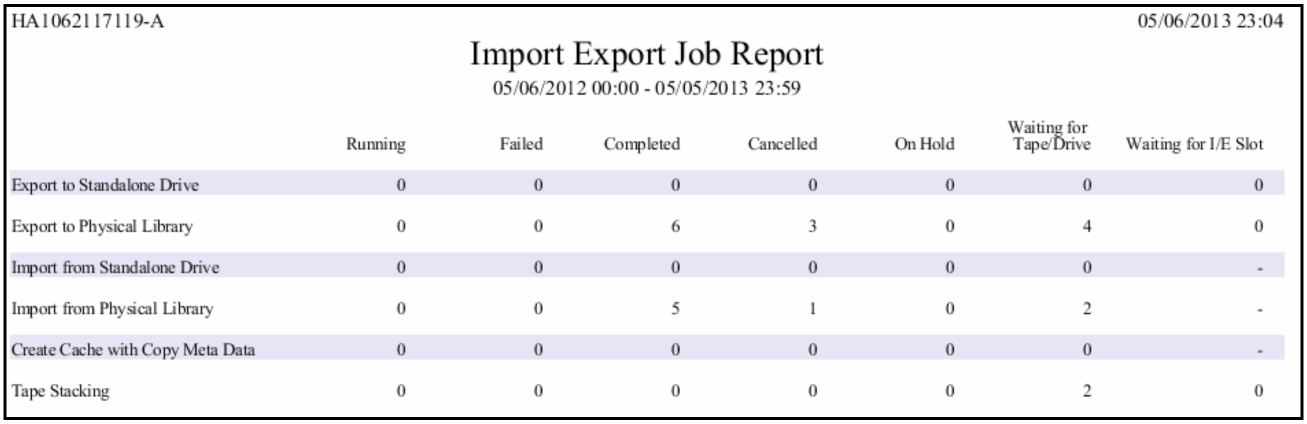

HA1062117119-A 05/06/2013 23:04

Import Export Job Report

05/06/2012 00:00 - 05/05/2013 23:59

	Running	Failed	Completed	Cancelled	On Hold	Waiting for Tape/Drive	Waiting for I/E Slot
Export to Standalone Drive	0	0	0	0	0	0	0
Export to Physical Library	0	0	6	3	0	4	0
Import from Standalone Drive	0	0	0	0	0	0	-
Import from Physical Library	0	0	5	1	0	2	-
Create Cache with Copy Meta Data	0	0	0	0	0	0	-
Tape Stacking	0	0	0	0	0	2	0

The detail pages display results based on the job type(s) and job status(es) you selected. Jobs are ordered by job ID. For each job, the report lists job ID, the job type, barcodes of source/destination tapes; locations of source/destination tapes, import/export mode, job status, start/end time of job, amount of data transferred, and job throughput.

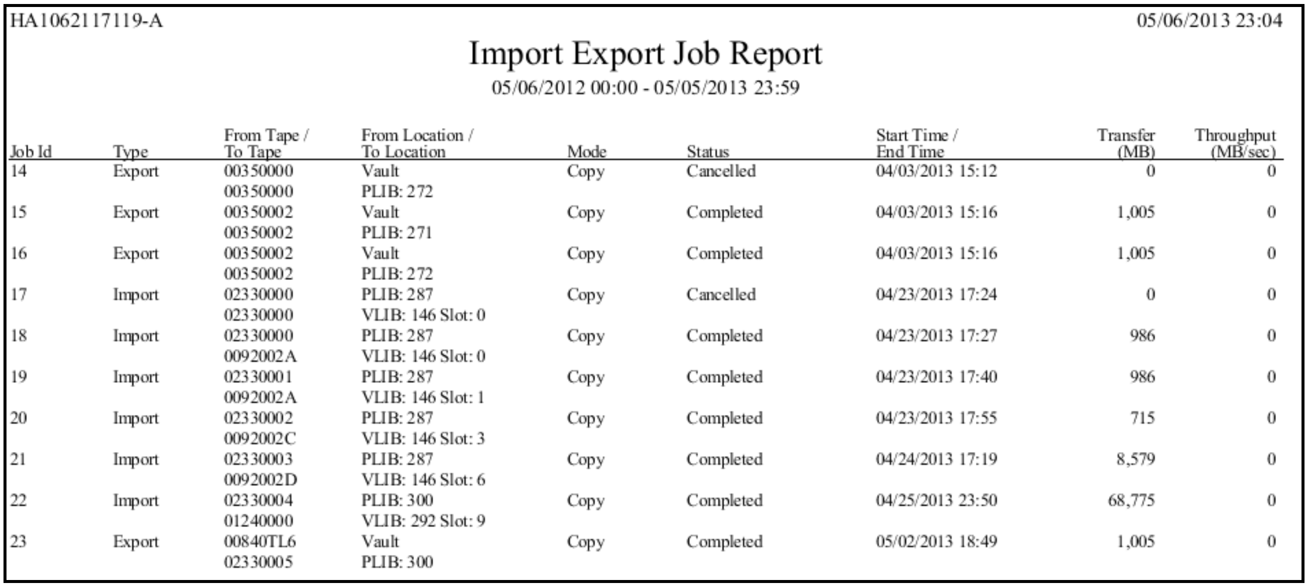

HA1062117119-A 05/06/2013 23:04

Import Export Job Report

05/06/2012 00:00 - 05/05/2013 23:59

Job Id	Type	From Tape / To Tape	From Location / To Location	Mode	Status	Start Time / End Time	Transfer (MB)	Throughput (MB/sec)
14	Export	00350000 00350000	Vault PLIB: 272	Copy	Cancelled	04/03/2013 15:12	0	0
15	Export	00350002 00350002	Vault PLIB: 271	Copy	Completed	04/03/2013 15:16	1,005	0
16	Export	00350002 00350002	Vault PLIB: 272	Copy	Completed	04/03/2013 15:16	1,005	0
17	Import	02330000 02330000	PLIB: 287 VLIB: 146 Slot: 0	Copy	Cancelled	04/23/2013 17:24	0	0
18	Import	02330000 0092002A	PLIB: 287 VLIB: 146 Slot: 0	Copy	Completed	04/23/2013 17:27	986	0
19	Import	02330001 0092002A	PLIB: 287 VLIB: 146 Slot: 1	Copy	Completed	04/23/2013 17:40	986	0
20	Import	02330002 0092002C	PLIB: 287 VLIB: 146 Slot: 3	Copy	Completed	04/23/2013 17:55	715	0
21	Import	02330003 0092002D	PLIB: 287 VLIB: 146 Slot: 6	Copy	Completed	04/24/2013 17:19	8,579	0
22	Import	02330004 01240000	PLIB: 300 VLIB: 292 Slot: 9	Copy	Completed	04/25/2013 23:50	68,775	0
23	Export	00840TL6 02330005	Vault PLIB: 300	Copy	Completed	05/02/2013 18:49	1,005	0

Replication Status

This history report displays information about virtual tapes enabled for replication and for virtual tape replicas, during the selected period of time. The wizard provides the following report options:

For virtual tapes

- *Sort by target server name* - For each target server, the report lists all virtual tape replicas and for each replica, the log dates and times of all replication activity.
- *Sort by log date and time* - The report lists log dates and times of all replication activity in ascending order. Details are arranged by virtual tape replica names for each target server.

For virtual tape replicas

- *Sort by primary server name* - For each primary server, the report lists all primary virtual tapes and for each tape, the log dates and times of all replication activity.
- *Sort by log date and time* - The report lists log dates and times of all replication activity in ascending order. Details are arranged by virtual tape names for each primary server.

Report results always identify the primary and target server, the name of the primary virtual tape and virtual tape replica, and the associated policy name and its replication options. Log information always incudes the log time, current replication activity status, start and end time, the amount of data analyzed, the percentage of data analyzed, the trigger, and comments. The sample below shows typical report layout, irrespective of sorting options.

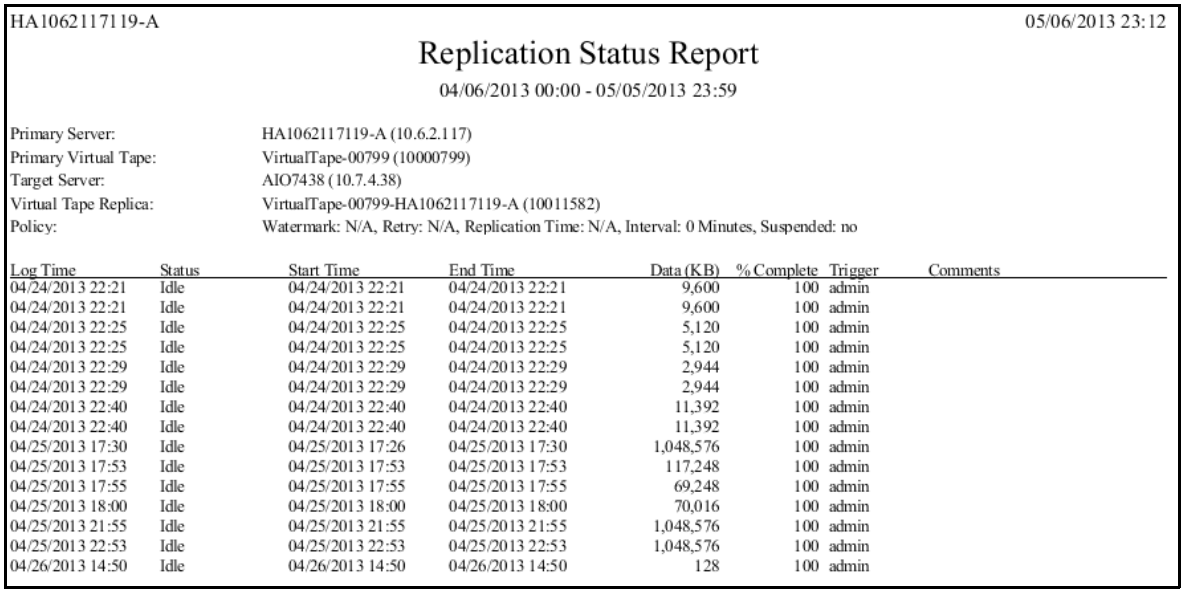

HA1062117119-A 05/06/2013 23:12

Replication Status Report

04/06/2013 00:00 - 05/05/2013 23:59

Primary Server:	HA1062117119-A (10.6.2.117)
Primary Virtual Tape:	VirtualTape-00799 (10000799)
Target Server:	AIO7438 (10.7.4.38)
Virtual Tape Replica:	VirtualTape-00799-HA1062117119-A (10011582)
Policy:	Watermark: N/A, Retry: N/A, Replication Time: N/A, Interval: 0 Minutes, Suspended: no

Log Time	Status	Start Time	End Time	Data (KB)	% Complete	Trigger	Comments
04/24/2013 22:21	Idle	04/24/2013 22:21	04/24/2013 22:21	9,600	100	admin	
04/24/2013 22:21	Idle	04/24/2013 22:21	04/24/2013 22:21	9,600	100	admin	
04/24/2013 22:25	Idle	04/24/2013 22:25	04/24/2013 22:25	5,120	100	admin	
04/24/2013 22:25	Idle	04/24/2013 22:25	04/24/2013 22:25	5,120	100	admin	
04/24/2013 22:29	Idle	04/24/2013 22:29	04/24/2013 22:29	2,944	100	admin	
04/24/2013 22:29	Idle	04/24/2013 22:29	04/24/2013 22:29	2,944	100	admin	
04/24/2013 22:40	Idle	04/24/2013 22:40	04/24/2013 22:40	11,392	100	admin	
04/24/2013 22:40	Idle	04/24/2013 22:40	04/24/2013 22:40	11,392	100	admin	
04/25/2013 17:30	Idle	04/25/2013 17:26	04/25/2013 17:30	1,048,576	100	admin	
04/25/2013 17:53	Idle	04/25/2013 17:53	04/25/2013 17:53	117,248	100	admin	
04/25/2013 17:55	Idle	04/25/2013 17:55	04/25/2013 17:55	69,248	100	admin	
04/25/2013 18:00	Idle	04/25/2013 18:00	04/25/2013 18:00	70,016	100	admin	
04/25/2013 21:55	Idle	04/25/2013 21:55	04/25/2013 21:55	1,048,576	100	admin	
04/25/2013 22:53	Idle	04/25/2013 22:53	04/25/2013 22:53	1,048,576	100	admin	
04/26/2013 14:50	Idle	04/26/2013 14:50	04/26/2013 14:50	128	100	admin	

Note: Because the report wizard is not designed to identify a server as a primary server or target server, report results will not be generated if you run the report on a target server and select *Virtual Tapes* or if you run the report on a source server and select *Virtual Tape Replicas*.

Virtual Library and Drive Assignment

This status report displays virtual tape library and drive assignments for all clients on the system, for the current server date. Results are presented from four different points of view: the Tape Library Summary, Drive Summary, and Client Summary.

The Tape Library Summary lists all virtual tape libraries on the system by name and ID and for each library displays its product ID, serial number, and assigned client, the initiator and target WWPNs assigned to the client, and the number of drives it includes.

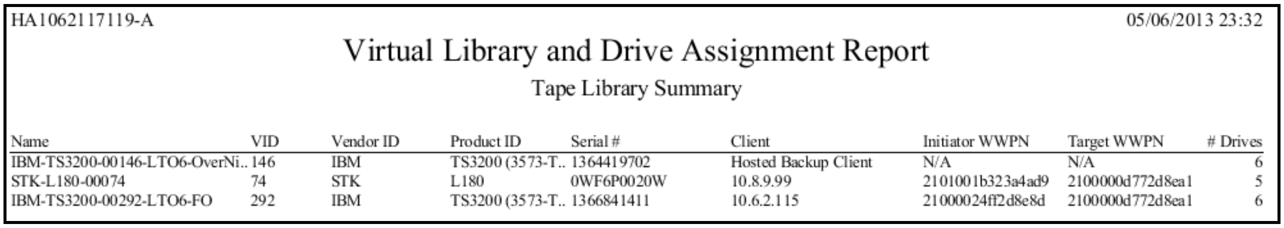

HA1062117119-A 05/06/2013 23:32

Virtual Library and Drive Assignment Report

Tape Library Summary

Name	VID	Vendor ID	Product ID	Serial #	Client	Initiator WWPN	Target WWPN	# Drives
IBM-TS3200-00146-LTO6-OverNi..	146	IBM	TS3200 (3573-T..	1364419702	Hosted Backup Client	N/A	N/A	6
STK-L180-00074	74	STK	L180	0WF6P0020W	10.8.9.99	2101001b323a4ad9	2100000d772d8ea1	5
IBM-TS3200-00292-LTO6-FO	292	IBM	TS3200 (3573-T..	1366841411	10.6.2.115	21000024ff2d8e8d	2100000d772d8ea1	6

The Tape Drive Summary lists all tape drives on the system by name and ID and for each tape drive displays its product ID, serial number, assigned client, and initiator and target WWPNs assigned to the client.

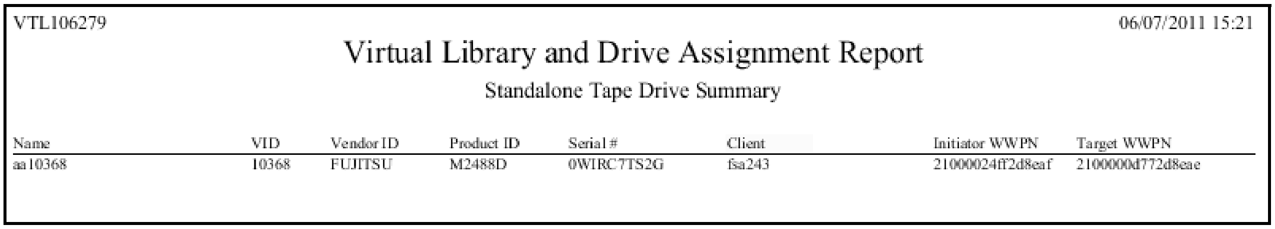

VTL106279 06/07/2011 15:21

Virtual Library and Drive Assignment Report

Standalone Tape Drive Summary

Name	VID	Vendor ID	Product ID	Serial #	Client	Initiator WWPN	Target WWPN
aa10368	10368	FUJITSU	M2488D	0WIRC7TS2G	fsa243	21000024ff2d8eaf	2100000d772d8eae

Summary information for each client (including the Hosted Backup Client) lists the name of all devices on the system and for each, displays the type of device (library or drive), ID, vendor, product ID, serial number, and initiator and target WWPNs assigned to the client.

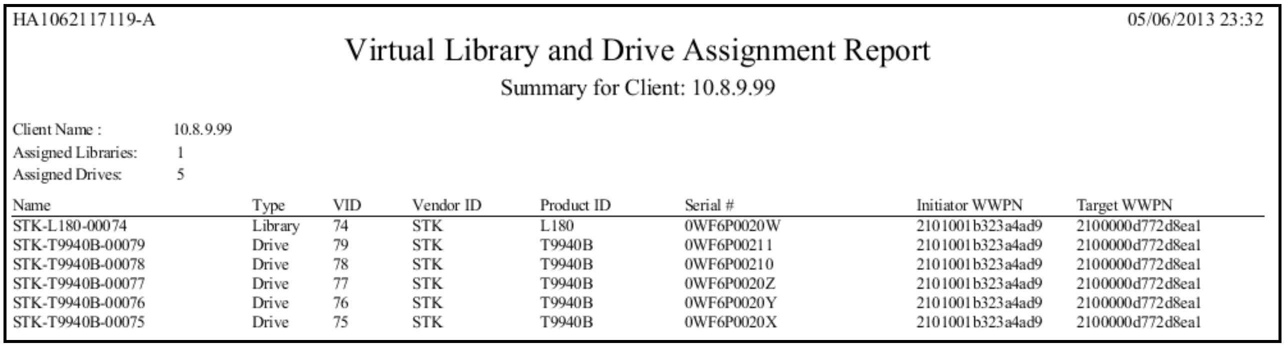

HA1062117119-A 05/06/2013 23:32

Virtual Library and Drive Assignment Report

Summary for Client: 10.8.9.99

Client Name : 10.8.9.99
Assigned Libraries: 1
Assigned Drives: 5

Name	Type	VID	Vendor ID	Product ID	Serial #	Initiator WWPN	Target WWPN
STK-L180-00074	Library	74	STK	L180	0WF6P0020W	2101001b323a4ad9	2100000d772d8ea1
STK-T9940B-00079	Drive	79	STK	T9940B	0WF6P00211	2101001b323a4ad9	2100000d772d8ea1
STK-T9940B-00078	Drive	78	STK	T9940B	0WF6P00210	2101001b323a4ad9	2100000d772d8ea1
STK-T9940B-00077	Drive	77	STK	T9940B	0WF6P0020Z	2101001b323a4ad9	2100000d772d8ea1
STK-T9940B-00076	Drive	76	STK	T9940B	0WF6P0020Y	2101001b323a4ad9	2100000d772d8ea1
STK-T9940B-00075	Drive	75	STK	T9940B	0WF6P0020X	2101001b323a4ad9	2100000d772d8ea1

Virtual Library Information

This status report displays information about each virtual tape library on the system, including the physical library it emulates, the amount of storage it occupies, and information about its drives, tapes, and slots. Results are displayed in a pie chart as well as in tabular format.

In the pie chart, the amount of disk space occupied by each library, including space occupied by the Virtual Vault, is displayed as a percentage of all available storage.

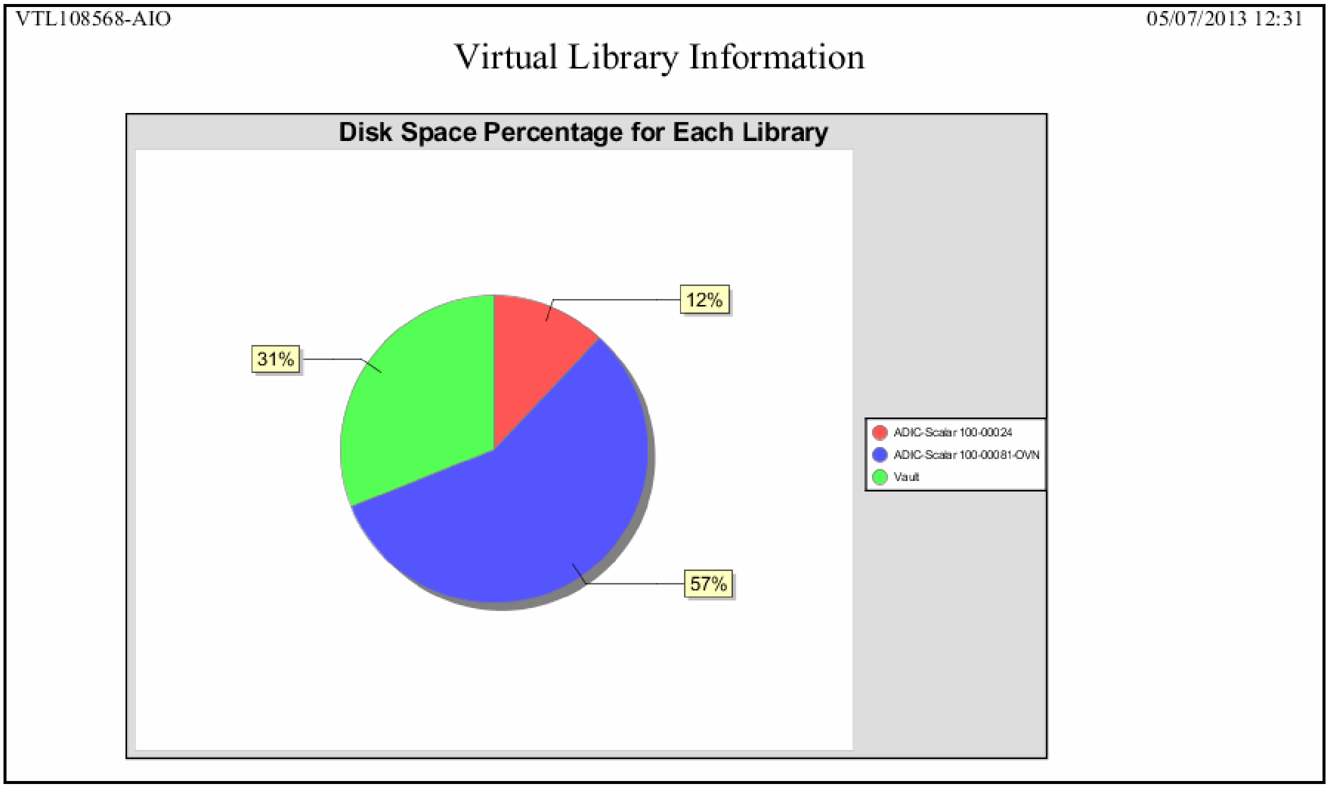

VTL108568-AIO 05/07/2013 12:31

Virtual Library Information

The results table lists all configured VTL libraries and for each one displays its name and virtual ID; the library and drive vendor/product it emulates; the number of drives, tapes, slots, and IE slots, the amount of storage it occupies, whether or not the library has encryption enabled, and assigned clients.

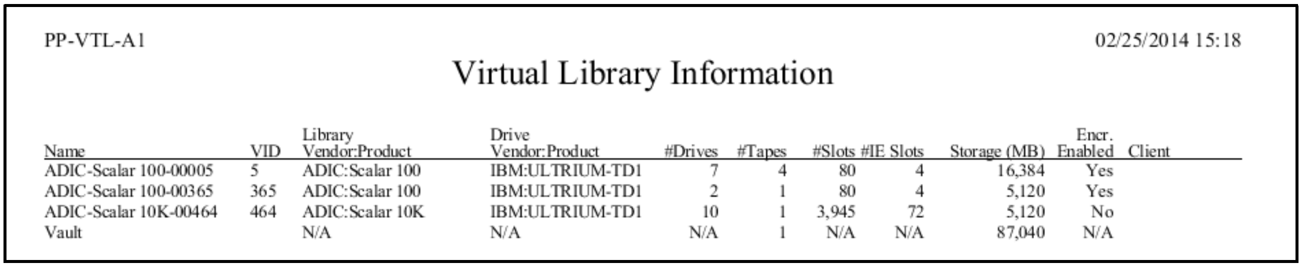

PP-VTL-A1 02/25/2014 15:18

Virtual Library Information

Name	VID	Library Vendor:Product	Drive Vendor:Product	#Drives	#Tapes	#Slots	#IE Slots	Storage (MB)	Encr. Enabled	Client
ADIC-Scalar 100-00005	5	ADIC:Scalar 100	IBM:ULTRIUM-TD1	7	4	80	4	16,384	Yes	
ADIC-Scalar 100-00365	365	ADIC:Scalar 100	IBM:ULTRIUM-TD1	2	1	80	4	5,120	Yes	
ADIC-Scalar 10K-00464	464	ADIC:Scalar 10K	IBM:ULTRIUM-TD1	10	1	3,945	72	5,120	No	
Vault		N/A	N/A	N/A	1	N/A	N/A	87,040	N/A	

This report does not include tape segments used for some types of internal tracking and processing when calculating used or allocated space.

Virtual Tape Activity

This history report shows activity for all virtual tapes within the specified period of time for three types of operations: Backup, Tape Import (*WRITE* operations), and Tape Export (*READ* operations).

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. Note that when you specify a date range larger than a week, it may take a considerable length of time (e.g., up to half an hour) to display due to the large amount data the report includes.

Regardless of the operation, the information for each report item includes the start and end time of the job, tape barcode, data compression ratio, job duration, and the speed of the operation in megabytes per second.

The data compression ratio is calculated as the total amount of user data divided by the actual amount of data saved to the tape when compression is enabled. This ratio includes space used for metadata, such as the tape header.

Typical displays include:

- For backup operations,

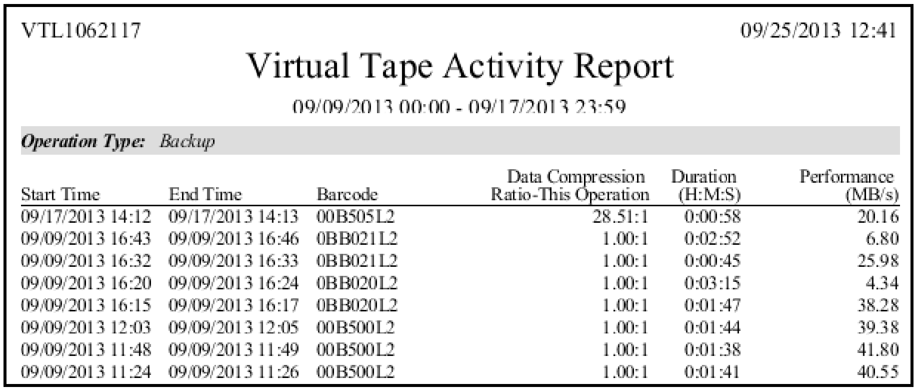

VTL1062117 09/25/2013 12:41

Virtual Tape Activity Report

09/09/2013 00:00 - 09/17/2013 23:59

Operation Type: *Backup*

Start Time	End Time	Barcode	Data Compression Ratio-This Operation	Duration (H:M:S)	Performance (MB/s)
09/17/2013 14:12	09/17/2013 14:13	00B505L2	28.51:1	0:00:58	20.16
09/09/2013 16:43	09/09/2013 16:46	0BB021L2	1.00:1	0:02:52	6.80
09/09/2013 16:32	09/09/2013 16:33	0BB021L2	1.00:1	0:00:45	25.98
09/09/2013 16:20	09/09/2013 16:24	0BB020L2	1.00:1	0:03:15	4.34
09/09/2013 16:15	09/09/2013 16:17	0BB020L2	1.00:1	0:01:47	38.28
09/09/2013 12:03	09/09/2013 12:05	00B500L2	1.00:1	0:01:44	39.38
09/09/2013 11:48	09/09/2013 11:49	00B500L2	1.00:1	0:01:38	41.80
09/09/2013 11:24	09/09/2013 11:26	00B500L2	1.00:1	0:01:41	40.55

- For export operations,

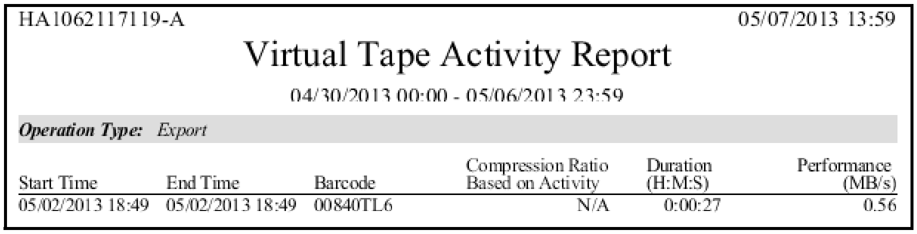

HA1062117119-A 05/07/2013 13:59

Virtual Tape Activity Report

04/30/2013 00:00 - 05/06/2013 23:59

Operation Type: *Export*

Start Time	End Time	Barcode	Compression Ratio Based on Activity	Duration (H:M:S)	Performance (MB/s)
05/02/2013 18:49	05/02/2013 18:49	00840TL6	N/A	0:00:27	0.56

- For import operations,

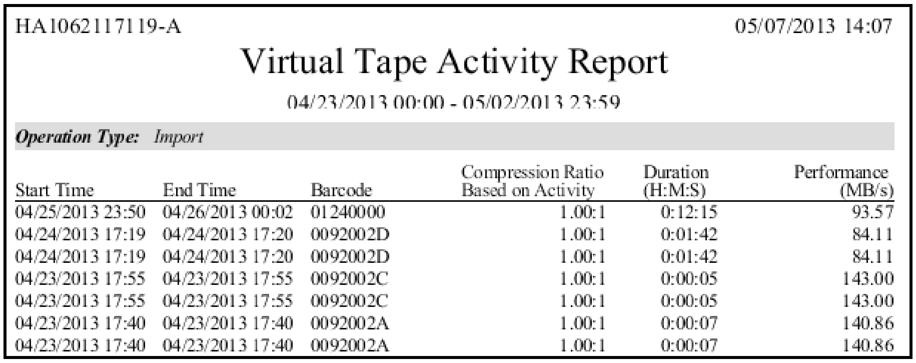

HA1062117119-A 05/07/2013 14:07

Virtual Tape Activity Report

04/23/2013 00:00 - 05/02/2013 23:59

Operation Type: *Import*

Start Time	End Time	Barcode	Compression Ratio Based on Activity	Duration (H:M:S)	Performance (MB/s)
04/25/2013 23:50	04/26/2013 00:02	01240000	1.00:1	0:12:15	93.57
04/24/2013 17:19	04/24/2013 17:20	0092002D	1.00:1	0:01:42	84.11
04/24/2013 17:19	04/24/2013 17:20	0092002D	1.00:1	0:01:42	84.11
04/23/2013 17:55	04/23/2013 17:55	0092002C	1.00:1	0:00:05	143.00
04/23/2013 17:55	04/23/2013 17:55	0092002C	1.00:1	0:00:05	143.00
04/23/2013 17:40	04/23/2013 17:40	0092002A	1.00:1	0:00:07	140.86
04/23/2013 17:40	04/23/2013 17:40	0092002A	1.00:1	0:00:07	140.86

Virtual Tape Information

This report displays the current status of all virtual tapes. The wizard lets you select the virtual tape libraries and deduplication policies that you want to include; all libraries policies are selected by default. 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.

Because of the amount of information available for virtual tapes, multiple sub-reports, referred to as *views*, present information related to a single VTL feature. In addition to the *Overall Summary*, you can choose the following: Deduplication View, Tape Caching View, Replica Resources View, Vault View, and Detailed Tape View.

This report does not include tape segments used for some types of internal tracking and processing when calculating used or allocated space.

Overall Summary View

This view is selected by default. For each tape, the report displays the bar code; amount of data written; whether or not the tape is full; tape caching status; whether migration, deduplication, and replication are required; whether remote export has been configured; tape location; and the deduplication policy (if any) to which the tape belongs.

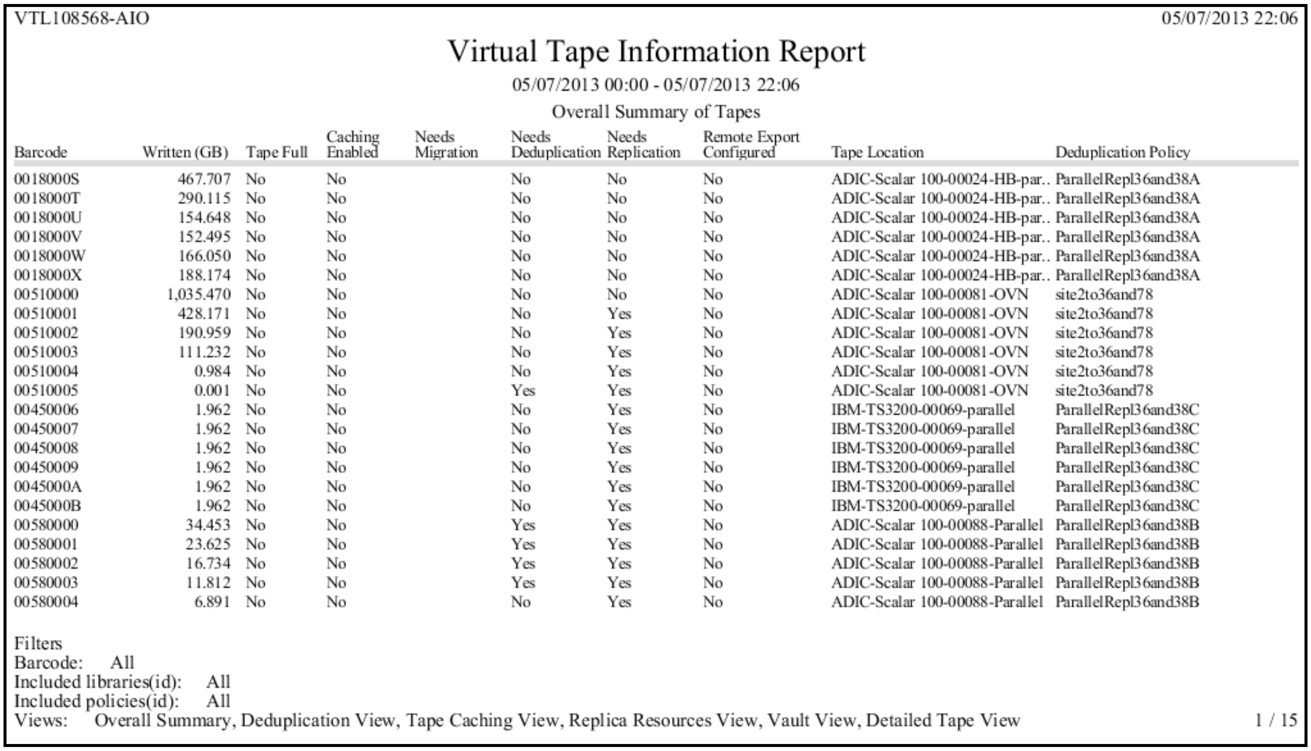

VTL108568-AIO 05/07/2013 22:06

Virtual Tape Information Report

05/07/2013 00:00 - 05/07/2013 22:06

Overall Summary of Tapes

Barcode	Written (GB)	Tape Full	Caching Enabled	Needs Migration	Needs Deduplication	Needs Replication	Remote Export Configured	Tape Location	Deduplication Policy
0018000S	467.707	No	No		No	No	No	ADIC-Scalar 100-00024-HB-par..	ParallelRepl36and38A
0018000T	290.115	No	No		No	No	No	ADIC-Scalar 100-00024-HB-par..	ParallelRepl36and38A
0018000U	154.648	No	No		No	No	No	ADIC-Scalar 100-00024-HB-par..	ParallelRepl36and38A
0018000V	152.495	No	No		No	No	No	ADIC-Scalar 100-00024-HB-par..	ParallelRepl36and38A
0018000W	166.050	No	No		No	No	No	ADIC-Scalar 100-00024-HB-par..	ParallelRepl36and38A
0018000X	188.174	No	No		No	No	No	ADIC-Scalar 100-00024-HB-par..	ParallelRepl36and38A
00510000	1,035.470	No	No		No	No	No	ADIC-Scalar 100-00081-OVN	site2to36and78
00510001	428.171	No	No		No	Yes	No	ADIC-Scalar 100-00081-OVN	site2to36and78
00510002	190.959	No	No		No	Yes	No	ADIC-Scalar 100-00081-OVN	site2to36and78
00510003	111.232	No	No		No	Yes	No	ADIC-Scalar 100-00081-OVN	site2to36and78
00510004	0.984	No	No		No	Yes	No	ADIC-Scalar 100-00081-OVN	site2to36and78
00510005	0.001	No	No		Yes	Yes	No	ADIC-Scalar 100-00081-OVN	site2to36and78
00450006	1.962	No	No		No	Yes	No	IBM-TS3200-00069-parallel	ParallelRepl36and38C
00450007	1.962	No	No		No	Yes	No	IBM-TS3200-00069-parallel	ParallelRepl36and38C
00450008	1.962	No	No		No	Yes	No	IBM-TS3200-00069-parallel	ParallelRepl36and38C
00450009	1.962	No	No		No	Yes	No	IBM-TS3200-00069-parallel	ParallelRepl36and38C
0045000A	1.962	No	No		No	Yes	No	IBM-TS3200-00069-parallel	ParallelRepl36and38C
0045000B	1.962	No	No		No	Yes	No	IBM-TS3200-00069-parallel	ParallelRepl36and38C
00580000	34.453	No	No		Yes	Yes	No	ADIC-Scalar 100-00088-Parallel	ParallelRepl36and38B
00580001	23.625	No	No		Yes	Yes	No	ADIC-Scalar 100-00088-Parallel	ParallelRepl36and38B
00580002	16.734	No	No		Yes	Yes	No	ADIC-Scalar 100-00088-Parallel	ParallelRepl36and38B
00580003	11.812	No	No		Yes	Yes	No	ADIC-Scalar 100-00088-Parallel	ParallelRepl36and38B
00580004	6.891	No	No		No	Yes	No	ADIC-Scalar 100-00088-Parallel	ParallelRepl36and38B

Filters
Barcode: All
Included libraries(id): All
Included policies(id): All
Views: Overall Summary, Deduplication View, Tape Caching View, Replica Resources View, Vault View, Detailed Tape View 1 / 15

Deduplication View

For each tape, the report displays the bar code; whether deduplication or replication is required; amount of data written; start/finish time for most recent deduplication job, amount of data processed, unique data, VIT size, deduplication ratio, throughput, and bandwidth utilization.

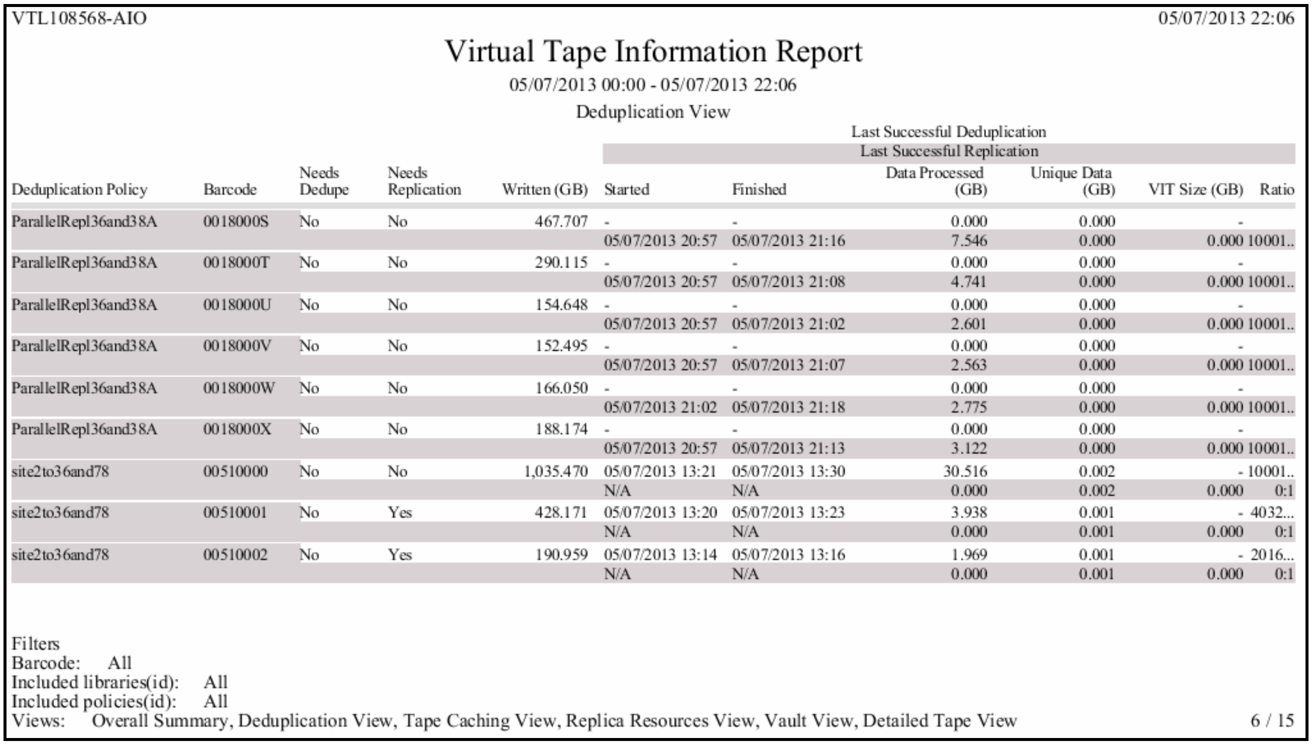

VTL108568-AIO 05/07/2013 22:06

Virtual Tape Information Report

05/07/2013 00:00 - 05/07/2013 22:06

Deduplication View

					Last Successful Deduplication Last Successful Replication					
Deduplication Policy	Barcode	Needs Dedupe	Needs Replication	Written (GB)	Started	Finished	Data Processed (GB)	Unique Data (GB)	VIT Size (GB)	Ratio
ParallelRepl36and38A	0018000S	No	No	467.707	-	-	0.000	0.000	-	
					05/07/2013 20:57	05/07/2013 21:16	7.546	0.000	0.000	10001..
ParallelRepl36and38A	0018000T	No	No	290.115	-	-	0.000	0.000	-	
					05/07/2013 20:57	05/07/2013 21:08	4.741	0.000	0.000	10001..
ParallelRepl36and38A	0018000U	No	No	154.648	-	-	0.000	0.000	-	
					05/07/2013 20:57	05/07/2013 21:02	2.601	0.000	0.000	10001..
ParallelRepl36and38A	0018000V	No	No	152.495	-	-	0.000	0.000	-	
					05/07/2013 20:57	05/07/2013 21:07	2.563	0.000	0.000	10001..
ParallelRepl36and38A	0018000W	No	No	166.050	-	-	0.000	0.000	-	
					05/07/2013 21:02	05/07/2013 21:18	2.775	0.000	0.000	10001..
ParallelRepl36and38A	0018000X	No	No	188.174	-	-	0.000	0.000	-	
					05/07/2013 20:57	05/07/2013 21:13	3.122	0.000	0.000	10001..
site2to36and78	00510000	No	No	1,035.470	05/07/2013 13:21	05/07/2013 13:30	30.516	0.002	-	10001..
					N/A	N/A	0.000	0.002	0.000	0:1
site2to36and78	00510001	No	Yes	428.171	05/07/2013 13:20	05/07/2013 13:23	3.938	0.001	-	4032...
					N/A	N/A	0.000	0.001	0.000	0:1
site2to36and78	00510002	No	Yes	190.959	05/07/2013 13:14	05/07/2013 13:16	1.969	0.001	-	2016...
					N/A	N/A	0.000	0.001	0.000	0:1

Filters
Barcode: All
Included libraries(id): All
Included policies(id): All
Views: Overall Summary, Deduplication View, Tape Caching View, Replica Resources View, Vault View, Detailed Tape View 6 / 15

Tape Caching View

This view displays information related to tape caching. For each tape, the report displays the library to which it belongs, its bar code, amount of data written, whether or not the tape is full, whether migration or deduplication is required, and the deduplication policy (if any) to which it belongs.

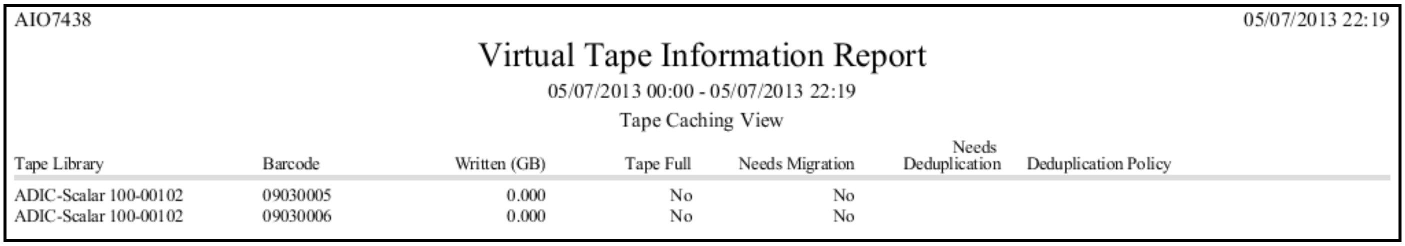

AIO7438 05/07/2013 22:19

Virtual Tape Information Report

05/07/2013 00:00 - 05/07/2013 22:19

Tape Caching View

Tape Library	Barcode	Written (GB)	Tape Full	Needs Migration	Needs Deduplication	Deduplication Policy
ADIC-Scalar 100-00102	09030005	0.000	No	No		
ADIC-Scalar 100-00102	09030006	0.000	No	No		

Vault View

This view is similar to the *Overall Summary View*, with the exception that instead of *Tape Location*, this view identifies the library to which each tape belongs. For each tape, the report displays the bar code, amount of data written, whether or not the tape is full, whether tape caching is configured, whether migration, deduplication, or replication is required, whether remote export is configured, the parent library, and the deduplication policy (if any) to which it belongs.

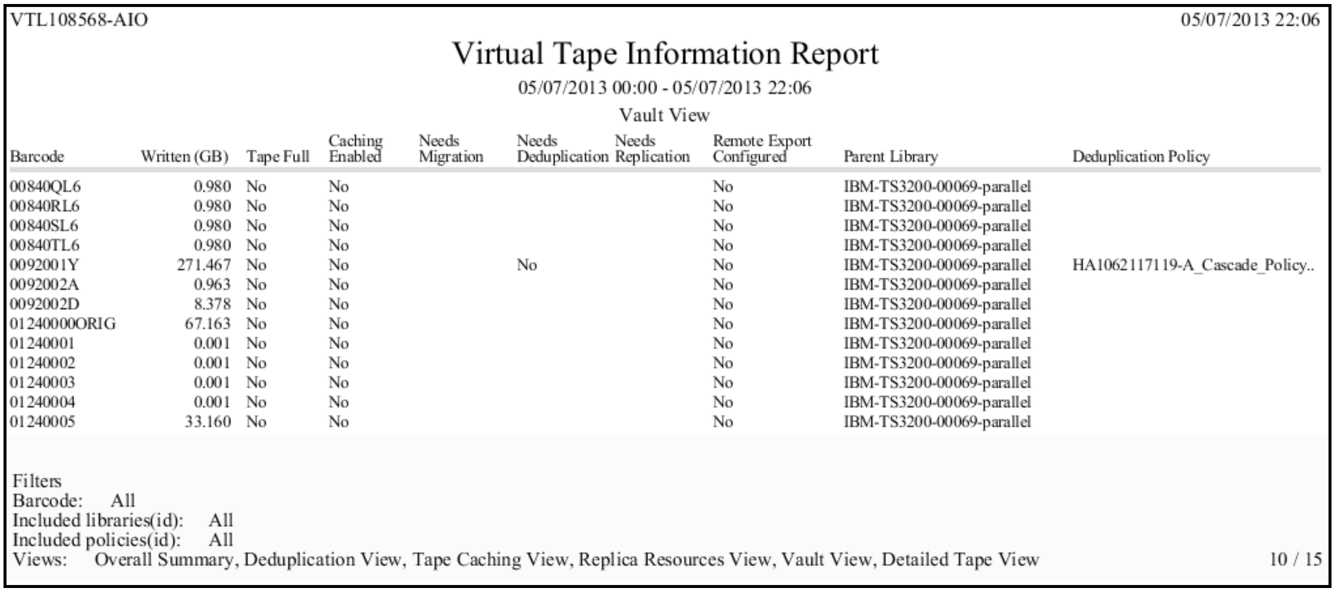

VTL108568-AIO 05/07/2013 22:06

Virtual Tape Information Report

05/07/2013 00:00 - 05/07/2013 22:06

Vault View

Barcode	Written (GB)	Tape Full	Caching Enabled	Needs Migration	Needs Deduplication	Needs Replication	Remote Export Configured	Parent Library	Deduplication Policy
00840QL6	0.980	No	No				No	IBM-TS3200-00069-parallel	
00840RL6	0.980	No	No				No	IBM-TS3200-00069-parallel	
00840SL6	0.980	No	No				No	IBM-TS3200-00069-parallel	
00840TL6	0.980	No	No				No	IBM-TS3200-00069-parallel	
0092001Y	271.467	No	No		No		No	IBM-TS3200-00069-parallel	HA1062117119-A_Cascade_Policy..
0092002A	0.963	No	No				No	IBM-TS3200-00069-parallel	
0092002D	8.378	No	No				No	IBM-TS3200-00069-parallel	
01240000ORIG	67.163	No	No				No	IBM-TS3200-00069-parallel	
01240001	0.001	No	No				No	IBM-TS3200-00069-parallel	
01240002	0.001	No	No				No	IBM-TS3200-00069-parallel	
01240003	0.001	No	No				No	IBM-TS3200-00069-parallel	
01240004	0.001	No	No				No	IBM-TS3200-00069-parallel	
01240005	33.160	No	No				No	IBM-TS3200-00069-parallel	

Filters
Barcode: All
Included libraries(id): All
Included policies(id): All
Views: Overall Summary, Deduplication View, Tape Caching View, Replica Resources View, Vault View, Detailed Tape View 10 / 15

Replica Resources View

This view presents information for tapes that are displayed when you select the *Replica Resources* object in the console. For each tape, the report displays the bar code, allocation size, whether or not the tape is full, the source VTL server, the tape ID, whether the tape is a FVIT or LVIT, and whether remote export has been configured.

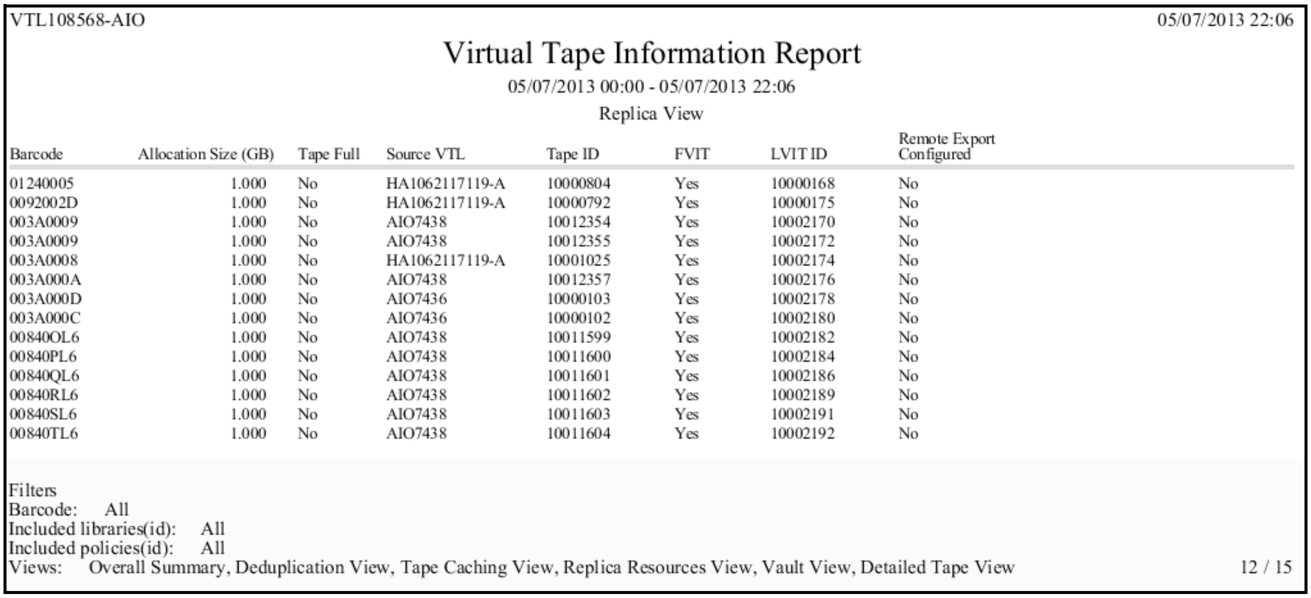

VTL108568-AIO 05/07/2013 22:06

Virtual Tape Information Report

05/07/2013 00:00 - 05/07/2013 22:06

Replica View

Barcode	Allocation Size (GB)	Tape Full	Source VTL	Tape ID	FVIT	LVIT ID	Remote Export Configured
01240005	1.000	No	HA1062117119-A	10000804	Yes	10000168	No
0092002D	1.000	No	HA1062117119-A	10000792	Yes	10000175	No
003A0009	1.000	No	AIO7438	10012354	Yes	10002170	No
003A0009	1.000	No	AIO7438	10012355	Yes	10002172	No
003A0008	1.000	No	HA1062117119-A	10001025	Yes	10002174	No
003A000A	1.000	No	AIO7438	10012357	Yes	10002176	No
003A000D	1.000	No	AIO7436	10000103	Yes	10002178	No
003A000C	1.000	No	AIO7436	10000102	Yes	10002180	No
00840OL6	1.000	No	AIO7438	10011599	Yes	10002182	No
00840PL6	1.000	No	AIO7438	10011600	Yes	10002184	No
00840QL6	1.000	No	AIO7438	10011601	Yes	10002186	No
00840RL6	1.000	No	AIO7438	10011602	Yes	10002189	No
00840SL6	1.000	No	AIO7438	10011603	Yes	10002191	No
00840TL6	1.000	No	AIO7438	10011604	Yes	10002192	No

Filters
Barcode: All
Included libraries(id): All
Included policies(id): All
Views: Overall Summary, Deduplication View, Tape Caching View, Replica Resources View, Vault View, Detailed Tape View 12 / 15

Tape Detail View

This view presents information for tapes in a specific library, including bar codes, tape ID, location (library), whether or not the tape is full, current allocation, maximum capacity, used size, amount of data written, and number of segments used.

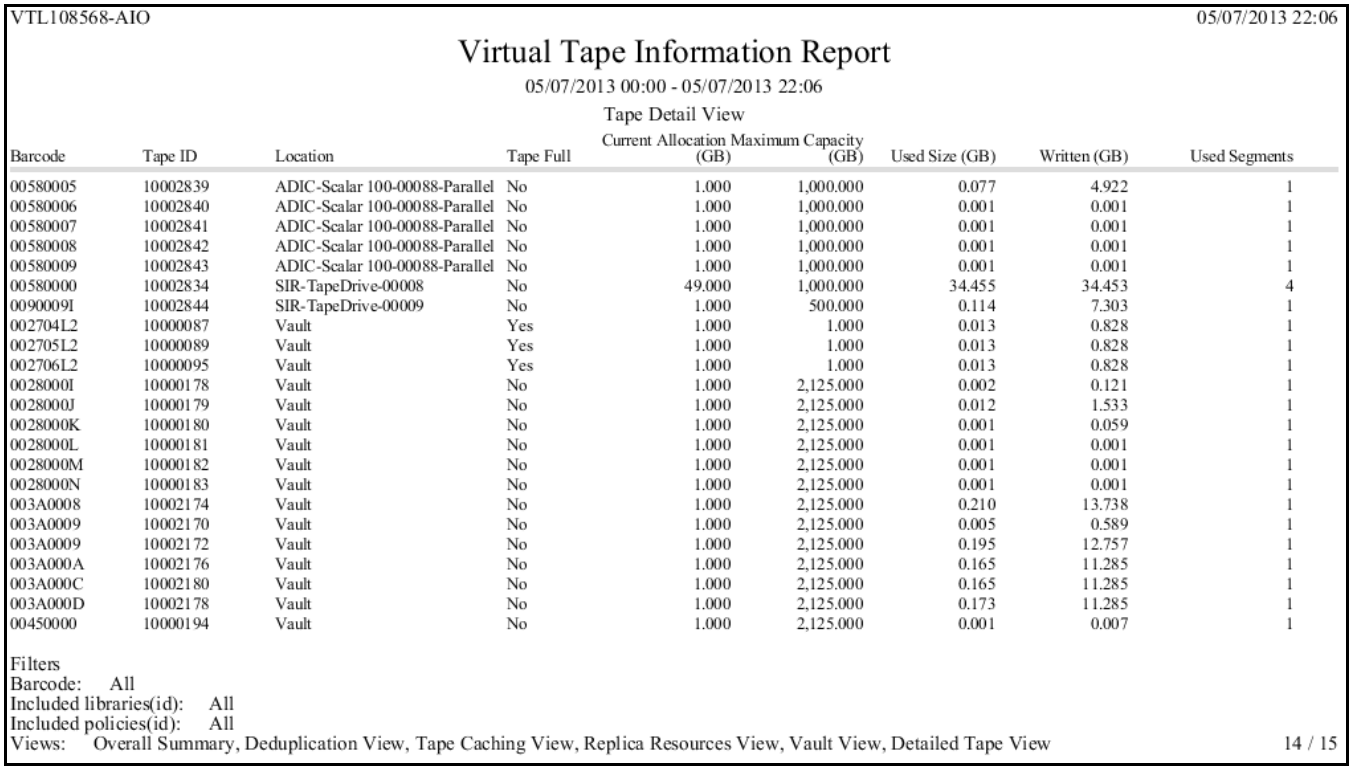

VTL108568-AIO 05/07/2013 22:06

Virtual Tape Information Report

05/07/2013 00:00 - 05/07/2013 22:06

Tape Detail View

				Current Allocation	Maximum Capacity			
Barcode	Tape ID	Location	Tape Full	(GB)	(GB)	Used Size (GB)	Written (GB)	Used Segments
00580005	10002839	ADIC-Scalar 100-00088-Parallel	No	1.000	1,000.000	0.077	4.922	1
00580006	10002840	ADIC-Scalar 100-00088-Parallel	No	1.000	1,000.000	0.001	0.001	1
00580007	10002841	ADIC-Scalar 100-00088-Parallel	No	1.000	1,000.000	0.001	0.001	1
00580008	10002842	ADIC-Scalar 100-00088-Parallel	No	1.000	1,000.000	0.001	0.001	1
00580009	10002843	ADIC-Scalar 100-00088-Parallel	No	1.000	1,000.000	0.001	0.001	1
00580000	10002834	SIR-TapeDrive-00008	No	49.000	1,000.000	34.455	34.453	4
0090009I	10002844	SIR-TapeDrive-00009	No	1.000	500.000	0.114	7.303	1
002704L2	10000087	Vault	Yes	1.000	1.000	0.013	0.828	1
002705L2	10000089	Vault	Yes	1.000	1.000	0.013	0.828	1
002706L2	10000095	Vault	Yes	1.000	1.000	0.013	0.828	1
0028000I	10000178	Vault	No	1.000	2,125.000	0.002	0.121	1
0028000J	10000179	Vault	No	1.000	2,125.000	0.012	1.533	1
0028000K	10000180	Vault	No	1.000	2,125.000	0.001	0.059	1
0028000L	10000181	Vault	No	1.000	2,125.000	0.001	0.001	1
0028000M	10000182	Vault	No	1.000	2,125.000	0.001	0.001	1
0028000N	10000183	Vault	No	1.000	2,125.000	0.001	0.001	1
003A0008	10002174	Vault	No	1.000	2,125.000	0.210	13.738	1
003A0009	10002170	Vault	No	1.000	2,125.000	0.005	0.589	1
003A0009	10002172	Vault	No	1.000	2,125.000	0.195	12.757	1
003A000A	10002176	Vault	No	1.000	2,125.000	0.165	11.285	1
003A000C	10002180	Vault	No	1.000	2,125.000	0.165	11.285	1
003A000D	10002178	Vault	No	1.000	2,125.000	0.173	11.285	1
00450000	10000194	Vault	No	1.000	2,125.000	0.001	0.007	1

Filters
Barcode: All
Included libraries(id): All
Included policies(id): All
Views: Overall Summary, Deduplication View, Tape Caching View, Replica Resources View, Vault View, Detailed Tape View 14 / 15

NAS Statistics Summary

This report displays summary information and settings for NAS replication sessions. You can include data on replication clients and/or replication servers.

This report can also be configured as a scheduled report.

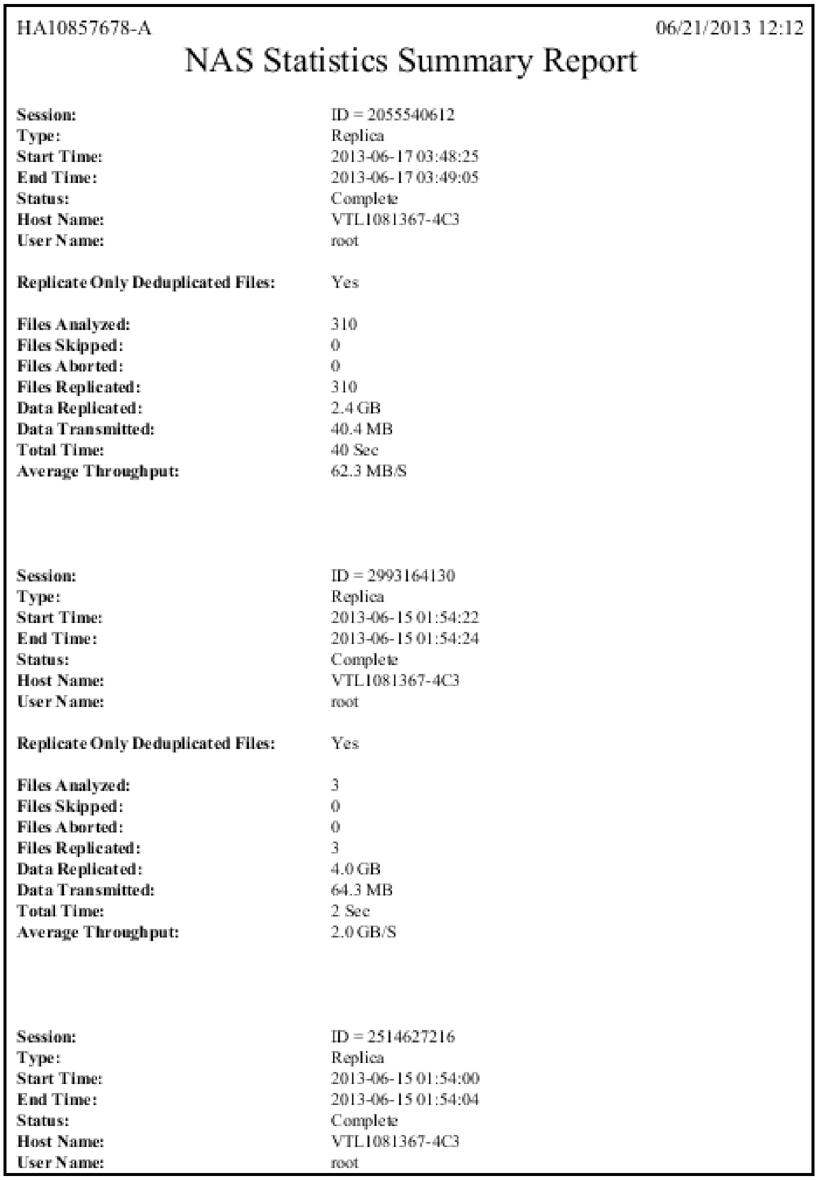

HA10857678-A 06/21/2013 12:12

NAS Statistics Summary Report

```
Session:                              ID = 2055540612
Type:                                 Replica
Start Time:                           2013-06-17 03:48:25
End Time:                             2013-06-17 03:49:05
Status:                               Complete
Host Name:                            VTL1081367-4C3
User Name:                            root

Replicate Only Deduplicated Files:    Yes

Files Analyzed:                       310
Files Skipped:                        0
Files Aborted:                        0
Files Replicated:                     310
Data Replicated:                      2.4 GB
Data Transmitted:                     40.4 MB
Total Time:                           40 Sec
Average Throughput:                   62.3 MB/S



Session:                              ID = 2993164130
Type:                                 Replica
Start Time:                           2013-06-15 01:54:22
End Time:                             2013-06-15 01:54:24
Status:                               Complete
Host Name:                            VTL1081367-4C3
User Name:                            root

Replicate Only Deduplicated Files:    Yes

Files Analyzed:                       3
Files Skipped:                        0
Files Aborted:                        0
Files Replicated:                     3
Data Replicated:                      4.0 GB
Data Transmitted:                     64.3 MB
Total Time:                           2 Sec
Average Throughput:                   2.0 GB/S



Session:                              ID = 2514627216
Type:                                 Replica
Start Time:                           2013-06-15 01:54:00
End Time:                             2013-06-15 01:54:04
Status:                               Complete
Host Name:                            VTL1081367-4C3
User Name:                            root
```

Usage reports

Deduplication - Tape Usage

This status report provides statistics for tapes in all deduplication policies. For each policy, the report lists tape names, barcodes, VIT status (yes, mixed, or no), capacity of each tape, amount of data written, and the deduplication ratio. Other information includes:

- New - Refers to data on tapes that have not yet been processed since the previous deduplication job.
- In SIR - Refers to the amount of processed data stored in the deduplication repository.
- Unique Data - Does not reflect any data previously stored in the deduplication repository; does not reflect the total amount of data stored in the deduplication repository.
- CB - Indicates if the tape contains at least one backup session in which data was compressed by backup software. If Symantec NetBackup or IBM Tivoli Storage Manager backup applications are being used, the report includes information about compressed data. For example, this flags shows *Yes* for tapes holding OST images that have been compressed by NetBackup.
- Client Name - Indicates the source of data on the tape. If the list of clients is too long to be displayed, you can export the report to any available format in order to see the complete list.

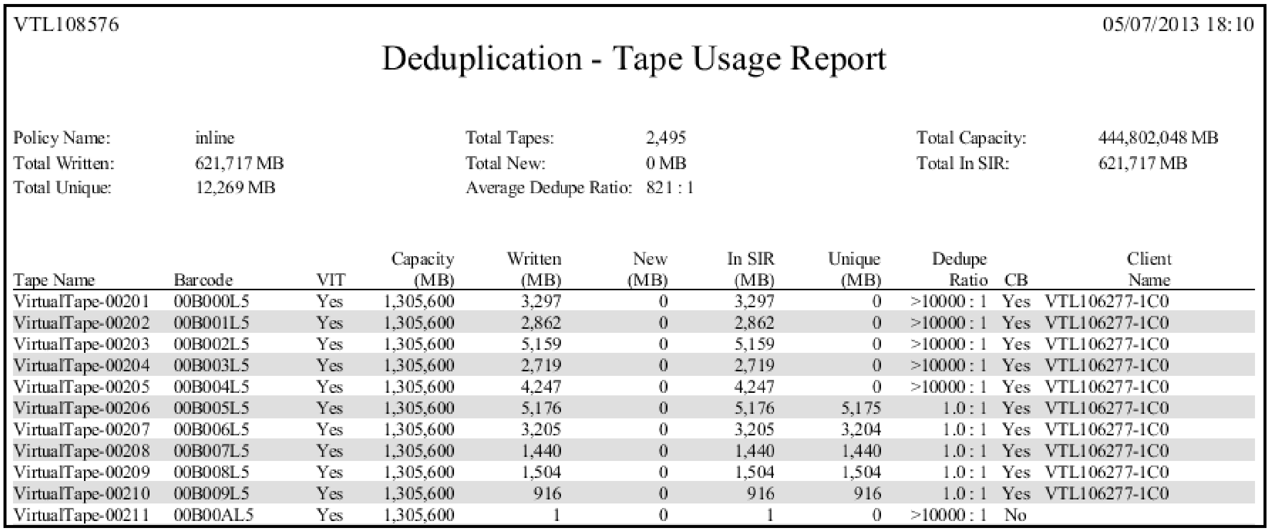

VTL108576 05/07/2013 18:10

Deduplication - Tape Usage Report

Policy Name:	inline	Total Tapes:	2,495	Total Capacity:	444,802,048 MB
Total Written:	621,717 MB	Total New:	0 MB	Total In SIR:	621,717 MB
Total Unique:	12,269 MB	Average Dedupe Ratio:	821 : 1		

			Capacity	Written	New	In SIR	Unique	Dedupe		Client
Tape Name	Barcode	VIT	(MB)	(MB)	(MB)	(MB)	(MB)	Ratio	CB	Name
VirtualTape-00201	00B000L5	Yes	1,305,600	3,297	0	3,297	0	>10000 : 1	Yes	VTL106277-1C0
VirtualTape-00202	00B001L5	Yes	1,305,600	2,862	0	2,862	0	>10000 : 1	Yes	VTL106277-1C0
VirtualTape-00203	00B002L5	Yes	1,305,600	5,159	0	5,159	0	>10000 : 1	Yes	VTL106277-1C0
VirtualTape-00204	00B003L5	Yes	1,305,600	2,719	0	2,719	0	>10000 : 1	Yes	VTL106277-1C0
VirtualTape-00205	00B004L5	Yes	1,305,600	4,247	0	4,247	0	>10000 : 1	Yes	VTL106277-1C0
VirtualTape-00206	00B005L5	Yes	1,305,600	5,176	0	5,176	5,175	1.0 : 1	Yes	VTL106277-1C0
VirtualTape-00207	00B006L5	Yes	1,305,600	3,205	0	3,205	3,204	1.0 : 1	Yes	VTL106277-1C0
VirtualTape-00208	00B007L5	Yes	1,305,600	1,440	0	1,440	1,440	1.0 : 1	Yes	VTL106277-1C0
VirtualTape-00209	00B008L5	Yes	1,305,600	1,504	0	1,504	1,504	1.0 : 1	Yes	VTL106277-1C0
VirtualTape-00210	00B009L5	Yes	1,305,600	916	0	916	916	1.0 : 1	Yes	VTL106277-1C0
VirtualTape-00211	00B00AL5	Yes	1,305,600	1	0	1	0	>10000 : 1	No	

This report is available only as a one-time report.

Deduplication Repository - Memory and Space Usage

The *Dashboard Summary* tab (displayed for a deduplication server) and the *Deduplication Repository* tab (displayed for a VTL-S server) display information about repository capacity and the usage of index cache, data disks, folder disks, and index disks. This status report is an alternate way to view this information.

An association between the VTL server and a deduplication server must exist in order for this report to be available. (For VTL-S, this association is created automatically as part of the process of enabling deduplication.) If a VTL server has been disassociated from a deduplication server, reports are unaffected - you can still view, export, or delete them.

The report range can be specified as the current server date (the default), yesterday, the past 7 days, the past 30 days, the past year, or a specific date range. In addition, you can choose the amount of time (the interval) represented by displayed bars/ table rows in report results to be an hour, a day, a week, a month, or a quarter (depending on the selected report dates).

Note: If data cannot be collected for an interval, the report will derive the value based on the next good value.

In the bar charts, each bar represents the total amount of repository space or index cache used by all nodes in the deduplication cluster as of the beginning of the interval.

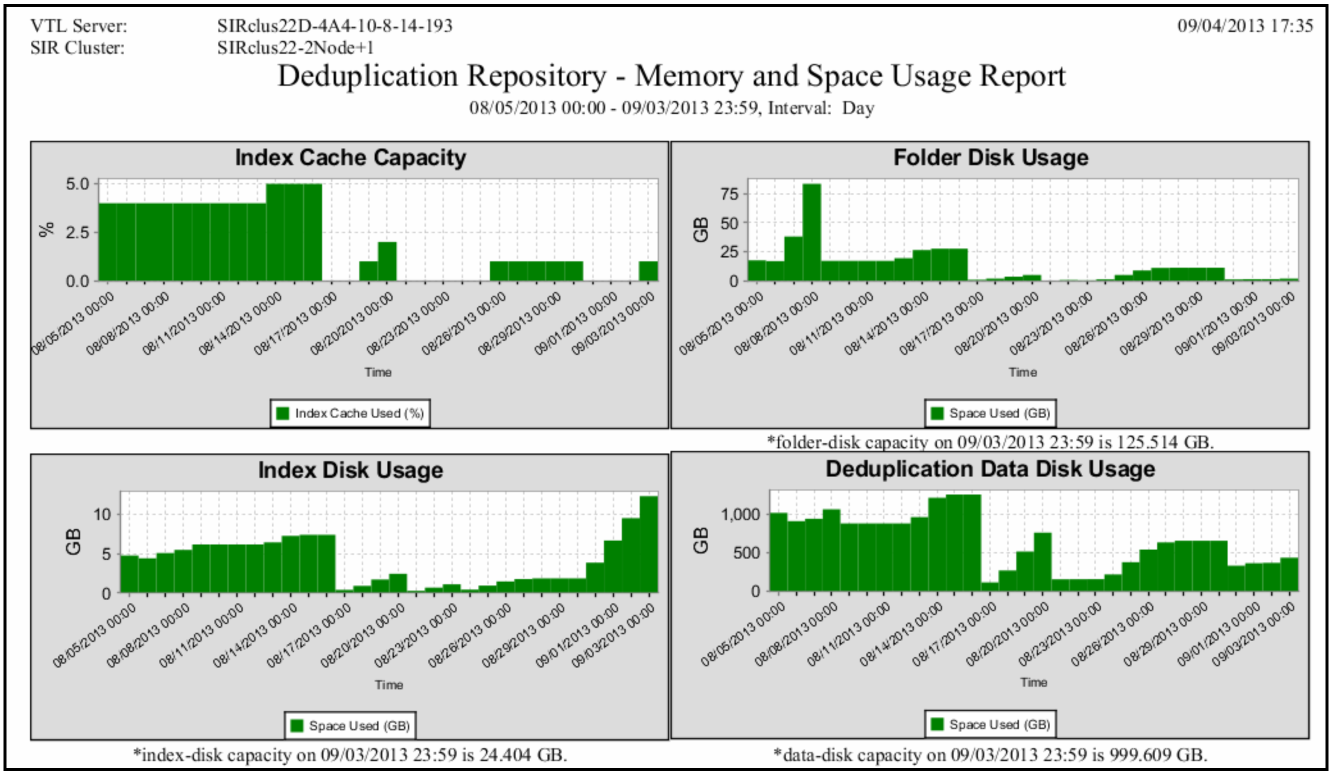

VTL Server: SIRclus22D-4A4-10-8-14-193
SIR Cluster: SIRclus22-2Node+1

09/04/2013 17:35

Deduplication Repository - Memory and Space Usage Report

08/05/2013 00:00 - 09/03/2013 23:59, Interval: Day

*folder-disk capacity on 09/03/2013 23:59 is 125.514 GB.

*index-disk capacity on 09/03/2013 23:59 is 24.404 GB.

*data-disk capacity on 09/03/2013 23:59 is 999.609 GB.

The second page displays the amount of repository space or memory used by the cluster as of the beginning of each interval.

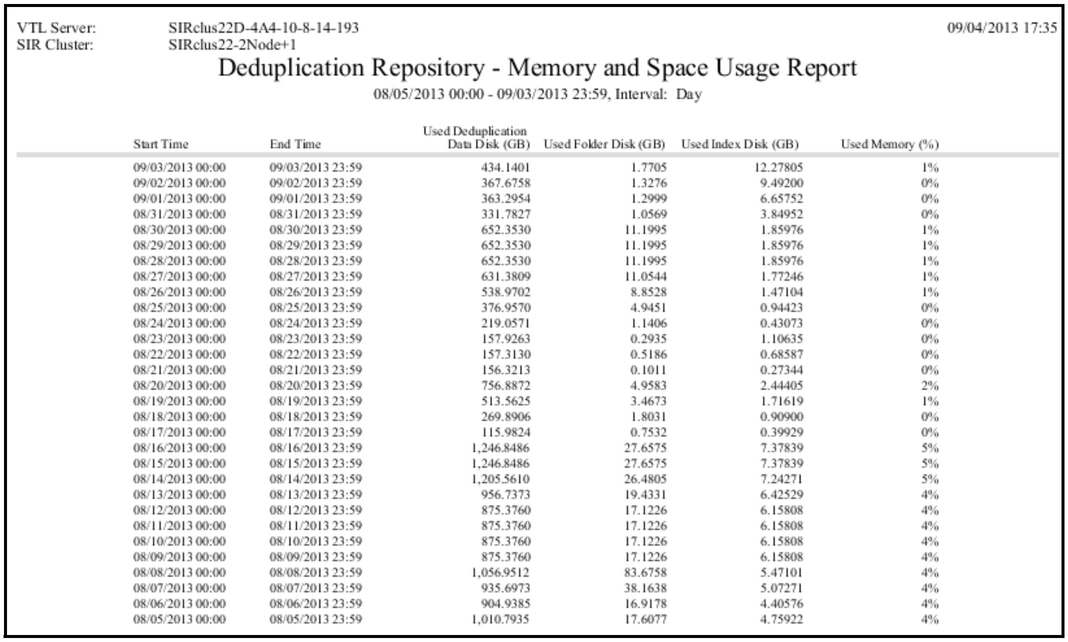

VTL Server: SIRclus22D-4A4-10-8-14-193 09/04/2013 17:35
SIR Cluster: SIRclus22-2Node+1

Deduplication Repository - Memory and Space Usage Report

08/05/2013 00:00 - 09/03/2013 23:59, Interval: Day

Start Time	End Time	Used Deduplication Data Disk (GB)	Used Folder Disk (GB)	Used Index Disk (GB)	Used Memory (%)
09/03/2013 00:00	09/03/2013 23:59	434.1401	1.7705	12.27805	1%
09/02/2013 00:00	09/02/2013 23:59	367.6758	1.3276	9.49200	0%
09/01/2013 00:00	09/01/2013 23:59	363.2954	1.2999	6.65752	0%
08/31/2013 00:00	08/31/2013 23:59	331.7827	1.0569	3.84952	0%
08/30/2013 00:00	08/30/2013 23:59	652.3530	11.1995	1.85976	1%
08/29/2013 00:00	08/29/2013 23:59	652.3530	11.1995	1.85976	1%
08/28/2013 00:00	08/28/2013 23:59	652.3530	11.1995	1.85976	1%
08/27/2013 00:00	08/27/2013 23:59	631.3809	11.0544	1.77246	1%
08/26/2013 00:00	08/26/2013 23:59	538.9702	8.8528	1.47104	1%
08/25/2013 00:00	08/25/2013 23:59	376.9570	4.9451	0.94423	0%
08/24/2013 00:00	08/24/2013 23:59	219.0571	1.1406	0.43073	0%
08/23/2013 00:00	08/23/2013 23:59	157.9263	0.2935	1.10635	0%
08/22/2013 00:00	08/22/2013 23:59	157.3130	0.5186	0.68587	0%
08/21/2013 00:00	08/21/2013 23:59	156.3213	0.1011	0.27344	0%
08/20/2013 00:00	08/20/2013 23:59	756.8872	4.9583	2.44405	2%
08/19/2013 00:00	08/19/2013 23:59	513.5625	3.4673	1.71619	1%
08/18/2013 00:00	08/18/2013 23:59	269.8906	1.8031	0.90900	0%
08/17/2013 00:00	08/17/2013 23:59	115.9824	0.7532	0.39929	0%
08/16/2013 00:00	08/16/2013 23:59	1,246.8486	27.6575	7.37839	5%
08/15/2013 00:00	08/15/2013 23:59	1,246.8486	27.6575	7.37839	5%
08/14/2013 00:00	08/14/2013 23:59	1,205.5610	26.4805	7.24271	5%
08/13/2013 00:00	08/13/2013 23:59	956.7373	19.4331	6.42529	4%
08/12/2013 00:00	08/12/2013 23:59	875.3760	17.1226	6.15808	4%
08/11/2013 00:00	08/11/2013 23:59	875.3760	17.1226	6.15808	4%
08/10/2013 00:00	08/10/2013 23:59	875.3760	17.1226	6.15808	4%
08/09/2013 00:00	08/09/2013 23:59	875.3760	17.1226	6.15808	4%
08/08/2013 00:00	08/08/2013 23:59	1,056.9512	83.6758	5.47101	4%
08/07/2013 00:00	08/07/2013 23:59	935.6973	38.1638	5.07271	4%
08/06/2013 00:00	08/06/2013 23:59	904.9385	16.9178	4.40576	4%
08/05/2013 00:00	08/05/2013 23:59	1,010.7935	17.6077	4.75922	4%

Physical Tape Usage

This status report displays information about how each physical tape in the physical tape database is mapped to a virtual tape, including barcode, serial number, creation date, the number of virtual tapes configured from the physical tape, and the total amount of data on the tape. The date/time of the most recent update to the tape is also displayed. Tapes without data are not included in report results.

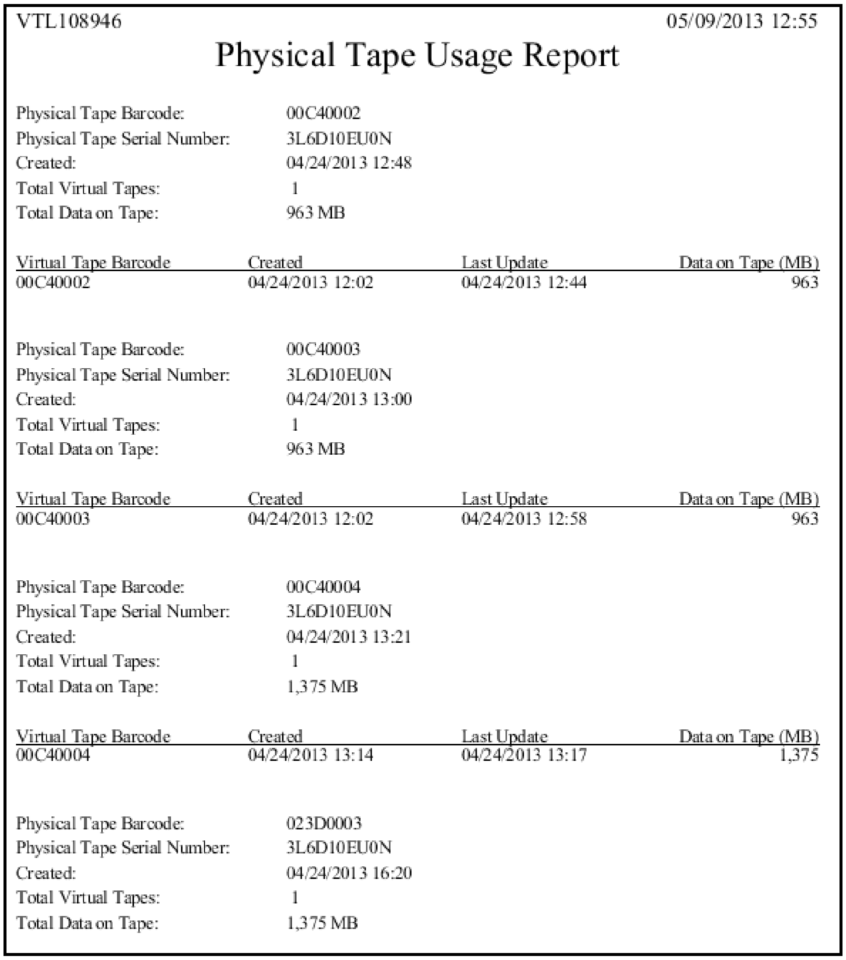

VTL108946 05/09/2013 12:55

Physical Tape Usage Report

Physical Tape Barcode:	00C40002
Physical Tape Serial Number:	3L6D10EU0N
Created:	04/24/2013 12:48
Total Virtual Tapes:	1
Total Data on Tape:	963 MB

Virtual Tape Barcode	Created	Last Update	Data on Tape (MB)
00C40002	04/24/2013 12:02	04/24/2013 12:44	963

Physical Tape Barcode:	00C40003
Physical Tape Serial Number:	3L6D10EU0N
Created:	04/24/2013 13:00
Total Virtual Tapes:	1
Total Data on Tape:	963 MB

Virtual Tape Barcode	Created	Last Update	Data on Tape (MB)
00C40003	04/24/2013 12:02	04/24/2013 12:58	963

Physical Tape Barcode:	00C40004
Physical Tape Serial Number:	3L6D10EU0N
Created:	04/24/2013 13:21
Total Virtual Tapes:	1
Total Data on Tape:	1,375 MB

Virtual Tape Barcode	Created	Last Update	Data on Tape (MB)
00C40004	04/24/2013 13:14	04/24/2013 13:17	1,375

Physical Tape Barcode:	023D0003
Physical Tape Serial Number:	3L6D10EU0N
Created:	04/24/2013 16:20
Total Virtual Tapes:	1
Total Data on Tape:	1,375 MB

Disk Space Usage History

This status report shows the peak amount of disk space available/used during the specified date range. Available intervals are based on the range: for single days, disk usage is shown for each 60-minute period; for a week, usage is shown for each four-hour period; for a 30-day period (as in the example below), usage is shown for each day. Categories showing an asterisk (*) indicate that the displayed value is based on the first available data in the interval.

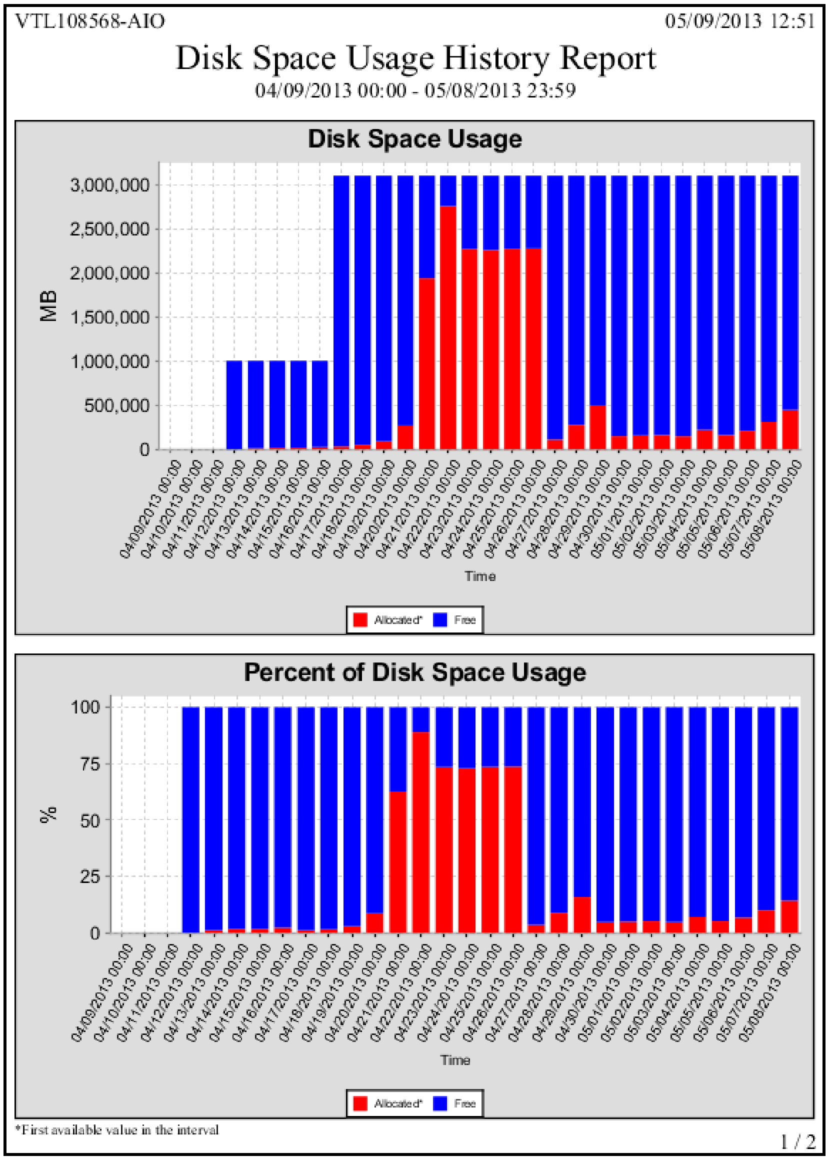

VTL108568-AIO 05/09/2013 12:51

Disk Space Usage History Report

04/09/2013 00:00 - 05/08/2013 23:59

*First available value in the interval

1 / 2

The results table includes a row of data for each interval in the graph: the start and stop time, the total capacity for all disks and the total allocated/free capacity expressed as a percentage of total capacity.

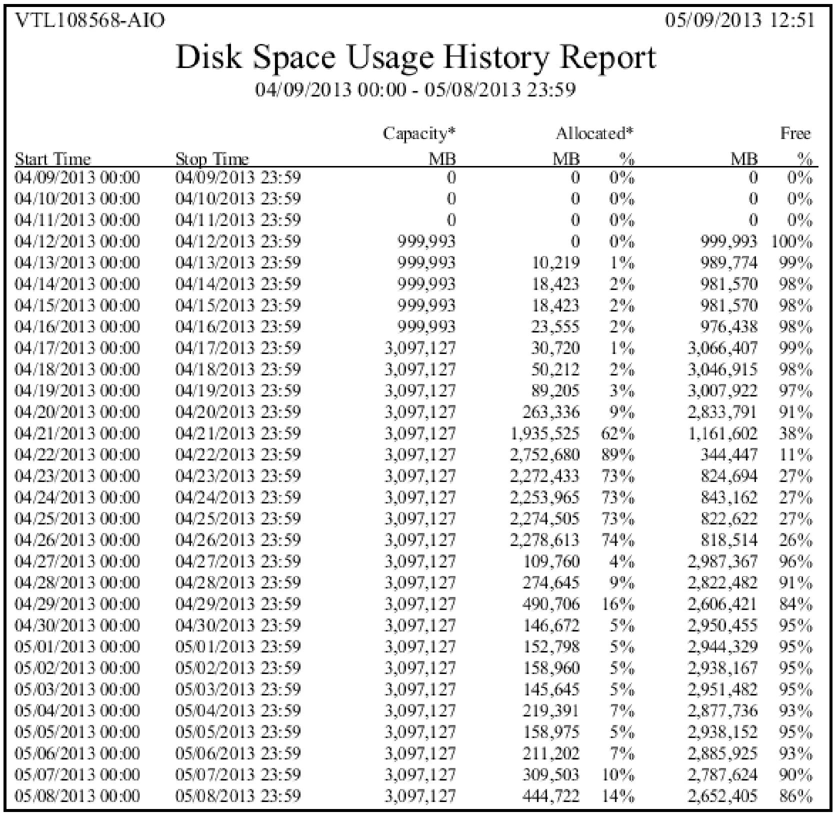

VTL108568-AIO 05/09/2013 12:51

Disk Space Usage History Report

04/09/2013 00:00 - 05/08/2013 23:59

		Capacity*	Allocated*		Free	
Start Time	Stop Time	MB	MB	%	MB	%
04/09/2013 00:00	04/09/2013 23:59	0	0	0%	0	0%
04/10/2013 00:00	04/10/2013 23:59	0	0	0%	0	0%
04/11/2013 00:00	04/11/2013 23:59	0	0	0%	0	0%
04/12/2013 00:00	04/12/2013 23:59	999,993	0	0%	999,993	100%
04/13/2013 00:00	04/13/2013 23:59	999,993	10,219	1%	989,774	99%
04/14/2013 00:00	04/14/2013 23:59	999,993	18,423	2%	981,570	98%
04/15/2013 00:00	04/15/2013 23:59	999,993	18,423	2%	981,570	98%
04/16/2013 00:00	04/16/2013 23:59	999,993	23,555	2%	976,438	98%
04/17/2013 00:00	04/17/2013 23:59	3,097,127	30,720	1%	3,066,407	99%
04/18/2013 00:00	04/18/2013 23:59	3,097,127	50,212	2%	3,046,915	98%
04/19/2013 00:00	04/19/2013 23:59	3,097,127	89,205	3%	3,007,922	97%
04/20/2013 00:00	04/20/2013 23:59	3,097,127	263,336	9%	2,833,791	91%
04/21/2013 00:00	04/21/2013 23:59	3,097,127	1,935,525	62%	1,161,602	38%
04/22/2013 00:00	04/22/2013 23:59	3,097,127	2,752,680	89%	344,447	11%
04/23/2013 00:00	04/23/2013 23:59	3,097,127	2,272,433	73%	824,694	27%
04/24/2013 00:00	04/24/2013 23:59	3,097,127	2,253,965	73%	843,162	27%
04/25/2013 00:00	04/25/2013 23:59	3,097,127	2,274,505	73%	822,622	27%
04/26/2013 00:00	04/26/2013 23:59	3,097,127	2,278,613	74%	818,514	26%
04/27/2013 00:00	04/27/2013 23:59	3,097,127	109,760	4%	2,987,367	96%
04/28/2013 00:00	04/28/2013 23:59	3,097,127	274,645	9%	2,822,482	91%
04/29/2013 00:00	04/29/2013 23:59	3,097,127	490,706	16%	2,606,421	84%
04/30/2013 00:00	04/30/2013 23:59	3,097,127	146,672	5%	2,950,455	95%
05/01/2013 00:00	05/01/2013 23:59	3,097,127	152,798	5%	2,944,329	95%
05/02/2013 00:00	05/02/2013 23:59	3,097,127	158,960	5%	2,938,167	95%
05/03/2013 00:00	05/03/2013 23:59	3,097,127	145,645	5%	2,951,482	95%
05/04/2013 00:00	05/04/2013 23:59	3,097,127	219,391	7%	2,877,736	93%
05/05/2013 00:00	05/05/2013 23:59	3,097,127	158,975	5%	2,938,152	95%
05/06/2013 00:00	05/06/2013 23:59	3,097,127	211,202	7%	2,885,925	93%
05/07/2013 00:00	05/07/2013 23:59	3,097,127	309,503	10%	2,787,624	90%
05/08/2013 00:00	05/08/2013 23:59	3,097,127	444,722	14%	2,652,405	86%

LUNs

This status report displays all virtual tapes that are currently allocated on all or specified LUNs. In the wizard, click *Individual* and select the device(s) you want to include in the report, or click *All* to include all devices (device selection is not necessary).

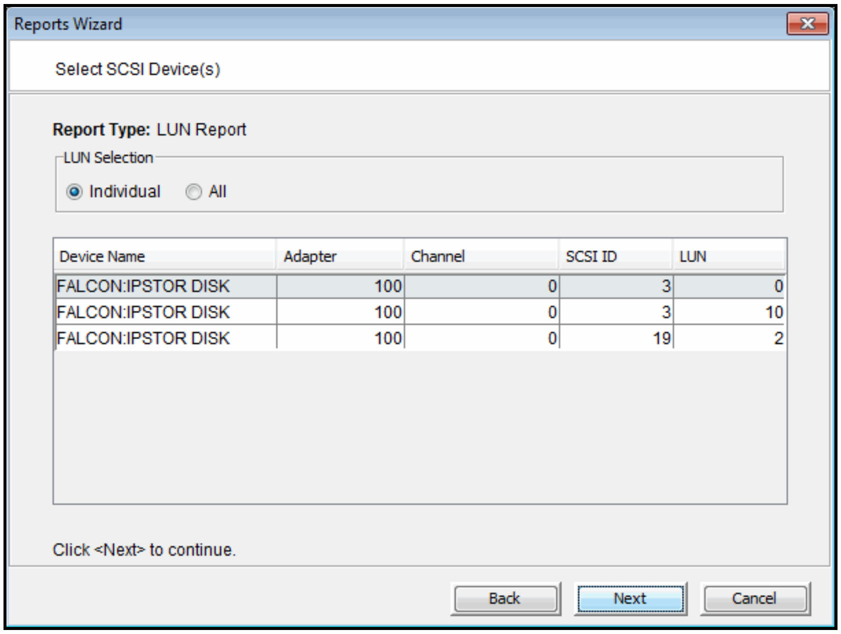

Report Type: LUN Report

LUN Selection

● Individual ○ All

Device Name	Adapter	Channel	SCSI ID	LUN
FALCON:IPSTOR DISK	100	0	3	0
FALCON:IPSTOR DISK	100	0	3	10
FALCON:IPSTOR DISK	100	0	19	2

Click <Next> to continue.

Results include the tape name and barcode, the library to which it belongs (including whether the tape is a replica resource or is in the vault), its current location, and assigned clients.

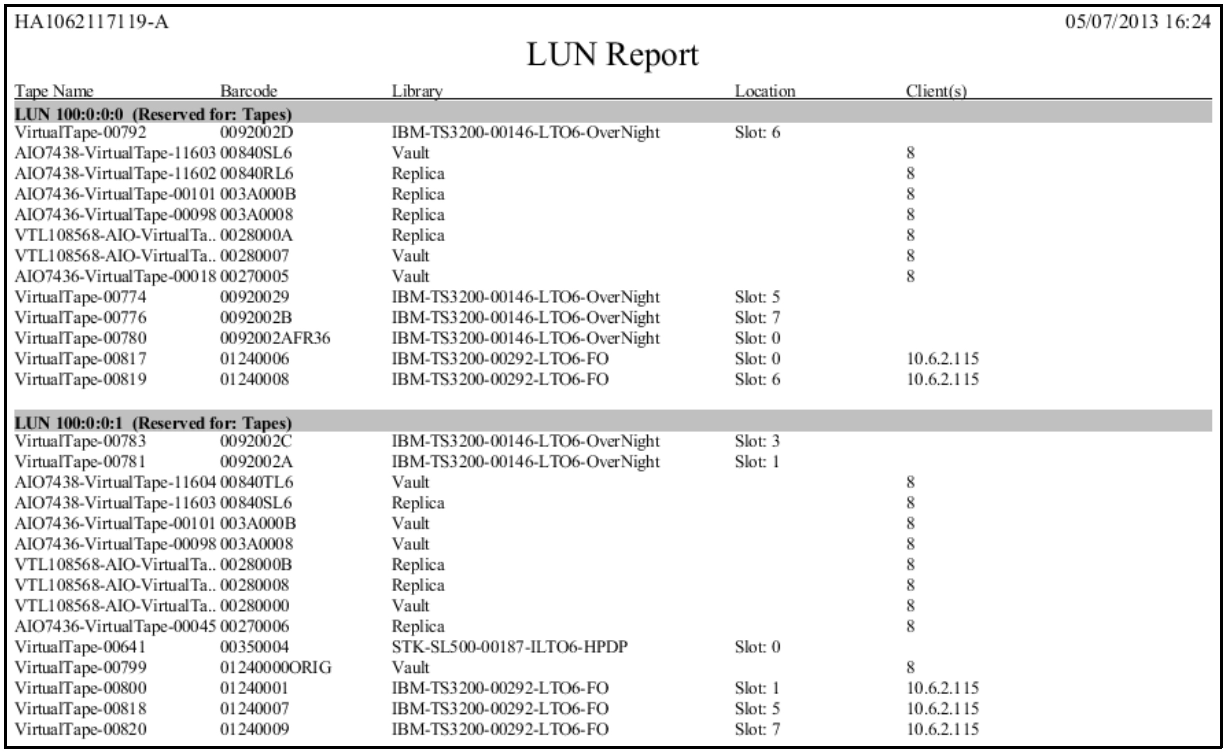

HA1062117119-A

05/07/2013 16:24

LUN Report

Tape Name	Barcode	Library	Location	Client(s)
LUN 100:0:0:0 (Reserved for: Tapes)				
VirtualTape-00792	0092002D	IBM-TS3200-00146-LTO6-OverNight	Slot: 6	
AIO7438-VirtualTape-11603	00840SL6	Vault		8
AIO7438-VirtualTape-11602	00840RL6	Replica		8
AIO7436-VirtualTape-00101	003A000B	Replica		8
AIO7436-VirtualTape-00098	003A0008	Replica		8
VTL108568-AIO-VirtualTa..	0028000A	Replica		8
VTL108568-AIO-VirtualTa..	00280007	Vault		8
AIO7436-VirtualTape-00018	00270005	Vault		8
VirtualTape-00774	00920029	IBM-TS3200-00146-LTO6-OverNight	Slot: 5	
VirtualTape-00776	0092002B	IBM-TS3200-00146-LTO6-OverNight	Slot: 7	
VirtualTape-00780	0092002AFR36	IBM-TS3200-00146-LTO6-OverNight	Slot: 0	
VirtualTape-00817	01240006	IBM-TS3200-00292-LTO6-FO	Slot: 0	10.6.2.115
VirtualTape-00819	01240008	IBM-TS3200-00292-LTO6-FO	Slot: 6	10.6.2.115
LUN 100:0:0:1 (Reserved for: Tapes)				
VirtualTape-00783	0092002C	IBM-TS3200-00146-LTO6-OverNight	Slot: 3	
VirtualTape-00781	0092002A	IBM-TS3200-00146-LTO6-OverNight	Slot: 1	
AIO7438-VirtualTape-11604	00840TL6	Vault		8
AIO7438-VirtualTape-11603	00840SL6	Replica		8
AIO7436-VirtualTape-00101	003A000B	Vault		8
AIO7436-VirtualTape-00098	003A0008	Vault		8
VTL108568-AIO-VirtualTa..	0028000B	Replica		8
VTL108568-AIO-VirtualTa..	00280008	Replica		8
VTL108568-AIO-VirtualTa..	00280000	Vault		8
AIO7436-VirtualTape-00045	00270006	Replica		8
VirtualTape-00641	00350004	STK-SL500-00187-ILTO6-HPDP	Slot: 0	
VirtualTape-00799	01240000ORIG	Vault		8
VirtualTape-00800	01240001	IBM-TS3200-00292-LTO6-FO	Slot: 1	10.6.2.115
VirtualTape-00818	01240007	IBM-TS3200-00292-LTO6-FO	Slot: 5	10.6.2.115
VirtualTape-00820	01240009	IBM-TS3200-00292-LTO6-FO	Slot: 7	10.6.2.115

NAS Resource Usage

This history report shows disk space usage, total disk space, and percentage used for a specific file system or all file systems on a specific date or range of dates.

This report can also be configured as a scheduled report.

The graph shows the total size and total used size of NAS resources based on the specified intervals.

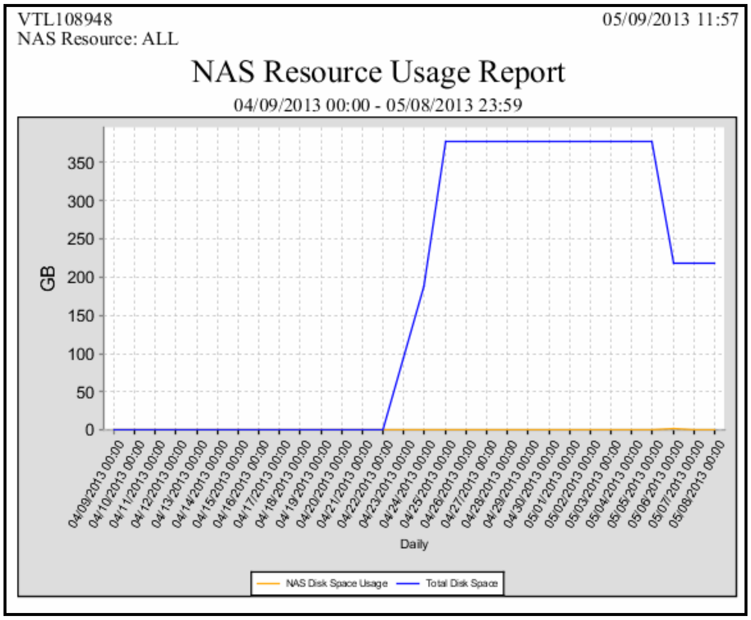

VTL108948
NAS Resource: ALL

05/09/2013 11:57

NAS Resource Usage Report

04/09/2013 00:00 - 05/08/2013 23:59

Information in the report table is based on the specified intervals and includes total size, percentage used, and total used size.

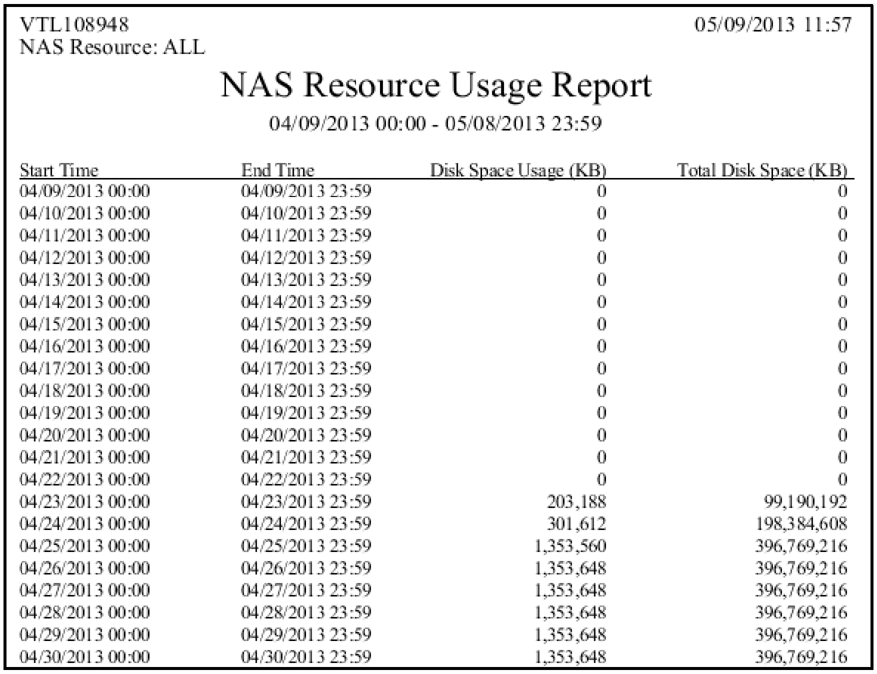

VTL108948
NAS Resource: ALL

05/09/2013 11:57

NAS Resource Usage Report

04/09/2013 00:00 - 05/08/2013 23:59

Start Time	End Time	Disk Space Usage (KB)	Total Disk Space (KB)
04/09/2013 00:00	04/09/2013 23:59	0	0
04/10/2013 00:00	04/10/2013 23:59	0	0
04/11/2013 00:00	04/11/2013 23:59	0	0
04/12/2013 00:00	04/12/2013 23:59	0	0
04/13/2013 00:00	04/13/2013 23:59	0	0
04/14/2013 00:00	04/14/2013 23:59	0	0
04/15/2013 00:00	04/15/2013 23:59	0	0
04/16/2013 00:00	04/16/2013 23:59	0	0
04/17/2013 00:00	04/17/2013 23:59	0	0
04/18/2013 00:00	04/18/2013 23:59	0	0
04/19/2013 00:00	04/19/2013 23:59	0	0
04/20/2013 00:00	04/20/2013 23:59	0	0
04/21/2013 00:00	04/21/2013 23:59	0	0
04/22/2013 00:00	04/22/2013 23:59	0	0
04/23/2013 00:00	04/23/2013 23:59	203,188	99,190,192
04/24/2013 00:00	04/24/2013 23:59	301,612	198,384,608
04/25/2013 00:00	04/25/2013 23:59	1,353,560	396,769,216
04/26/2013 00:00	04/26/2013 23:59	1,353,648	396,769,216
04/27/2013 00:00	04/27/2013 23:59	1,353,648	396,769,216
04/28/2013 00:00	04/28/2013 23:59	1,353,648	396,769,216
04/29/2013 00:00	04/29/2013 23:59	1,353,648	396,769,216
04/30/2013 00:00	04/30/2013 23:59	1,353,648	396,769,216

NAS CIFS Share Usage

This history report shows disk space usage for a specific CIFS share on a specific date or range of dates.

This report can also be configured as a scheduled report.

The graph displays how much space was used on a given share over time.

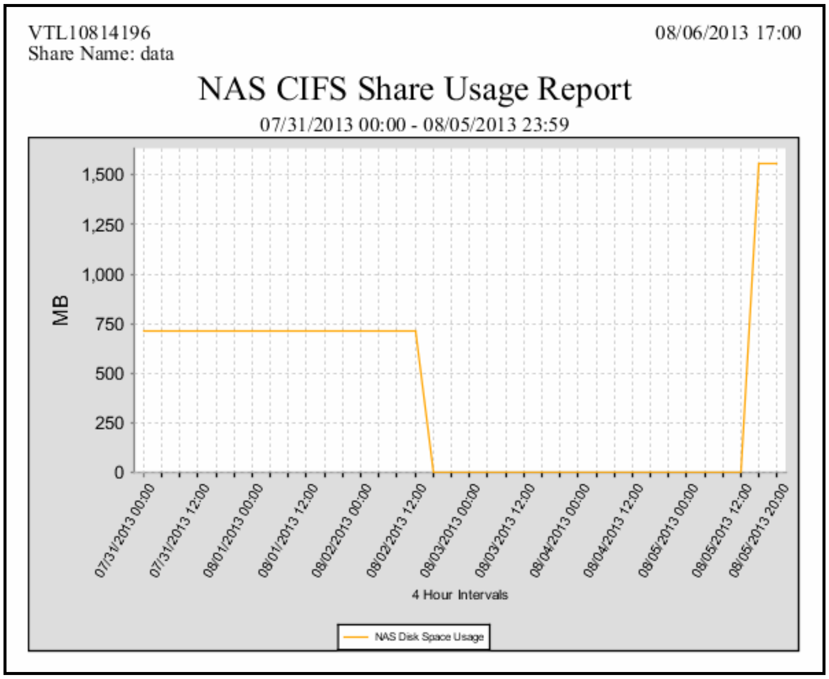

Information in the report table shows how much space is used on the share during each interval within the report time frame.

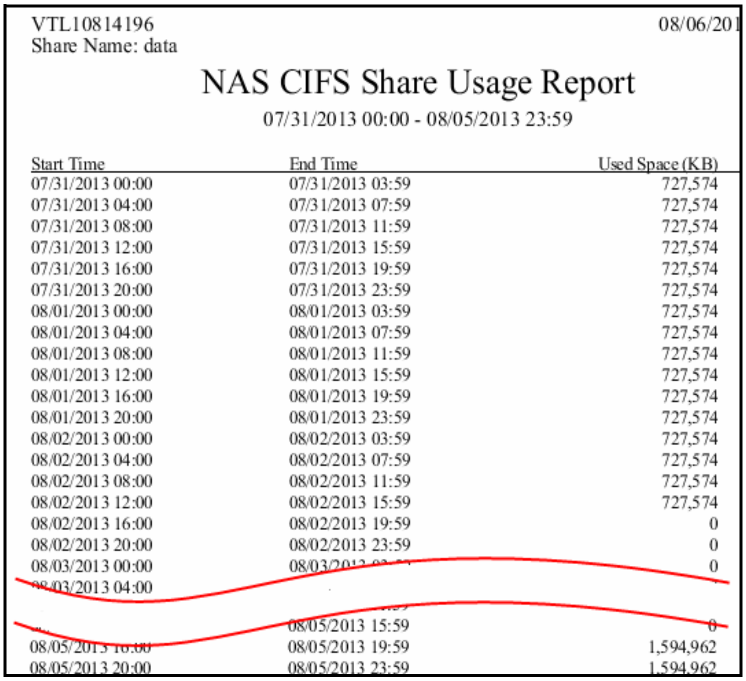

VTL10814196
Share Name: data

08/06/201

NAS CIFS Share Usage Report

07/31/2013 00:00 - 08/05/2013 23:59

Start Time	End Time	Used Space (KB)
07/31/2013 00:00	07/31/2013 03:59	727,574
07/31/2013 04:00	07/31/2013 07:59	727,574
07/31/2013 08:00	07/31/2013 11:59	727,574
07/31/2013 12:00	07/31/2013 15:59	727,574
07/31/2013 16:00	07/31/2013 19:59	727,574
07/31/2013 20:00	07/31/2013 23:59	727,574
08/01/2013 00:00	08/01/2013 03:59	727,574
08/01/2013 04:00	08/01/2013 07:59	727,574
08/01/2013 08:00	08/01/2013 11:59	727,574
08/01/2013 12:00	08/01/2013 15:59	727,574
08/01/2013 16:00	08/01/2013 19:59	727,574
08/01/2013 20:00	08/01/2013 23:59	727,574
08/02/2013 00:00	08/02/2013 03:59	727,574
08/02/2013 04:00	08/02/2013 07:59	727,574
08/02/2013 08:00	08/02/2013 11:59	727,574
08/02/2013 12:00	08/02/2013 15:59	727,574
08/02/2013 16:00	08/02/2013 19:59	0
08/02/2013 20:00	08/02/2013 23:59	0
08/03/2013 00:00	08/03/2013 03:59	0
08/03/2013 04:00		
	08/05/2013 15:59	0
08/05/2013 16:00	08/05/2013 19:59	1,594,962
08/05/2013 20:00	08/05/2013 23:59	1,594,962

Allocation reports

Disk Space Allocation for Virtual Tapes in Libraries

This report can generate the current status of space for use by tapes that are currently in all virtual tape libraries. It can also generate a historical view of space allocated for use by tapes in all or specified virtual tape libraries. In this report, the concept of *used space* is defined as *total allocated space*, which refers to storage consumed by virtual tapes, mixed VITs, and VITs for each virtual tape library.

Results include tape segments used for some types of internal tracking and processing when calculating used or allocated space; however, these segments do not affect the maximum capacity available for tapes. Disks devoted to deduplication, the VTL repository, and standalone tape drives are not represented.

Status report

To display a status report, choose the *Current disk space allocation* option in the report wizard. By default, the report will include data for the current server date. Several results are displayed:

- Pie charts representing disk space allocated to tapes
- A table and bar graphs showing information for each LUN
- A table and bar graphs representing disk space allocated to each library

The *Total Disk Space Allocation chart* shows total disk space, the amount of space that has been allocated for tapes on all included disks, and free (not allocated) space. The *Total Tape Space Usage Summary* chart focuses on the distribution of the space allocated for tapes and shows where tapes are located: in libraries, in the virtual vault, and in replica resources.

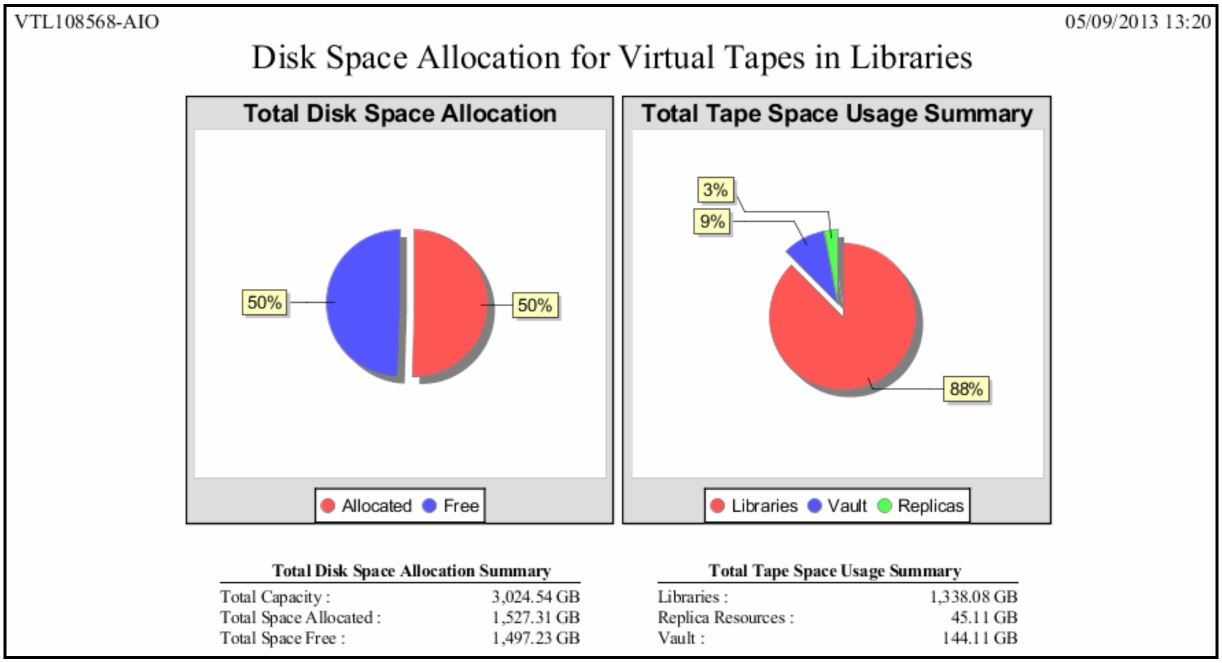

VTL108568-AIO

05/09/2013 13:20

Disk Space Allocation for Virtual Tapes in Libraries

Total Disk Space Allocation Summary	
Total Capacity :	3,024.54 GB
Total Space Allocated :	1,527.31 GB
Total Space Free :	1,497.23 GB

Total Tape Space Usage Summary	
Libraries :	1,338.08 GB
Replica Resources :	45.11 GB
Vault :	144.11 GB

LUN data on the next page of the report includes the SCSI address, vendor ID, and product ID of each LUN, plus the LUN’s total capacity and allocated/free space. A bar chart further represents the total space on each LUN that is allocated to virtual tapes.

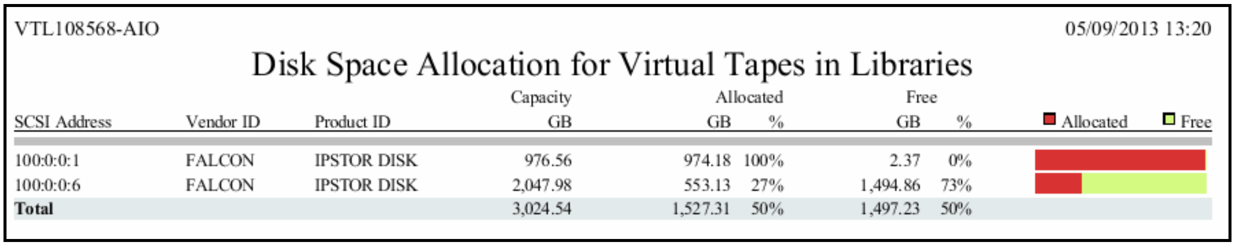

VTL108568-AIO 05/09/2013 13:20

Disk Space Allocation for Virtual Tapes in Libraries

			Capacity	Allocated		Free		🟥 Allocated 🟩 Free
SCSI Address	Vendor ID	Product ID	GB	GB	%	GB	%	
100:0:0:1	FALCON	IPSTOR DISK	976.56	974.18	100%	2.37	0%	
100:0:0:6	FALCON	IPSTOR DISK	2,047.98	553.13	27%	1,494.86	73%	
Total			3,024.54	1,527.31	50%	1,497.23	50%	

Library data includes tape space allocation for each selected library. A bar chart indicates the percentage of disk allocated to tapes.

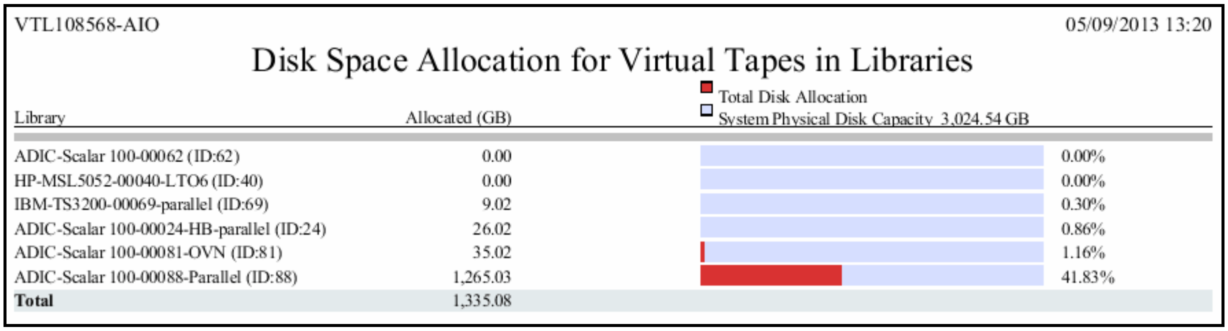

VTL108568-AIO 05/09/2013 13:20

Disk Space Allocation for Virtual Tapes in Libraries

Library	Allocated (GB)	🟥 Total Disk Allocation □ System Physical Disk Capacity 3,024.54 GB	
ADIC-Scalar 100-00062 (ID:62)	0.00		0.00%
HP-MSL5052-00040-LTO6 (ID:40)	0.00		0.00%
IBM-TS3200-00069-parallel (ID:69)	9.02		0.30%
ADIC-Scalar 100-00024-HB-parallel (ID:24)	26.02		0.86%
ADIC-Scalar 100-00081-OVN (ID:81)	35.02		1.16%
ADIC-Scalar 100-00088-Parallel (ID:88)	1,265.03		41.83%
Total	1,335.08		

History report

To display a history report, choose the *Historical library space allocation* option in the report wizard, then select the virtual tape libraries you want to include in the results. You can choose a report period of up to one year. Available interval(s) between data points depend on the period you select.

Each data point represents the allocation value for that point in time, which is the first recorded data from each data interval. Results include a line graph for each library included in the report, showing allocation over time for virtual tapes, mixed VITs, and VITs.

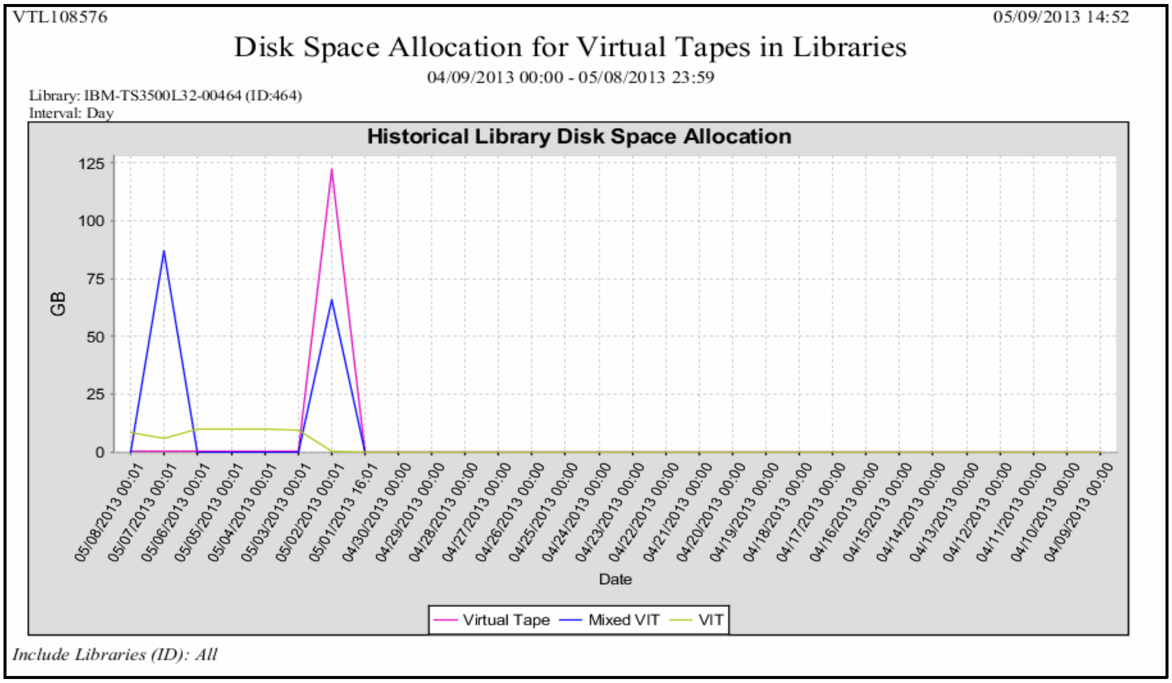

VTL108576 05/09/2013 14:52

Disk Space Allocation for Virtual Tapes in Libraries

04/09/2013 00:00 - 05/08/2013 23:59

Library: IBM-TS3500L32-00464 (ID:464)
Interval: Day

Include Libraries (ID): All

Physical Resource Allocation

This status report displays all virtual devices that have been allocated from all or selected virtualized LUNs. LUNs devoted to use by direct devices are excluded; SCSI aliases are not displayed. Results for each LUN include its name, SCSI address, device type, the category for which it is used (such as for virtual devices), its total capacity, the amount and percentage of allocated space, the amount and percentage of free space, the number of segments on the device. For a resource on a backup server, the report includes the number of virtual tapes allocated on the device.

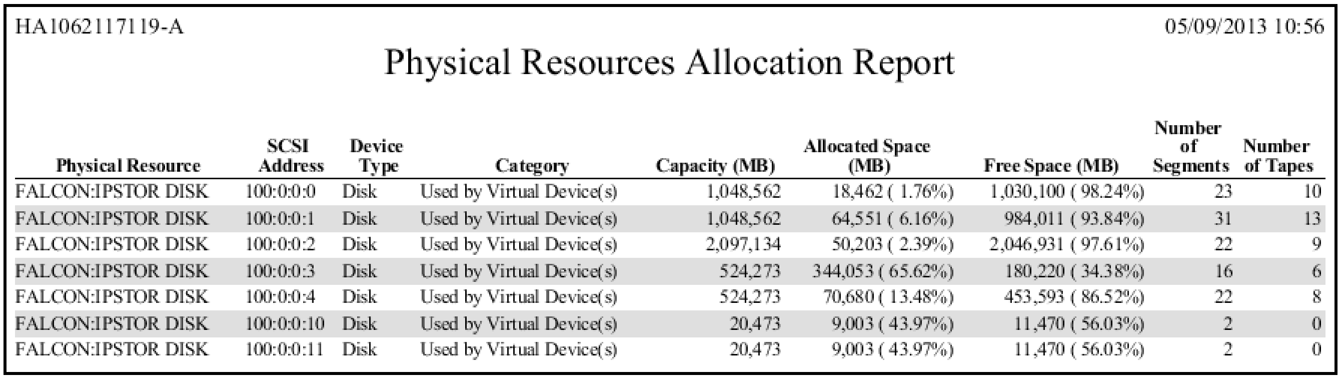

HA1062117119-A 05/09/2013 10:56

Physical Resources Allocation Report

Physical Resource	SCSI Address	Device Type	Category	Capacity (MB)	Allocated Space (MB)	Free Space (MB)	Number of Segments	Number of Tapes
FALCON:IPSTOR DISK	100:0:0:0	Disk	Used by Virtual Device(s)	1,048,562	18,462 (1.76%)	1,030,100 (98.24%)	23	10
FALCON:IPSTOR DISK	100:0:0:1	Disk	Used by Virtual Device(s)	1,048,562	64,551 (6.16%)	984,011 (93.84%)	31	13
FALCON:IPSTOR DISK	100:0:0:2	Disk	Used by Virtual Device(s)	2,097,134	50,203 (2.39%)	2,046,931 (97.61%)	22	9
FALCON:IPSTOR DISK	100:0:0:3	Disk	Used by Virtual Device(s)	524,273	344,053 (65.62%)	180,220 (34.38%)	16	6
FALCON:IPSTOR DISK	100:0:0:4	Disk	Used by Virtual Device(s)	524,273	70,680 (13.48%)	453,593 (86.52%)	22	8
FALCON:IPSTOR DISK	100:0:0:10	Disk	Used by Virtual Device(s)	20,473	9,003 (43.97%)	11,470 (56.03%)	2	0
FALCON:IPSTOR DISK	100:0:0:11	Disk	Used by Virtual Device(s)	20,473	9,003 (43.97%)	11,470 (56.03%)	2	0

Configuration reports

Fibre Channel Adapters Configuration

On a VTL or deduplication server, this status report shows the World Wide Port Name (WWPN) and port information for all Fibre Channel adapters; this report is useful for matching up WWPNs with clients.

On a deduplication server, the report also shows mode (initiator vs target) information.

This report is available only as a one-time report.

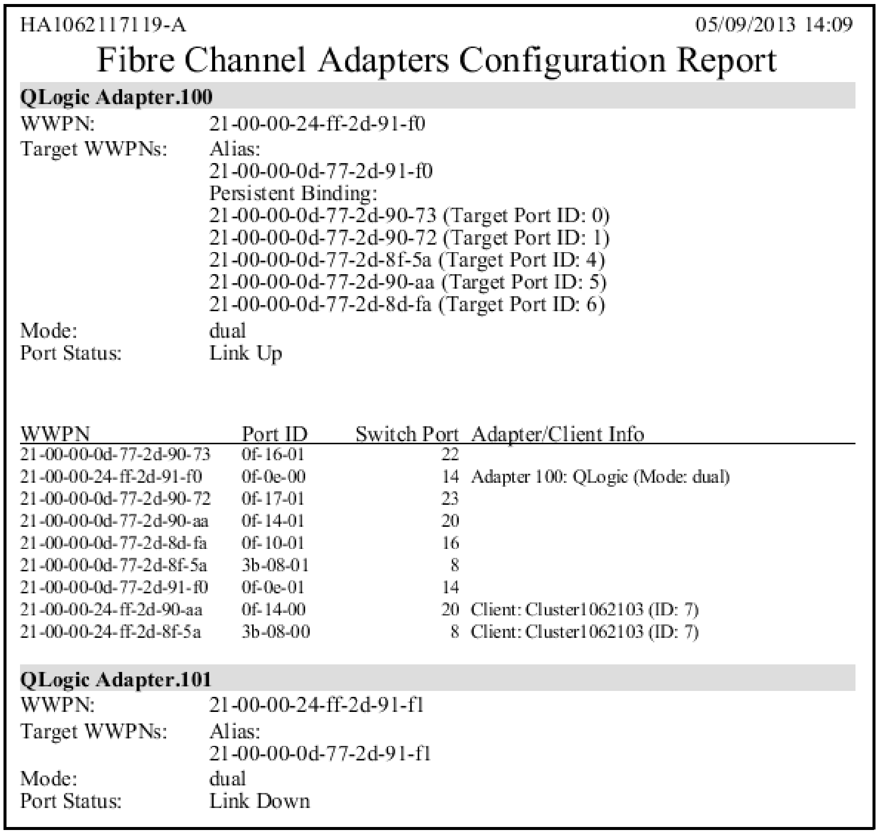

HA1062117119-A 05/09/2013 14:09

Fibre Channel Adapters Configuration Report

QLogic Adapter.100

WWPN: 21-00-00-24-ff-2d-91-f0

Target WWPNs: Alias:
21-00-00-0d-77-2d-91-f0
Persistent Binding:
21-00-00-0d-77-2d-90-73 (Target Port ID: 0)
21-00-00-0d-77-2d-90-72 (Target Port ID: 1)
21-00-00-0d-77-2d-8f-5a (Target Port ID: 4)
21-00-00-0d-77-2d-90-aa (Target Port ID: 5)
21-00-00-0d-77-2d-8d-fa (Target Port ID: 6)

Mode: dual

Port Status: Link Up

WWPN	Port ID	Switch Port	Adapter/Client Info
21-00-00-0d-77-2d-90-73	0f-16-01	22	
21-00-00-24-ff-2d-91-f0	0f-0e-00	14	Adapter 100: QLogic (Mode: dual)
21-00-00-0d-77-2d-90-72	0f-17-01	23	
21-00-00-0d-77-2d-90-aa	0f-14-01	20	
21-00-00-0d-77-2d-8d-fa	0f-10-01	16	
21-00-00-0d-77-2d-8f-5a	3b-08-01	8	
21-00-00-0d-77-2d-91-f0	0f-0e-01	14	
21-00-00-24-ff-2d-90-aa	0f-14-00	20	Client: Cluster1062103 (ID: 7)
21-00-00-24-ff-2d-8f-5a	3b-08-00	8	Client: Cluster1062103 (ID: 7)

QLogic Adapter.101

WWPN: 21-00-00-24-ff-2d-91-f1

Target WWPNs: Alias:
21-00-00-0d-77-2d-91-f1

Mode: dual

Port Status: Link Down

Physical Resources Configuration

For VTL tape and NAS file resources, this status report displays details per LUN for all physical adapters on the server. For each adapter, the report shows information about each physical device that has been configured to the adapter, including its vendor, product name, SCSI ID, LUN number, disk type and size, and category - system disk, unassigned, or reserved for virtual device. The *Reserved for* column will show if the resource has been reserved for the VTL Configuration Repository, tapes, or NAS resources.

For a deduplication server, the report lists all of the physical resources, including physical resources and physical devices available to backup servers and system drives. The *Reserved for* column will show whether a resource has been reserved for the Deduplication Repository.

This report is available only as a one-time report.

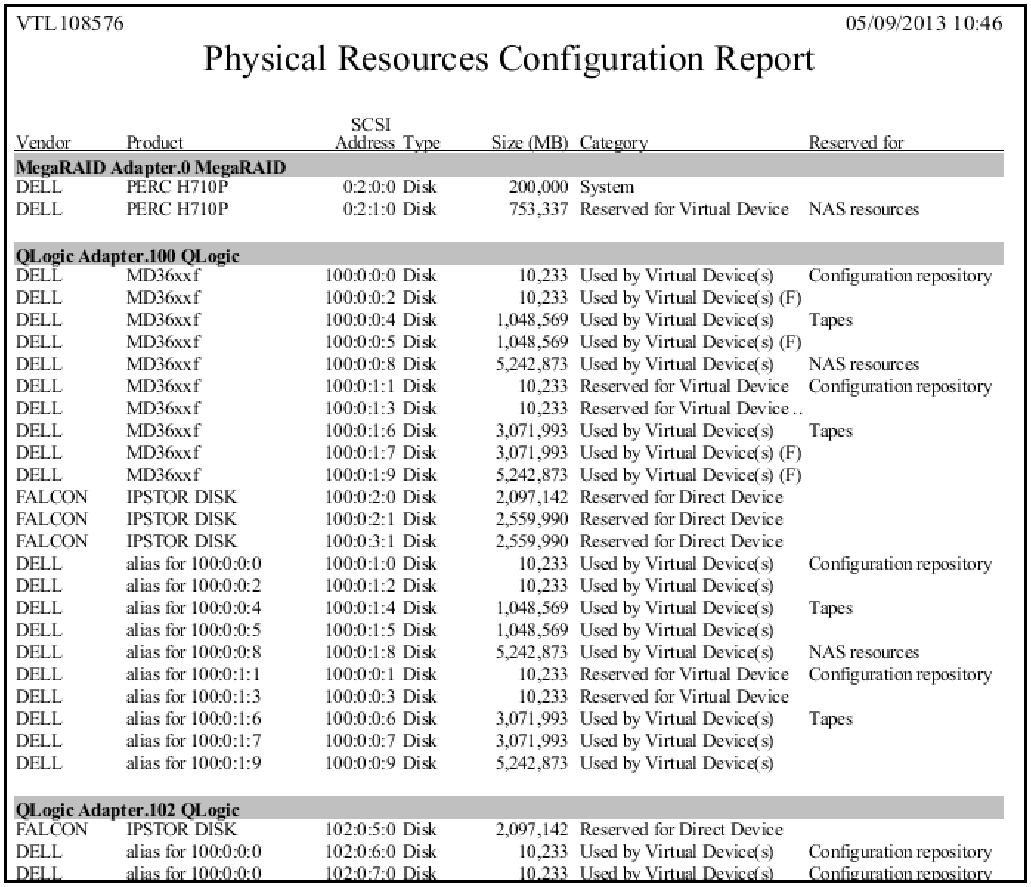

VTL 108576 05/09/2013 10:46

Physical Resources Configuration Report

Vendor	Product	SCSI Address	Type	Size (MB)	Category	Reserved for
MegaRAID Adapter.0 MegaRAID						
DELL	PERC H710P	0:2:0:0	Disk	200,000	System	
DELL	PERC H710P	0:2:1:0	Disk	753,337	Reserved for Virtual Device	NAS resources
QLogic Adapter.100 QLogic						
DELL	MD36xxf	100:0:0:0	Disk	10,233	Used by Virtual Device(s)	Configuration repository
DELL	MD36xxf	100:0:0:2	Disk	10,233	Used by Virtual Device(s) (F)	
DELL	MD36xxf	100:0:0:4	Disk	1,048,569	Used by Virtual Device(s)	Tapes
DELL	MD36xxf	100:0:0:5	Disk	1,048,569	Used by Virtual Device(s) (F)	
DELL	MD36xxf	100:0:0:8	Disk	5,242,873	Used by Virtual Device(s)	NAS resources
DELL	MD36xxf	100:0:1:1	Disk	10,233	Reserved for Virtual Device	Configuration repository
DELL	MD36xxf	100:0:1:3	Disk	10,233	Reserved for Virtual Device ..	
DELL	MD36xxf	100:0:1:6	Disk	3,071,993	Used by Virtual Device(s)	Tapes
DELL	MD36xxf	100:0:1:7	Disk	3,071,993	Used by Virtual Device(s) (F)	
DELL	MD36xxf	100:0:1:9	Disk	5,242,873	Used by Virtual Device(s) (F)	
FALCON	IPSTOR DISK	100:0:2:0	Disk	2,097,142	Reserved for Direct Device	
FALCON	IPSTOR DISK	100:0:2:1	Disk	2,559,990	Reserved for Direct Device	
FALCON	IPSTOR DISK	100:0:3:1	Disk	2,559,990	Reserved for Direct Device	
DELL	alias for 100:0:0:0	100:0:1:0	Disk	10,233	Used by Virtual Device(s)	Configuration repository
DELL	alias for 100:0:0:2	100:0:1:2	Disk	10,233	Used by Virtual Device(s)	
DELL	alias for 100:0:0:4	100:0:1:4	Disk	1,048,569	Used by Virtual Device(s)	Tapes
DELL	alias for 100:0:0:5	100:0:1:5	Disk	1,048,569	Used by Virtual Device(s)	
DELL	alias for 100:0:0:8	100:0:1:8	Disk	5,242,873	Used by Virtual Device(s)	NAS resources
DELL	alias for 100:0:1:1	100:0:0:1	Disk	10,233	Reserved for Virtual Device	Configuration repository
DELL	alias for 100:0:1:3	100:0:0:3	Disk	10,233	Reserved for Virtual Device	
DELL	alias for 100:0:1:6	100:0:0:6	Disk	3,071,993	Used by Virtual Device(s)	Tapes
DELL	alias for 100:0:1:7	100:0:0:7	Disk	3,071,993	Used by Virtual Device(s)	
DELL	alias for 100:0:1:9	100:0:0:9	Disk	5,242,873	Used by Virtual Device(s)	
QLogic Adapter.102 QLogic						
FALCON	IPSTOR DISK	102:0:5:0	Disk	2,097,142	Reserved for Direct Device	
DELL	alias for 100:0:0:0	102:0:6:0	Disk	10,233	Used by Virtual Device(s)	Configuration repository
DELL	alias for 100:0:0:0	102:0:7:0	Disk	10,233	Used by Virtual Device(s)	Configuration repository

Performance reports

Deduplication Repository - Performance

This cluster-level report analyzes deduplication performance and CPU usage for the entire cluster for the specified period of time. It is generated from a VTL server on behalf of the associated deduplication server.

You can set the interval between data points to be an hour, a day, a week, a month, or a quarter (depending on the selected date range). Each data point (bar) represents the total performance or usage during the interval.

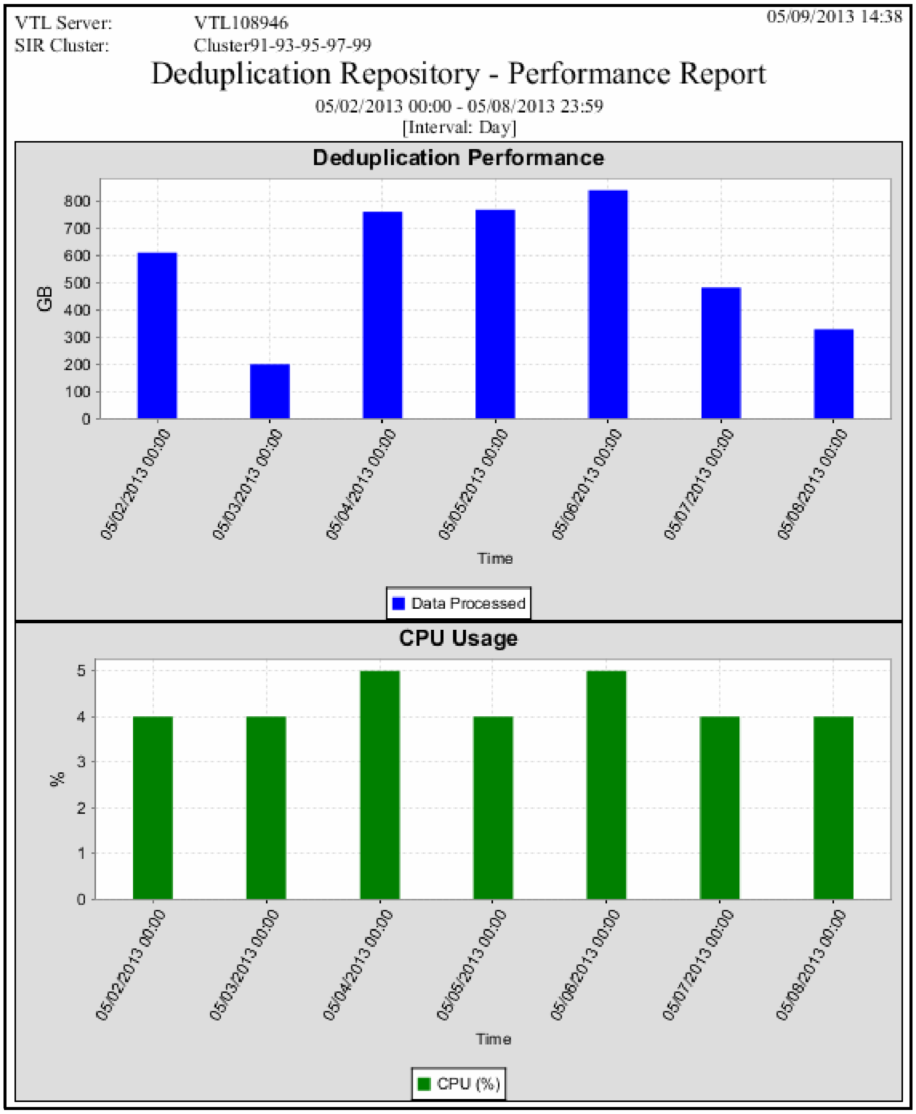

VTL Server: VTL108946
SIR Cluster: Cluster91-93-95-97-99
05/09/2013 14:38

Deduplication Repository - Performance Report
05/02/2013 00:00 - 05/08/2013 23:59
[Interval: Day]

Information in the results table shows the start and end time of the interval, the name of the deduplication cluster or node in the cluster, total data processed, the deduplication ratio, and CPU usage during deduplication as a percentage of system CPU activity. Results for the cluster as a whole are displayed first, followed by a separate results section for each node in the cluster.

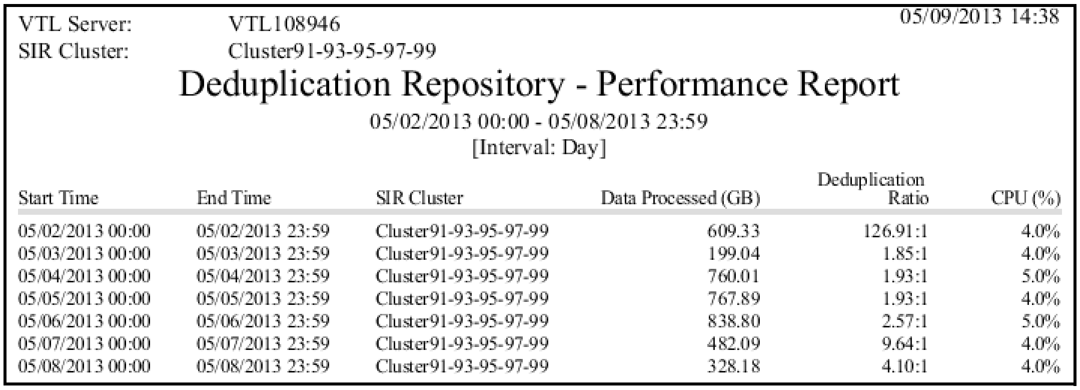

VTL Server: VTL108946 05/09/2013 14:38
SIR Cluster: Cluster91-93-95-97-99

Deduplication Repository - Performance Report

05/02/2013 00:00 - 05/08/2013 23:59
[Interval: Day]

Start Time	End Time	SIR Cluster	Data Processed (GB)	Deduplication Ratio	CPU (%)
05/02/2013 00:00	05/02/2013 23:59	Cluster91-93-95-97-99	609.33	126.91:1	4.0%
05/03/2013 00:00	05/03/2013 23:59	Cluster91-93-95-97-99	199.04	1.85:1	4.0%
05/04/2013 00:00	05/04/2013 23:59	Cluster91-93-95-97-99	760.01	1.93:1	5.0%
05/05/2013 00:00	05/05/2013 23:59	Cluster91-93-95-97-99	767.89	1.93:1	4.0%
05/06/2013 00:00	05/06/2013 23:59	Cluster91-93-95-97-99	838.80	2.57:1	5.0%
05/07/2013 00:00	05/07/2013 23:59	Cluster91-93-95-97-99	482.09	9.64:1	4.0%
05/08/2013 00:00	05/08/2013 23:59	Cluster91-93-95-97-99	328.18	4.10:1	4.0%

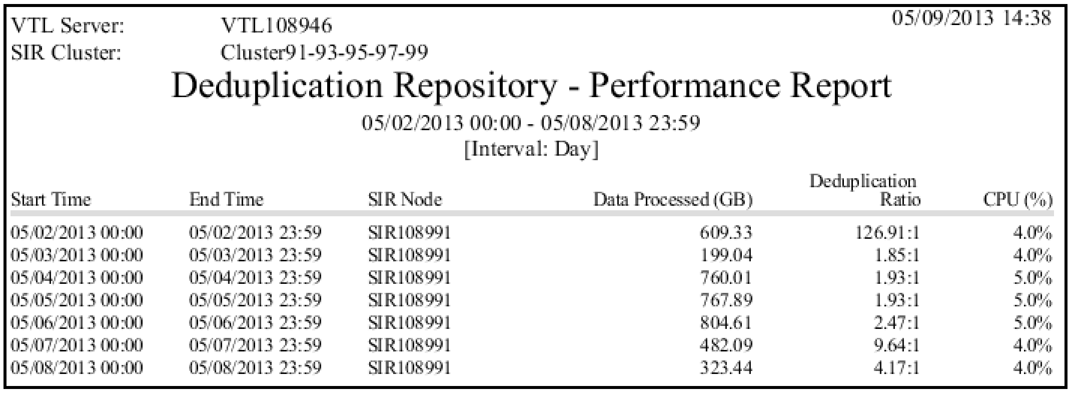

VTL Server: VTL108946 05/09/2013 14:38
SIR Cluster: Cluster91-93-95-97-99

Deduplication Repository - Performance Report

05/02/2013 00:00 - 05/08/2013 23:59
[Interval: Day]

Start Time	End Time	SIR Node	Data Processed (GB)	Deduplication Ratio	CPU (%)
05/02/2013 00:00	05/02/2013 23:59	SIR108991	609.33	126.91:1	4.0%
05/03/2013 00:00	05/03/2013 23:59	SIR108991	199.04	1.85:1	4.0%
05/04/2013 00:00	05/04/2013 23:59	SIR108991	760.01	1.93:1	5.0%
05/05/2013 00:00	05/05/2013 23:59	SIR108991	767.89	1.93:1	5.0%
05/06/2013 00:00	05/06/2013 23:59	SIR108991	804.61	2.47:1	5.0%
05/07/2013 00:00	05/07/2013 23:59	SIR108991	482.09	9.64:1	4.0%
05/08/2013 00:00	05/08/2013 23:59	SIR108991	323.44	4.17:1	4.0%

VTL Performance

This history report displays the average CPU/memory usage and total amount of I/O data for each time interval for the entire VTL system, including all (the default) or selected adapters, LUNs, clients, and virtual tape libraries. If compression is enabled, total I/O data is computed as the compressed value(s), except for Virtual Tape Libraries, for which both uncompressed and compressed data are displayed.

You can set the interval between data points to be an hour, a day, a week, a month, or a quarter (depending on the selected report dates). In the graphs, each data point represents the total throughput/usage during the interval since the previous data point.

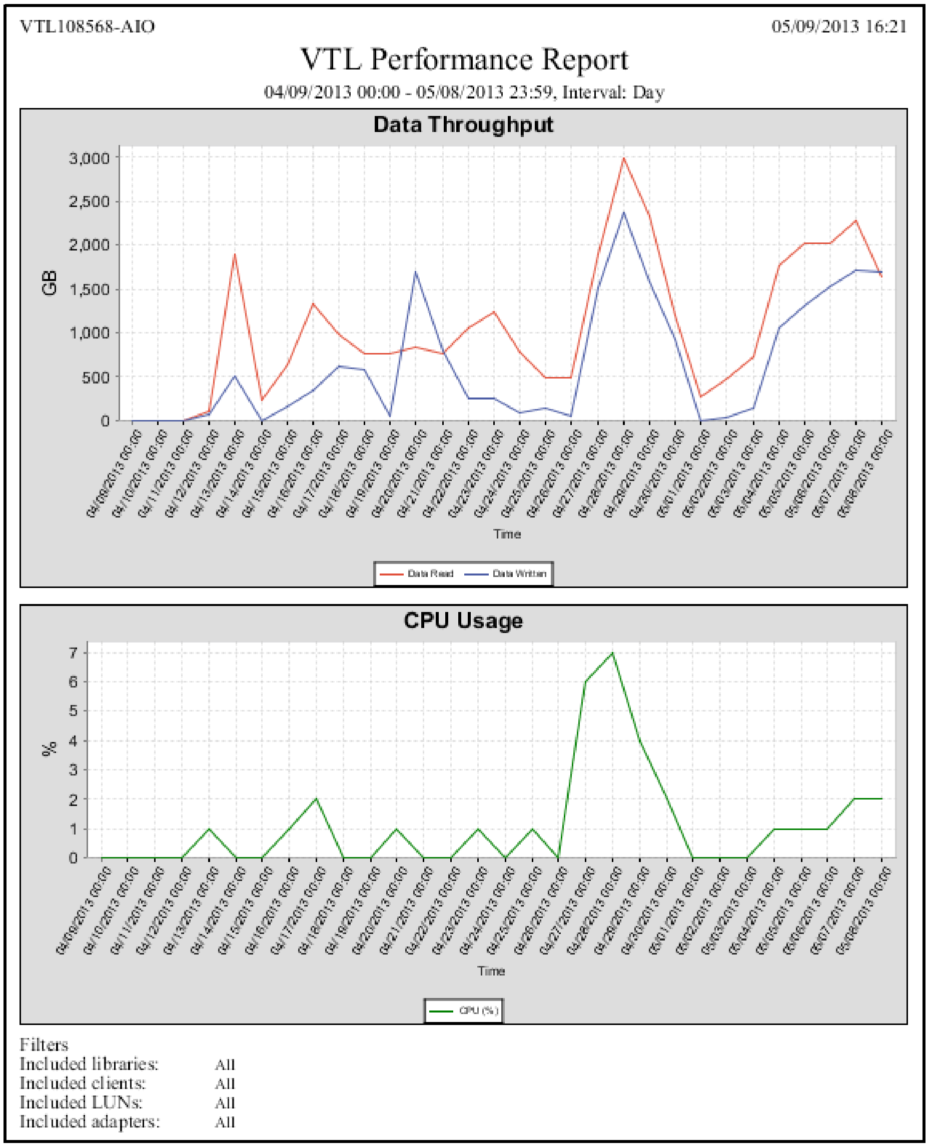

VTL108568-AIO 05/09/2013 16:21

VTL Performance Report

04/09/2013 00:00 - 05/08/2013 23:59, Interval: Day

Filters

Included libraries:	All
Included clients:	All
Included LUNs:	All
Included adapters:	All

Results for the server, adapters, LUNs, clients, and virtual tape libraries appear in dedicated sections of the report.

- Performance information for the VTL server, each adapter, each LUN, and each client device includes the start and end time of the interval, the amount of data read, and the amount of data written.

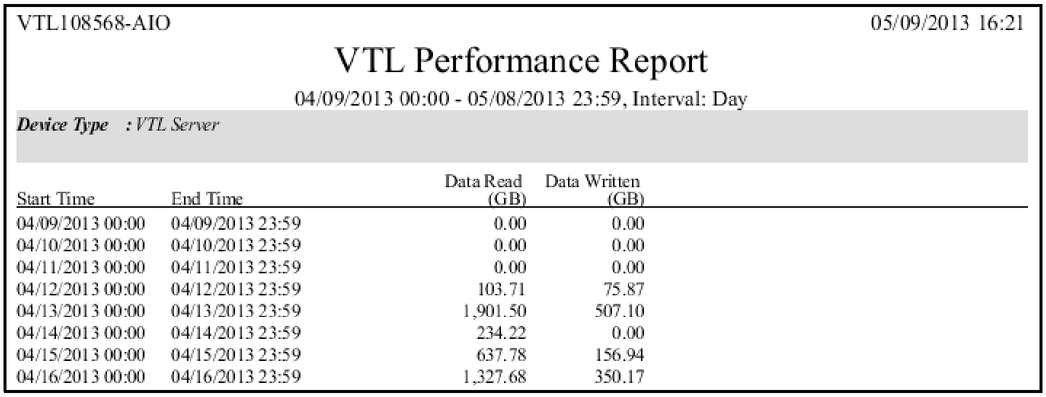

VTL108568-AIO 05/09/2013 16:21

VTL Performance Report

04/09/2013 00:00 - 05/08/2013 23:59, Interval: Day

***Device Type** : VTL Server*

Start Time	End Time	Data Read (GB)	Data Written (GB)
04/09/2013 00:00	04/09/2013 23:59	0.00	0.00
04/10/2013 00:00	04/10/2013 23:59	0.00	0.00
04/11/2013 00:00	04/11/2013 23:59	0.00	0.00
04/12/2013 00:00	04/12/2013 23:59	103.71	75.87
04/13/2013 00:00	04/13/2013 23:59	1,901.50	507.10
04/14/2013 00:00	04/14/2013 23:59	234.22	0.00
04/15/2013 00:00	04/15/2013 23:59	637.78	156.94
04/16/2013 00:00	04/16/2013 23:59	1,327.68	350.17

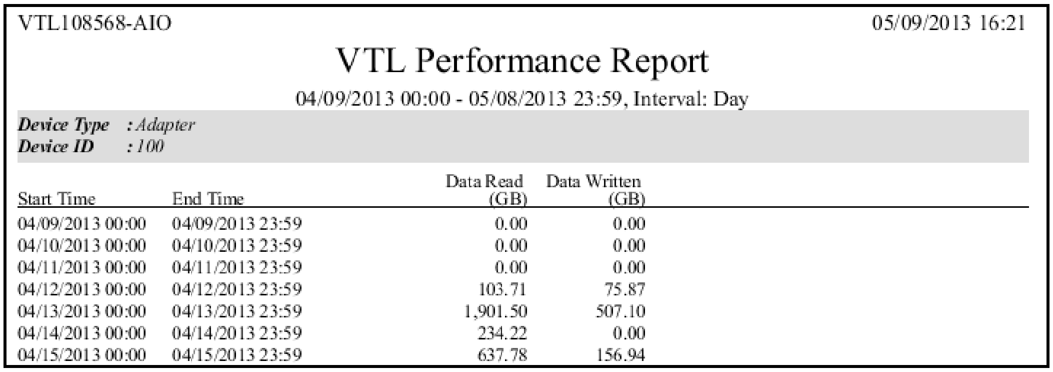

VTL108568-AIO 05/09/2013 16:21

VTL Performance Report

04/09/2013 00:00 - 05/08/2013 23:59, Interval: Day

***Device Type** : Adapter*
***Device ID** : 100*

Start Time	End Time	Data Read (GB)	Data Written (GB)
04/09/2013 00:00	04/09/2013 23:59	0.00	0.00
04/10/2013 00:00	04/10/2013 23:59	0.00	0.00
04/11/2013 00:00	04/11/2013 23:59	0.00	0.00
04/12/2013 00:00	04/12/2013 23:59	103.71	75.87
04/13/2013 00:00	04/13/2013 23:59	1,901.50	507.10
04/14/2013 00:00	04/14/2013 23:59	234.22	0.00
04/15/2013 00:00	04/15/2013 23:59	637.78	156.94

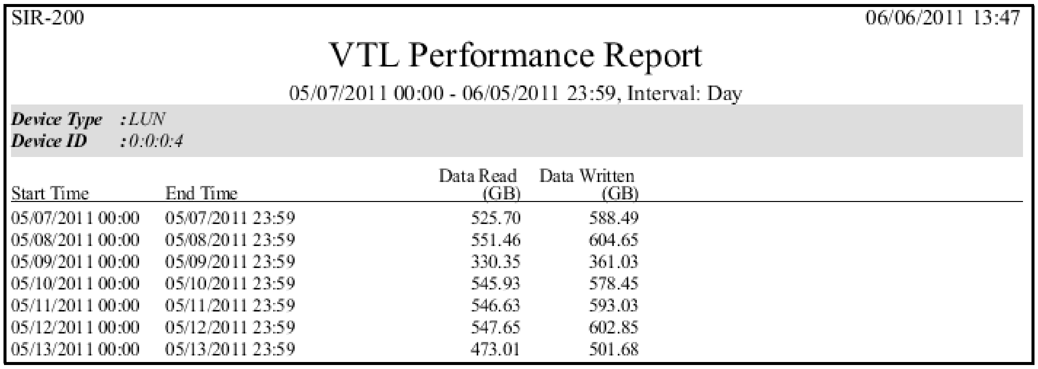

SIR-200 06/06/2011 13:47

VTL Performance Report

05/07/2011 00:00 - 06/05/2011 23:59, Interval: Day

***Device Type** : LUN*
***Device ID** : 0:0:0:4*

Start Time	End Time	Data Read (GB)	Data Written (GB)
05/07/2011 00:00	05/07/2011 23:59	525.70	588.49
05/08/2011 00:00	05/08/2011 23:59	551.46	604.65
05/09/2011 00:00	05/09/2011 23:59	330.35	361.03
05/10/2011 00:00	05/10/2011 23:59	545.93	578.45
05/11/2011 00:00	05/11/2011 23:59	546.63	593.03
05/12/2011 00:00	05/12/2011 23:59	547.65	602.85
05/13/2011 00:00	05/13/2011 23:59	473.01	501.68

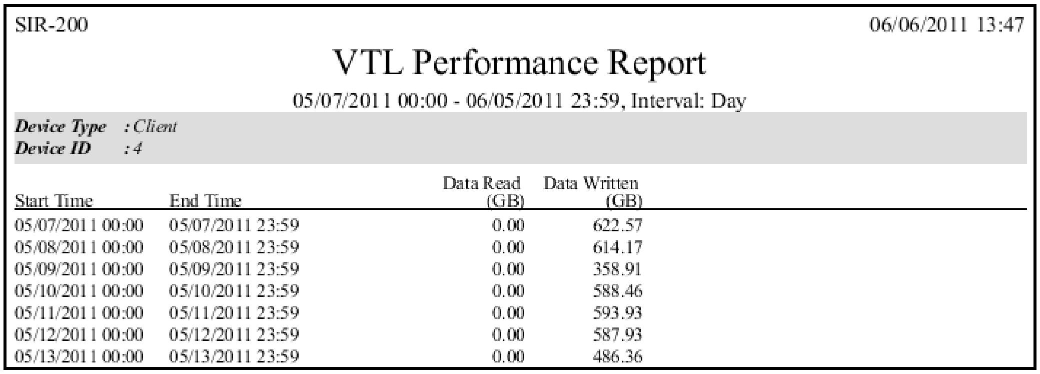

SIR-200 06/06/2011 13:47

VTL Performance Report

05/07/2011 00:00 - 06/05/2011 23:59, Interval: Day

Device Type *: Client*
Device ID *: 4*

Start Time	End Time	Data Read (GB)	Data Written (GB)
05/07/2011 00:00	05/07/2011 23:59	0.00	622.57
05/08/2011 00:00	05/08/2011 23:59	0.00	614.17
05/09/2011 00:00	05/09/2011 23:59	0.00	358.91
05/10/2011 00:00	05/10/2011 23:59	0.00	588.46
05/11/2011 00:00	05/11/2011 23:59	0.00	593.93
05/12/2011 00:00	05/12/2011 23:59	0.00	587.93
05/13/2011 00:00	05/13/2011 23:59	0.00	486.36

- Performance information for a Virtual Tape Library includes the start and end time of each interval, the amount of uncompressed data read and written, and the amount of compressed data read and written.

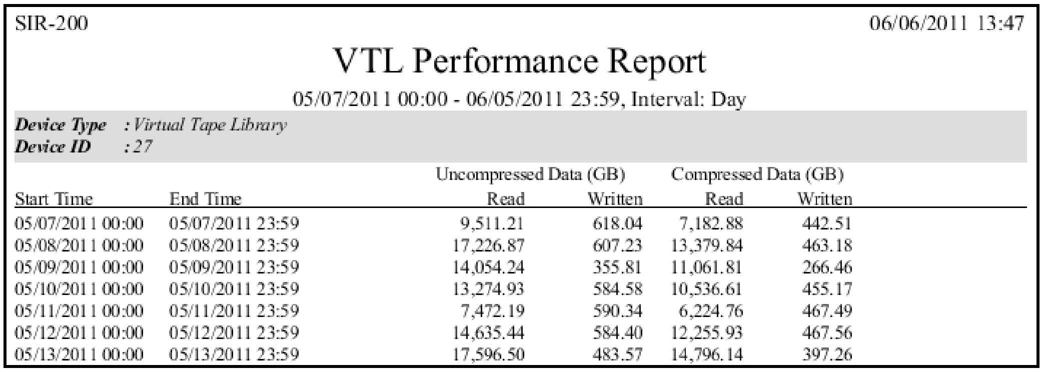

SIR-200 06/06/2011 13:47

VTL Performance Report

05/07/2011 00:00 - 06/05/2011 23:59, Interval: Day

Device Type *: Virtual Tape Library*
Device ID *: 27*

		Uncompressed Data (GB)		Compressed Data (GB)	
Start Time	End Time	Read	Written	Read	Written
05/07/2011 00:00	05/07/2011 23:59	9,511.21	618.04	7,182.88	442.51
05/08/2011 00:00	05/08/2011 23:59	17,226.87	607.23	13,379.84	463.18
05/09/2011 00:00	05/09/2011 23:59	14,054.24	355.81	11,061.81	266.46
05/10/2011 00:00	05/10/2011 23:59	13,274.93	584.58	10,536.61	455.17
05/11/2011 00:00	05/11/2011 23:59	7,472.19	590.34	6,224.76	467.49
05/12/2011 00:00	05/12/2011 23:59	14,635.44	584.40	12,255.93	467.56
05/13/2011 00:00	05/13/2011 23:59	17,596.50	483.57	14,796.14	397.26

Email Alerts

Email Alerts is a unique customer support utility that proactively identifies and diagnoses potential system or component failures and automatically notifies system administrators via email.

With *Email Alerts*, the performance and behavior of servers can be monitored so that system administrators are able to take corrective measures within the shortest amount of time, ensuring optimum service uptime and IT efficiency.

Using pre-configured scripts (called *triggers*), Email Alerts monitors a set of pre-defined, critical system components (memory, disk, server modules, etc.) and system log messages. With its open architecture, administrators can easily register new elements to be monitored by these scripts.

Configure Email Alerts

Email Alerts should be enabled on each VTL server.

> **Note:** If you are going to use failover, it is very important to enable Email Alerts on each server so that an alert will be sent if failure occurs on any server. This should be done before configuring failover.

1. In the console, right-click your server and select *Options* --> *Enable Email Alerts*.

2. Enter general information for your Email Alerts configuration.

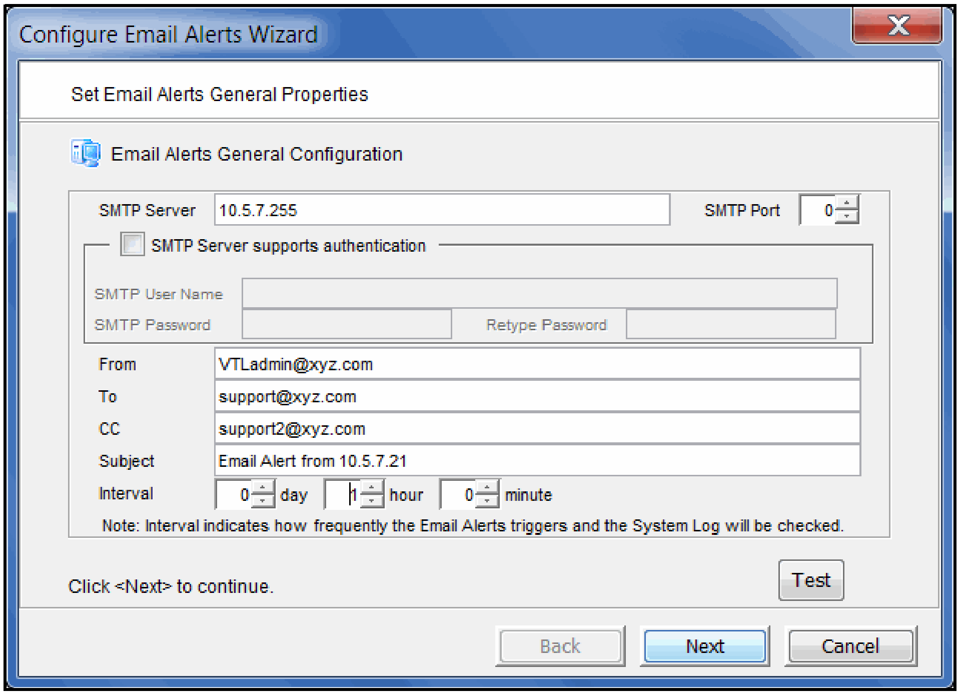
Configure Email Alerts Wizard
Set Email Alerts General Properties
Email Alerts General Configuration
Test
Back
Next
Cancel

SMTP Server: 10.5.7.255   SMTP Port: 0

☐ SMTP Server supports authentication

SMTP User Name:

SMTP Password: Retype Password:

From: VTLadmin@xyz.com

To: support@xyz.com

CC: support2@xyz.com

Subject: Email Alert from 10.5.7.21

Interval: 0 day 1 hour 0 minute

Note: Interval indicates how frequently the Email Alerts triggers and the System Log will be checked.

Click <Next> to continue.

SMTP Server - Specify the mail server that Email Alerts should use to send out notification emails. You can enter an IP address or hostname consisting of alphabet letters, numbers, "_", "-", or ".". The maximum length is 255 characters.

SMTP Port - Specify the mail server port that Email Alerts should use.

SMTP Server supports authentication - Indicate if the SMTP server supports authentication.

SMTP Username/Password - If you enabled the authentication option on the SMTP server, specify the user account that will be used by Email Alerts to log into the mail server. Email Alerts may not work if the SMTP username and password are set without authentication.

From - Specify the email account that will be used in the “From” field of emails sent by Email Alerts.

To - Specify the email address of the account that will receive emails from Email Alerts. This will be used in the “To” field of emails sent by Email Alerts. Separate multiple email addresses with semicolons.

CC - Specify any other email accounts that should receive emails from Email Alerts.

Subject - Specify the text that should appear in the subject line. The general subject defined during setup will be followed by the server name and the trigger-specific subject. If the email is sent based on event severity, the event ID will be appended to the general email subject.

Interval - Specify how frequently the Email Alerts triggers should be checked.

Test - Click to test the configuration by sending a test email to the address defined in the *s* field.

3. Enter the contact information that should appear in each Email Alerts email.

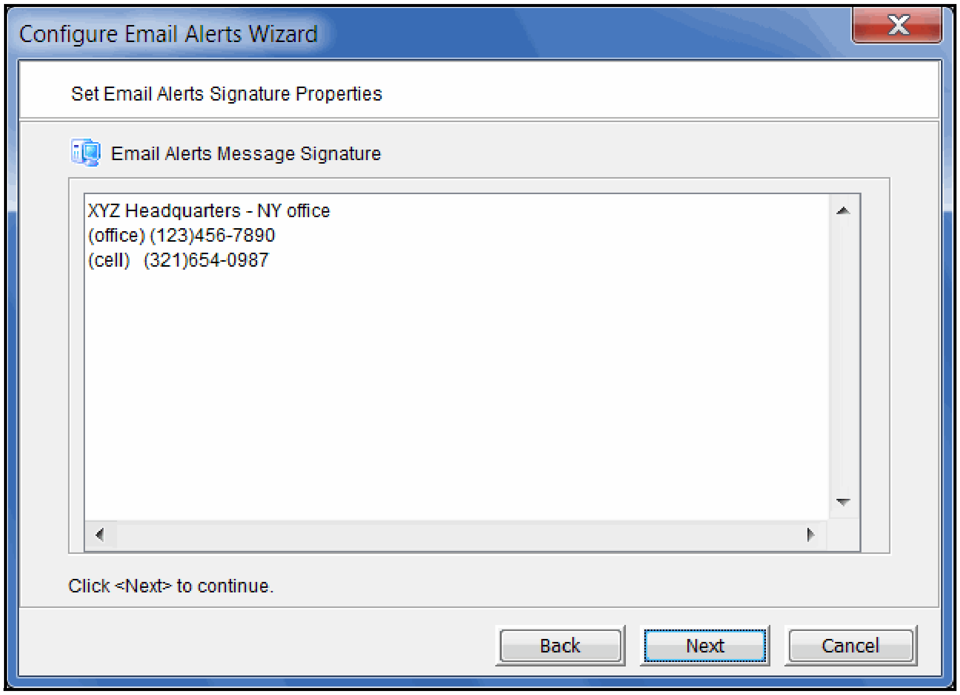

Set Email Alerts Signature Properties

Email Alerts Message Signature

XYZ Headquarters - NY office
(office) (123)456-7890
(cell) (321)654-0987

Click <Next> to continue.

4. Set the triggers that will cause Email Alerts to send an email.

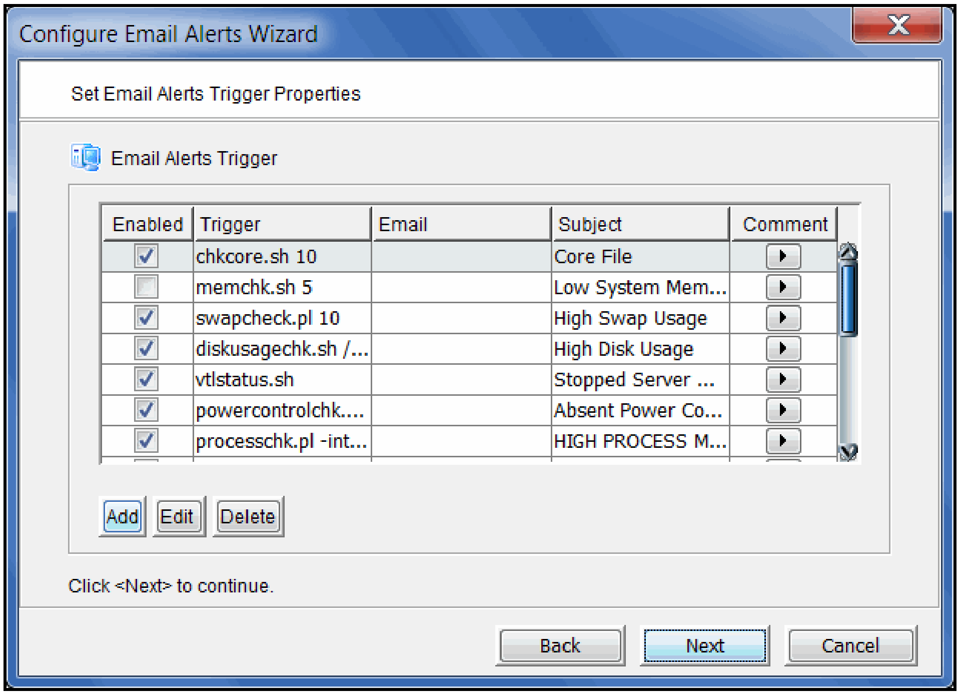

Set Email Alerts Trigger Properties

Email Alerts Trigger

Enabled	Trigger	Email	Subject	Comment
☑	chkcore.sh 10		Core File	▸
□	memchk.sh 5		Low System Mem...	▸
☑	swapcheck.pl 10		High Swap Usage	▸
☑	diskusagechk.sh /...		High Disk Usage	▸
☑	vtlstatus.sh		Stopped Server ...	▸
☑	powercontrolchk....		Absent Power Co...	▸
☑	processchk.pl -int...		HIGH PROCESS M...	▸

Click <Next> to continue.

Triggers are the scripts/programs that perform various types of error checking. By default, DSI includes triggers that check for low system resources, low disk space, and relevant new entries in the system log.

The following scripts are pre-defined:

chkcore.sh 10 (Core file check) - Sends an alert if a new core file has been generated in $ISHOME/var/xray/. If a core file is found, it compresses the file, checks the specified number of core files to keep (default 10), and if the limit has been reached, deletes old core files.

memchk.sh 5 (System memory check) - Sends an alert if the available system memory is below the specified percentage (default 5%).

swapcheck.pl 10 (Swap disk usage check) - Sends an alert if the available swap disk is below the specified percentage (default 10%).

diskusagechk.sh / 95 (Disk usage check) - Sends an alert if the available disk space usage on the specified file system (default is the "/" root file system) is over the specified percentage (default 95). To check usage for multiple disks, append multiple "mount point/threshold" parameters. For example, "diskusagechk.sh / 95 /usr 80" will check "/" and "/usr" with thresholds of 95 and 80, respectively.

vtlstatus.sh (VTL status check) - Sends an alert if a server module has stopped.

powercontrolchk.pl -interval 1440 (Power control check) - Sends an alert if the power control option is missing in a failover configuration. The default is to check once a day (every 1440 minutes).

processchk.pl -interval 60 (System process check) - Sends an alert if a process uses more than 1 GB of memory or more than 90% of usage. The default is to check every hour.

zombiechk.pl 10 -interval 1440 (Defunct process check) - Sends an alert if the number of defunct processes is over the specified value (default 10). The default is to check once a day (every 1440 minutes).

neterrorchk.pl -interval 60 (Network error check) - Sends an email alert for network errors, overruns, dropped frames, or network collisions. The default is to check every hour.

netconfchk.pl -interval 1440 (Network configuration check) - Sends an email alert for inactive network interfaces and invalid broadcast IP addresses. The default is to check once a day (every 1440 minutes).

fcchk.pl -interval 60 (QLogic HBA check) - Sends an alert if the status of a QLogic Fibre Channel initiator port (to storage) is not *Online* or if a Fibre Channel link is down. The default is to check every hour.

nasusagechk.sh 200 1024 (NAS resource usage check) - (VTL server only) Sends an alert if free disk space on a NAS resource is below the specified threshold (default 200 GB) for any resource larger than the specified size (default 1024 GB).

promisecheck.pl 10.x.x.x administrator password -interval 10 (Promise storage check) - Sends an alert for Promise storage hardware errors. This trigger needs to be enabled on-site and requires the IP address and user/ password account to access the storage via *ssh*. The *ssh* service must be enabled and started on the Promise storage. The default is to check every 10 minutes.

SIRmonitor.sh (SIR repository usage check) - (SIR server only) Sends an alert if repository data disk usage is over 90%, index disk usage is over 90%, or memory usage is over 75%.

scsitimeoutchk.pl -interval 60 (Storage connection check) - Sends an alert if a SCSI connection has timed out. The default is to check every hour.

reportheartbeat.pl -interval 1440 (Heartbeat check) - Sends an email to indicate that the server is alive. The default is to check once a day (every 1440 minutes).

chknewpatch.pl -interval 1440 (New patch check) - Sends an alert if new patches are detected in the $ISHOME/newpatches directory. The default is to check every day.

syslogchk.pl (System log check) - Sends an alert if a message in the system log matches a pattern specified in the syslog.check file (set on the next dialog/ tab). This script looks at the last 20 MB of messages. If the system log is rotated prior to the Email Alerts checking interval, the previous log is checked as well as the current log.
To limit the number of email alerts, you can use the -memorize parameter to set the timeframe (in minutes) to remember each event. Refer to 'Limit repetitive emails' for more information.
Some of the more common events checked in the system log are:

- Failover events
- Replication failure
- Storage path failure
- Mirror failure
- Mirror swap
- SCSI error
- Abandoned commands
- FC pending commands, busy FC, or FC loop down
- Unplugged HBA or missing HBA target
- Storage logout or offline device
- iSCSI client reset
- Kernel error, stack, lock up, or segmentation fault
- Out of memory condition
- Machine reboot
- Hardware or file system error

If you need to modify an existing script or create a new script/program, refer to 'Script/program trigger information' for more information. You cannot delete the predefined triggers.

5. View the message patterns that are tracked in the system log by Email Alerts.

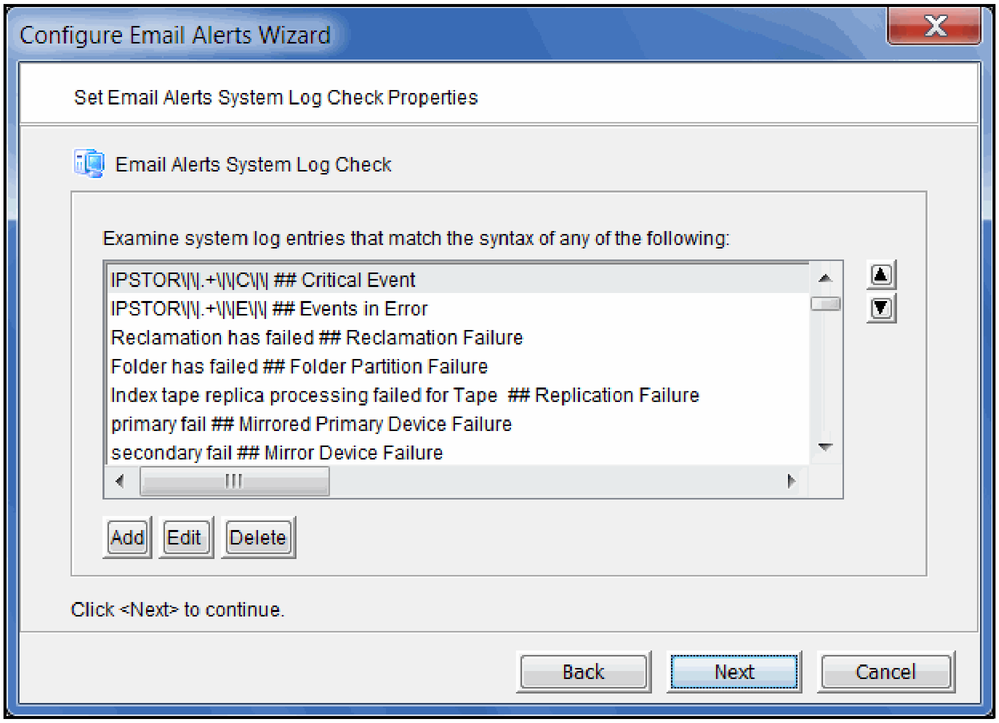

Configure Email Alerts Wizard
Set Email Alerts System Log Check Properties
Email Alerts System Log Check

Examine system log entries that match the syntax of any of the following:

```
IPSTOR\|\|.+\|\|C\|\| ## Critical Event
IPSTOR\|\|.+\|\|E\|\| ## Events in Error
Reclamation has failed ## Reclamation Failure
Folder has failed ## Folder Partition Failure
Index tape replica processing failed for Tape  ## Replication Failure
primary fail ## Mirrored Primary Device Failure
secondary fail ## Mirror Device Failure
```

Click <Next> to continue.

The system log records important events or errors that occur in the system, including those generated by VTL.

Each line is a regular expression. The regular expression rules follow the pattern for AWK (a standard Unix utility).

If needed, you can temporarily add, edit, or delete a pattern.

6. Indicate the severity level of event log messages that should be sent as email alerts.

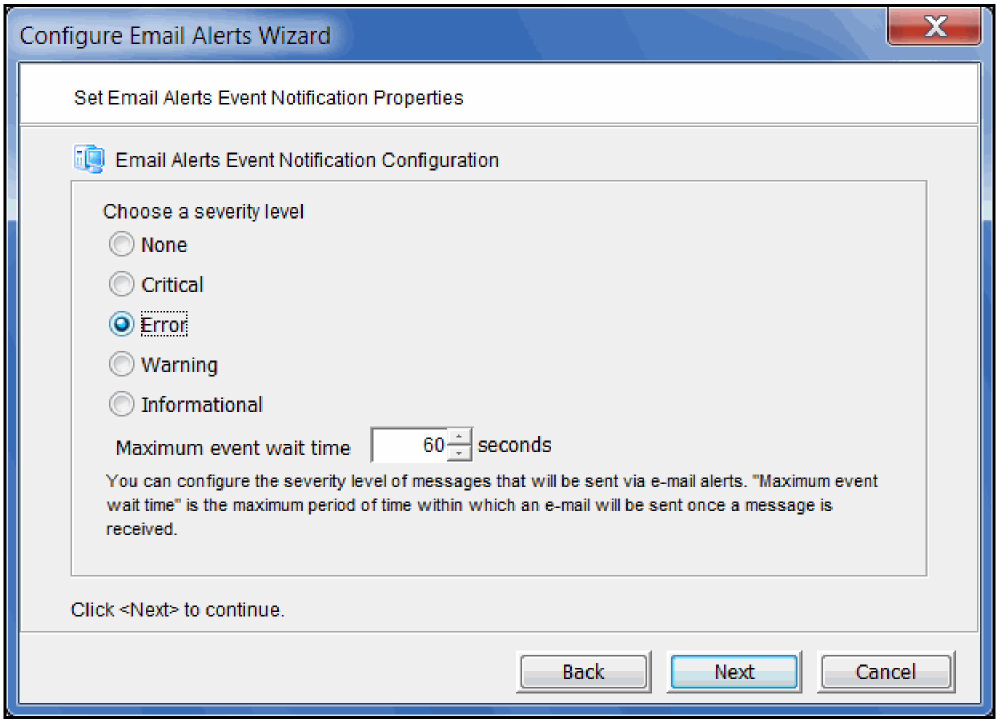

Configure Email Alerts Wizard
Set Email Alerts Event Notification Properties
Email Alerts Event Notification Configuration
Back
Next
Cancel

Choose a severity level

○ None
○ Critical
● Error
○ Warning
○ Informational

Maximum event wait time 60 seconds

You can configure the severity level of messages that will be sent via e-mail alerts. "Maximum event wait time" is the maximum period of time within which an e-mail will be sent once a message is received.

Click <Next> to continue.

You can select one of the following severity levels:

- Critical - sends only critical system log messages
- Error - sends error and critical system log messages.
- Warning - sends warning and higher system log messages.
- Informational - sends system log messages of all severity levels.

If you select them here, critical and error alerts will be sent based on the interval set on the *Email Alerts General Configuration* tab and on the *Maximum event wait time* set below. Warnings/informational messages will only be sent based on the *Maximum event wait time*.

If you add warnings and/or informational messages on the *Email Alerts System Log Check* tab and select *None* here, alerts will only be sent based on the interval set on the *Email Alerts General Configuration* tab.

Maximum event wait time is the maximum period of time within which an email will be sent once a system log event occurs.

7. Confirm all information and click *Finish* to enable Email Alerts.

Email format

The email subject will contain the general subject defined during setup followed by the server name and the trigger-specific subject. If the email is sent based on event severity, the event ID will be appended to the general email subject

The email body will contain the messages returned by the triggers. The alert text starts with the category followed by the actual message coming from the system log. The first 30 lines are displayed. If the email body is more than 16 KB, it will be compressed and sent as an attachment to the email. The signature defined during setup appears at the end of the email body.

Modify Email Alerts properties

Once Email Alerts is enabled, you can modify the information by right-clicking on your server and selecting *Email Alerts*.

Click the appropriate tab to update the desired information.

Limit repetitive emails

You have several options to limit repetitive emails.

Interval to trigger scripts

To override the global Email Alerts interval and run a specific trigger less frequently, you can use the -interval parameter with any trigger. Adding this parameter to a trigger indicates how frequently (in minutes) you want a trigger to be run.

Memorize events

You can limit the number of email alerts for the same event. By using the -memorize parameter for the *syslogchk.pl* trigger, you can have the Email Alerts module memorize events and timestamps of events for which an alert is sent.

If an event is detected several times during the current timeframe, only the first occurrence is reported in the email and the number of repetitions is indicated at the end of the email body with the last occurrence of the message.

The default value is the same as the Email Alerts interval that was set on the first dialog (or the *General* tab if Email Alerts is already configured).

Customize the email for a specific trigger

You can specify an email address to override the default *To* address or a text subject to override the default *Subject*. To do this:

1. Right-click your server and select *Email Alerts* --> *Trigger* tab*.*

2. Highlight the trigger and click *Edit*.

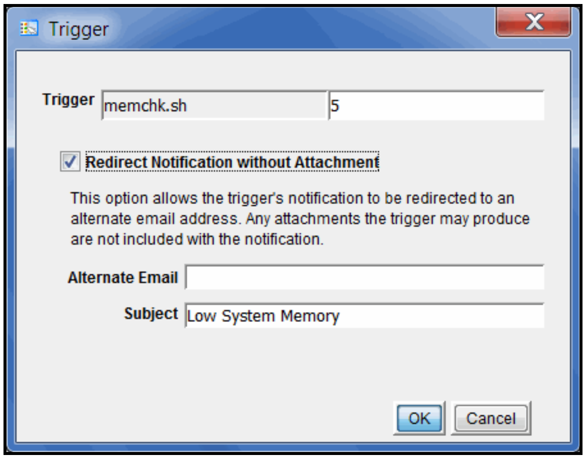

Trigger memchk.sh 5

☑ Redirect Notification without Attachment

This option allows the trigger's notification to be redirected to an alternate email address. Any attachments the trigger may produce are not included with the notification.

Alternate Email

Subject Low System Memory

3. Check the *Redirect Notification Without Attachment* checkbox.

4. Enter the alternate email address or subject.

   The alternate email address and subject are saved to the $ISHOME/etc/callhome/trigger.conf.

   > **Note:** If you specify an email address, it overrides the return code. Therefore, no attachment will be sent, regardless of the return code.

Script/program trigger information

Email Alerts uses script/program triggers to perform various types of error checking. By default, DSI includes several scripts/programs that check for low system memory, changes to the server XML configuration file, and relevant new entries in the system log.

Add a new script

The trigger must be an executable shell script with an .sh extension. If you create a new script, you must add it in the console so that Email Alerts knows of its existence.

To do this:

1. Right-click your server and select *Email Alerts*.
2. Select the *Trigger* tab.
3. Click *Add*.
4. Click *Browse* to locate the shell script/program.
5. If required, enter an argument for the trigger.

   You can also enter a comment for the trigger and specify alternate email information.

Return codes

Return codes determine what happens as a result of the script's execution. The following return codes are valid:

- 0: No action is required and no email is sent.
- 1: Email Alerts sends an email without any attachments.
- 2: Email Alerts attaches all files in $ISHOME/etc and $ISHOME/log to the email.
- 3: Email Alerts sends the X-ray file as an attachment (which includes all files in $ISHOME/etc and $ISHOME/log). Because of its size (minimum of 2 MB), it is recommended that you do not attach the X-ray file to the notification email sent for a trigger.

The $ISHOME/etc directory contains a configuration file (containing virtual device, physical device, HBA, etc. information). The $ISHOME/log directory contains Email Alerts logs (containing events and output of triggers).

Output from trigger

In order for a trigger to send useful information in the email body, it must redirect its output to the environment variable $IPSTORCLHMLOG.

Sample script

The following is the content of the VTL status check trigger, vtlstatus.sh:

```
#!/bin/sh
RET=0
if [ -f /etc/.is.sh ]
then
    . /etc/.is.sh
else
   echo Installation is not complete. Environment profile is missing in
/etc.
    echo
    exit 0 # don't want to report error here so have to exit with error
code 0
fi
$ISHOME/bin/vtl status | grep STOPPED >> $IPSTORCLHMLOG
if [ $? -eq 0 ] ; then
        RET=1
fi
exit $RET
```

If any VTL module has stopped, this trigger generates a return code of 1 and sends an email.

Command Line

Virtual Tape Library (VTL) provides a simple utility that allows you to perform some of the more common functions at a command line instead of through the DSI Management Console. You can use this command line utility to automate many tasks on servers in the VTL system, as well as integrate VTL with your existing management tools.

Use the command line utility

Type `iscon` at the command line to display a list of commands. Each command must be combined with the appropriate long or short arguments (ex. Long: `--server-name`   Short: `-s servername`) that are described in this chapter.

If you type the command name (for example, `c:\iscon importtape`), a list of arguments will be displayed for that command.

Commands

On the following pages are groups of commands you can use to perform functions from the command line on servers with various roles.

- ‘Common commands’ (VTL/NAS, SIR, VTL-S)
- ‘Virtual Tape Library commands’ (VTL/NAS, VTL-S)
- ‘Deduplication Repository commands’ (SIR)
- ‘NAS commands’ (NAS)

You should be aware of the following as you enter commands:

- Type each command on a single line, separating arguments with a space.
- You can use either the short or long arguments.
- Variables are listed in <> after each argument.
- Arguments listed in brackets [ ] are optional.
- The order of the arguments is irrelevant.
- Arguments separated by | are choices. Only one can be selected.
- For a value entered as a literal, it is necessary to enclose the value in quotes (double or single) if it contains special characters such as *, <, >, ?, |, %, $, or space. Otherwise, the system will interpret the characters with a special meaning before it is passed to the command.
- Literals cannot contain leading or trailing spaces. Leading or trailing spaces enclosed in quotes will be removed before the command is processed.

Common arguments

The following arguments are used by many commands. For each, a long and short variation is included. You can use either one. The short arguments **ARE** case sensitive. For arguments that are specific to each command, refer to the section for that command.

Short Argument	Long Argument	Value/Description
-s	--server-name	VTL server name (hostname or IP address)
-u	--server-username	VTL server username
-p	--server-password	VTL server user password
-c	--client-name	VTL client name
-v	--vdevid	VTL virtual device ID

Note: You only need to use the `--server-username (-u)` and `--server-password (-p)` arguments when you log into a server. You do not need them for subsequent commands on the same server during your current session.

Common commands

The commands in this section can be used on a server with any role: VTL/NAS, SIR, or VTL-S.

Server login/logout

Log in to the server

```
iscon login [-s <server-name> -u <username> -p <password> | -e] [-X <rpc-timeout>]

iscon login [--server-name=<server-name> --server-username=<username>
--server-password=<password> | --environment] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command stores the provided set of credentials for the specified server in a secure location. Those credentials will be used in order to authenticate future commands until the logout command is executed.

In order to use the -e (--environment) parameter, you must set the following three environment variables:

- ISSERVERNAME
- ISUSERNAME
- ISPASSWORD

After setting these variables, the environment parameter can be used in the login command in place of -s <server-name>, -u <user-name> and -p <password>.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Note: To set environment variables in the bash shell, you must set three variables as follows:

- export ISSERVERNAME=*10.1.1.1*
- export ISUSERNAME=*root*
- export ISPASSWORD=*password*

Log out from the server

```
iscon logout -s <server-name> [-X <rpc-timeout>]

iscon logout --server-name=<server-name> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command destroys the credentials information stored by the login command for the specified server. Subsequent commands will require the authentication information to be provided at the time of execution.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Note: If, when this command is issued, you are not logged in to the server or you have already logged out, error 0x0902000f will be returned.

Server info

Get server info

```
iscon getserverinfo -s <server-name> [-u <username> -p <password>]
[-X <rpc-timeout>]

iscon getserverinfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command queries information about the specified server and returns server version, operating system version, kernel version, and installed patches.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get server version

```
iscon getserverversion -s <server-name> [-u <username> -p <password>]
[-X <rpc-timeout>]

iscon getserverversion --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command queries information about the specified server and returns the server version and software build number.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Physical devices

Get physical device information

```
iscon getpdevinfo -s <server-name> [-u <username> -p <password>]
[-F [-M | -C <category>] | [-a] [-A] [-I <ACSL>] ] [-o <output-format>]
[-X <rpc-timeout>]

iscon getpdevinfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--config [--include-system-info | --category=<category>] |
[--allocated-list] [--available-list] [--scsiaddress=<ACSL>] ]
[--output-format=<output-format>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves information about the physical devices detected by the specified server. By default, the command displays the allocation information for the virtualized disks and the assigned physical libraries and drives owned by the server. The "allocated" and "available" options work as filters for the default execution. The "config" option displays information about all the physical devices that are detected by the server.

-F (--config) is an option to get the physical device configuration information. The default is to exclude the system device information.

-M (--include-system-info) is an option to include the system device information when -F (--config) is used.

-C (--category) is an option to be used as a filter to get the configuration information for the specified category in one of the values, when -F (--config) is used: *virtual* (default), *service-enabled*, or *direct*.

-M (--include-system-info) and -C (--category) options are mutually exclusive.

-o (--output-format) is an option to specify the output format. The <output-format> for the -F (--config) option is one of the following values: *list* or *detail* or *guid* or *scsi*.

-a (--allocated-list) is an option to get the allocated physical device information.

-A (--available-list) is an option to get the available physical device information.

-I (--scsiaddress) is an option to specify the SCSI address as a device filter in the following format: <ACSL>=#:#:#:# (adapter:channel:id:lun)

The <output-format> for the -a (--allocated-list) and the -A (--available-list) options is one of the following values: *list* or *detail* or *size-only*.

-F (--config), and -a (--allocated-list) and/or -A (--available-list) are mutually exclusive. You can either get the configuration information or get the allocation information. When getting the allocation information, you can specify either -a (--allocated-list), or -A (--available-list) or both. The default is to display both the device allocation and availability information if none of the options is specified.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get adapter info

```
iscon getadapterinfo -s <server-name> [-u <username> -p <password>]
[-a <adapter>] [-N] [-B][-o <output-format>]
[-X <rpc-timeout>]

iscon getadapterinfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--adapter=<adapter>] [--sns-info] [--binding-info]
[--output-format=<output-format>][--rpc-timeout=<rpc-timeout>]
```

Description:

This command displays information for all adapters on the specified server.

-a (--adapter) is an option to display the information for the specified adapter number only.

-N (--sns-info) is an option to get SNS information.

-B (--binding-info) is an option to get persistent binding information.

-o (--output-format) is an option for output format in one of the following values: *list* (default), *detail.*

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Rescan physical devices

```
iscon rescandevices -s <server-name> [-u <username> -p <password>]
[-a <adapter-range>] [-i <scsi-range>] [-l <lun-range>] [-L] [-X <rpc-timeout>]

iscon rescandevices --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--adapter-range=<adapter-range>] [--scsi-range=<scsi-range>] [--lun-range=<lun-range>]
[--sequential] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command rescans the existing physical devices on the specified server and updates the VTL system with the new configuration.

-a (--adapter-range) is the adapter or adapter range to be rescanned. The default is to rescan all the adapters if it is not specified. For example, *-a 5* or *-a 5-10* or *-a auto*.

-i (--scsi-range) is the starting SCSI ID and ending SCSI ID to be rescanned. The default is to rescan all the SCSI IDs if the range is not specified. For example, -i 0-5

-l (--lun-range) is the starting LUN and ending LUN to be rescanned. The default is not to rescan any LUN if it is not specified. For example, -l 0-10

If you want the system to rescan the device sequentially, you can specify the –L (--sequential) option. The default is not to rescan sequentially.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Prepare disk

```
iscon preparedisk -s <server-name> [-u <username> -p <password>]
[-U <target-username> -P <target-password>]
-i <guid> | -I <ACSL> -C <category> [-l <lun-reservation>]
[-X <rpc-timeout>]

iscon preparedisk --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--target-username=<username> --target-password=<password>]
--scsiaddress=<ACSL> | --guid=<guid> --category=<category>
[--lun-reservation=<lun-reservation>][--rpc-timeout=<rpc-timeout>]
```

Description:

This command changes the category for the specified physical device. You must use care with this command; the system will not check the current category and LUN reservation setting before making the specified change.

<guid> is the unique identifier of the physical device.

<ACSL> is the SCSI address of the physical device in the following format: #:#:#:# (adapter:channel:scsi id:lun).

Either -i (--guid) or -I (--scsiaddress) has to be specified for the disk to be prepared.

-C (--category) is required to specify the new category for the physical device in one of the following values: *unassigned*, *virtual*, *direct*.

-l (--lun-reservation) is needed if the category is set to "virtual". The LUN reservation determines what kind of resources can be created on this device. The accepted values are:

- None
- ConfigurationRepository
- DeduplicationRepository
- Tapes
- NASResources

DeduplicationRepository is only valid for SIR and VTL-S servers.

Tapes and *NASResources* are only valid for VTL and VTL-S servers.

If *None* is selected, the device will not be allocated.

If the server is set up for failover, the failover partner has to be rescanned after the disk preparation. <target-username> and <target-password> options specify the user name and password for the failover partner.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Change LUN reservation

```
iscon changelunreservation -s <server-name> [-u <username> -p <password>]
-i <guid> | -I <#:#:#:#> -l <lun-reservation> -f
[-X <rpc-timeout>]

iscon changelunreservation --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--guid=<guid> | --scsiaddress=<#:#:#:#>
--lun-reservation=<lun-reservation> --force
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command changes the LUN reservation for a virtualized device. The new reservation type must be compatible with the existing resources on the device or the device must be empty.

Either -i (--guid) or -I (--scsiaddress) must be used in order to identify the device.

-l (--lun-reservation) is required and determines what kind of resources can be created on this device. The accepted values are:

- None
- ConfigurationRepository
- DeduplicationRepository
- Tapes
- NASResources

DeduplicationRepository is only valid for SIR and VTL-S servers.

Tapes and *NASResources* are only valid for VTL and VTL-S servers.

If *None* is selected, the device will not be allocated.

-f (--force) is required in order to confirm the operation.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Rename physical device

```
iscon renamephysicaldevice -s <server-name> [-u <username> -p <password>]
-i <guid> | -I <ACSL> -n <new-name> [-X <rpc-timeout>]

iscon renamephysicaldevice --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--guid=<guid> | --scsiaddress=<ACSL> --name=<new-name> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command renames a physical device.

-i (guid) specifies the unique identifier of the physical device.

-I (--scsiaddress) is the SCSI address of the physical device in the following format: #:#:#:# (adapter:channel:scsi id:lun)

Either -i (--guid) or -I (--scsiaddress) can be specified for the physical device.

-n (--new-name) is required for the new physical device name. The maximum length for the name is 64. The following characters are invalid for the name: <>"&$/\'

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete physical device

```
iscon deletephysicaldevice -s <server-name> [-u <username> -p <password>]
-i <guid> | -I <#:#:#:#> -f
[-X <rpc-timeout>]

iscon deletephysicaldevice --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--guid=<guid> | --scsiaddress=<#:#:#:#> --force
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes a physical device from server configuration. The device must be offline. The *force* argument must be used in order to confirm the operation.

Either -i (--guid) or -I (--scsiaddress) must be used in order to identify the device to be deleted.

-f (--force) is required in order to confirm the operation.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Restore system preferred path

```
iscon restoresystempreferredpath -s <server-name>
[-u <username> -p <password>]
[-X <rpc-timeout>]

iscon restoresystempreferredpath --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command restores the system preferred path configuration for multi-path Fibre Channel devices.

This command may cause storage to trespass.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Virtual devices

Get virtual device list

```
iscon getvdevlist -s <server-name> [-u <username> -p <password>]
[-l [-v <vdevid> | -n <vdevname> | -B <barcode>] [-A] [-C] [-M <output-delimiter>] ]
[-X <rpc-timeout>]

iscon getvdevlist --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--longlist [--vdevid=<vdevid> | --vdevname=<vdevname> | --barcode=<barcode>]
[--long-physical-layout] [--long-client-list]
[--output-delimiter=<output-delimiter>] ] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command lists all the virtual tape libraries, drives, tapes, and replica tapes on the specified server.

–l (--longlist) is an option to display detailed information in property=value format. Additional options can be specified along with the –l (--longlist) option.

-v (--vdevid) is an option to query and report a single device by its virtual ID. Cannot be combined with -B (--barcode).

-n (--vdevname) is an option to query and report a single device, other than a virtual tape, by its name. Cannot be combined with -B (--barcode).

-B (--barcode) is an option to query and report virtual tapes by barcode. The format for this argument is a list of barcodes separated by commas.

-A(--long-physical-layout) is an option to display the physical layout associated with the device.

-C (--long-client-list) displays the assigned client list when -l (--longlist) option is specified.

-M (--output-delimiter) can be specified when -l is specified to replace the linefeed with the specified delimiter. The maximum length of the delimiter is 8.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Clients

Get client virtual device list

```
iscon getclientvdevlist -s <server-name> [-u <username> -p <password>]
-c <client-name> [-t <client-type>] [-l [-M <output-delimiter>] ]
[-X <rpc-timeout>]

iscon getclientvdevlist --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> [--client-type=<client-type>]
[--longlist [--output-delimiter=<output-delimiter>] ] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves and displays information about all virtual devices assigned to the client from the specified server. The default output format is a list with heading.

-c (--client-name) is required to specify a client name or * for all clients.

-t (client-type) is the type of the client to be retrieved with one of the following values: *FC* or *ISCSI*. The client type will only take effect when the client name is *. Be aware that in some platforms you are required to enclose the "*" in double quote to take it as a literal.

-l(--longlist) is an option to display the long format.

-M (--output-delimiter) can be specified when -l is specified to replace the linefeed with the specified delimiter. The maximum length of the delimiter is 8.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Add Fibre Channel client

```
iscon addclient -s <server-name> [-u <username> -p <password>]
-c <client-name>
[-I <initiator-wwpns>] [[-a on] [-A on]] | [-C on]
[-X <rpc-timeout>]

iscon addclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> [--initiator-wwpns=<initiator-wwpns>]
[[--enable-VSA=on] [--enable-iSeries=on]] | [--enable-Celerra=on]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds a Fibre Channel client to the specified server.

-c (--client-name) is a unique client name for the client to be created. The maximum length of the client name is 32. The following characters are invalid for a client name: <>"&$/\'

-I (--initiator-wwpns) is an option to set the initiator WWPNs. An initiator WWPN is a 16-byte Hex value. Separate initiator WWPNs with commas if more than one initiator WWPN is specified. For example: 13af35d2f4ea6fbc,13af35d2f4ea6fad

-a (--enable-VSA) is an option to enable Volume Set Addressing.

-A (--enable-iSeries) is an option to enable IBM iSeries Server support.

-C (--enable-Celerra) is an option to enable Celerra support.

The Celerra option cannot be combined with VSA or iSeries options.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete client

```
iscon deleteclient -s <server-name> [-u <username> -p <password>]
-c <client-name> [-X <rpc-timeout>]

iscon deleteclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes a client from the specified server.

-c (--client-name) is the name of the client to be deleted.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Rename client

```
iscon renameclient -s <server-name> [-u <username> -p <password>]
-c <client-name> -n <new-name> [-X <rpc-timeout>]

iscon renameclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> --name=<new-name> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command changes the display name for the specified client. For compatibility purposes, the initial name of the client when it was added to the system will be preserved. Either the display name or the initial client name can be used as the client name argument in related commands.

-c (--client-name) is required to specify either the client name or the current alias.

-n (--name) is required to specify the new alias. The maximum length of the alias is 32. The following characters are invalid for the alias: <>"&$/\'

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get client properties

```
iscon getclientprop -s <server-name> [-u <username> -p <password>]
-c <client-name> [-X <rpc-timeout>]

iscon getclientprop --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command returns the current configuration of the specified client.

-c (--client-name) is required to specify the client name.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Mirroring

Create a mirror

```
iscon createmirror -s <server-name> [-u <username> -p <password>]
-v <vdevid> -I <acsl>
[-X <rpc-timeout>]

iscon createmirror --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> --scsiaddress=<acsl>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a mirror of the specified virtual device on the specified physical device. The data will be automatically synchronized.

-v (--vdevid) is required to specify the ID of the virtual device.

-I (--scsi-address) is required to specify the LUN address of the physical device that will contain the mirror. For repository devices, the argument can be a list of scsi addresses separated with commas. For maximum redundancy, the mirror should be on a separate physical device from the primary (preferably on different controllers). The mirror can be defined with disks that are not necessarily identical to each other in terms of vendor, type, or even interface (SCSI, FC, iSCSI).

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get mirror status

```
iscon getmirrorstatus -s <server-name> [-u <username> -p <password>]
-v <vdevid>
[-X <rpc-timeout>]

iscon getmirrorstatus --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command returns mirroring status for the specified virtual device. If mirroring is active, the command will also include synchronization progress and the estimated time to completion.

-v (--vdevid) is required to specify the ID of the mirrored virtual device.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Remove a mirror

```
iscon removemirror -s <server-name> [-u <username> -p <password>]
-v <vdevid>
[-X <rpc-timeout>]

iscon removemirror --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command cancels any active mirror synchronization job for the specified virtual device and removes the mirror.

-v (--vdevid) is required to specify the ID of the virtual device.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Swap a mirror

```
iscon swapmirror -s <server-name> [-u <username> -p <password>]
-v <vdevid>
[-X <rpc-timeout>]

iscon swapmirror --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command swaps the specified primary device with its mirrored copy.

-v (--vdevid) is required to specify the ID of the virtual device.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Sync a mirror

```
iscon syncmirror -s <server-name> [-u <username> -p <password>]
-v <vdevid>
[-X <rpc-timeout>]

iscon syncmirror --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command synchonizes the specified virtual device with its mirroring device. The command does not wait for the operation to finish.

-v (--vdevid) is required to specify the ID of the virtual device to be synchronized.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Replication

Set up network throttling

```
iscon setupreplthrottling -s <server-name> [-u <username> -p <password>]
-V <throttle-value>
[-X <rpc-timeout>]

iscon setupreplthrottling --server-name=<server-name>
[--server-username=<username>
--server-password=<password>]
--throttle-value=<throttle-value>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command sets a maximum data transmission rate during data replication and data resolving.

-V (--throttle-value) is required to provide a throttle value between 10 and 1000000 [KB/s] in order to enable the feature or change the current value. Use 0 to disable the feature.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Reports

The reports below can be generated through the command line interface.

Fibre Channel adapters configuration report

```
iscon createfcaconfreport -s <server-name> [-u <username> -p <password>] [-o <filename>]
[-X <rpc-timeout>]

iscon createfcaconfreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>] [--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the fibre channel adapters configuration. The report can be viewed, printed, emailed, or exported to other formats from the console.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: FCAdaptersConfig-server-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Physical resource allocation report

```
iscon createphyresourceallocreport -s <server-name> [-u <username> -p <password>]
-I <ACSL> [-o <filename>] [-X <rpc-timeout>]

iscon createphyresourceallocreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>] [--scsiaddress=<ACSL>]
[--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the physical resource allocation of the specified device. The report can be viewed, printed, emailed, or exported to other formats from the console.

-I <ACSL> (--scsiaddress) is the LUN address of the device.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: PhysicalResourceAllocation-server-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Physical resources configuration report

```
iscon createphyresourcesconfreport -s <server-name> [-u <username> -p <password>]
[-o <filename>] [-X <rpc-timeout>]

iscon createphyresourcesconfreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>] [--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the physical resources configuration. The report can be viewed, printed, emailed, or exported to other formats from the console.

This command creates a report that lists all physical adapters for a specific server. For each adapter, the report shows all information about each physical device that has been configured to the adapter.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: PhysicalResourcesConfiguration-server-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Users

These commands apply to all users.

Add user

```
iscon adduser -s <server-name[-u <username-p <password>]
-t <user-type> -N <new-username> -W <new-password>
[-X <rpc-timeout>]

iscon adduser --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--user-type=<user-type> --username=<new-username> --password=<new-password>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a new user on the specified server. You must log in as "root" in order to perform this operation.

-t (--user-type) is required to specify the user type:

- A (for VTL Administrator)
- S (for VTL Read-only User)
- I (for VTL iSCSI User)

-N (--username) is required to specify the username.

-W (--password) is required to specify the password to authenticate the user. The password must conform the password security policy of your organization.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete user

```
iscon deleteuser -s <server-name[-u <username-p <password>]
-N <username-to-be-deleted>
[-X <rpc-timeout>]

iscon deleteuser --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--username=<username-to-be-deleted>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes the specified user. Note that this may prevent a backup server from accessing VTL. You must log in as "root" in order to perform this operation.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

List users

```
iscon listusers --server-name=<server-name> [--server-username=<username>
--server-password=<password>]
[--rpc-timeout=<rpc-timeout>]

iscon listusers -s <server-name[-u <username-p <password>]
[-X <rpc-timeout>]
```

Description:

This command displays usernames and user types for user accounts on the specified server. You must log in as "root" in order to perform this operation.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

Set user password

```
iscon setuserpassword -s <server-name[-u <username-p <password>]
-W <new-password> [-X <rpc-timeout>]

iscon setuserpassword --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--password=<new-password> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command allows the connecting user to change his/her own password.

-W (--password) is required to specify the new password to authenticate the user. The password must conform the password security policy of your organization.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Reset user password

```
iscon resetuserpassword -s <server-name[-u <username-p <password>]
-N <username> -W <new-password> [-X <rpc-timeout>]

iscon resetuserpassword --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--username=<username> --password=<new-password>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command changes an account password without requiring entry of the existing password. You must log in as "root" in order to perform this operation.

-N (--username) is required to specify the account name.

-W (--password) is required to specify the new password to authenticate the user. The password must conform the password security policy of your organization.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Licensing

Get license keycode information

```
iscon getlicense -s <server-name> [-u <username> -p <password>] [-l]
[-X <rpc-timeout>]

iscon getlicense --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--longlist] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command gets license keycode information (including license type, description, and registration information) for the specified server.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Add a license keycode

```
iscon addlicense -s <server-name> [-u <username> -p <password>] -k <license-keycode>
[-X <rpc-timeout>]

iscon addlicense --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --license=<license-keycode>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds a license keycode.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Remove a license keycode

```
iscon removelicense -s <server-name> [-u <username> -p <password>] -k <license-keycode>
[-X <rpc-timeout>]

iscon removelicense --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --license=<license-keycode>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command removes a license keycode.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Register a license keycode

```
iscon registerlicense -s <server-name> [-u <username> -p <password>] -k <license-keycode>
[-X <rpc-timeout>]

iscon registerlicense --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--license=<license-keycode> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command registers a specific license key code.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Data encryption

Unlock data encryption

```
iscon unlockdataencryptionoption -s <server-name> [-u <username> -p <password>]
[-X <rpc-timeout>]

iscon unlockdataencryptionoption --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command unlocks the virtual tape encryption option or the deduplication repository encryption option on the specified server, based on the server role. You must log in as "root" in order to perform this operation.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get data encryption information

```
iscon getdataencryptioninfo -s <server-name> [-u <username> -p <password>]
[-X <rpc-timeout>]

iscon getdataencryptioninfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves data encryption information (including encryption and activation status) for the specified server.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Activate data encryption

```
iscon activateencryption -s <server-name> [-u <username> -p <password>]
-W <activation password> [-X <rpc-timeout>]

iscon activateencryption --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--password=<activation password> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command activates encryption, allowing access to data stored on encrypted virtual tapes and replicas and on an encrypted deduplication repository.

-W (--password) is required to provide the data encryption activation password.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Change encryption password

```
iscon changeencryptionactivationpassword -s <server-name> [-u <username> -p <password>]
-O <old password> -W <new password> -C <new password> [-H <password hint>]
[-X <rpc-timeout>]

iscon changeencryptionactivationpassword --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--old-password=<old password> --password=<new password> --confirm-password=<new password>
[--password-hint=<password hint>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command changes the password set for encryption activation. If the specified server is part of a cluster, the password will be changed for all nodes in the cluster.

-O (--old-password) is required to provide the current password.

-W (--password) is required to provide the new password. The password length is between 10 and 16 characters.

-C (--confirm-password) is required to provide the new password again. The two new password arguments must match.

-H (--password-hint) is optional text that can provide password clues. The text is up to 32 characters and it is shown whenever other commands fail due to a password mismatch. In order to replace the old hint, use -H " ". If the argument is not provided, the old hint is preserved.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Miscellaneous

Get X-ray

```
iscon getxray -s <server-name> [-u <username> -p <password>]
[-o <filename>] [-f] [-O <additional options>]
[-m <FTP>] [-fs <ftp server-name> -fo <ftp port> -fd <ftp target directory>
-fu <ftp username> -fp <ftp password>]
[-X <rpc-timeout>]

iscon getxray --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--output-file=<filename>] [--force]
[--options=<additional options] [--method=<FTP>]
[--ftp-server=<ftp server-name> --ftp-port=<ftp port>
--ftp-directory=<ftp target directory> --ftp-user=<ftp username>
--ftp-password=<ftp password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command generates an X-ray file for the specified server and saves it locally, or remotely via the FTP protocol.

-o (--output-file) is an option to specify the X-ray file name. Unless the FTP method is selected, this can include the full path. The default output file name format is: <hostname>-xray-<YYMMDD-HHMMSS>-build<#>.tar.gz

-f (--force) is an option to overwrite the existing file when the output file already exists. Otherwise, an error will be returned. This argument cannot be used with the FTP method.

-O (--options) is an option to add core files and/or detailed log files to the X-ray file using one or both of the following values, separated with a comma:
CORE   LOG

-m (--method) is an option to transfer the X-ray file to a remote server via the FTP protocol. The only accepted value is FTP.

The following ftp arguments are required when the FTP method is used:

-fs (--ftp-server) is the server to which the X-ray file should be transferred.

-fo (--ftp-port) is the ftp port used to transfer the X-ray file.

-fd (--ftp-directory) is the target directory where the X-ray should be stored.

-fu (--ftp-username) is the ftp user name used to authenticate the transfer.

-fp (--ftp-password) is the ftp password used to authenticate the transfer.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get Event Log

```
iscon geteventlog -s <server-name> [-u <username> -p <password>]
[-D <date-range>] [-F <fileFormat>] [-o <filename>] [-H] [-f] [-X <rpc-timeout>]

iscon geteventlog --server-name=<server-name>
[--server-username=<username> --server-password=<password>] [--date-range=<date-range>]
[--file-format=<fileFormat>] [--include-heading] [--output-file=<filename>] [--force]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves the event log messages recorded between specified dates.

-D (--date-range) is the starting date/time and ending date/time in the following format: YYYYMMDDhhmmss-YYYYMMDDhhmmss. The starting time must precede the ending time.

-F (--fileFormat) is one of the following formats: *csv* (default) or *txt*.

-H (--include-heading) is an option to include the event log data heading.

-o (--output-file) is the full path of the file name to save the event log data. If the output filename is not specified, the default filename is: eventlogYYYY-MM-DD-hh-mm-<servername>[.#]

[.#] is the additional suffix when there is a duplicate.

-f (--force) is an option to overwrite the existing file if the output file already exists.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get attention required information

```
iscon getattentionrequired -s <server-name> [-u <username> -p <password>]
[-X <rpc-timeout>]

iscon getattentionrequired --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This commands displays the attention required messages.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Deduplication

Start reclamation

```
iscon startsirreclamation -s <server-name> [-u <username> -p <password>]
-T <SPACE | INDEX> [-f] [-X <rpc-timeout>]

iscon startsirreclamation --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--type=<SPACE | INDEX> [--force] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command triggers reclamation on the associated SIR cluster or server. On VTL servers, the command always starts a new reclamation job regardless of the reclamation type.

On VTL-S servers with deduplication enabled, the command does not start a new space reclamation job unless the force argument is used. The *force* argument can be used only on VTL-S servers, but cannot be used for index reclamation.

-T (--type) is required to specify the reclamation type. Use *SPACE* for space reclamation or *INDEX* for index pruning.

-f (--force) is an option for VTL-S servers with deduplication enabled.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Virtual Tape Library commands

The commands in this section can only be used on a server with the VTL/NAS or VTL-S role.

Options

Enable Hosted Backup

```
iscon enablehostedbackup -s <server-name> [-u <username> -p <password>]
[-X <rpc-timeout>]

iscon enablehostedbackup --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command enables the Hosted Backup option.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Disable Hosted Backup

```
iscon disablehostedbackup -s <server-name> [-u <username> -p <password>]
[-X <rpc-timeout>]

iscon disablehostedbackup --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command disables the Hosted Backup option.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Physical devices

Show storage allocation

```
iscon showstorageallocation -s <server-name> [-u <username> -p <password>]
[-o <csv|list>] [-X <rpc-timeout>]

iscon showstorageallocation --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--output-format=<csv|list>][--rpc-timeout=<rpc-timeout>]
```

Description:

This command displays information about how your storage is allocated.

-o (--output-format) is an option to choose one of the following formats for the output: *csv* (default) or *list.*

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Virtual devices

Get VTL info

```
iscon getvtlinfo -s <server-name> [-u <username> -p <password>]
[-T <vtl-info_type> [-L <tape-library-vid>]] [-F <vtl-info-filter>] [-l [-A] [-M]]
[-X <rpc-timeout>]

iscon getvtlinfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--vtl-info-type=<vtl-info-type> [--tape-library-vid=<tape-library-vid>] ]
[--vtl-info-filter=<vtl-info-filter>]
[--longlist [--long-physical-layout] [--output-delimiter=<output-delimiter>] ]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command lists all the virtual tape libraries, drives, and tapes on the specified server.

-T (--vtl-info-type) is the VTL information type with one of the following values: *VLIBS* or *VDRIVES* or *VAULT* or *PLIBS* or *PDRIVES*.

- VLIBS   = display virtual tape libraries only.
- VDRIVES = display standalone virtual tape drives only
- VAULT   = display virtual tape vault only.
- PLIBS   = display physical tape libraries only.
- PDRIVES = display standalone physical tape drives only.

The default is to display all the information.

-L (--tape-library-vid) is an option to specify the virtual tape library when VLIBS is specified, or to specify the physical tape library when PLIBS is specified.

-F (--vtl-info-filter) is an additional filter that can be combined using the following values separated with commas: *library* or *drive* or *tape*.

- library = include physical and/or virtual library information.
- drive = include physical and/or virtual drive information.
- tape = include physical and/or virtual tape information.

For example: -F "library,drive,tape" or --vtl-info-filter="library,drive,tape"

The default is to display all of the information that applies. There will be an error if <vtl-info-type> is specified and the <vtl-info-filter> specified does not apply. For example, "library" does not apply to "VDRIVES".

-l (--longlist) is an option to display detailed information.

-A (--long-physical-layout) is an option to display the physical layout associated with the device, if applicable. The argument is ignored if -l (--long-list) is not specified.

-M (--output-delimiter) can be specified when -l is specified to replace the linefeed with the specified delimiter. The maximum length of the delimiter is 8.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Assign virtual device

```
iscon assignvdev -s <server-name> [-u <username> -p <password>]
-v <vdevid> -c <client-name> [-y]
[-I <initiatorWWPN|*>] [-T <targetWWPN|*>
[-X <rpc-timeout>]

iscon assignvdev --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --vdevid=<vdevid>
--client-name=<client-name> [--vlib-only]
[--initiatorWWPN=<initiatorWWPN|*>] [--targetWWPN=<targetWWPN|*>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command prepares and assigns a virtual device to an existing Fibre Channel client on the specified server.

-v (--vdevid) is required to specify the virtual device ID of the virtual tape library or virtual tape drive to be assigned.

-c (--client-name) is required to specify the client to which the virtual tape library or drive will be assigned.

-y (--vlib-only) is an option that assigns the virtual tape library to the client without assigning all of the virtual tape drives in the library. The default is to assign all of the virtual tape drives in the library.

-I (--initiatorWWPN) and -T (--targetWWPN) are options for Fibre Channel clients. The initiator WWPN or target WWPN is a 16-byte hex value or "*" for all. For example, 13af35d2f4ea6fbc. The default is "*" if it is -I or the -T option is not specified.

-l (--lun) is another option for Fibre Channel clients. The range is between 0 and 15. The next available LUN will be assigned if is it is not specified.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Assign virtual library or drive to an iSCSI client

```
iscon assignvdevtoiscsiclient -s <server-name> [-u <username> -p <password>]
-v <vdevid> -c <client-name> -r <iscsi-target-id> [-y] [-l <lun>] [-X <rpc-timeout>]

iscon assignvdevtoiscsiclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> --client-name=<client-name> --iscsi-target-id=<iscsi-target-id>
[--vlib-only] [--lun=<lun>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command attaches a virtual library or drive to an iSCSI client.

-v (--vdevid) is required to specify the virtual device ID of the virtual tape library or virtual tape drive to be assigned.

-c (--client-name) is required to specify the client name to assign the virtual tape library or drive to.

-r (--iscsi-target-id) is required to provide the iSCSI target ID.

-y (--vlib-only) is an option for virtual tape library assignment. The default is to assign all of the virtual tape drives in the library. This option assigns the virtual tape library to the client without assigning all of the virtual tape drives in the library.

-l (--lun) is an option to specify LUN ID.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Unassign virtual device

```
iscon unassignvdev -s <server-name> [-u <username> -p <password>]
-v <vdevid> -c <client-name> [-y] [-X <rpc-timeout>]

iscon unassignvdev --server-name=<server-name> [--server-username=<username>]
[--server-password=<password>] --vdevid=<vdevid> --client-name=<client-name>
[--vlib-only] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command unassigns a virtual tape library or drive from the specified client.

-v (--vdevid) is required to specify the virtual device ID of the virtual tape library or drive to be unassigned.

-c (--client-name) is required to specify the client name from which to unassign the library or drive.

-y (--vlib-only) is an option that unassigns the virtual tape library to the client without unassigning all of the virtual tape drives in the library. The default is to unassign all of the virtual tape drives in the library.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Rename virtual device

```
iscon renamevirtualdevice -s <server-name> [-u <username> -p <password>]
-v <vdevid> -n <vdevname>
[-X <rpc-timeout>]

iscon renamevirtualdevice --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> --vdevname=<vdevname>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command can be used to rename any virtual device except a device used for the Deduplication Repository.

-v (--vdevid) is required to specify the virtual device ID of the device to be renamed.

-f (--vdevname) is required to specify the new device name. The name can include a maximum of 64 characters. The following characters are invalid: <>"&$/\ .

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Clients

Add iSCSI client

```
iscon addiscsiclient -s <server-name> [-u <username> -p <password>]
-c <client-name> -I <initiator-name-list>
[-a <user-name-list>] [-X <rpc-timeout>]

iscon addiscsiclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> --initiator-name-list=<initiator-name-list>
[--user-name-list=<user-name-list>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds an iSCSI client to the specified server.

-c (--client-name) is a unique client name for the client to be created. The maximum length of the client name is 64. The following characters are invalid for the client name: <>"&$/\'

-I (--initiator-name-list) is required to provide at least one valid initiator name. Multiple names must be separated with commas.

-a (--user-name-list) is an option to limit client access to specified iSCSI users. There must be an existing iSCSI user account for each name specified. Multiple names must be separated with commas. By default, the client will allow unauthenticated access.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create an iSCSI target

```
iscon createiscsiclienttarget -s <server-name> [-u <username> -p <password>]
-c <client-name> -I <ip-address> -R <iscsi-target-name> [-l <lun>] [-X <rpc-timeout>]

iscon createiscsiclienttarget --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name>
--ip-address=<ip-address> --iscsi-target-name=<iscsi-target-name>
[--lun=<lun>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates and assigns a new iSCSI target for the specified client.

-c (--client-name) is required to specify the client name. The client type must be iSCSI.

-I (--ip-address) is required to specify the IP address of the target.

-R (--iscsi-target-name) is required to specify the target name. Valid characters are: a to z, 0 to 9, and .

-l (--lun) is an option to specify the starting LUN ID.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Assign an iSCSI target

```
iscon assigniscsiclienttarget -s <server-name> [-u <username> -p <password>]
-c <client-name> -r <iscsi-target-id>
[-X <rpc-timeout>]

iscon assigniscsiclienttarget --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> --iscsi-target-id=<iscsi-target-id>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates and assigns a new iSCSI target for the specified client.

-c (--client-name) is required to specify the client name.

-r (--iscsi-target-id) is required to specify the target ID.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete an iSCSI target

```
iscon deleteiscsiclienttarget -s <server-name> [-u <username> -p <password>]
-c <client-name> -r <iscsi-target-id> [-X <rpc-timeout>]

iscon deleteiscsiclienttarget --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name>
--iscsi-target-id=<iscsi-target-id> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes the specified iSCSI target. All virtual devices must be unassigned from the target prior to running the command.

-c (--client-name) is required to specify the client name.

-r (--iscsi-target-id) is required to specify the target ID.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Virtual libraries and drives

Get supported virtual tape libraries

```
iscon getsupportedvlibs -s <server-name> [-u <username> -p <password>]
[-l [-t <vlib-type>] [-c][-M <output-delimiter>] ] [-X <rpc-timeout>]

iscon getsupportedvlibs --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--longlist [--vlib-type=<vlib-type>] [--compatible-drive-list]
[--output-delimiter=<output-delimiter>] ] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves information about all supported virtual tape libraries.

-l (--longlist) can be specified to get the supported library information in a long format. The default is to display the information in a list format.

-t (--vlib-type) is an option with the -l (--longlist) option to get the detail library information for a specific library. The format for the <vlib-type> is: <vendorID>:<productID>. For example, ADIC:Scalar 100

-c (--compatible-drive-list) is an option to display the compatible drives in a tabular format instead of the default long format.

-M (--output-delimiter) can also be specified with the -l (--longlist) option to replace the linefeed with the specified delimiter. The maximum length of the delimiter is 8.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get supported virtual drives

```
iscon getsupportedvdrives -s <server-name> [-u <username> -p <password>]
[-l [-M <output-delimiter>] ] [-X <rpc-timeout>]

iscon getsupportedvdrives --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--longlist [--output-delimiter=<output-delimiter>] ] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves information about all supported virtual tape drives.

-l (--longlist) can be specified to get the supported drive information in a long format. The default is to display the information in a list format.

-M (--output-delimiter) can be specified when -l is specified to replace the linefeed with the specified delimiter. The maximum length of the delimiter is 8.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create virtual tape library

```
iscon createvirtuallibrary -s <server-name> [-u <username> -p <password>]
-t <vlib-type> [-n <vlib-name>] -d <vdrive-type> [-r <vdrive-name-prefix>]
[-R <num-of-drives>] [-A <auto-archive-mode> [-Y <days>] [-J] | -N <auto-repl-mode>
-S <target-name> [-U <target-username> -P <target-password>] [-M <#[D|H|M]>] ]
[-B <barcode-range>] [-T <num-of-slots>] [-E <import-export-slots>]
[-D -I <initial-size> -C <increment-size>] [-m <max-capacity>] [-L <on|off>] [-f]
[-K <data-key-name> -G <data-key-password>] [-k <key-name> -W <key-password>]
[-X <rpc-timeout>]

iscon createvirtuallibrary --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vlib-type=<vlib-type> [--vlib-name=<vlib-name>] --vdrive-type=<vdrive-type>
[--vdrive-name-prefix=<vdrive-name-prefix>] [--num-of-drives=<num-of-drives>]
[--auto-archive-mode=<auto-archive-mode> [--delay-delete-days=<days>]
[--auto-eject-to-ie] | --auto-replication=<auto-repl-mode> --target-name=<target-name>
[--target-username=<target-username> --target-password=<target-password>]
[--delay-delete-time=<#[D|H|M]>] ] [--barcode-range=<barcode-range>]
[--num-of-slots=<num-of-slots>] [--import-export-slots=<import-export-slots>]
[--capacity-on-demand --initial-size=<initial-size> --increment-size=<increment-size>]
[--max-capacity=<max-capacity>] [--auto-loader=<on|off>] [--force]
[--data-key-name=<key-name> --data-key-password=<key-password>]
[--key-name=<key-name> --key-password=<key-password>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a virtual tape library with the specified configuration.

-t (--vlib-type) is required in the following format: “<vendorID>:<productID>”

-n (--vlib-name) is optional. A default name will be provided in the format of <vendorID>-<productID>-<vid> if it is not specified.

-d (--vdrive-type) is required to specify the type of tape drive to be created in the library. The format of <vdrive-type> is as follows: “<vendorID>:<productID>”

-r (--vdrive-name-prefix) is an option to specify the prefix of the virtual drive. The default prefix is in the format of <drive-vendorID>-<drive-productID>-<vid>.

-R (--num-of-drives) is an option to create the specified number of drives, up to the maximum allowed by the library. By default, the library will be created with 1 drive. Use -f (--force) to override the default maximum value for the specified library in order to create up to 256 drives.

-A (--auto-archive-mode) is an option with one of the following values: *copy* or *move*.

-Y (--delay-delete-days) is an option for *move* mode to specify the number of days to wait before deletion. The maximum is 365 days. The default value is 365 days.

-J (--auto-eject-to-ie) is an option to be specified with -A (--auto-archive-mode) to eject the tape to the import/export (IE) slot after the export job.

-N (--auto-replication) is an option with one of the following values: *replication* or *remotemove*.

-S (--target-name) is the target server name for auto-replication. It is required for auto-replication.

-U (--target-username) and -P (--target-password): Target credentials are required if the target server was not already connected to with the login command or if they are not the same as the primary server. If they are not provided, the primary server credentials will be used.

-M (--delay-delete-time) is an option for *remotemove* mode to specify a time to wait before deletion. It can be specified in days(D), hours(H) or minutes(M). For example, 2D, 10H, 150M

-B (--barcode-range) can be specified in the following format: <barcodeB>-<barcodeE>
Barcode is an alphanumeric value with a length of 4 to 12. <barcodeB> and <barcodeE> have to be the same length. <barcodeE> has to be greater than <barcodeB>. A default <barcode-range> will be generated if it is not specified.

-T (--num-of-slots) and -E (--import-export-slots) are optional.

The (--num-of-slots) can exceed the maximum number of slots supported by the specified library type, but it is limited to 64,000.

The (--import-export-slots) cannot exceed the maximum number of IE slots supported by the specified library type. The default is to use the maximum number of slots supported by the specified library type.

-D (--capacity-on-demand) is an option to expand the virtual tape when needed. The default is to create the virtual tape with the maximum capacity if it is not specified.

-I (--initial-size) and -C (--increment-size) are options to be specified with <capacity-on-demand> option. The default value for both options is 5 GB. The (--increment-size) cannot be less than 5 GB.

-m (--max-capacity) is an option to set the maximum capacity of the virtual tapes, up to the maximum value allowed by the library. Use -f (--force) to override the default maximum value for the specified library in order to set the value up to 1,800 GB.

The unit of <max-capacity>, <initial-size>, and <increment-size> are all in GB.

-L (--auto-loader) is an option to set the auto-loader for those libraries that support the feature. The default value is *off*.

-f (--force) is an option to override the maximum default values for the specified library and allow up to a maximum of 256 drives and 1,800 GB of tape capacity.

-K (--data-key-name) -G (--data-key-password) are options for tape data encryption on disk storage. All newly created tapes in this library will be encrypted using this key.

-k (--key-name) and -W (--key-password) are options for tape encryption support to be set in conjunction with Auto-Archive Mode. Specify the key name and key password of the encryption key if you wish to encrypt the data when exporting the virtual tape to the physical tape.

A virtual device ID will be assigned to the virtual library when it is created successfully.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete virtual tape library

```
iscon deletevirtuallibrary -s <server-name> [-u <username> -p <password>]
-v <vdevid> [-d]
[-X <rpc-timeout>]

iscon deletevirtuallibrary --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> [--delete-virtual-tapes]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes a virtual tape library if there are no clients currently connected to it.

-v (--vdevid) is required to specify the device virtual ID.

-d (--delete-virtual-tapes) is an option to delete all of the existing virtual tapes from the virtual tape library. If not specified, the virtual tapes are moved to the vault.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Set virtual tape library duplication

```
iscon setvirtuallibrarytapeduplication -s <server-name> [-u <username> -p <password>]
-v <vdevid> -Z <on|off> [-Q <num-of-copies> -q <A:B:C>][-X <rpc-timeout>]

iscon setvirtuallibrarytapeduplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> --tape-duplication=<on|off>
[--num-of-copies=<num-of-copies> --prefix=<A:B:C>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command sets the Tape Duplication property for a virtual tape library.

-v (--vdevid) is required in order to identify the virtual library.

-Z (--tape-duplication) is required in order to enable or disable the Tape Duplication property: *on* (enable) or *off* (disable).

-Q (--num-of-copies) is an option to specify the number of copies to be made using the same barcode if the tape duplication option is enabled. The maximum value is 5. The default value is 1.

-q (--prefix) is an option to add a single character prefix to the barcode of each copy, in the following format: A:B:C

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Add virtual tape drive

```
iscon addvirtualdrive -s <server-name> [-u <username> -p <password>]
-L <tape-library-vid> [-r <vdrive-name-prefix>] [-R <num-of-drives>] [-X <rpc-timeout>]

iscon addvirtualdrive --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-library-vid=<tape-library-vid> [--vdrive-name-prefix=<vdrive-name-prefix>]
[--num-of-drives=<num-of-drives>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds a virtual tape drive to a specific virtual tape library.

-L (--tape-library-vid) is required to specify the virtual tape library to add the virtual tape drive(s).

-r (--vdrive-name-prefix) is an option to specify the prefix of the virtual tape drive. The default prefix is in the format of <drive-vendorID>-<drive-productID>-<vid>.

-R (--num-of-drives) is optional, the default is 1 if it is not specified.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create standalone tape drive

```
iscon createstandalonedrive -s <server-name> [-u <username> -p <password>]
-d <vdrive-type> [-r <vdrive-name-prefix>] [-R <num-of-drives>]
[-D -I <initial-size> -C <increment-size>] [-m <max-capacity>] [-X <rpc-timeout>]

iscon createstandalonedrive --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdrive-type=<vdrive-type> [--vdrive-name-prefix=<vdrive-name-prefix>]
[--num-of-drives=<num-of-drives>] [--capacity-on-demand --initial-size=<initial-size>
--increment-size=<increment-size>] [--max-capacity=<max-capacity>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a standalone virtual tape drive.

-d (--vdrive-type) is required to specify the type of tape drive to be created in the following format: <vendorID>:<productID>

-r (--vdrive-name-prefix) is an option to specify the prefix of the virtual drive. The default prefix is in the format of <drive-vendorID>-<drive-productID>-<vid>.

-R (--num-of-drives) can be specified to create multiple drives of the same type. The default is 1 if it is not specified. The maximum number of drives is 10.

-D (--capacity-on-demand) is an option to expand the virtual tape when needed. The default is to create the virtual tape with the maximum capacity if it is not specified.

-I (--initial-size) and -C (--increment-size) are options to be specified with <capacity-on-demand> option.

-m (--max-capacity) is an option to specify the maximum capacity of the virtual tape. The maximum capacity configured for the specified type of virtual drive will be used if it is not specified.

The unit of <max-capacity>, <initial-size> and <increment-size> are all in GB.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete virtual tape drive

```
iscon deletevirtualdrive -s <server-name> [-u <username> -p <password>]
-v <vdevid> [-d]
[-X <rpc-timeout>]

iscon deletevirtualdrive --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> [--delete-virtual-tapes]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes a standalone virtual tape drive or a virtual tape drive from a library, if the following conditions are satisfied:

- there are no clients connected to the drive.
- the specified virtual device is not the only existing virtual tape drive in the parent virtual tape library.
- the virtual tape drive has the highest element number in the parent virtual tape ibrary.

-v (--vdevid) is required to specify the device virtual ID.

-d (--delete-virtual-tapes) is an option to delete the loaded virtual tape from the virtual tape drive. If not specified and the specified device is a standalone virtual tape drive, the virtual tape is moved to the vault, otherwise the tape is moved back to the parent virtual tape library.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Virtual tapes

Get virtual tape information

```
iscon getvirtualtapeinfo -s <server-name> [-u <username> -p <password>]
[-L <parent-library-id] [-B <barcode>] [-X <rpc-timeout>]

iscon getvirtualtapeinfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--tape-library-vid=<tape-library-vid>] [--barcode=<barcode>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command displays information about the specified virtual tapes, in CSV format. By default, the command reports all virtual tapes found in the VTL system that are located in virtual libraries.

-L (--tape-library-vid) is an option to choose a single virtual library to be queried and report on only those virtual tapes that are located in this library.

-B (--barcode) is an option to only display information about the virtual tape that is specified by this barcode. The tape must be in a virtual tape library.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create virtual tape

```
iscon createvirtualtape -s <server-name> [-u <username> -p <password>] -v <parent-vid>
[ [-g <#(GB)> [-I <ACSL>] ] [-n <vdevname>] [-B <barcode | barcode-range>] -t <count>]
[-A -l <plib-vid> -b <physical-tape-barcode> [-J] | -N [-S <target-name>]
[-U <target-username> -P <target-password>] [-X <rpc-timeout>]

iscon createvirtualtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--parent-vid=<parent-vid> [ [--size-gb=<#(GB)>] [--scsiaddress=<ACSL>] ]
[--vdevname=<vdevname>] [--barcode=<barcode | barcode-range>] [--count=<count>]
[--enable-auto-archive --plib-vid=<plib-vid>
--physical-tape-barcode=<physical-tape-barcode>
[--auto-eject-to-ie] | --enable-auto-remotecopy
[--enable-auto-replication
--target-name=<target-name> [--target-username=<target-username>
--target-password=<target-password>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a virtual tape with the specified configuration.

-v (--parent-vid) is the virtual device id of the virtual library or standalone drive.

-g (--size-gb) is an option to specify the size in GB. The size of the virtual tape will be the size configured in the properties of the virtual library or virtual drive if it is not specified.

-I (--scsiaddress) is an option to specify preferred physical devices for creating a virtual device. It can be a list of ACSLs separated by a comma or a file enclosed in <> containing an ACSL on each line.
ACSL=#:#:#:# (adapter:channel:id:lun)

-n (--vdevname) is an option to specify the virtual tape name or prefix when creating more than one tape. The maximum length of the virtual device name is 64. Leading and trailing spaces will be removed. Enclose the name in double quotes to ensure the proper name. The following characters are invalid for the name: <>"&$/\'

-B (--barcode) is an option to either set the virtual tape with the provided barcode or create virtual tapes in batch mode configured with barcodes form the specified barcode range. The argument must be within the barcode range configured for the library and must not contain used barcodes. When provided as a barcode range, the option creates a virtual tape for each barcode in the range.

-t (--count) is an option to create multiple virtual tapes having the barcode automatically chosen from within the barcode range configured at library level. The library must have the required number of free slots available. If combined, "count" and "barcode" options must agree in number.

If the parent library has the auto-archive/remotecopy property enabled, use the following options to provide additional information for virtual tape creation:

-A (--enable-auto-archive) is an option when the parent library is enabled with auto-archive option.

-l (--plib-vid) is required when <auto-archive-mode> is specified. It is the physical tape library where the tape will be exported to automatically.

-b (--physical-tape-barcode) is required to specify the list of physical tape barcode(s) when auto-archive option is specified. Separate multiple barcodes with commas. For example, -b 00010001,00010009,0001000A

-J (--auto-eject-to-ie) is optional when <auto-archive-mode> is specified.

-N (--enable-auto-replication) is an option when the parent library is enabled with the auto-replication option.

-S (--target-name) can be specified when auto-replication option is specified. The default remote server from the parent library will be used if it is not specified.

The *count* and *barcode* options cannot be specified when the -A (--enable-auto-archive) option is specified because the number of tapes will be obtained from the list of barcodes specified with -b (--physical-tape-barcode) option.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Set tape properties

```
iscon settapeproperty -s <server-name> [-u <username> -p <password>]
-v <vdevid> [-B <barcode>] [-f] [-F] [-w <on|off>] [-A <auto-archive-mode> [-Y <days>]
[-J <on|off>] | -N <auto-repl-mode> -S <target-name>
[-U <target-username> -P <target-password>] [-M <#[D|H|M]>] ]
[-k <key-name> -W <key-password> | -d] [-Z <on|off> -Q <num-of-copies>]
[-X <rpc-timeout>]

iscon settapeproperty --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --vdevid=<vdevid>
[--barcode=<barcode>] [--force] [--full-capacity] [--tape-write-protect=<on|off>]
[--auto-archive-mode=<auto-archive-mode> [--delay-delete-days=<days>]
[--auto-eject-to-ie] | --auto-replication=<auto-replication-mode>
--target-name=<target-name>
[--server-username=<username> --server-password=<password>]
[--delay-delete-time=<#[D|H|M]>] ]
[--key-name=<key-name> --key-password=<key-password> | --disable-key]
[--tape-duplication=<on|off> --num-of-copies=<num-of-copies>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command configures tape properties for the specified virtual tape. The tape must be located in a virtual tape library slot. If the specified virtual tape is in the vault, only the write protection property can be configured.

-v (--vdevid) is required to specify the ID of the virtual tape to set the properties.

-B (--barcode) is an option to specify the new barcode for the tape. -f (--force) option is required if the new barcode is not in the barcode range specified for the parent library. Barcode is an alphanumerical value in the length of 4 to 12.

-F (--full-capacity) is an option to expand the tape to the maximum capacity and turn off the <capacity-on-demand> option if it is enabled for the virtual tape.

-w (--tape-write-protect) is an option to turn on and off the tape write protection with the following values: *on* (enable) or *off* (disable).

-A (--auto-archive-mode) is an option with one of the following values: *copy* or *move* or *inherited* or *none*.

- "none" is the value to turn off the auto-archive mode if the virtual tape is enabled with the auto-archive option.
- "inherited" can only be specified when the parent library is enabled with the auto-archive option.

-Y (--delay-delete-days) is an option for auto-archive *move* mode to specify up to 365 days to wait before the deletion. The default value is 365 days.

-J (--auto-eject-to-ie) is an option for auto-archive mode to eject the physical tape to the IE slot after a successful archive job: *on* (enable) or *off* (disable).

-N (--auto-replication) is an option in one of the following values: *localcopy*, *localmove*, *remotecopy, remotemove*, or *none*.

-S (--target-name) is the remote server name for auto-replication. It is required for auto-replication.

-U (--target-username) and -P (--target-password) are options to specify a different user ID and password to log in to the remote server.

-M (--delay-delete-time) is an option for auto-replication move mode to specify up to 30 days to wait before deletion. The default value is 1 day. The value can be specified in days(D), hours(H) or minutes(M). For example, 2D, 10H, 150M

-A (--auto-archive-mode) and -N (--auto-replication) cannot be specified if replication is enabled for the tape.

-k (--key-name), -W (--key-password) and -d (--disable-key) are options for tape encryption support to be set in conjunction with Auto-Archive Mode. Specify the key name and key password of the encryption key if you wish to encrypt the data when exporting the virtual tape to physical tape. Specify -d (--disable-key) if you wish to disable tape encryption for this tape.

-Z (--tape-duplication) is an option to set the Tape Duplication property with one of the following values: *on* (enable), *off* (disable), or *inherited*.

-Q (--num-of-copies) is an option to specify the number of copies to be made using the same barcode if the tape duplication option is enabled. The maximum value is 5. The default value is 1.

At least one of the above properties has to be specified.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete virtual tape

```
iscon deletevirtualtape -s <server-name> [-u <username> -p <password>]
[-v <vdevid> ] | [-B <barcode> -l <lib/sa_drive ID | 0 (Vault)>]
[-f]
[-X <rpc-timeout>]

iscon deletevirtualtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--vdevid=<vdevid>] |
[--barcode=<barcode> --from-location-id=<lib/sa_drive ID | 0 (Vault)>]
[--force]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes a virtual tape.

To delete a virtual tape, specify either the -v (--vdevid) or the -B (--barcode) of the tape, as they are mutually exclusive. You can also specify the -l (--from-location-id) option.

-v (--vdevid) is an option to specify the tape virtual ID.

-B (--barcode) is an option to specify the barcode of the virtual tape. By default, the command queries all libraries, drives, and the vault. The barcode must be unique. If you have duplicate barcodes, use -l (--from-location-id) to narrow the search. If the tape's -v (--vdevid) is provided, the barcode and location ID options are ignored.

-l (--from-location-id) is an option to specify the virtual ID of the library or standalone drive where the virtual tape is located when you use the -B (--barcode) option. If the tape is located in the vault, use 0 for the location ID.

-f (--force) is an option to force the deletion of a virtual tape configured for replication. The corresponding virtual tape replica will not be deleted or promoted.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Move virtual tape

```
iscon movevirtualtape -s <server-name> [-u <username> -p <password>]
-v <vdevid> | -B <barcode> [-i]
[-L <tape-library-vid> | -D <tape-drive-vid> | -l <slot-no>] [-X <rpc-timeout>]

iscon movevirtualtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> | --barcode=<barcode> [--include-filter]
[--tape-library-vid=<tape-library-vid> | --tape-drive-vid=<tape-drive-vid> |
--slot-no=<slot-no>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command moves a virtual tape to a different location.

-v (--vdevid) or -B (--barcode) is required to identify the virtual tape to be moved to a different location.

-i (--include-filter) is an optional filter that can be used to uniquely identify a virtual tape when multiple tapes have the same barcode and -B (--barcode) is used. This option can be one of the following values:

- TapeName="*"
- Location=*
- ParentID=#

"Location" is the current location: the library ID if the tape is in a slot, the drive ID, or Vault. "ParentID" is the ID of the last library that hosted the tape and it is preserved when the tape is moved to the vault. If the tape cannot be uniquely identified, the command will fail.

-L (--tape-library-vid) is the virtual library to move to. It is not required if the virtual tape is moved within the same library.

-D (--tape-drive-vid) is the virtual drive in a library or the standalone drive to move to.

-l (--slot-no) is the slot in a library to move to.

If none of the above locations are specified, the vault will be assumed to be the new location.

If the tape is in a slot in a library, it can be moved to a different slot or a drive in the library, or it can be moved to the vault.

- Vlib Slot -> Tape drive (in the library only)
- Vlib Slot -> Slots in same library
- Vlib Slot -> Vault

If it is in a drive in the library, it can be moved to an available slot in the library or to the vault.

- Vlib Drive -> Slots in same library
- Vlib Drive -> Vault

If the tape is in a standalone drive, it can only be moved to the vault.

- Standalone Tape Drive -> Vault

If the tape is in the vault, it can be moved to an available slot in a library, or an available standalone drive.

- Vault -> Vlib (First available slot)
- Vault -> Standalone Tape Drive

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Note: If you are moving virtual tapes from within a script, be sure to include the appropriate delays, as it can take several seconds to complete the move. During this time, the tape is still considered as being in its original slot.

Tape copy

```
iscon tapecopy -s <server-name> [-u <username> -p <password>]
-v <source-vdevid> -S <target-name> [-U <target-username> -P <target-password>] | [-h]
[-L <tape-library-vid> | -D <tape-drive-vid>] [-n <vdevname>] [-f]
[-X <rpc-timeout>]

iscon tapecopy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--source-vdevid=<source-vdevid> --target-name=<target-name>
[--target-username=<target-username> --target-password=<target-password>] | [--local]
[--tape-library-vid=<tape-library-vid> | --tape-drive-vid=<tape-drive-vid>]
[--vdevname=<vdevname>] [--force] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a copy of the specified virtual tape. The data is transferred through a replication job.

-v (--source-vdevid) is required to specify the ID of the virtual tape to be copied from.

-S (--target-name) is required to specify the target server name where the remote tape copy will be created and copied to. If the replication is local, use the -h (--local) option.

-U (--target-username) and -P (--target-password) are optional for connection and login to the target server if the target server was not logged in with login command.

-h (--local) is an option to create a local tape copy. Target server information and credentials are not required when using this option and are ignored if they are specified.

-L <tape-library-vid> and -D <tape-drive-vid> are options to move the tape copy to the virtual tape library or virtual tape drive when the copy is completed.

-n (--vdevname) is an option to specify the virtual tape name of the tape copy. The maximum length of the virtual device name is 64. Leading and trailing spaces will be removed. Enclose the name in double quotes. The following characters are invalid for the name: <>"&$/\'

A default name with the primary server and source virtual tape name will be generated if it is not specified.

-f (--force) option is required when the tape is scheduled to be deleted. The deletion schedule for the virtual tape will be removed and the replication will be configured.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Shred virtual tape

```
iscon shredvirtualtape -s <server-name> [-u <username> -p <password>]
-B <barcode> | -v <vid> [-d] [-X <rpc-timeout>]

iscon shredvirtualtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--barcode=<barcode> | --vdevid=<vid> [--delete-virtual-tapes]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes the data stored on the specified virtual tapes located in the vault. Either the barcode or the virtual tape ID can be used in order to identify the tapes. When barcode identification is used, the command will shred all of the virtual tapes that share the same barcode. The format for the identification arguments is a list of items separated by commas.

-B (--barcode) can be used to specify the virtual tapes by barcode.

-v (--vdevid) can be used to specify the virtual tapes by ID.

-d (--delete-virtual-tapes) is an option to delete the virtual tapes after the shredding operation is executed.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Physical libraries and drives

Inventory physical tape library

```
iscon plibinventory -s <server-name> [-u <username> -p <password>]
[-l <physical-tape-library-vid>] [-X <rpc-timeout>]

iscon plibinventory --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--plib-vid=<tape-library-vid>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command refreshes the physical tape information for the specified physical library if the information is out of sync.

-l (--plib-vid) is an option to specify the physical tape library to perform the inventory.

Inventory operation will be performed for all the physical tape libraries if -l (--plib-vid) is not specified.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Mark physical library or drive enabled/disabled

```
iscon markphysicallibdrvstate -s <server-name> [-u <username> -p <password>]
-v <plib-or-pdrive-vid> [-E <on | off>] [-X <rpc-timeout>]

iscon markphysicallibdrvstate --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--plib-or-pdrive-vid=<plib-or-pdrive-vid> [--state=<on | off>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command marks a physical tape library or drive as enabled or disabled.

-v (--plib-or-pdrive-vid) is required to specify the ID of the device.

-E (--state) is required to specify the new state for the device. If the argument is not specified, the command retrieves the current state for the device.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Physical tapes

Move physical tape

```
iscon movephysicaltape -s <server-name> [-u <username> -p <password>]
-m <move-operation> -L <physical-tape-library-vid>
-B <physical-tape-barcode> | -l <from-location-id> -t <to-location-id>
[-X <rpc-timeout>]

iscon movephysicaltape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--move-operation=<move-operation> --tape-library-vid=<physical-tape-library-vid>
--physical-tape-barcode=<barcode> | --from-location-id=<from-location-id>
--to-location-id=<to-location-id> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command moves a physical tape to a new location.

-l (--plib-vid) is the ID of the physical tape library.

-m(--move-operation) is one of the following operations:

- DriveToSlot
- SlotToSlot
- SlotToDrive
- IESlotToSlot
- SlotToIESlot

-L(--tape-library-vid) is the physical library virtual ID where the tape is located.

-B(--physical-tape-barcode) identifies the physical tape to be moved. If barcode is not provided, the current tape location must be provided accordingly to the requested operation.

-l(--from-location-d) is the current slot or import/export (IE) slot number, or the physical drive virtual ID.

-t(--to-location-id) is the destination slot or IE slot number or the physical drive virtual ID.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Eject physical tape

```
iscon ejectphysicaltape -s <server-name> [-u <username> -p <password>]
-L <physical-tape-library-vid> -B <physical-tape-barcode-list>
[-A <acs-lsm-cap>] [-X <rpc-timeout>]

iscon ejectphysicaltape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-library-vid=<physical-tape-library-vid>
--tape-barcode-list=<physical-tape-barcode-list> | [--acs-lsm-cap=<acs-lsm-cap>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command ejects physical tapes from the specified library.

-L(--tape-library-vid) is the physical library virtual ID where the tapes are located.

-B(--tape-barcode-list) identifies the physical tapes to be ejected. This argument can be a list of barcodes separated with commas. The list should be enclosed in quotes.

-A <--acs-lsm-cap> is an optional argument representing the Cartridge Access Port for the Automated Cartridge System Library Software libraries. The format of the argument is acs:lsm:cap

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Prepare physical tape for stacking

```
iscon preparestackingptape -s <server-name> [-u <username> -p <password>]
-v <standalone-physical-tape-drive-vid> -b <physical-tape-barcode> -O [-X <rpc-timeout>]

iscon preparestackingptape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--pdrive-vid=<physical drive id> --physical-tape-barcode=<barcode>
--overwrite-mode [--rpc-timeout=<rpc-timeout>]
```

Description:

This command prepares a physical tape for stacking. If the tape is not blank, all existing information will be deleted in the process.

-v (--pdrive-vid) is required to provide the ID of the physical tape drive where the physical tape is loaded.

-b (--physical-tape-barcode) is required to provide a barcode for the physical tape. This barcode will replace the old barcode. The barcode must be between 4 and 12 characters.

-O (--overwrite-mode) is required in order to confirm that the existing data will be erased.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Scan stacked physical tape

```
iscon retrieveptapebarcode -s <server-name> [-u <username> -p <password>]
-v [-X <rpc-timeout>]

iscon retrieveptapebarcode --server-name=<server-name>
[--server-username=<sername> --server-password=<password>]
--pdrive-vid=<physical drive id> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command scans a stacked physical tape loaded in the specified standalone physical tape drive in order to retrieve the barcode from the tape header. The command does not wait for the operation to complete. Use "getvdevlist" with "-v" and "-l" to get the barcode displayed.

-v (--pdrive-vid) is required to provide the ID of the physical tape drive where the physical tape is loaded.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Virtual tape caching

Set tape caching

```
iscon settapecaching -s <server-name> [-u <username> -p <password>]
-L <library-vid> -t <tape-caching-enable> -m <delete-cache | data-deduplication>
-T <# of days> [-H <hours>] [-S <start-time>] [-W <day-of-the-week>][-b <and-or>]
[-d <# of hours>][-c][[-e][-f]] [[-I | -M | -R <# of days> | -N] [-X <rpc-timeout>]

iscon settapecaching --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-library-vid=<library-vid> --tape-caching-enable=<tape-caching-enable>
[--reclamation-method=<delete-cache|data-deduplication>]
[--turn-into-directlink=<# of days>
[--hourly <hours>] [--start-time=<start-time>] [--day-of-the-week=<day-of-the-week>]
[--trigger-combine=<and-or>] [--retention-hours=<# of hours>]
[--migration-threshold] [[--tape-ejected-to-slot] [--tape-full]]
[--immediately | --reclamation-threshold | --retention-days=<# of days> | --never]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command can be used in order to enable, disable, or change the Automated Tape Caching policy for a virtual tape library. The reclamation method used is "delete cache".

-L (--tape-library-vid) is the virtual device ID of the virtual tape library to be set.

Set -t (--tape-caching-enable) to 1 for enable or 0 to disable. If the *disable* option is used, all other arguments will be ignored. The *enable* option must be used in order to set or change the tape caching policy.

-m (--reclamation-method) is an option to select from the following reclamation methods: delete-cache (default method) and data-deduplication.

Time-based data migration triggers: Time Based

-H (--hourly) is an option to start migration every specified number of hours, up to 23. This option cannot be combined with any other migration trigger.

-S (--start-time) alone can be used to start daily migrations at the time specified. When combined with other data migration triggers, the *-S* option will delay the migration execution to the specified time.

-W (--day-of-the-week) is the default time based trigger that can be used to start weekly migrations on the specified day at 00:00(am): Sunday: 0, Monday: 1, ..., Saturday: 6. This option is ignored if Policy Based triggers are used.

Policy-based data migration triggers:

-b (--trigger-combine) tells how trigger policies are combined (specified by -d, -c, -e). 1 -- and; 0 -- or. The default value is 1 (and).

-d (--retention-hours) triggers the data migration after the data was retained on the disk for the specified number of hours, up to a year.

-c (--migration-threshold) triggers the data migration when the disk usage percentage is above the global disk space threshold.

-e (--tape-ejected-to-slot) triggers the data migration when unloading a virtual tape from a drive that had data written to it.

-f (--tape-full) applies to the --tape-ejected-to-slot trigger. The data is migrated only if the tape becomes full.

Reclamation triggers:

-I (--immediately) triggers the data reclamation immediately after the data migration completes.

-M (--reclamation-threshold) triggers the data reclamation when the disk usage percentage is above the global disk space threshold.

-R (--retention-days) triggers the data reclamation after the specified number of days, up to 2000.

-N (--never) means that virtual tapes will never be reclaimed.

The reclamation triggers are mutually exclusive.

The command always overwrites the current policy. Therefore, all of the desired properties must be provided each time the command is executed. The following properties are set by default if no triggers are provided:

- Migration trigger: daily at 00:00(am)
- Reclamation trigger: immediately

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get physical tape list

```
iscon getphysicaltapelist -s <server-name> [-u <username> -p <password>]
-l <physical-tape-library-vid> | -v <standalone-physical-tape-drive-vid>
[-F <filter>] [-S] [-X <rpc-timeout>]

iscon getphysicaltapelist --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--plib-vid=<physical-tape-library-vid> | --pdrive-vid=<standalone-physical-tape-drive-vid>
[--ptape-filter=<filter>] [--stacking-info]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command displays a list of physical tapes located in the specified physical tape library or standalone physical tape drive.

-l (--plib-vid) is the ID of the physical tape library.

-v (--pdrive-vid) is the ID of the standalone physical tape drive where the physical tape is located.

-F (--ptape-filter) is an option to show only the physical tapes having the specified property:

- SYNC (eligible for sync operation, applies to library only) OR
- STACKED (stacked tape).

-S <--stacking-info> is an option to retrieve and display the stacking information of the stacked physical tapes when the stacked tape filter is used.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Sync physical tapes

```
iscon syncphysicaltape -s <server-name> [-u <username> -p <password>]
-l <plib-vid> -b <physical-tape-barcode> -L <virtual-tape-library-id>
-t <virtual-tape-slot-no> [-M <sync-mode>] [-k <key-name> -W <key-password>]
[-I <ACSL list>] [-n <vdevname>] [-g <#(GB)>] [-X <rpc-timeout>]

iscon syncphysicaltape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--plib-vid=<physical-tape-library-vid> --physical-tape-barcode=<physical-tape-barcode>
--tape-library-vid=<virtual-tape-library-id>
--virtual-tape-slot-no=<virtual-tape-slot-no> [--sync-mode=<sync-mode>]
[--key-name=<key-name> --key-password=<key-password>] [--scsiaddress=<ACSL list>]
[--vdevname=<vdevname>] [--size-gb=<#(GB)>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a synchronized virtual tape for each physical tape provided. The physical tapes must be from the specified physical tape library and the virtual tape will be created in the specified virtual tape library. The virtual tape library must have the tape caching feature enabled.

-l (--plib-vid) is the virtual ID of the physical tape library where the physical tapes are located.

-b (--physical-tape-barcode) is the barcode of the physical tape. If the barcode contains leading or trailing spaces, it must be enclosed in double quotes. For batch mode, the argument can be a list of barcodes separated by commas or a file name enclosed in "< >" (i.e., "<file>") containing a barcode on each line. Do not use quotes inside the file. The file must be located in the same folder as the command line utility or a full path is required. The virtual tape(s) will be created with the same barcode as the physical tape(s). The barcode(s) must not be in use by any other virtual tape in the system.

-L (--tape-library-vid) is the ID of the virtual tape library where the virtual tapes will be created.

-t (--virtual-tape-slot-no) is an option to provide an empty destination slot for the virtual tape. Not for "-M cache" mode. For batch mode, virtual tapes will be created starting with the specified slot number.

-M (--sync-mode) is an option to select the synchronization mode from one of the following values (default is “cache"):

- cache (create cache)
- metadata (create cache and copy meta data)
- directlink (create direct link)

-k (--key-name) and -W (--key-password) are options for tape encryption support. If the tape to be synchronized was encrypted through the system, you need to specify the key name and the key password of the encryption key to decrypt the data.

The following three options can be selected for create cache mode only:

-I (--scsiaddress) is an option to specify preferred physical devices for creating the virtual device. It can be a list of ACSLs separated with commas. ACSL=#:#:#:# (adapter:channel:id:lun)

-n (--vdevname) is an option to specify the virtual tape name or prefix when creating more than one tape. The maximum length of the virtual device name is 64. Leading and trailing spaces will be removed. Enclose the name in double quotes to ensure the proper name is used. The following characters are invalid for the name: <>"&$/\'

-g (--size-gb) is an option to specify the initial size, in GB, of the virtual tapes, if the capacity-on-demand property for the virtual tape library is enabled. The default is 1 GB.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Migrate virtual tapes

```
iscon migratevirtualtapes -s <server-name> [-u <username> -p <password>]
-T <tape-vid-list> [-f] [-X <rpc-timeout>]

iscon migratevirtualtapes --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-vid-list=<tape-vid-list> [--tape-full] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command migrates the specified virtual tapes to the physical libraries they are synchronized with.

-T (--tape-vid-list) is a list of virtual tape ID(s) separated with commas.

-F (--tape-full) is an option to force full tape migration. By default, the migration operation is incremental.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Reclaim disk space

```
iscon reclaimtapes -s <server-name> [-u <username> -p <password>]
-T <tape-vid-list> [-X <rpc-timeout>]

iscon reclaimtapes --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-vid-list=<tape-vid-list> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command reclaims the disk space occupied by the specified migrated virtual tapes.

-T (--tape-vid-list) is required to specify the ID of the virtual tapes to be reclaimed, separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Renew cache

```
iscon renewcache -s <server-name> [-u <username> -p <password>]
-v <vdevid> [-M <metadata>] [-k <key-name> -W <key-password>] [-I <ACSL>] [-n <vdevname>]
[-g <#(GB)>] [-X <rpc-timeout>]

iscon renewcache --server-name=<server-name>
[--server-username=<username> --server-password=<password>] [--import-mode=<metadata>]
[--key-name=<key-name> --key-password=<key-password>] [--scsiaddress=<ACSL>]
[--vdevname=<vdevname>] [--size-gb=<#(GB)>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command converts a virtual stub tape into a virtual cache tape.

-v (--vdevid) is required to specify the ID of the virtual direct access tape.

-M (--import-mode) is an option to specify that the header area should be copied from the physical tape to the new virtual tape cache. The value of this option must be: *metadata*.

-k (--key-name) and -W (--key-password) are options for tape encryption support. If the tape to be renewed was encrypted through the system, you need to specify the key name and the key password of the encryption key to decrypt the data.

The following properties of the virtual cache tape can be set if the "-M" option is not specified:

-I (--scsiaddress) is an option to specify preferred physical devices for creating the virtual device. It can be a list of ACSLs separated with commas or a file enclosed in <> containing an ACSL on each line.
ACSL=#:#:#:# (adapter:channel:id:lun)

-n (--vdevname) is an option to specify the virtual tape name or prefix when creating more than one tape. The maximum length of the virtual device name is 64. Leading and trailing spaces will be removed. Enclose the name in double quotes to ensure it is properly parsed and interpreted. The following characters are invalid for the name:
<>"&$/\'

-g (--size-gb) is an option to specify the size in GB. The size of the virtual tape will be the size configured in the properties of the virtual tape library or virtual tape drive if it is not specified. This option cannot be specified if the capacity-on-demand option is not enabled at library level.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get tape caching info

```
iscon gettapecachinginfo -s <server-name> [-u <username> -p <password>]
[-X <rpc-timeout>]

iscon gettapecachinginfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command reports a list of virtual tapes that are candidates for tape migration and summarizes the information about the space used by those tapes.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Replication

Create a replica

```
iscon createreplication -s <server-name> [-u <username> -p <password>]
-v <source-vdevid> [-S <target-name> [-U <target-username> -P <target-password>]] | [-h]
[-w <watermark(MB)> | [-d <YYYYMMDDHHMM> -i <#[H|M]>]] [-r <on>]
[[-t <timeout>] [-I <retry-in]>] [-C <retry-for>]] [-c <on|off>] [-e <on|off>]
[-L <#:#:#:#>] [-n <replica-vdev-name>] [-X <rpc-timeout>]

iscon createreplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--source-vdevid=<source-vdevid> --target-name=<target-name>
[--target-username=<target-username> --target-password=<target-password>]] | [--local]
[--watermark=<watermark(MB)> | [--date=<YYYYMMDDHHMM> --interval=<#[H|M]>]] |
[--repl-first <on>] [[--replication-timeout=<timeout>]
[--replication-retry-interval=<retry-in>] [--replication-retry-count=<retry-for]]
[--compression=<on|off>] [--encryption=<on|off>] [--preferred-lun=#:#:#:#]
[--vdevname=<replica-name>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command sets up a tape replication configuration.

-v (--source-vdevid) is required to specify the ID of the virtual tape to be configured for replication.

-S (--target-name) is an option to specify the target server name where the tape replica will be created and replicated to. If the replication is local, use -H (--local) option.

-U (--target-username) and -P (--target-password) are optional for connection and login to the target server if the target server was not logged in with a login command.

-h (--local) is an option to create a local replica. Target server information and credentials are not required when using this option and are ignored if they are specified.

The replication configuration requires a trigger policy to be set. If no trigger policy is specified, the command will automatically apply the appropriate default policy based on the tape caching property of the specified virtual tape.

Any combination of the following two options can be used in order to set up a replication trigger policy for a virtual tape with the tape caching property disabled. The default policy is 1024 MB watermark.

-w (--watermark) is a data size based trigger in MB. The watermark is checked when the tape is unloaded from the tape drive and the replication is triggered if the amount of new data on the tape has reached the specified watermark.

-d (--date) combined with -i (--interval) is a time based trigger. The replication is triggered at the time specified by date and then repeated every interval. -d (--date) format is YYYYMMDDHHMM and -i (--interval) format is a number followed by H for hours or M for minutes (e.g. -i 2H or --interval=120M). The default value for interval is 1H (one hour).

For virtual tapes with tape caching enabled, replication is triggered based on the tape caching policy:

-r (--repl-first) is an option to replicate the virtual tape before it is migrated. Use *on* in order to enable this policy or *off* to have tape migration executed first. The default policy is to replicate the virtual tape after it is migrated.

Replication is retried based on the timeout policy:

- -t (--replication-timeout) in seconds (default 60).
- -I (--replication-retry-interval) in seconds (default 60).
- -C (--replication-retry-count) retry count (default 1).

-c (--compression) is an option for remote replication only and applies to compression of data during network transmission. Possible values are: *on* or *off*. This option cannot be used for encrypted virtual tapes.

-e (--encryption) is an option for remote replication only and applies to encryption of data during network transmission. Possible values are: *on* or *off.* This option cannot be used for encrypted virtual tapes.

-L (--preferred-lun) is an option to specify preferred physical devices for creating the virtual device. The format for this option is: #:#:#:# (adapter:channel:id:lun)

-n (--vdevname) is an option to specify the replica tape name.The maximum length of the device name is 64. Leading and trailing spaces will be removed. Enclose the name in double quotes to ensure proper parsing. The following characters are invalid for the name: <>"&$/\'"

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Promote a replica

```
iscon promotereplica -s <server-name> -v <vdevid> | -S <target-name> -V <replicaid>
[-u <username> -p <password>] [-U <target-username> -P <target-password>] [-f]
[-X <rpc-timeout>]
```

```
iscon promotereplica --server-name=<server-name> --vdevid=<vdevid> |
--target-name=<target-name> --replicaid=<replicaid> [--server-username=<username>
--server-password=<password>] [--target-username=<target-username>
--target-password=<target-password>] [--force] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command promotes a replica to a regular virtual device if the primary disk is available and the replica disk is in a valid state.

Specify either the primary server and the source virtual tape ID or the target server and the tape replica ID. The user name and password must be provided for both servers, if the servers were not registered using the login command.

-v (--vdevid) is the ID of the source virtual tape and -V (--replicaid) is the ID of the tape replica.

If the source virtual tape is still valid and available, and the tape replica is in an invalid state, the tape replica can be promoted with the force option. But, it is recommended to synchronize the tape replica with the source virtual tape first unless the source virtual tape is physically defective or unavailable.

If the source virtual tape is no longer available, the tape replica can be promoted with the force option -f (--force) even when it is in invalid state if you are sure the data on the tape replica is useful.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Remove replication

```
iscon removereplication -s <server-name> -v <vdevid> | -S <target-name> -V <replicaid>
[-u <username> -p <password>] [-U <target-username> -P <target-password>] [-f]
[-X <rpc-timeout>]

iscon removereplication --server-name=<server-name> --vdevid=<vdevid> |
--target-name=<target-name> --replicaid=<replicaid> [--server-username=<username>
--server-password=<password>] [--target-username=<target-username>
--target-password=<target-password>] [--force] [--rpc-timeout=<rpc-timeout>]
```

This command removes replication configuration from the specified source virtual tape and deletes the replica tape from the target.

Specify either the primary server and the source virtual tape ID or the target server and the tape replica ID. The user name and password must be provided for both servers, if the servers were not registered using the login command.

-v (--vdevid) is the ID of the source virtual tape and -V (--replicaid) is the ID of the tape replica.

Either the primary server with the source virtual tape or the target server with the tape replica can be specified to remove the replication configuration, but not both.

If the target server no longer exists or cannot be connected to, only the replication configuration on the primary server will be removed.

If the primary server no longer exists or cannot be connected to, only the tape replica will be deleted.

-f (--force) option has to be specified when either the primary server or target server no longer exists or cannot be connected.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Suspend replication

```
iscon suspendreplication -s <server-name> [-u <username> -p <password>]
-v <vdevid> [-X <rpc-timeout>]

iscon suspendreplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --vdevid=<vdevid>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command suspends scheduled replication for a virtual device that will be triggered by your replication policy. It will not stop a replication that is currently in progress.

-v (--source-vdevid) is the ID of the source virtual tape on the primary server to be suspended.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Resume replication

```
iscon resumereplication -s <server-name> [-u <username> -p <password>]
-v <vdevid> [-X <rpc-timeout>]

iscon resumereplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --vdevid=<vdevid>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command resumes scheduled replication for a virtual device that was suspended by the *suspendreplication* command. The replication will then be triggered by the replication policy once it is resumed.

-v (--source-vdevid) is the ID of the source virtual tape on the primary server to be resumed.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Set replication properties

```
iscon setreplicationproperties -s <server-name> [-u <username> -p <password>]
-v <source-vdevid> [-w <watermark(MB)> | [-d <YYYYMMDDHHMM> -i <#[H|M]>]] | [-r <on|off>]
[[-t <timeout>] [-I <retry-in]>] [-C <retry-for>]] [-c <on|off>] [-e <on|off>]
[-X <rpc-timeout>]

iscon setreplicationproperties --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--source-vdevid=<source-vdevid> [--watermark=<watermark(MB)>
[--watermark=<watermark(MB)> | [--date=<YYYYMMDDHHMM> --interval=<#[H|M]>]] |
[--repl-first <on|off>] [[--replication-timeout=<timeout>]
[--replication-retry-interval=<retry-in>] [--replication-retry-count=<retry-for]]
[--compression=<on|off>] [--encryption=<on|off>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command changes the replication policy for the specified virtual tape.

-v (--source-vdevid) is required to specify the ID of the source virtual tape.

Any combination of the following two options can be used to set up a replication trigger policy for a virtual with the tape caching property disabled.

-w (--watermark) is a data size based trigger in MB. The watermark is checked when the tape is unloaded from the tape drive and the replication is triggered if the amount of new data on the tape has reached the specified watermark.

-d (--date) combined with -i (--interval) is a time based trigger. The replication is triggered at the time specified by date and then repeated every interval. -d (--date) format is YYYYMMDDHHMM and -i (--interval) format is a number followed by H for hours or M for minutes (e.g. -i 2H or --interval=120M).

To delete a watermark trigger specify 0 for the watermark. To delete a time based trigger specify NA for date. At least one trigger must remain active.

The date argument is not required if you are only changing the interval.

For virtual tapes having the tape caching property enabled, the replication is triggered based on the tape caching policy:

-r (--repl-first) is required to replicate the virtual tape before it is migrated. Use "on" in order to enable this policy or "off" to have tape migration executed first.

The replication retry policy can be changed using the following options:

- -t (--replication-timeout) in seconds (default 60).
- -I (--replication-retry-interval) in seconds (default 60).
- -C (--replication-retry-count) retry count (default 1).

-c (--compression) is an option to enable or disable compression with one of the values: *on* or *off*.

-e (--encryption) is an option to enable or disable encryption with one of the values: *on* or *off*.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get replication properties

```
iscon getreplicationproperties -s <server-name> [-u <username> -p <password>]
-v <source-vdevid> [-X <rpc-timeout>]

iscon getreplicationproperties --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--source-vdevid=<source-vdevid> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command shows the replication configuration for the specified virtual tape.

-v (--source-vdevid) is required to specify the ID of the source virtual tape.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get replication status

```
iscon getreplicationstatus -S <target-name> [-U <username> -P <password>]
-v <replicaid> [-X <rpc-timeout>]

iscon getreplicationstatus --target-name=<target-name>
[--target-username=<username> --target-password=<password>]
--replicaid=<replicaid> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command shows replication status for the specified virtual replica tape.

-S (--target-name) is the target server and -v (--replicaid) is ID of the tape replica, both of which are required.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Start replication

```
iscon startreplication -s <server-name> [-u <username> -p <password>]
-v <vdevid> [-X <rpc-timeout>]

iscon startreplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command starts replication on demand for a virtual device.

-v (--source-vdevid) is the ID of the source virtual tape on the primary server to start.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Stop replication

```
iscon stopreplication -s <server-name> [-u <username> -p <password>]
-v <vdevid> [-X <rpc-timeout>]

iscon stopreplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
-vdevid=<vdevid> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command stops the replication that is in progress for a virtual device.

-v (--source-vdevid) is the ID of the source virtual tape on the primary server to stop.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Promote replica in test mode

```
iscon testmodepromotereplica -S <replica-server-name> -V <replicaid>
[-U <replica-server-username> -P <replica-server-password>]
[-u <primary-server-username> -p <primary-server-password>] [-X <rpc-timeout>]

iscon testmodepromotereplica
--target-name=<replica-server-name> --replicaid=<replicaid>
[--target-username=<replica-server-username>
--target-password=<replica-server-password>]
[--server-username=<primary-server-username>
--server-password=<primary-server-password>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command promotes a tape replica in test mode and suspends the replication property for its virtual tape source.

Both, tape replica and its virtual tape source must be valid and available. The information identifying the virtual source tape is automatically retrieved from the tape replica properties. If not already logged in, the user name and password must be specified for both replica and source servers.

-V (--replicaid) is the ID of the tape replica.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Demote replica in test mode

```
iscon testmodedemotetape -S <testmode-server-name> -V <testmode-tape-id>
[-U <testmode-server-username> -P <testmode-server-password>]
[-u <primary-server-username> -p <primary-server-password>] [-X <rpc-timeout>]

iscon testmodedemotetape --target-name=<testmode-server-name>
--testmode-tape-id=<testmode-tape-id> [--target-username=<testmode-server-username> --
target-password=<testmode-server-password> [--server-username=<primary-server-username>
--server-password=<primary-server-password>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command demotes a test mode virtual tape to a replica and resumes the replication property for its virtual tape source. The test mode virtual tape must be in the virtual vault.

Both the test mode virtual tape and its source virtual tape must be valid and available. The information identifying the source virtual tape is automatically retrieved from the test mode virtual tape properties. If not already logged in, the user name and password must be specified for both servers holding the virtual tapes.

-V (--testmode-tape-id) is the test mode virtual tape ID.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Import/Export

Get import/export job status

```
iscon getimportexportjobstatus -s <server-name> [-u <username> -p <password>]
[-j <job-id-list>] [-T <job-type> -S <job_status> -D <date-range|date> -l]
[-X <rpc-timeout>]

iscon getimportexportjobstatus --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--job-id-list=<job-id-list>] | [--job-type=<job_type> --job_status=<job_status>]
--date-range=<date-range|date> --longlist] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command displays the status of the import/export jobs present in the queue. If no filters are specified, the command displays all the jobs that are in the queue.

-j <--job-id-list> is an optional list of job IDs separated with commas. The command displays the status of specified jobs only. All other filters are ignored.

-T <--job-type> is an optional job type based filter. The command displays those jobs matching the provided type. The accepted job type values are: IMPORT, EXPORT, or OTHER (such as scan).

-S <--job_status> is an optional job status based filter. The command displays those jobs matching the provided status. The accepted job status values are: FAILED, HOLD, READY, or OTHER (such as waiting for tape/drive or cancelled).

-D (--date-range) is an option to specify the date range for the report (future dates are ignored): YYYYMMDD-YYYYMMDD or YYYYMMDD.

-l (--longlist) is an option to display detailed information.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Import tape

```
iscon importtape -s <server-name> [-u <username> -p <password>]
[-M <import-mode>] -v <plib-or-pdrive-vid> [-B <barcode> | -l <slot-no>]
-L <tape-library-vid> [-b <virtual-tape-barcode>] -t <virtual-tape-slot-no>
[-j <job-description>] [-k <key-name> -W <key-password>] [-X <rpc-timeout>]

iscon importtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--import-mode=<import-mode>] --plib-or-pdrive-vid=<plib-or-pdrive-vid>
[--barcode=<barcode> | --slot-no=<slot-no>] --tape-library-vid=<tape-library-vid>
--virtual-tape-slot-no=<virtual-tape-slot-no>
[--virtual-tape-barcode=<virtual-tape-barcode>] [--job-description=<job-description>]
[--key-name=<key-name> --key-password=<key-password>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command imports a physical tape to a virtual tape. The virtual tape is created during the import operation.

-M (--import-mode) is an option in one of the following values: *copy* (default) or *direct-access* or *recycle*.

-v (--pdrive-or-pdrive-vid) is required to specify the virtual device ID of the physical tape library or physical tape drive from which the physical tape is to be imported.

If the physical tape is from a physical tape library, either <barcode> or <slot-no> of the physical tape must be specified with -B (--barcode) or -l (--slot-no) to identify the physical tape. If the barcode contains leading or trailing spaces, it must be enclosed in double quotes. Physical tape information is not required if the physical tape is imported from a standalone physical tape drive.

-L (--tape-library-vid) is the virtual device ID of the virtual tape library to which the physical tape is to be imported.

-t (--virtual-tape-slot-no) is required for the virtual tape location.

-b (--virtual-tape-barcode) is optional when the physical tape from a physical tape library contains a barcode. It is required if the physical tape does not have a barcode or when it is from a physical tape drive.

-j (--job-description) is an option to specify a description for the import job.

-k (--key-name) and -W (--key-password) are options for tape encryption support. If the tape to be imported was encrypted through the system, you need to specify the key name and the key password of the encryption key to decrypt the data.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Import stacked tape

```
iscon importstackedtape -s <server-name> [-u <username> -p <password>]
[-l <physical-tape-library-vid> -b | -B <physical-tape-barcode> | -o <slot-no>] |
[-v <standalone-physical-tape-drive-vid> [-b <physical-tape-barcode>]]
-L <virtual-tape-library-vid> [-k <key-name> -W <key-password>]
[-z <originalbarcode:newbarcode:slotnumber>] [-j <job-description>] [-X <rpc-timeout>]

iscon importstackedtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--plib-vid=<physical-tape-library-vid>
--physical-tape-barcode=<physical-tape-barcode> | --slot-number=<slot-no>] |
[-pdev-vid=<standalone-physical-tape-drive-vid>
[--physical-tape-barcode=<physical-tape-barcode>]]
--tape-library-vid=<virtual-tape-library-vid>
[--key-name=<key-name> --key-password=<key-password>]
[--vtape-info=<originalbarcode:newbarcode:slotnumber>]
[--job-description=<job-description>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command submits an import job in order to copy the content of the specified stacked physical tape to virtual tapes.

If the physical tape is located in a physical tape library:

- -l (--plib-vid) is required to specify the ID of physical tape library where the physical tape is located.
- Either -b | -B (--physical-tape-barcode) or -o (--slot-number) must be used in order to identify the physical tape. If the barcode contains leading or trailing space characters, it must be enclosed in double quotes.

If the physical tape is located in a standalone physical tape drive:

- -v (--pdrive-vid) is required to specify the ID of the standalone physical tape drive where the physical tape is located.
- -b | -B (--physical-tape-barcode) is an option to validate the physical tape by matching this barcode with the barcode in the tape header. If barcode is not provided, the command will import the physical tape without checking. If the barcodes do not match, the operation will fail.

The following arguments are common:

-L (--tape-library-vid) is required to specify the ID of the virtual tape library to which you want to import the stacked tape.

-k (--key-name) and -W (--key-password) are options for tape encryption support. If the tape to be imported was encrypted through the system, you need to specify the key name and the key password of the encryption key to decrypt the data.

-z (--vtape-info) is an option to select the virtual tapes that are to be imported from the specified stacked physical tape and specify how to create those tapes. By default, all virtual tapes found on the physical tape will be imported. This option contains a list of arguments separated with commas, in the following format: "originalbarcode:newbarcode:slotnumber,..."

Enclose the list in double quotes if the barcodes contain leading or trailing spaces.

In order to create the corresponding virtual tape in the first available slot in the specified virtual tape library, use the keyword "any": originalbarcode:newbarcode:any

In order to preserve the original barcode of the virtual tape, use the keyword "same": originalbarcode:same:any

-j (--job-description) is an option to specify a description for the import job.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Export virtual tape

```
iscon exportvirtualtape -s <server-name> [-u <username> -p <password>]
-v <vdevid> -L <tape-library-vid> -b | -B <barcode> | -l <slot-no>
[-Z <on> -Q <num-of-copies>] [-J]] | [-pd <standalone-physical-tape-drive-vid>
[-M <export-mode> [-Y <days>] ] [-j <job-description>] [-f]
[-k <key-name> -W <key-password>] [-X <rpc-timeout>]

iscon exportvirtualtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> --tape-library-vid=<tape-library-vid>
--same-barcode | --barcode=<barcode> | --slot-no=<slot-no>
[--tape-duplication=<on> --num-of-copies=<num-of-copies>]
[--auto-eject-to-ie]] [--pdrive-vid=<standalone-physical-tape-drive-vid>
[--export-mode=<export-mode> [--delay-delete-days=<days>]]
[--job-description=<job-description>] [--force]
[--key-name=<key-name> --key-password=<key-password>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command exports the information from a virtual tape to a physical tape.

-v (--vdevid) is required to specify the ID of the virtual tape to be exported to the physical tape.

If the physical tape is located in a physical tape library:
-L (--tape-library-vid) is required to specify the ID of the target physical tape library.

You must include one of the following arguments to select the physical tape:

- -b (--same-barcode) lets you select a physical tape with the same barcode of the virtual tape if a physical tape with the same barcode exists.
- -B (--barcode) lets you specify the barcode of an available physical tape in the physical tape library. If the barcode contains leading or trailing spaces, it must be enclosed in double quotes.
- -l (--slot-no) lets you specify the slot number of an available physical tape in the physical tape library.

-Z (--tape-duplication) is an option to enable Tape Duplication for this export job: *on* (enable); default is *off* (disabled).

-Q (--num-of-copies) is an option to specify the number of copies to be made using the same barcode if the tape duplication option is enabled. The maximum value is 5. The default value is 1.

-J (--auto-eject-to-ie) is an option to eject the tape to the IE slot after the export job.

If the physical tape is located in a standalone physical tape drive:
-pd (--pdrive-vid) is required to specify the ID of the standalone physical tape drive where the physical tape is located.

The following arguments are common:

-M (--export-mode) is an option with one of the following values: *copy* (default) or *move*.

-Y (--delay-delete-days) is an option for *move mode* to specify the number of days to wait before deletion. The maximum is 365 days. The default value is 365 days.

-j (--job-description) is an option to specify a description for the tape export job.

-f (--force) is required when the tape is scheduled to be deleted.

-k (--key-name) and -W (--key-password) are options for tape encryption support. Specify the key name and key password of the encryption key if you wish to encrypt the data when exporting the virtual tape to the physical tape.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Export virtual tapes with stacking

```
iscon exportvirtualtapeswithstacking -s <server-name> [-u <username> -p <password>]
-T <tape-vid-list> [-L <physical-tape-library-vid>
-B <physical-tape-barcode> | -l <slot-no>] |
[-v <standalone-physical-tape-drive-vid> -b | -B <physical-tape-barcode>]]
[-M <export-mode> [-Y <days>] [-j <job-description>] [-f] [-J] [-O]
[-k <key-name> -W <key-password>] [-X <rpc-timeout>]

iscon exportvirtualtapeswithstacking --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-vid-list=<tape-vid-list> [--tape-library-vid=<physical-tape-library-vid>
--physical-tape-barcode=<physical-tape-barcode> | --slot-no=<slot-no>] |
[--pdrive-vid=<standalone-physical-tape-drive-vid>
--same-barcode | --physical-tape-barcode=<physical-tape-barcode>]]
[--export-mode=<export-mode> [--delay-delete-days=<days>]]
[--job-description=<job-description>] [--force] [--auto-eject-to-ie] [--overwrite-mode]
[--key-name=<key-name> --key-password=<key-password>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command exports up to one thousand virtual tapes to a physical tape. The virtual tapes must be in the virtual vault.

-T <--tape-vid-list> is required to specify the ID of the virtual tapes to be exported. The argument can be a list of IDs or ranges of IDs, separated with commas, or a file name enclosed in "< >" (i.e., "<file>"). The file must contain a single tape ID or range on each line. The file must be located in the same folder as the command line utility or the full path is required. For example: -T 10000001,10000010-10000020 or -T "</tmp/exporttapes>"

If the physical tape is located in a physical tape library:

-L <--tape-library-vid> is required to specify the ID of the physical tape library hosting the physical tape.

Either -B (--physical-tape-barcode) or -l (--slot-number) must be used in order to identify the physical tape. If the barcode contains leading or trailing space characters, it must be enclosed in double quotes.

If the physical tape is located in a standalone physical tape drive:

-v (--pdrive-vid) is required to specify the ID of the standalone physical tape drive where the physical tape is located.

-b | -B (--physical-tape-barcode) is required to specify the physical tape barcode. The specified barcode must match the the barcode form the physical tape header or the command will fail.

The following arguments are common:

-M (--export-mode) is an option with one of the following values: *copy* (default) or *move*.

-Y (--delay-delete-days) is an option for *move* mode to specify the number of days to wait before deletion. The maximum is 365 days. The default value is 365 days.

-j (--job-description) is an option to specify a description for the tape export job.

-f (--force) is required when the tape is scheduled to be deleted.

-J (--auto-eject-to-ie) is an option to eject the tape to the IE slot after the export job.

-O (--overwrite-mode) is an option to overwrite the physical tape. By default, the data is appended to the physical tape.

-k (--key-name) and -W (--key-password) are options for tape encryption support. Specify the key name and key password of the encryption key if you wish to encrypt the data when exporting the virtual tape to the physical tape.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Scan physical tapes

```
iscon scanphysicaltapes -s <server-name> [-u <username> -p <password>]
-l <physical-tape-library-vid> [-b | -B <physical-tape-barcode-list>] |
[-v <standalone-physical-tape-drive-vid> -b | -B <physical-tape-barcode>]]
[-X <rpc-timeout>]

iscon scanphysicaltapes --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--plib-vid=<physical-tape-library-vid>
[--tape-barcode-list=<physical-tape-barcode-list>]] |
[--pdrive-vid=<standalone-physical-tape-drive-vid>
[--physical-tape-barcode=<physical-tape-barcode>]] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command submits individual scan jobs for each physical tape found in the specified location. The command does not wait for the scan jobs to finish.

If the location to be scanned is a physical tape library:

-l (--plib-vid) is the ID of the physical tape library.

-b | -B (--tape-barcode-list) is an option to scan only the specified physical tapes from the same library. Multiple barcodes must be separated with commas and the list must be enclosed in double quotes if the barcodes contain leading or trailing space characters.

If the location to be scanned is a standalone physical tape drive:

-v (--pdriveID) is the ID of the standalone physical tape drive.

-b | -B (--physical-tape-barcode) is an option to validate the physical tape by matching this barcode with the barcode in the tape header. If barcode is not provided, the command will scan the physical tape without checking.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Resume import/export jobs

```
iscon resumeimportexportjobs -s <server-name> [-u <username> -p <password>]
-j <job-id-list> [-X <rpc-timeout>]

iscon resumeimportexportjobs --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --job-id-list=<job-id-list>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command resumes the specified import/export jobs. The jobs must be in the import/export queue in a suspended state.

-j <--job-id-list> is a list of job IDs separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Restart import/export jobs

```
iscon restartimportexportjobs -s <server-name> [-u <username> -p <password>]
-j <job-id-list> [-X <rpc-timeout>]

iscon restartimportexportjobs --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--job-id-list=<job-id-list> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command restarts the specified import/export jobs. The jobs must be in the import/export queue and they must have either been cancelled or failed.

.-j <--job-id-list> is a list of job IDs separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete import/export jobs

```
iscon deleteimportexportjobs -s <server-name> [-u <username> -p <password>]
-j <job-id-list> [-X <rpc-timeout>]

iscon deleteimportexportjobs --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --job-id-list=<job-id-list>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes the specified import/export jobs. The jobs must be in the import/export queue.

-j <--job-id-list> is a list of job IDs separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Suspend import/export jobs

```
iscon suspendimportexportjobs -s <server-name> [-u <username> -p <password>]
-j <job-id-list> [-X <rpc-timeout>]

iscon suspendimportexportjobs --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --job-id-list=<job-id-list>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command suspends the specified import/export jobs. The jobs must be in the import/export queue and must be idle.

-j <--job-id-list> is a list of job IDs separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Cancel import/export jobs

```
iscon cancelimportexportjobs -s <server-name> [-u <username> -p <password>]
-j <job-id-list> [-X <rpc-timeout>]

iscon cancelimportexportjobs --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --job-id-list=<job-id-list>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command cancels the specified import/export jobs. The jobs must be in the import/export queue and must be running.

-j <--job-id-list> is a list of job IDs separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Reports

The reports below can be generated through the command line interface. Many of the reports allow you to select a date range. The definition of a day (midnight to midnight or noon to noon) is set in the console (right-click the *Reports* object and select *Properties* --> *Other* tab).

Deduplication policy status report

```
iscon creatededupepolicystatusreport -s <server-name> [-u <username> -p <password>]
[-z <report period>] | [-D <date-range>] [-o <filename>] [-X <rpc-timeout>]

iscon creatededupepolicystatusreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the current deduplication policies and the deduplication jobs executed by those policies. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is:
YYYYMMDD-YYYYMMDD or YYYYMMDD.

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The date option is applied to the local server time. The default value is: "-z t" (today).

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: DeduplicationPolicyStatus-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Deduplication tape activity report

```
iscon createdPedupetapeactivityreport -s <server-name> [-u <username> -p <password>]
[-i <"policyIDlist">] [-B <barcode-range>] [-S <job-status>]
[-z <report period>] | [-D <date-range>] [-O <additional options>] [-o <filename>]
[-X <rpc-timeout>]

iscon createdPedupetapeactivityreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--policyid=<"policyIDlist">] [--barcode-range=<barcode-range>]
[--job-status=<status> [--report-period=<report-period>] | [--date-range=<date-range>]
[--options=<aditional options>] [--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the deduplication history at tape level. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

The optional arguments can be combined in order to perform advanced queries. The default relationship for any optional argument combination is "and".

-i (--policyid) is an option to report the activity of the specified polices only. Multiple policy IDs must be separated with semicolons and no spaces are allowed. For example: -i 1;2;3

-B (--barcode-range) is an option to report the activity of the specified virtual tapes only. The format for this argument is a barcode range.

-S (job-status) is an option to report the activity based on the job status. The accepted values for this argument are: OK, FAILED, CANCELED, and NEW.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD.

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The date option is applied to the local server time. The default value is: "-z t" (today).

-O (--options) is an option to specify additional parameters for report generation, separated by commas:

- GJS[ET] - group by job status with optional sort by end time
- STL - display the information as a continuous tape list
- IAT - include active tapes
- LEJ - show only jobs from last policy run
- LCJ - show only the last completed job from policy

LEJ and LCJ are mutually exclusive values.

For example: -O GJS,IAT
This command will group jobs by status and will include active tapes.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: DeduplicationTapeActivity-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Deduplication tape usage report

```
iscon creatededupetapeusagereport -s <server-name> [-u <username> -p <password>]
[-o <filename>] [-X <rpc-timeout>]

iscon creatededupetapeusagereport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the current disk usage and data deduplication information for the existing polices. The information is shown at tape level. The report can be viewed, printed, emailed, or exported to other formats from the console.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: DeduplicationTapeUsager-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Deduplication replication status report

```
iscon creatededupereplicationstatusreport -s <server-name> [-u <username> -p <password>]
[-L <server-source-list>] [-d <RNZ>] [-o <filename>] [-X <rpc-timeout>]

iscon creatededupereplicationstatusreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--source-list=<server-source-list>] [--details=<RNZ>] [--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the replication coverage for all the servers that have deduplication policies set up to replicate to this server. The report can be viewed, printed, emailed or exported to other formats from the console.

-L (--source-list) is an option to include only the specified source servers in the report. This argument can be a list of server names separated by commas or the file name enclosed in "< >" (e.g. "<file>") of a text file containing the list in the first line. The file must be located in the same folder as the command line utility or the full path is required. The server name is case sensitive.

-d (--details) in an option to request additional information, using one or more of the following characters (e.g. -d RZ): "R", "N", and "Z".

- "R" lists the last successful replication job information for all virtual tapes that are in a deduplication policy and have their current status as resolved.

- "N" lists all virtual tapes that are in a deduplication policy and are not fully replicated.
- "Z" lists all virtual tapes that are in a deduplication policy and have no data. These tapes do not require replication.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: DeduplicationReplicationStatus-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Deduplication repository memory and space usage report

```
iscon creatededuperepositorymemspacereport -s <server-name> [-u <username> -p <password>]
[-z <report period>] | [-D <date-range>] [-d <interval>] [-o <filename>]
[-X <rpc-timeout>]

iscon creatededuperepositorymemspacereport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--data-points=<interval>] [--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the repository memory and space usage during the specified time period. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: : YYYYMMDD-YYYYMMDD or YYYYMMDD

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The default value is: "-z t" (today).

-d (--data-points) is an option to specify the time interval between the data points: "daily", "weekly", "monthly", "quarterly". The default values for the data points interval are:

- hourly - when reporting up to 3 days of data
- daily - when reporting between 4 and 60 days of data
- weekly - when reporting more than 60 days of data

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: DedupeRepositoryMemorySpaceUsage-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Deduplication repository performance report

```
iscon creatededuperepositoryperformancereport -s <server-name> [-u <username> -p
<password>]
[-z <report period>] | [-D <date-range>] [-d <interval>] [-o <filename>]
[-X <rpc-timeout>]

iscon creatededuperepositoryperformancereport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--data-points=<interval>] [--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the repository deduplication activity during the specified time period. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The default value is: "-z t" (today).

-d (--data-points) is an option to specify the time interval between data points when either the report period or date range argument is used. In order to limit the number of data points and prevent reports with a single data point, the accepted values are:

- hourly - when reporting fewer than 4 days of data
- daily - when reporting between 2 and 59 days of data
- weekly - when reporting more than 13 days of data
- monthly - when reporting more than 59 days of data
- quarterly - when reporting more than 121 days of data

The default values for the interval between data points are:

- hourly - when reporting up to 3 days of data
- daily - when reporting between 4 and 60 days of data
- weekly - when reporting more than 60 days of data

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: DedupeRepositoryPerformance-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Deduplication repository reclamation report

```
iscon creatededuperepositoryreclamationreport -s <server-name> [-u <username> -p
<password>] [-z <report period>] | [-D <date-range>] [-o <filename>] [-X <rpc-timeout>]

iscon creatededuperepositoryreclamationreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the deduplication reclamation activity for the specified time period. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD.

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The date option is applied to the server local time. The default value is: "-z t" (today).

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: DedupeRepositoryReclamation-MM-DD-YYYY-hh-mm-ss.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Disk space allocation for virtual tapes in libraries report

```
iscon creatediskspaceallocreport -s <server-name> [-u <username> -p <password>]
[-h [-R <resource-list> -z <report period> | -D <date-range> -d <interval>]]
[-o <filename>] [-X <rpc-timeout>]

iscon creatediskspaceallocreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--historical [--resource-list->resource-list>
--report-period=<report-period> | --date-range=<date-range> --data-points=<interval>]]
[--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side summarizing the disk space used by the allocated tapes. The report can be viewed, printed, emailed, or exported to other formats from the console. By default, the report presents the current disk allocation.

-h (--historical) is an option to create a historical report. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties.

The following four options can be used only when the historical report option is selected:

-R <--resource-list> in an option to report the status of the specified libraries only. The argument can be a list of virtual identifiers separated with commas, or the file name, enclosed in "< >", of a text file containing the list in the first line. The file must be located in the same folder as the command line utility or the full path is required. For example: -R 10,17 or -R "<lib_id_file.txt>"

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The default value is: "-z t" (today).

-d (--data-points) is an option to specify the time interval between the data points when either the rep[ort period or date range argument is used. In order to limit the number of data points and prevent reports with a single data point, the accepted values are:

- hourly - when reporting fewer than 4 days of data
- daily - when reporting between 2 and 59 days of data
- weekly - when reporting more than 13 days of data
- monthly - when reporting more than 59 days of data
- quarterly - when reporting more than 121 days of data

The default values for the interval between data points are:

- hourly - when reporting up to 3 days of data
- daily - when reporting between 4 and 60 days of data
- weekly - when reporting more than 60 days of data

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: DiskSpaceAllocationVirtualTapes-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Disk space usage history report

```
iscon creatediskspaceusagehistoryreport -s <server-name> [-u <username> -p <password>]
[-z <report period>] | [-D <date-range>] [-o <filename>] [-X <rpc-timeout>]

iscon creatediskspaceusagehistoryreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the disk space usage for the selected time period. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is:
YYYYMMDD-YYYYMMDD or YYYYMMDD.

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The date option is applied to the local server time. The default value is: "-z t" (today).

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: DiskSpaceUsageHistory-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Import export job report

```
iscon createimportexportjobreport -s <server-name> [-u <username> -p <password>]
[-z <report period>] | [-D <date-range>] [-i <filter>] [-o <filename>] [-X <rpc-timeout>]

iscon createimportexportjobreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--include-filter=<filter>] [--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing all import/export and tape caching jobs that were placed in the queue during the specified period of time, regardless of job status. The report can be viewed, printed, emailed, or exported to other formats from the console.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday

- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is:
YYYYMMDD-YYYYMMDD or YYYYMMDD

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The date option is applied to the local server time. The default value is: "-z t" (today).

-i (--include-filter) is an optional filter to include only the specified jobs. The following values are accepted. Multiple values must be separated with commas.

Job Type:

- ESD[C | M] - Export to standalone drive, Copy or Move
- EPL[C | M] - Export to physical library, Copy or Move
- ISD[C | R] - Import from standalone drive, Copy or Recycle
- IPL[C | R] - Import from physical library, Copy or Recycle
- CC - Create cache with copy meta data
- TS - Tape stacking jobs

The default is to include tapes that were exported with either Copy or Move or were imported with either Copy or Recycle.

Job Status:

- WTD - Waiting for tape/drive
- FAIL - Failed
- COMP - Completed
- CANC - Cancelled
- HOLD - On hold
- WIE - Waiting for IE slot
- RUN - Running

For example: -i EPL,COMP,CANC
This command will include only jobs with export to physical library, *copy* and *move*, *completed* and *cancelled*.

By default all jobs are included.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: ImportExportJobReport-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

LUN report

```
iscon createlunreport -s <server-name> [-u <username> -p <password>]
[-I <ACSL>] [-o <filename>] [-X <rpc-timeout>]

iscon createlunreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>] [--scsiaddress=<ACSL>]
[--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing information about the resources allocated per LUN. The report can be viewed, printed, emailed, or exported to other formats from the console.

-I <ACSL> (--scsiaddress) is an option to specify a single LUN address to be reported.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: LUNReport-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Physical tape usage report

```
iscon createphytapeusagereport -s <server-name> [-u <username> -p <password>]
[-o <filename>] [-X <rpc-timeout>]

iscon createphytapeusagereport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the information queried from the physical tape database. The report can be viewed, printed, emailed, or exported to other formats from the console.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: PhysicalTapeUsage-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Replication status report

```
iscon createreplicationstatusreport -s <server-name> [-u <username> -p <password>]
[-z <report period>] | [-D <date-range>] [-r <repl-resource-type> -R <resourceList>]
[-O <sorting>] [-o <outputFilename>] [-X <rpc-timeout>]

iscon createreplicationstatusreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--repl-resource-type=<repl-resource-type> --resource-list=<resourceList>]
[--options=<sorting>] [--output-file=<outputFilename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing replication job related information about the specified resources. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD.

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The default value is: "-z t" (today).

-r (--repl-resource-type) is an option to specify a generic resource type to be queried. It can be one of the following: TAPE or TAPEReplica. The default value is TAPE.

-R <--resource-list> in an option to report the status of the specified resources only. The argument can be a list of virtual identifiers separated with commas or the name of a file enclosed in <> containing the resource ID on each line. All the resources must be of the type specified by "-r".

- Example 1: -R 10000005,10000006
- Example 2: -R "<res_id_file.txt>"

-O (--options) specify the output sorting for report generation. The default value is SRV:

- SRV - sort the output by remote server name
- LOG - to sort the output by job start time

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: ReplicationStatus-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Virtual library and drive assignment report

```
iscon createvirtuallibdrvassignreport -s <server-name> [-u <username> -p <password>]
[-o <filename>] [-X <rpc-timeout>]

iscon createvirtuallibdrvassignreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing all the virtual tape libraries and drives assigned to different clients. The report can be viewed, printed, emailed, or exported to other formats from the console.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: LibraryDriveAssignment-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Virtual library information report

```
iscon createvirtuallibinforeport -s <server-name> [-u <username> -p <password>]
[-o <filename>] [-X <rpc-timeout>]

iscon createvirtuallibinforeport --server-name=<server-name>
[--server-username=<username> --server-password=<password>] [--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing information about each virtual tape library created on this server, including the physical library it emulates, the amount of storage occupied by its virtual tapes, the clients that this library is assigned to, and the number of drives, slots, and tapes. The report can be viewed, printed, emailed, or exported to other formats from the console.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: VirtualLibraryInfo-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Virtual tape activity report

```
iscon createvirtualtapeactivityreport -s <server-name> [-u <username> -p <password>]
[-z <report period>] | [-D <date-range>] [-i <filter>] [-o <filename>] [-X <rpc-timeout>]

iscon createvirtualtapeactivityreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--include-filter=<filter>] [--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing all virtual tape activity for all virtual tapes for three types of operations: Backup, Tape Import (*write* operations), and Tape Export (*read* operations). The displayed information includes the start time, end time, duration, job performance, the barcode of the virtual tape, and the compression rate if applicable. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console. The tape activity includes: backup, export and import jobs.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is the starting date and ending date in the following format (maximum 365 days): YYYYMMDD-YYYYMMDD or YYYYMMDD.

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The date option is applied to the server local time. The default value is: "-z t" (today).

-i (--include-filter) is an optional filter to include only the virtual tapes that match the barcode filter. This option can be one of the following values:

- BARCODEPREFIX=barcodePrefix,
- BARCODECONTAINS=pattern,
- BARCODERANGE=barcodeStart-barcodeEnd,

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: VirtualTapeActivity-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Virtual tape information report

```
iscon createvirtualtapeinforeport -s <server-name> [-u <username> -p <password>]
[-i <filter>] [-o <filename>] [-X <rpc-timeout>]

iscon createvirtualtapeinforeport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--include-filter=<filter>][--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing detailed information about each virtual tape and tape replica. Based on the selected views, ths information includes tape barcode, tape status, data size, action needed for tape, tape location, deduplication information, etc.The report can be viewed, printed, emailed, or exported to other formats from the console.

-i (--include-filter) is an optional filter to include only the specified virtual tapes. This option can be any combination of the following values, separated by commas. Multiple IDs of the same type must be separated by semicolons. The barcode filters are mutually exclusive.

- BARCODEPREFIX=barcodePrefix
- BARCODECONTAINS=pattern
- BARCODERANGE=barcodeStart-barcodeEnd
- LIBRARY=ID1;ID2 (virtual library ID list)
- POLICY=ID1;ID2 (deduplication policy ID list)

Additionally, the following values can be used to select from different output templates. Multiple views must be separated with semicolons. The default view is overall summary.

- VIEW=OS (overall summary)
- DE (deduplication view)
- TC (tape caching view)
- RR (replica resources view, includes all replica tapes)
- VV (vault view, includes all tapes from the vault)
- DT (detailed tape view)

The argument must be enclosed in double quotes. For example:
-i "LIBRARY=10;11,BARCODEPREFIX==00,VIEW=OS;DE;TC"

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: VirtualTapeInfo-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

VTL performance report

```
iscon createvtlperformancereport -s <server-name> [-u <username> -p <password>]
[-z <report period>] | [-D <date-range>] [-d <interval>] [-i] [-o <filename>]
[-X <rpc-timeout>]

iscon createvtlperformancereport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--data-points=<interval>] [--include-filter=<filter>] [--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the average CPU/memory usage for the entire VTL system, and the amount of data read/written for the server, adapters, LUNs, client devices, and virtual tape libraries during each interval in the specified period of time. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is:
YYYYMMDD-YYYYMMDD or YYYYMMDD.

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The date option is applied to the server local time. The default value is: "-z t" (today).

-d (--data-points) is an option to choose the time interval between data points when either the report period or date range argument is used. In order to limit the number of data points and prevent reports with a single data point, the accepted values are:

- "hourly" when the report period is less than 4 days
- "daily" when the report period is between 2 and 59 days
- "weekly" when the report period is more than 13 days
- "monthly" when the report period is more than 59 days
- "quarterly" when the report period is more than 121 days

The default values for the data points are:

- "hourly" when the report includes up to 3 days of data
- "daily" when the report includes between 4 and 60 days of data
- "weekly" when the report includes more than 60 days of data

If a report period is not specified, there is no need to use -d (--data-points).

-i (--include-filter) is an optional filter to include only the specified devices. This option can be any combination of the following values, separated by commas. Multiple IDs of the same type must be separated by semicolons. The argument must be enclosed in quotes.

- Storage HBAs - ADAPTER=all or ADAPTER_NO_1;ADAPTER_NO_2
- Storage devices - LUN=all or A:C:S:L(1);A:C:S:L(2)
- Clients - CLIENT=all or CLIENT_ID_1;CLIENT_ID_2

- Virtual libraries - LIBRARY=all or ID_1;ID_2

By default, the report lists the performance for the whole VTL server and for all devices mentioned above. Use "none" in order to filter out individual devices. For example:

- -i "ADAPTER=all,LUN=100:0:0:1;100:0:0:5,LIBRARY=100"
- -i "none"

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: VTLPerformance-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Failover

Get failover status

```
iscon getfailoverstatus -s <server-name> [-u <username> -p <password>] [-X <rpc-timeout>]

iscon getfailoverstatus --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command obtains the failover status from either the primary failover server or the secondary failover server.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Data encryption

Enable virtual tape encryption

```
iscon enablevirtualtapeencryption -s <server-name> [-u <username> -p <password>]
-W <activation password> -C <activation password> [-H <password hint>] [-X <rpc-timeout>]

iscon enablevirtualtapeencryption --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--password=<activation password> --confirm-password=<activation password>
[--password--hint<password hint>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command enables the virtual tape encryption option. In order to be able to access encrypted data, the server must have encryption activated. Virtual tapes inherit the encryption property from their parent virtual library for the lifetime of the virtual tape.

-W (password) is required to provide the 10 to 16 character password needed for encryption activation.

-C (--confirm-password) is required to provide the activation password again. The two new password arguments must match.

-H (--password-hint) is optional text that can provide password clues. The text is up to 32 characters and it is shown whenever other commands fail due to a password mismatch.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Deduplication

List deduplication policies

```
iscon dedupelistpolicies -s <server-name> [-u <username> -p <password>]
[-X <rpc-timeout>]

iscon dedupelistpolicies --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command displays the deduplication policies created on the specified server.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Start a deduplication policy

```
iscon dedupestartpolicy -s <server-name> [-u <username> -p <password>]
-I <"policyname"> [-X <rpc-timeout>]

iscon dedupestartpolicy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--policyname=<"policyname"> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command starts the execution of the specified policy.

-I (--policyname) is required to specify the policy name. Enclose the policy name in double quotes.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Stop a deduplication policy

```
iscon dedupestoppolicy -s <server-name> [-u <username> -p <password>]
-I <"policyname"> [-X <rpc-timeout>]

iscon dedupestoppolicy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--policyname=<"policyname">] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command stops the execution of the specified policy.

-I (--policyname) is required to specify the policy name. Enclose the policy name in double quotes.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Add a deduplication policy

```
iscon dedupeaddpolicy -s <server-name> [-u <username> -p <password>]
-I <"policyname"> [-t] [i] [-F <H@hh:mm>]
[-e [-m <size-mb> | -f]] [-Q <low> -M <retry-count> -V <retry-interval>]
[-N -n <replication-mode>
-T <target-server-ip> -U <targetusername> -P <targetpassword>
-T2 <target-server-ip> -U2 <targetusername> -P2 <targetpassword>
-L <vlib-id> -L2 <vlib-id>
-w <hh:mm-hh:mm> -A -c <on|off> -z <on|off> -R <retry-interval> -C <retry-count>
[-X <rpc-timeout>]

iscon dedupeaddpolicy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--policyname=<"policyname"> [--turbo] [--inline] [--frequency=<H@mm:mm>]
[--tape-ejected-to-slot [--size-mb=<size-mb> | --tape-full]]
[--policy-priority=<low> --max-retry=<retry-count> --retry-interval=<retry-interval>]
[--enable-replication
--replication-mode=<replication mode>
--target-server-ip=<target-server-ip>
--target-username=<targetusername> --target-password=<targetpassword>
--target-server-ip2=<target-server-ip>
--target-username2=<targetusername> --target-password2=<targetpassword>
--tape-library-vid=<vlib-id> --tape-library-vid2=<vlib-id>
--repl-window=<hh:mm-hh:mm> --auto-delete
--compression=<on|off> --encryption=<on|off>
--replication-retry-interval=<retry-interval> --replication-retry-count=<retry-count>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a deduplication policy configured with the specified arguments.

-I (--policyname) is required to specify the policy name. Enclose the policy name in double quotes.

-t (--turbo) is an option to enable the "turbo deduplication" feature. When enabled, this feature can improve the overall backup throughput and/or the deduplication performance. This option cannot be used with the "inline deduplication" trigger.

Deduplication triggers:

No Schedule (Manual) is the default trigger. The deduplication triggers are mutually exclusive.

-i (--inline) triggers deduplication process while backup is in progress.

-F (--frequency) is a time-based trigger with the following options:

- H@hh:mm triggers the policy execution hourly starting at the specified time
- D@hh:mm triggers the policy execution daily starting at the specified time
- Sunday@hh:mm triggers the policy execution weekly starting on the specified day and time.
  For example: -F Wednesday@23:00.

-e (--tape-ejected-to-slot) triggers deduplication for an individual tape in this policy whenever it is ejected to the slot, if the new data written to tape is greater than 1 MB.

-m (--size-mb) is an additional option for the eject trigger that activates the trigger only if the size of the new data on the tape is greater than the specified size.

-f (--tape-full) is an additional option for the eject trigger that activates the trigger only if the tape is full.

-Q (--policy-priority) is an option to prioritize the execution of the queued deduplication jobs. Default is "none". The accepted values are: "low", "medium" or "high".

-M (--max-retry) is an option to retry a failed deduplication job the specified number of times (0 to 99999). The default is 0.

-V (--retry-interval) is an option to specify the time between retries (1 to 60 minutes). The default is 30 minutes.

Replication options:

-N (--enable-replication) is an option to create a policy with replication of deduplicated data. The target information arguments are mandatory when replication is enabled.

-n (--replication-mode) is an option to enable advanced replication. The advanced replication values are:

- CASCADE is an option to replicate from server A to server B, then server B will replicate to server C;
- PARALLELC is an option to replicate from server A to server B and C concurrently;
- PARALLELS is an option to replicate from server A to server B and C sequentially.

-T (--target-server-ip) is the IP address of the VTL target server. When advanced replication is enabled, this information identifies server B.

-U (--target-username) and -P (--target-password) are required in order to access VTL target server information.

-T2 (--target-server-ip2) is the IP address of the VTL target server for advanced replication mode. This argument identifies server C.

-U2 (--target-username2) and -P2 (--target-password2) are required in order to access the target server specified by -T2.

-L (--tape-library-vid) is an option to move the LVIT tape to a virtual tape library on the remote server. The virtual library must support the same media type as the tapes in the policy. When advanced replication is enabled, this argument identifies a virtual tape library on server B.

-L2 (--tape-library-vid2) is an option for advanced replication to move the LVIT tape to a virtual tape library on the remote server C. The virtual library must support the same media type as the tapes in the policy.

-w (--repl-window) is an option to restrict replication to the specified time interval: hh:mm-hh:mm.

-A (--auto-delete) is an option to remove the replica when the tape is full.

-c (--compression) is an option to enable replication compression. The default value is "off".

-z (--encryption) is an option to enable replication encryption. The default value is "off".

-R (--replication-retry-interval) is an option to retry replication after the specified number of seconds when failure occurs. The default value is 60 seconds.

-C (--replication-retry-count) is an option to specify the the number of replication retries. The default value is 1.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete a deduplication policy

```
iscon dedupedelpolicy -s <server-name> [-u <username> -p <password>]
-I <"policyname"> [-U <targetusername> -P <targetpassword>]
[-X <rpc-timeout>]

iscon dedupedelpolicy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--policyname=<"policyname">
[--target-username=<targetusername> --target-password=<targetpassword>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes the specified deduplication policy.

-I (--policyname) is required to specify the policy name. Enclose the policy name in double quotes.

-U (--target-username) and -P (--target-password) - When executing this command on the primary server (server A) in order to also delete the primary policy in a cascaded replication setup, these arguments are required in order to delete the policy from the first target server (server B) if either of the following conditions applies:

- You are not already logged into the target server (using the login command).
- Credentials for server B are not the same as credentials for server A.

If these arguments are not provided, credentials for the primary server will be used. If the command fails to delete the cascaded policy, run the command on server B in order to delete the corresponding policy.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Add a tape to a deduplication policy

```
iscon dedupeaddtapetopolicy -s <server-name> [-u <username> -p <password>]
-T <tapevidlist> -I <"policyname"> [-X <rpc-timeout>]

iscon dedupeaddtapetopolicy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-vid-list=<tapevidlist> --policyname=<"policyname"> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds virtual tapes to an existing policy.

-T (--tape-vid-list) is required to specify the ID of the virtual tapes to be added to the policy as a list of numbers separated by commas.

-I (--policyname) is required to specify an existing name. Enclose the policy name in double quotes.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Remove a tape from a deduplication policy

```
iscon deduperemovetapefrompolicy -s <server-name> [-u <username> -p <password>]
-T <tapevidlist> -I <"policyname"> [-X <rpc-timeout>]

iscon deduperemovetapefrompolicy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-vid-list=<tapevidlist> --policyname=<"policyname"> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command removes virtual tapes from an existing policy.

-T (--tape-vid-list) is required to specify the ID of the virtual tapes to be removed from the policy as a list of numbers separated by commas.

-I (--policyname) is required to specify the policy name. Enclose the policy name in double quotes.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get deduplication tape activity

```
iscon dedupetapeactivityinfo -s <server-name> [-u <username> -p <password>]
[-I <"policynamelist">] [-T <tapevidlist>] [-S <job-status>]
[-D <YYYYMMDDhhmmss-YYYYMMDDhhmmss>] [-w <hh:mm-hh:mm>] [-x] [-d] [-O] [-Z] [-l]
[-M <delim>] [-X <rpc-timeout>]

iscon dedupetapeactivityinfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--policyname=<"policyname">] [--tape-vid-list=<tapevidlist>]
[--job-status=<job-status>] [--date-range=<YYYYMMDDhhmmss-YYYYMMDDhhmmss>]
[--backup-window=<hh:mm-hh:mm>] [--last-run] [--skip-deleted][--order-by-status]
[--extra-filter][--longlist] [--output-delimiter=<delim>][--rpc-timeout=<rpc-timeout>]
```

Description:

This command reports the deduplication history for tapes on the specified server. The optional arguments can be combined in order to perform advanced queries. The default relationship for any optional argument combination is "and".

-I (--policyname) is an option to report the activity of the specified policy only. Multiple names must be separated by commas and the whole argument must be enclosed in double quotes: e.g. "Policy 1,Policy 2,Policy 3".

-T (--tape-vid-list) is an option to report the activity of the specified virtual tapes only. The format for this argument must be a list of numbers separated by commas.

-S (job-status) is an option to report activity based on job status. The accepted values for this argument are: *OK*, *FAILED, CANCELED*, or *NEW*.

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDDhhmmss-YYYYMMDDhhmmss.

-w (--backup-window) is an option to report the activity for the specified time interval only. This option can be combined with -D (--date-range) to generate the report for a specific interval over multiple days. For example,
-D 20131201000000-20131231235959 -w 01:00-04:00
would generate the report for the hours of 1:00am to 4:00am for the 31 days specified.

-x (--last-run) is an option to report the last execution for each tape per policy.

-d (--skip-deleted) is an option to filter out the records for the tapes that were deleted or moved from policies.

-O (--order--by-status) is an option to order the records for each policy by execution status.

-Z (--extra-filter) is an option to add an additional filter to skip the records for the complete execution if there is no scanned or replicated data. This argument is ignored if the detailed format output is requested.

-l (--longlist) is an option to display the detailed report in the format "Label=Value".

-M (--output-delimiter) is an option to display the report using the specified string as the field delimiter. The delimiter can be up to 8 characters long.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Deduplication Repository commands

The commands in this section can only be used on a SIR server.

Reports

Get report data

```
iscon getreportdata -s <server-name> [-u <username> -p <password>]
[-d <date>] [-I #[D|W|M]]
[-C <clientList>] [-R <resourceList>]
[-N]
[-o <outputFilename>] [-F <fileFormat>] [-H] [-f]
[-X <rpc-timeout>]

iscon getreportdata --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--date=<date>] [--units=#[<D|W|M>]]
[--client-list=<clientList>]
[--resource-list=<resourceList>]
[--no-system-info]
[--output-file=<outputFilename>]
[--file-format=<fileFormat>]
[--include-heading] [--force]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command generates a report listing deduplication server total I/O activity, average system memory, and CPU usage for each interval between data points of the selected unit.

The date format is YYYYMMDD and the default is today's date.

The time unit can be specified in three different ways:

- D = Day
- W = Week
- M = Month

-R (--resourceList) can be any of the following:

- a list of resource ids separated by comma
- "*" for all the resources
- a filename enclosed in <> containing the resource ID on each line.

For example:

- -R 1,3,5,10
- -R "<res_id_file.txt>"

-C (--clientList) can be any of the following:

- a list of client names separated by comma

- “*" for all clients
- a filename enclosed in <> containing client name in each line.

For example:

- -C client1,client2
- -C "<client_file.txt>"

System information, including memory and CPU usage information, is included by default.

-N (--no-system-info) is an option to exclude system memory and CPU usage information.

If an output filename is not specified, the default filename is “ServerThroughputYYYY-MM-DD[.#]”

[.#] is added to the filename automatically if a report with the requested filename already exists.

-F (--fileFormat) can be either “csv” (the default) or “txt”.

-H (--include-heading) is the option to include a data header.

-f (--force) overwrites the existing file when the output file already exists.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

NAS commands

The commands in this section can only be used on a server on which NAS has been enabled.

Virtual devices

Create NAS device

```
iscon createnasdevice -s <server-name> [-u <username> -p <password>]
-I <ACSL list>
[-O] [-F <file-system>
[-fo <format-options>]
[-mo <mount-options>]
[-n <vdevname>]
[-m <#(MB)>]
[-X <rpc-timeout>]

iscon createnasdevice --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--scsiaddress=<ACSL list>
[--ost]
[--file-system=<file-system>
[--format-options=<format-options>]
[--mount-options=<mount-options>]
[--vdevname=<vdevname>]
[--size-mb=<#(MB)>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a NAS virtual device.

-I (--scsiaddress) is required to specify a physical device to be used to create the virtual device. The argument can be a list of ACSLs separated by commas. #:#:#:#,#:#:#:# (adapter:channel:id:lun)

-O (--ost) is an option used to create a device for OpenStorage.

-F (--file-system) is an option to choose the file system type (e.g. EXT4). The default file system type will be determined by the NAS configuration.

-fo (--format-options) is an option to specify the format options; e.g. -F -I 512 -m 0 -v -j -E resize=16383G -J size=128 -b 4096

-mo (--mount-options) is an option to specify the mount options; e.g. rw,nosuids,user_attr

-m (--size-mb) is an option to specify the size in MB for the allocated space. If this option is not specified, the total available space will be allocated. The minimum size for the NAS resource is 1000.

-n (--vdevname) is an option to specify the virtual device name. The maximum length is 64. Leading and trailing spaces will be removed. Enclose the name in double quotes to ensure proper parsing. These characters are invalid for NAS Resource name: <>"&$/\'()%# :;|*?`

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Expand NAS device

```
iscon expandnasdevice -s <server-name> [-u <username> -p <password>]
-v <vdevid> -I <ACSL list> [-m <#(MB)>]
[-W <wait-for-device-status>]
[-X <rpc-timeout>]

iscon expandnasdevice --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid>
--scsiadress=<ACSL list>
[--additional-mb=<#(MB)>]
[--wait-for-device-status=<wait-for-device-status>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This commmand expands a NAS resource. It is recommended that you detach all clients from the resource before executing this command.

-v (--vdev-id) is required to provide the resource ID.

-I (--scsiaddress) is required to specify the SCSI address of the luns to be used for expansion, separated by commas; e.g., #:#:#:#,#:#:#:#

-m (--additional-mb) is an option to specify how much space (in MB) to use for expansion. If this option is not specified, the total available space will be allocated.

-W (--wait-for-device-status) is an option to wait for the operation to complete, between 3 and 3600 seconds. By default, the command does not wait for the expansion result.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

Delete NAS device

```
iscon deletenasdevice -s <server-name> [-u <username> -p <password>]
[-v <vdevid> ]
[-X <rpc-timeout>]

iscon deletenasdevice --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--vdevid=<vdevid>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes a NAS resource.

-v (--vdevid) is required - specify the device ID.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

NAS resources

Format NAS resource

```
iscon formatnas -s <server-name> [-u <username> -p <password>]
-v <vdevid>
[-mo <mount-options>] [-fo <format-options>]
[-X <rpc-timeout>]

iscon formatnas --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid>
[--mount-options=<mount-options>]
[--format-options=<format-options>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command formats a NAS resource file system.

-v (--vdevid) is required to specify the virtual device ID.

-mo (--mount-options) specifies options for mounting the file system. Refer to the Linux mount command man page for all available file system mount options.

-fo (--format-options) specifies options for formatting the file system. Refer to the Linux mke4fs command man page for details. Options must be entered inside double quotes; e.g., “option1 option2 ...”.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create a share

```
iscon createshare -s <server-name> [-u <username> -p <password>]
-v <vdevid> [-x <export-path>]
-f <folder-name>
-W <cifs-share-name>
[-C <connection-limit>] [-c <share-comment>]
-A <cifs-account-name> [-a <cifs-account-access>] |
-cn <client-name> | -m <client-machine>
-ca <client-access-mode> [-E] [-Q <squash>]
[-X <rpc-timeout>]

iscon createshare --server-name=<server-name>
[--server-username=<username--server-password=<password>]
--vdevid=<vdevid> [--export-path=<export-path>] --folder-name=<folder-name>
--cifs-share-name=<cifs-share-name> [--connection-limit=<connection-limit>]
[--share-comment=<share-comment>]
--cifs-account-name=<cifs-account-name> [--cifs-account-access=<cifs-account-access>] |
--client-name=<client-name> | --client-machine=<client-machine>
--access-mode=<client-access-mode> [--insecure] [--squash=<squash>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a folder and shares it with an existing NFS and/or CIFS client.

-v (--vdevid) is required to specify the device ID of the NAS resource.

-f (--folder-name) is required to specify the folder name.

-x (--export-path) is an option to specify the parent path of the folder if it is not under the root of the NAS resource. For example, to create a share named 'share-folder-1' under '/nas/NAS-00005/fds/account-depart', use the following arguments:

- --vdevid=5
- --folder-name=share-folder-1
- --export-path=account-depart

Share the folder with an existing client.

CIFS client:

-W (--cifs-share-name) is required.

-C (--connection-limit) is an option to limit the number of connections to the share. Specify 0 for “unlimited” or a positive number as the maximum number of connections allowed.

-c (--share comment) is an option to add a comment for the share. The maximum length is 255 characters.

-A (--cifs-account-name) is required.

-a (--cifs-account-access) is an option to specify access mode: *ReadOnly* or *ReadWrite* (the default).

Example for a CIFS share:

```
iscli createshare --server-name=h41-93 --vdevid=10 --folder-name=shrFldTest --cifs-share-name=shrFldTest
--window-account-name=winTest --window-account-access=ReadOnly
```

NFS client:

Either -cn (--client name) or -m (--client-machine) is required.

-cn (--client-name) is the name of the NFS client.

-m (--client-machine) defines the client machine. Enter either the IP address or full hostname. On certain subnetworks, this is defined using network IP/netmask.

- Netmask can be specified using the IP address; e.g., "255.255.0.0" or CIDR notation (a prefix length); e.g., "16". Either specification has the same meaning in the command.
- Whether you use the IP address or CIDR format to specify the netmask, the network IP should have the same number of characters as netmask; e.g., 10.1.0.0/255.255.0.0 or 172.30.123.0/24.

-ca <client-access-mode> is an option to specify client access. This can be *ReadOnly* or *ReadWrite*.

-E (--insecure) is a security option. The default is “secure”. Enter -E (--insecure) if the client’s operating system does not use a reserved port for NFS.

-Q (--squash) is an option that maps user IDs. The default is “root”, which remaps UID=0 (the root user) to “nfsnobody:nasgrp”*.* Enter -Q (--no) to prevent mapping UIDs or -Q (--all) to map all UIDs to “nfsnobody:nasgrp”.

Example for an NFS share:

```
iscli createshare --server-name=h41-93 --vdevid=10 --folder-name=shrFldTest --client-name=h41-93 --access-
mode=ReadOnly
```

Delete a share

```
iscon deleteshare -s <server-name> [-u <username> -p <password>]
-v <vdevid> -x <export-path> | -W <cifs-share-name> [-X <rpc-timeout>]

iscon deleteshare --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> --export-path=<export-path> | --cifs-share-name=<cifs-share-name>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes the specified share.

-v (--vdevid) is required to specify the device id of the NAS resource.

-x (--export-path) or -W (--cifs-share-name) is required.

-x (--export-path) is required to identify an NFS share and must include the full path; e.g., /nas/NAS-00005/fds/account-depart/share-folder-1

Refer to ‘Create a share’ for definitions of other parameters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

Assign a share to a CIFS client

```
iscon assignshare -s <server-name> [-u <username> -p <password>]
-W <cifs-share-name> -A <cifs-account-name> [-a <cifs-account-access>][-G]
[-X <rpc-timeout>]

iscon assignshare --server-name=<server-name>
[--server-username=<username--server-password=<password>]
--cifs-share-name=<cifs-share-name>
--cifs-account-name=<cifs-account-name>
[--cifs-account-access=<cifs-account-access>]
[--group]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command assigns an existing CIFS share to an existing CIFS account.

-W (--cifs-share-name) and -A (cifs-account-name) are required.

-a (--cifs-account-access) is optional. Permitted value is *ReadOnly* or *ReadWrite* (default).

-G (--group) is optional for a group account.

Refer to ‘Create a share’ for definitions of other parameters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Unassign a share from a CIFS client

```
iscon unassignshare -s <server-name> [-u <username> -p <password>]
-W <cifs-share-name> -A <cifs-account-name> [-X <rpc-timeout>][-G]

iscon unassignshare --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--cifs-share-name=<cifs-share-name> --cifs-account-name=<cifs-account-name>
[--group]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command unassigns a share from a CIFS account.

-W (--cifs-share-name) and -A (cifs-account-name) are required.

-G (--group) is optional for a group account.

Refer to 'Create a share' for definitions of other parameters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Assign a share to an NFS client

```
iscon assignnfs -s <server-name> [-u <username> -p <password>]
-v <vdevid> -x <export-path>
-cn <client-name> | -m <client-machine>
-ca <client-access-mode> [-E] [-Q <squash>]
[-X <rpc-timeout>]

iscon assignnfs --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> --export-path=<export-path>
--client-name=<client-name> | --client-machine=<client-machine>
[--access-mode=<client-access-mode>][--insecure] [--squash=<squash>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command assigns an existing NFS share to an existing NFS client.

-v (--vdevid) is required to specify the device ID.

-x (--export-path) is required.

Either -cn (--client name) or -m (--client-machine) is required in order to assign an NFS share.

Refer to 'Create a share' for definitions of other parameters.

-ca <client-access-mode> is an option to specify client access. This can be *ReadOnly* or *ReadWrite*.

-E (--insecure) is a security option. The default is “secure”. Enter -E (--insecure) if the client’s operating system does not use a reserved port for NFS.

-Q (--squash) is an option that maps user IDs. The default is “root”, which remaps UID=0 (the root user) to “nfsnobody:nasgrp”. Enter -Q (--no) to prevent mapping UIDs or -Q (--all) to map all UIDs to “nfsnobody:nasgrp”.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Unassign a share from an NFS client

```
iscon unassignnfs -s <server-name> [-u <username> -p <password>]
-v <vdevid> -x <export-path> -cn <client-name> | -m <client-machine>
[-X <rpc-timeout>]

iscon unassignnfs --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<vdevid> --export-path=<export-path>
--client-name=<client-name> | --client-machine=<client-machine>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command unassigns a share from an existing NFS client.

-v (--vdevid) is required to specify the device ID.

-x (--export-path) is required.

Either -cn (--client name) or -m (--client-machine) is required.

Refer to ‘Create a share’ for definitions of other parameters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Clients

Add a CIFS client

```
iscon addcifsclient -s <server-name[-u <username> -p <password>]
-c <client-name> -t user -W <user-password> | -t group
[-d <description>]
[-X <rpc-timeout>]

iscon addcifsclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> --client-type=user
--password=<user-password> | --client-type=group
[--description=<description>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds a CIFS client on a server with authentication mode set to USER.

-c (--client-name) is the user/group name.

-t (--client-type) is required with one of the following values: *user* (default) or *group*.

-W (--user-password) is required if the client type is *user*.

-d (--description) can be added for client type of user or group. The maximum length is 128 characters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete a CIFS client

```
iscon deletecifsclient -s <server-name> [-u <username> -p <password>]
-c <client-name> -t <client-type>
[-X <rpc-timeout>]

iscon deletecifsclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> --client-type=<client-type>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes a CIFS client on a server with authentication mode set to USER.

-c (--client-name) is the user/group name.

-t (--client-type) is required with one of the following values: *user* (default) or *group*.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Add an NFS client

```
iscon addnfsclient -s <server-name> [-u <username> -p <password>]
-cn <client-name> -m <client-machine> [-d <description>]
[-X <rpc-timeout>]

iscon addnfsclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> --client-machine=<client-machine>
[--description=<description>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds an NFS client, which can be a single host or a subnetwork.

-cn (--client name) and -m (--client-machine) are required in order to add an NFS client.

-cn (--client-name) is the name of the NFS client.

-m (--client-machine) defines the client machine. Enter either the "IP address" or “Full Hostname". On certain subnetworks, this is defined using "Network IP/Netmask".

- "Netmask" can be specified using the IP address; e.g., "255.255.0.0" or CIDR notation (a prefix length); e.g., "16". Either specification has the same meaning in the command.
- Whether you use the "IP address" or "CIDR” format to specify the Netmask, the Network IP should have the same number of characters as Netmask; e.g., 10.1.0.0/255.255.0.0 or 172.30.123.0/24.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete an NFS client

```
iscon deletenfsclient -s <server-name> [-u <username> -p <password>]
-cn <client-name> [-X <rpc-timeout>]

iscon deletenfsclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes an NFS client.

-cn (--client-name) is the name of the NFS client.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Set security mode

```
iscon setsecuritymode -s <server-name> [-u <username> -p <password>]
{ -D -S <domain-controller> -N <domain-username> -P <domain-password>
[-B <bind-point>] [-O <OU>,<...>]
[-a <administrative-username> -w <administrative-password>] |
-U [-W <workgroup>] }
[-I <RANGE>,<...>] [-G <RANGE>,<...>] [-f] [-c <comment>]
[-X <rpc-timeout>]

iscon setsecuritymode --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
{ --domain-mode --domain-controller=<domain-controller>
--domain-username=<domain-username> --domain-password=<domain-password>
[--bind-point=<bind-point>] [--organizational-units=<OU>,<...>]
[--administrative-login=<administrative-login> --login-password=<login-password >] |
--user-mode [--workgroup=<workgroup>] }
[--uid-ranges=<RANGE>,<...>] [--gid-ranges=<RANGE>,<...>]
[--force] [--comment=<comment>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command is used to set up CIFS user security mode in domain mode or user mode.

Domain mode:

-D (--domain-mode) is required to set up domain mode.

-S (--domain-controller) is required to indicate the domain controller(s). Up to two controllers can be specified in the format “-S DC1,DC2” or “--domain-controller=DC1,DC2”.

-N (--domain-username) -P (--domain-password) are required to define the account VTL will use to log into Active Directory on the domain server.

-B (--bind-point) is optional to direct VTL to a starting point or a single OU tree on the domain controller for retrieving organizational units. The default value is "/", which is the root of the OU tree.

-O (--organizational-units) is optional to specify one or more organizational units (OU) to which you will offer NAS shares, in the format "OU1,OU2,OU/OU3,...". Invalid OUs will be omitted. The process will be terminated if no OUs are valid.

- When OU(s) are not specified and the bind point is set to "/", first-level OUs on the domain controller will be used.
- When OU(s) are not specified but a bind point was specified, the bind point will be used.

-a (--administrative-login) -w (--login-password) An account with administrative privileges to allow this NAS server to become a trusted member within the domain realm. If these are empty, the default values will be domain-username and domain-password. These are optional and will not be saved.

User mode:

-U (--user-mode) indicates to set up user mode.

-W (--workgroup) is optional to specify the workgroup name.

Common arguments:

-I (--uid-ranges) -G (--gid-ranges) These are optional to specify id ranges preserved for CIFS users or groups. The format of an id range is firstID-lastID, e.g., 12345-23456. It can be a list of id ranges separated by commas, "RANGE1,RANGE2,...". If no specified id range list when changing mode, the default range is "45435-65436". When assigning new id range list without changing mode, the given range value will be validated and insert to list if is valid.

-f (--force) will force to remove all CIFS shared assignment(s) when changing security mode.

-c (--comment) is the description of CIFS service. The maximum length is up to 128 characters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Reports

The reports below can be generated through the command line interface. These reports allow you to select a date range. The definition of a day (midnight to midnight or noon to noon) is set in the console (right-click the *Reports* object and select *Properties* --> *Other* tab).

Create NAS Resource usage report

```
iscon createnasresourceusagereport -s <server-name> [-u <username> -p <password>]
-n <nas-id>
[-z <report period> | -D <date-range>]
[-o <filename>]
[-X <rpc-timeout>]

iscon createnasresourceusagereport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--nas-id=<nas-id>] [--report-period=<report-period> | --date-range=<date-range>]
[--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing NAS device information. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

The default report is for the current day.

-n (--nas-id) is an option to specify a single device.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The default value is: "-z t" (today).

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is “IndividualNasResourceUsage-MM-DD-YYYY--hh-mm-ss”.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create NAS CIFS share usage report

```
iscon createnascifsshareusagereport -s <server-name> [-u <username> -p <password>]
-S <share-name>
[-z <report period> | -D <date-range>]
[-o <filename>]
[-X <rpc-timeout>]

iscon createnascifsshareusagereport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--share-name=<share-name>
[--report-period=<report-period> | --date-range=<date-range>]
[--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing information about the specified NAS CIFS share. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-S (--share-name) is required to specify the name of the share.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The default value is: "-z t" (today).

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: NASShareUsage-MM-DD-YYYY--hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create NAS statistics summary report

```
iscon createnasstatisticssummaryreport -s <server-name> [-u <username> -p <password>]
[-d] [-v] [-c]
[-z <report period> | -D <date-range>]
[-o <filename>]
[-X <rpc-timeout>]

iscon createnasstatisticssummaryreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--nas-isdedup] [--nas-isrepserver] [--nas-isrepclient]
[--report-period=<report-period> | --date-range=<date-range>]
[--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing NAS statistics. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-d (--nas-isdedup) is an option to include deduplication information in the report.

-v (--nas-isrepserver) is an option to include replication server information in the report.

-c (--nas-isrepclient) is an option to include replication client information in the report.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the range of dates the report should cover. There is no limit to the date range, but future dates are ignored. Valid formats are:

- date range: YYYYMMDD-YYYYMMDD
- specific day: YYYYMMDD

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The default value is: "-z t" (today).

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the output file name is not specified, the default file name is: NASStatisticsSummary-MM-DD-YYYY--hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

NAS information

Get deduplication statistics

```
iscon getdedupestatsinfo -s <server-name> [-u <username> -p <password>]
[-T <disp-type>] [-l]
[-X <rpc-timeout>]

iscon getdedupestatsinfo --server-name=<server-name>
[--server-username=<username>
--server-password=<password>]
[--disp-type=<disp-type>] [--longlist]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves repository usage statistics and the global deduplication ratio from the associated deduplication (SIR) cluster or server.

-T (--disp-type) is the option to specify the display type:

- “Usage” displays repository usage statistics.
- “Ratio” displays the global deduplication ratio only.

-l (--longlist) is the option to display the long format.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get realtime performance information

```
iscon getrealtimeperformanceinfo -s <server-name> [-u <username> -p <password>]
-T <disp-type>
[-b [-n <number>] [-d <delay-time>]]
[-S <source-server>] [-v <vdevid>] [-l]
[-X <rpc-timeout>]

iscon getrealtimeperformanceinfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--disp-type=<disp-type>
[--batch-mode [--number=<number>] [--delay-time=<delay-time>]]
[--source-server=<source-server>] [--vdevid=<vdevid>] [--longlist]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves real-time performance statistics for the specified operation.

-T (--disp-type) is required in order to specify the type of performance for which to display statistics:

- “Dedup” - display deduplication performance
- “RepliOut” - display outgoing replication performance
- “RepliIn” - display incoming replication performance
- “FileSystem” - display file system performance

-b (--batch-mode) is the option to include performance of batch mode operations.

-n (--number) is the option to specify the number of iterations of batch mode operations.

-d (--delay-time) is the option to specify the delay in seconds between two iterations of batch mode operations.

-S (--source-server) is the option to specify the name of the source server for incoming replication.

-v (--vdevid) is the option to return performance statistics for a specific virtual device ID.

-l (--longlist) is the option to display the long format.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get folder details

```
iscon getfolderdetailedinfo -s <server-name> [-u <username> -p <password>]
-fp <folder-path>]
[-fn <file-name>] [-l]
[-X <rpc-timeout>]

iscon getfolderdetailedinfo --server-name=<server-name>
[--server-username=<username>
--server-password=<password>]
--folder-path=<folder-path>
[--file-name=<file-name>] [--longlist]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command returns information on all files in the specified folder.

-fp (--folder-path) is required - specify the folder path.

-fn (--file-name) is the option to only display information only for a specified file on the folder. For example, to display information for a file whose full path is “/nas/NAS-00001/fds/myfolder/myfile.dat”, specify
“-fp /nas/NAS-00001/fds/myfolder -fn myfile.dat”.

-l (--longlist) is the option to display the long format.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

NAS utilities

Start Fast Copy

```
iscon startfastcopy -s <server-name> [-u <username> -p <password>]
-fs <source-path> -ft <target-path> [-f]
[-X <rpc-timeout>]

iscon startfastcopy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--source-path=<source-path> --target-path=<target-path> [--force]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command sets up a copy procedure between NAS resources within the same server. If source files have been deduplicated, the stub file will be copied so that the file does not have to be restored.

-fs (--source-path) is required to specify the source directory.

-ft (--target-path) is required to specify the target directory; this can be a new or existing directory.

-f (--force) is an option to overwrite the existing target directory.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

Example: #iscon startfastcopy -s xxx.x.x.x -fs/nas/NAS-xxx/fds/source -ft /nas/NAS-xxx/fds/target/8

The command returns the process identifier (pid) and status:

Fastcopy process started with pid 18313

Command: startfastcopy executed successfully.

Stop Fast Copy

```
iscon stopfastcopy -s <server-name> [-u <username> -p <password>]
-fp <pid>
[-X <rpc-timeout>]

iscon stopfastcopy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--pid=<pid>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command stops an in-progress copy procedure between NAS resources. If you do not have the process identifier (pid), run *ps aux* | *grep "fdscli fastcopy"* on the server to obtain it.

-fp (--pid) is required to specify the process to stop.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

Get Fast Copy info

```
iscon getfastcopystatinfo -s <server-name> [-u <username> -p <password>]
-fp <pid>
[-X <rpc-timeout>]

iscon getfastcopystatinfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--pid=<pid>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves information about the specified Fast Copy process, including the directory being copied, its size, and percentage of completion. If you do not have the process identifier (pid), run *ps aux* **|** *grep "fdscli fastcopy"* on the server to obtain it.

-fp (--pid) is required to specify the process to retrieve.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

SNMP Integration

VTL provides Simple Network Management Protocol (SNMP) support to integrate VTL management into an existing enterprise management solution, such as HP OpenView, CA Unicenter, IBM Tivoli NetView, or BMC Patrol.

VTL can send different types of information to your SNMP manager:

- **Event Log messages** - By default, Event Log messages (informational, warnings, errors, and critical errors) are not sent, but you may want to configure VTL to send certain types of messages as traps to your SNMP manager. Refer to 'Server properties' to configure this.
- **MIBs** - Each MIB (Management Information Base) monitors information and processes in VTL. You will need to compile the VTL MIBs into your SNMP manager. The procedure to do this will vary by SNMP manager. You will need to compile the following VTL MIB files which can be found in `$ISHOME/etc/snmp/mibs`:
    - falconstor-all.mib
    - falc-vtl-trap.mib

OID

Each MIB object has a unique ID, an OID (object ID). This OID is comprised of a fixed number, followed by the FalconStor ID (7368) and a unique product/object number. For example, 1.3.6.1.4.1.7368.1.1.0.1001.

Community name

By default, VTL uses `falcon` as the community name for MIB browsing. This is different than the community name used for SNMP traps that is set through *Server properties*.

If you need to change the community name used for MIB browsing, you need to modify the setting in the *Access Control* section of the `$ISHOME/etc/snmp/snmpd.conf` file.

VirtualTapeLibraryMIB

The falcVtlMonitor MIB has multiple tables that display different capacity, usage, and performance statistics.

falcVtlMonitorMIB - falcVtlMonCapacity

Displays VTL capacity and usage statistics.

falcVtlCapCacheGeneralInfo

Displays VTL capacity statistics.

Object	Description
BackupCacheCapacityAvailable	Available cache capacity, in GB
BackupCacheCapacityTotal	Total cache capacity, in GB
CacheCapacityUsed	Cache used space, in GB
falcUnassignedVTLCacheCapacity	Unassigned cache space, in GB
CacheCapacityPercentFree	Free cache space percentage
CacheCapacityPercentUsed	Used cache space percentage
CacheCapacityPercentUnassigned	Unassigned cache space percentage

LibCacheUsage

Displays cache usage by all tapes in each individual library.

Object	Description
falcVtlLibID	Virtual tape library ID
falcVtlLibUsedByAllVtapes	Space used by virtual tapes, in MB
falcVtlLibUsedByAllMixedVIT	Space used by mixed tapes, in MB
falcVtlLibUsedByAllVIT	Space used by VIT tapes, in MB
falcVtlLibPendingDedup	Available space pending deduplication of virtual tapes and mixed tapes, in MB

PolicyCacheUsage

Displays cache usage by all tapes in each policy.

Object	Description
CacheIndex	Index key
CacheName	Policy name
UsedByAllVtapes	Space used by virtual tapes, in MB
UsedByAllMixedVIT	Space used by mixed tapes, in MB
UsedByAllVIT	Space used by VIT tapes, in MB
PendingDedup	Available space pending deduplication of virtual tapes and mixed tapes, in MB

falcVtlMonitorMIB - falcVtlMonPerformance

Displays VTL performance and statistics.

falcVtlPerfOneDayIntervalDataInfo

Displays VTL performance for the past 24 hours. Note that no data will be available for a failover server that has been taken over.

Object	Description
DataWritten	Amount of data written for all job types (backup, deduplication, import), in MB
CacheSpaceUsed	Amount of space used after compression, in MB
DataRead	Amount of total uncompressed data read, in MB
DataCompressedRead	Amount of compressed data read, in MB
ReplRawDataTx	(VTL-S servers only) Raw replication data transferred, in MB
ReplActualDataTx	(VTL-S servers only) Actual amount of replication data transferred, in MB
ReplTotalTapesTx	(VTL-S servers only) Number of tapes replicated

falcVtlMonPerformanceInfo

Displays VTL performance.

Object	Description
AvgDedupRatio	(VTL-S servers only) Average tape total deduplication ratio
CompressRatio	Average tape compression ratio

falcVtlHistoryMIB

The falcVtlHistoryMIB has multiple tables that display historical information about activity on the server, dashboard, tape history, and deduplication status.

Activity

Displays tape activity history, including operations related to client backup and restore, import, export, and deduplication:

Object	Description
StartTime	Start time for the tape operation
EndTime	End time for the tape operation
VlibraryID	Virtual library in which the tape was present when the operation took place
VDriveID	Virtual drive used by the operation
VTapeID	Virtual tape used for the operation
Operation	Operation performed
WriteDataMB	Uncompressed data written, in MB
WriteCompressDataMB	Compressed data written, in MB
ReadDataMB	Uncompressed data read, in MB
ReadCompressDataMB	Compressed data read, in MB
StartEODMB	Location on the tape that marks the end of data at the start of the operation (in MB to denote the location as offset from the beginning of the tape)
EndEODMB	Location on the tape that marks the end of data at the end of the operation (in MB to denote the location as offset from the beginning of the tape)
Barcode	Tape barcode

DashStatistics

Displays the following dashboard activity history:

Object	Description
TimeStamp	Date and time the data was collected
DiskSpaceTotal	Total disk space used, in GB
DiskSpaceAvailable	Available disk space, in GB

Object	Description
PerformanceRead	Read performance, in KB/sec (calculated at the HBA level)
PerformanceWrite	Write performance, in KB/sec (calculated at the HBA level)

TapeHistory

Displays the following statistics history for each tape in a deduplication policy:

Object	Description
TapeID	Tape ID
TapeName	Tape name
TapeBarcode	Tape barcode
PolicyID	Policy ID
DriveSerialNumber	Drive serial number
TapeOperation	Operation conducted on tape: Scan Tape, Write Tape, Resolve Tape, Empty Tape, Upgrade Tape, Verify Tape, Inline Parsing Tape
Parser	Data type/parser used: ArcServe, Atempo, Bacula, Commvault, DataProtector, IBMiSeries, Legato, MicrosoftTapeFormat, NDMP, Netbackup, NetbackupFalconstorOpenStorage, Netvault, OracleSecureBackup, Syncsort, TSM, Virbak, Unknown
VTLServer	VTL server
Result	Result of operation: Success, Failed
ErrorCode	Error code
Message	Error message
DedupeStatus	Deduplication status: Total Operation Running, Total Operation Finished, Total Operation Paused, Total Operation Failed, Total Operation Preparing, Total Operation Stopped, Queued, Running, Finished, Paused, Failed, Preparing, Stopped, Unknown
TapeSize	Tape size, in GB
ScannedData	Amount of data processed, in GB
NumberOfTapeBlocks	Physical allocation tape blocks (divided by 1024)
NumberOfFiles	Number of files
Data	Total amount of data on tape, in GB
UniqueData	Unique data, in GB
VVTapeData	Used size of a VIT (the amount of data written to a VIT after deduplication) in GB

Object	Description
WrittenVITData	Same as VVTapeData
ScanStartTime	Scan start time
ScanEndTime	Scan end time
RunTime	Run time, in seconds
Throughput	Throughput, in MB per second
WriteVITTime	VIT write time, in seconds
WriteVITThroughput	VIT write throughput, in MB per second
ReplicationStart	Replication start time
ReplicationEnd	Replication end time
ReplicationThroughput	Replication throughput, in MB per second
DataReplicated	Amount of data replicated, in MB
ReplicationStatus	Replication status: New, Idle, In Progress, Replication Failed, Scanning, Scan Failed
ResolverStartTime	Resolver start time
ResolverEndTime	Resolver end time
ResolverThroughput	Resolver throughput, in MB per second
ResolvedVITData	Amount of resolved VIT data, in GB
ResolveUniqueData	Unique data resolved, in GB
ResolveStatus	Resolver status: Total Operation Running, Total Operation Finished, Total Operation Paused, Total Operation Failed, Total Operation Preparing, Total Operation Stopped, Queued, Running, Finished, Paused, Failed, Preparing, Stopped, Unknown
ResolveErrorCode	Resolver error code
PolicyStartTime	Policy start time
PolicyEndTime	Policy end time
RunType	Run type (manual, scheduled, tape caching)
RunStatus	Status: Complete, Incomplete

falcVtlServer

Displays information about the VTL server and the options configured.

Processor

Displays information about all of the processors in the VTL server.

Object	Description
falcVtlProcessorInfo	Processor type

NetInterface

Displays information about all of the network interfaces in the VTL server.

Object	Description
falcVtlNetInterfaceInfo	Network interface and maximum transfer unit of each IP packet (MTU)

FailoverInfo

Displays failover information for the configured failover node.

Object	Description
Configuration Type	Type of failover (active-passive or mutual)
Failover Partner	Name of failover partner
Failover Configuration Properties	VTL servers in failover group Self check interval Heartbeat interval Recovery setting Current failover state

FCInfo

Displays failover information about all of the Fibre Channel HBAs in the VTL server.

Object	Description
WWPN	World Wide Port Name
Mode	Mode: Initiator, Target, Standby, Null.

falcVtlServerOptionsInfo

Displays VTL server options.

Object	Description
falcVTLServerOptionsInfo	Displays VTL server options (Fibre Channel, iSCSI, VTL database, Email Alerts, Hosted Backup, NDMP) and informs whether each is enabled or disabled.

falcVtlServerInfo

Displays VTL server information.

Object	Description
ServerName	Server hostname
LoginMachineName	Server IP
ServerVersion	Server version
OsVersion	Server operating system version
KernelVersion	Server kernel version
Memory	Amount of server memory
SwapSpace	Amount of server swap space

falcVtlLibrarySystem

Displays information about virtual and physical libraries, clients, and physical resources. It also contains deduplication policy information, if applicable.

VirtualLibrary / falcVtlVirtualLibsInfo

Displays information about the configuration and properties of each virtual tape library.

Object	Description
falcVtlVirtLibID	Virtual tape library ID
falcVtlVirtLibName	Virtual tape library name
falcVtlVirtLibVendorID	Vendor ID
falcVtlVirtLibProductID	Product ID
falcVtlVirtLibRev	Firmware version
falcVtlVirtLibNumSlots	Number of slots
falcVtlVirtLibNumDrives	Number of drives
falcVtlVirtLibBarcodeBegin	First barcode in range for library
falcVtlVirtLibBarcodeEnd	Last barcode in range for library
falcVtlVirtLibTapeCapacityOnDemand	Indicates if tape capacity on demand (COD) is enabled
falcVtlVirtLibInitAllocSize	Initial tape size, in MB (if tape COD is enabled)
falcVtlVirtLibIncrementSize	Incremental amount, in MB (if tape COD is enabled)
falcVtlVirtLibMaxCapacity	Maximum tape capacity, in MB (if tape COD is enabled)
falcVtlVirtLibAutoArchive	Indicates if auto archive is enabled
falcVtlVirtLibMediaType	Media type
falcVtlVirtLibNumTapes	Number of tapes in library
falcVtlVirtLibSerialNum	Serial number of library
falcVtlVirtLibAutoReplication	Indicates if auto replication is enabled
falcVtlVirtLibAutoTapeCaching	Indicates if tape caching is enabled
falcVtlVirtLibTapeDuplication	Indicates if tape duplication is enabled
falcVtlVirtualLibsInfo	Total number of virtual tape libraries

VirtualDrive / falcVtlVirtualDrivesInfo

Displays information about the configuration and properties of each virtual tape drive (inside a virtual tape library, standalone, and used for deduplication).

Object	Description
falcVtlVirtDriveID	Virtual tape drive ID
falcVtlVirtDriveName	Virtual tape drive name
falcVtlVirtDriveVendorID	Vendor ID (“FALCON” for deduplication tape drives)
falcVtlVirtDriveProductID	Product ID (“SIR” for deduplication tape drives)
falcVtlVirtDriveRevision	Firmware version
falcVtlVirtDriveMediaType	Media type (“SIR001” for deduplication tape drives)
falcVtlVirtDriveLocationType	Virtual tape drive location (virtual tape library or standalone)
falcVtlVirtDriveLocationID	ID of the virtual tape library where the virtual tape drive resides
falcVtlVirtDriveGBRead	GB read from this tape drive
falcVtlVirtDriveGBWritten	GB written to this tape drive
falcVtlVirtDriveCompression	Indicates if compression is enabled
falcVtlVirtDriveStatus	Operational status: unknown, empty, loaded, ejected (but not removed), offline, passthrough (all commands will be sent), becoming ready, unloading, online
falcVtlVirtualDrivesInfo	Total number of virtual tape drives

VirtualTape / falcVtlVirtualTapesInfo

Displays information about the configuration and properties of each virtual tape.

Object	Description
falcVtlVirtTapeID	Virtual tape ID
falcVtlVirtTapeName	Virtual tape name
falcVtlVirtTapeTotalSize	Size, in MB
falcVtlVirtTapeStatus	Status (online, offline)
falcVtlVirtTapeGUID	Globally Unique Identifier (GUID)
falcVtlVirtTapeUsedSize	Used size, in MB
falcVtlVirtTapeBarcode	Barcode
falcVtlVirtTapeMediaType	Media type
falcVtlVirtTapeCapacityOnDemand	Indicates if tape capacity on demand is enabled

Object	Description
falcVtlVirtTapeAutoArchive	Indicates if auto archive is enabled or disabled
falcVtlVirtTapeWriteProtection	Indicates if the tape is write protected
falcVtlVirtTapeLocationType	Tape location (library slot, drive, or vault)
falcVtlVirtTapeLocationID	ID of the virtual device where the tape resides (-1, not applicable, for vault)
falcVtlVirtTapeLocationSlot	Slot number of virtual tape library that tape resides in (-1, not applicable, for vault and drive)
falcVtlVirtTapeInitAllocSize	Initial tape size, in MB (if tape COD is enabled)
falcVtlVirtTapeIncrementSize	Incremental amount, in MB (if tape COD is enabled)
falcVtlVirtTapeMaxCapacity	Maximum tape capacity, in MB (if tape COD is enabled)
falcVtlVirtTapeRdeTapeType	Type of tape (regular virtual tape, mixed VIT, pure VIT)
falcVtlVirtTapeRdeEndOfVITData	Location on the tape that marks the end of VIT data (in MB to denote the location as offset from the beginning of the tape)
falcVtlVirtTapeRdeUniqueData	Unique data on the tape, in MB
falcVtlVirtTapeRdeDedupedData	Deduplicated data on the tape, in MB
falcVtlVirtualTapesInfo	Total number of virtual tapes

PhysicalLibrary / falcVtlPhysicalLibsInfo

Displays information about the configuration and properties of each physical tape library.

Object	Description
falcVtlPhyLibID	Physical tape library ID
falcVtlPhyLibName	Physical tape library name
falcVtlPhyLibAllocationType	Same as falcVtlPhyLibName
falcVtlPhyLibVendorID	Vendor ID
falcVtlPhyLibProductID	Product ID
falcVtlPhyLibStatus	Status (online, offline)
falcVtlPhyLibGUID	Globally Unique Identifier
falcVtlPhyLibSerialNumber	Serial number
falcVtlPhyLibNumSlots	Number of slots
falcVtlPhysicalLibsInfo	Total number of physical tape libraries

PhysicalDrive / falcVtlPhysicalDrivesInfo

Displays information about the configuration and properties of each physical tape drive (inside a physical tape library and standalone).

Object	Description
falcVtlPhyDriveID	Physical tape ID
falcVtlPhyDriveName	Physical tape drive name
falcVtlPhyDriveVendorID	Vendor ID
falcVtlPhyDriveProductID	Product ID
falcVtlPhyDriveStatus	Status (online, offline)
falcVtlPhyDriveGUID	Globally Unique Identifier
falcVtlPhyDriveSerialNumber	Serial number
falcVtlPhyDriveState	Operational status: empty, loaded, ejected (but not removed from drive), offline, passthrough (all commands will be sent)
falcVtlPhysicalDrivesInfo	Total number of physical tape drives

PhysicalTape / falcVtlPhysicalTapesInfo

Displays information about the configuration and properties of each physical tape.

Object	Description
TapeSlot	Physical slot the tape is in
TapeBarcode	Barcode
falcVtlPhysicalTapesInfo	Total number of physical tapes

VtlJob / falcVtlJobQueueInfo

Displays information about the tape import/export job queue.

Object	Description
falcVTLJobID	Job ID
falcVTLJobType	Job type: • exportToSADriveAndCopy - Export with copy on a standalone physical tape drive • exportToSADriveAndMove - Export with move on a standalone physical tape drive • exportToPhyLibAndCopy - Export with copy to a physical tape library • exportToPhyLibAndMove - Export with move to a physical tape library • importFromSADriveAndCopy - Import in copy mode from a standalone physical tape drive • importFromSADriveAndRecycle - Import in recycle mode from a standalone physical tape drive • importFromPhyLibAndCopy - Import in copy mode from a physical library • importFromPhyLibAndRecycle - Import in recycle mode from a physical library • createFromCacheMetaDataModeAndCopy - Create cache in copy meta data mode • moveTapeIESlot - Moving tape to an IE slot • importFromPhyLibStackAndCopy - Import from physical library with stacking in copy mode • exportToSADriveStackAndCopy - Export to standalone physical drive with stacking in move mode • exportToPhyLibStackAndCopy - Export to physical library with stacking in copy mode • moveTapeIESlotStack - Move tape to an IE slot with stacking • scanPhyTapeStack - Scan physical tape • exportToPhyLibstackAndMove - Export to physical library with stacking in move mode
falcVTLJobPhysicalLibName	Physical library used for import/export
falcVTLJobPhysicalTapeBarcode	Physical tape barcode used for import/export
falcVTLJobPhysicalSlot	Physical slot number used for import/export
falcVTLJobVirtualLibName	Virtual library used for import/export
falcVTLJobTapeName	Virtual tape used for import/export
falcVTLJobTapeBarcode	Virtual tape barcode used for import/export
falcVTLJobVirtualSlot	Library virtual slot number used for import/export
falcVTLJobStatus	Job status: ready, running, completed, cancelled, or failed falcVTLJobDescription - Job description
falcVtlJobQueueInfo	Total number of import/export drives

ReplicaResource

Displays information about replica resources.

Object	Description
VirtualID	Replica resource ID
VirtualName	Resource name
AllocationType	Resource type: Null, Virtual Device, Direct Device, System, Reserved, Service Enabled Disk, Unknown
TotalSize	Size, in MB
ConfigurationStatus	Status: Online or offline
GUID	Globally Unique Identifier
PrimaryVirtualID	Source server and device in the format <hostname of source>:<virtual device ID>.
ReplicationStatus	Current status of the replication schedule - Replication Failed, New, Idle, Merging, Unknown (stopped)
LastStartTime	Last replication start time

ReplicaPhyAllocationLayout

Displays information about the physical devices which were used to create replica resources.

Object	Description
VirtualID	Replica resource ID
VirtualName	Resource name
Name	Physical device name
Type	Type (primary or mirror) of the physical layout
SCSIAddress	SCSI address of the replica resource in the format <Adapter:Channel:SCSI:LUN>
FirstSector	First sector of the physical device used for this resource
LastSector	Last sector of the physical device used for this resource
Size	Size allocated for the replica resource, in MB

ReplicationPolicy / falcVtlReplicaResourcesInfo

Displays information about tapes with replication policies.

Object	**Description**
ResourceID	Virtual tape ID
ResourceName	Virtual tape name
Option	Replication status (enabled or disabled)
ReplicaServer	Target replica server name
ReplicaDeviceID	Target replica device ID
Schedule	Current status of the schedule: On Schedule, Suspended, or N/A
Watermark	Watermark set to trigger replication
WatermarkRetry	Retry interval if replication fails
Time	Time when replication is scheduled to occur
Interval	Time interval between replication jobs
falcVtlNumOfReplica	Total number of replica resources

DeduplicationPolicy / falcVtlDeduplicationPoliciesInfo

Displays information about deduplication policies.

Object	**Description**
PolicyName	Policy name
PolicyID	Policy ID
NumberOfTapes	Number of tapes in the policy
TriggerType	Trigger type: • endOfBackupTapeFull • endOfBackup • noSchedule • scheduleHourly • scheduleDaily • scheduleSunday • scheduleMonday • scheduleTuesday • scheduleWednesday • scheduleThursday • scheduleFriday • scheduleSaturday • tapeCaching

Object	Description
SirCluster	Deduplication cluster to which data will be deduplicated
ReplicationStatus	Indicates if replication is enabled or disabled
TurboDeduplication	Indicates if turbo deduplication is enabled or disabled
ReplicationMode	Replication mode
falcVtlDeduplicationPoliciesInfo	Total number of deduplication policies.

falcVtlPhysicalResources

Displays information about physical resources.

Storage HBAs

Displays information about each storage HBA.

Object	Description
Number	SCSI adapter number
Info	Model/type
WWPN	World wide port name
Mode	Mode: Target or initiator
AliasWWPN	Alias WWPN, if dual mode
AliasMode	Mode of alias: Target or initiator
GBRead	Amount of data read, in GB, for target ports
GBWrite	Amount of data written, in GB, for target ports

StorageDevices

Information about the hardware specifications and characteristics of each SCSI device (physical tape library, physical tape drive, and storage device).

Object	Description
DeviceType	Access type for the attached device (Direct-Access, Sequential-Access, Medium Changer)
VendorID	Product vendor
ProductID	Product model
FirmwareRev	Firmware version
AdapterNo	SCSI adapter number
ChannelNo	SCSI channel
ScsiID	SCSI ID
LUN	SCSI LUN
TotalSectors	Number of sectors or blocks
SectorSize	Number of bytes in each sector or block
TotalSize	Total size of the device, in MB

Object	Description
ConfigStatus	Status (online or offline)
UsedSize	Space used, in MB
FreeSize	Free space, in MB
StorageOwner	Owner of the device. Storage devices will show the server name. Media changers and drives will be displayed as local owner.

falcVtlPhysicalResourcesInfo

Total number of devices.

Object	Description
Adapters	Number of adapters
Devices	Number of physical devices

falcVtlSanClients

Displays information about clients (backup servers).

SANClient

Displays information about each backup server.

Object	Description
falcVtlSanClientID	Client ID
falcVtlSanClientName	Client name
falcVtlSanClientType	Client type: Fibre Channel, iSCSI, HostedBackup

FCClientResource

Displays information about virtual resources assigned to each FC client.

Object	Description
ResourceID	Virtual resource ID
ResourceName	Virtual resource name
ClientID	Client ID
ClientName	Client name
ResourceAllocType	Resource type: Direct Device Virtual Library, Direct Device Virtual Drive
LUN	SCSI LUN of client
InitatorWWPN	WWPN of the client's initator HBA
TargetWWPN	WWPN of the client's target HBA
Access	Read/write access mode: Null, Read-only, Read/Write, Read/Write Non-Exclusive, Undefined

ISCSIClientResource

Displays information about virtual resources assigned to each iSCSI client.

Object	Description
ResourceID	Virtual resource ID
ResourceName	Virtual resource name
ClientID	Client ID

Object	Description
ClientName	Client name
ResourceAllocType	Resource type: Direct Device Virtual Library, Direct Device Virtual Drive
LUN	SCSI LUN of client
IPAddress	Client IP address
TargetID	Client target ID
TargetName	Client target name

falcVtlSanClientsInfo

Total number of backup servers.

Object	Description
falcVtlNumOfSANClients	Total number of backup servers

Deduplication Repository MIBs

falcSirMonitorMIB

Displays deduplication threshold, capacity and performance related information.

falcSirMonSwapMemoryInfo

Displays swap memory usage information.

Object	Description
Total	Total swap size, in MB
Used	Used swap size, in MB

falcSirMonCapacityInfo

Displays usage information for the cluster.

Object	Description
Name	Deduplication cluster name
DataDiskAvailable	Repository data storage available, in MB
DataDiskTotal	Total repository data storage, in MB
DataDiskAvailablePercentage	Percentage of total repository data storage that is available
IndexDiskAvailable	Index storage available, in MB
IndexDiskTotal	Total index storage, in MB
IndexDiskAvailablePercentage	Percentage of total index storage that is available
RepositoryObjectRetainedPercentage	Percentage of index cache capacity retained by reclamation
RepositoryObjectUsed	Percentage of index cache capacity that is used
RepositoryObjectAvailablePercentage	Percentage of index cache capacity that is available
FolderDiskAvailable	Folder space available, in MB
FolderDiskTotal	Total folder space, in MB
FolderDiskAvailablePercentage	Percentage of total folder space that is available
ReclamationLastRunTime	Last date and time data reclamation was run
ReclamationDataSaved	Data space saved after reclamation, in MB
IndexReclamationLastRunTime	Last date and time index pruning was run
ReclamationIndexDataSaved	Index space saved after reclamation, in MB

falcSirDedupeRatioRangesInfo

Displays deduplication ratio information for the cluster. The statistics are updated every hour; these are not real-time statistics. By default, the statistics are updated on the hour.

Object	Description
AvgTapeDedupeRatioIn1Hour	Average deduplication ratio (data scanned / data stored) for the last hour
AvgVITReplicaDedupeRatioIn1Hour	Average VIT replication deduplication ratio (data replicated / unique data replicated) for the last hour
AvgSIRDedupeRatioIn1Hour	Average deduplication ratio (total data / total compressed) for the last hour
VITsWith1-2DedupeRatio	The number of tapes with a deduplication ratio between 1.0 and < 2
VITsWith2-4DedupeRatio	The number of tapes with a deduplication ratio between 2.0 and < 4
VITsWith4-8DedupeRatio	The number of tapes with a deduplication ratio between 4.0 and < 8
VITsWith8-16DedupeRatio	The number of tapes with a deduplication ratio between 8.0 and < 16
VITsWithGreaterThan16DedupeRatio	The number of tapes with a deduplication ratio greater than 16

falcSirServerPerformanceInfo

Displays information about total data deduplicated. The statistics are updated every day; these are not real-time statistics. By default, the statistics are updated at midnight.

Object	Description
DataDeduplicatedIn24Hour	Amount of data deduplicated by the deduplication server in the past 24 hours, in MB
NumberofTapesDeduplicatedIn24 Hour	Number of tapes deduplicated by the deduplication server in the past 24 hours, in MBs
DataReplicatedIn24Hour	Amount of data replicated by the deduplication target replication server in the past 24 hours, in MBs
AverageDedupeRatioIn24Hour	Average deduplication ratio of data processed by the deduplication server in the past 24 hours, in MBs
AverageCompressionRatioIn24Hour	Average compression ratio of unique blocks processed by the deduplication server in the past 24 hours, in MBs

falcSirMonNodePerformanceInfo

Displays performance information for the cluster. The statistics are updated every day; these are not real-time statistics. By default, the statistics are updated at midnight.

Object	Description
DataDeduplicated	Amount of data deduplicated by the deduplication cluster in the past 24 hours, in MB
OverallDedupeRatio	Overall deduplication ratio by the deduplication cluster
AverageDedupeRatio	Average deduplication ratio for the deduplication cluster in the past 24 hours
AverageCompressionRatio	Average compression ratio for the deduplication cluster of unique blocks in the past 24 hours
ReplicationDataTx	Amount of actual data replicated in the past 24 hours
LastIndexLoadTime	Number of minutes required to load the index into RAM the last time it was loaded

falcSirCluster - falcSirClusterConf

Displays information about the deduplication server related to deduplication cluster configuration

falcSirCCGeneralInfo

Displays general cluster information.

Object	Description
ClusterName	Name of the deduplication cluster
FailoverOption	Indicates if failover is enabled or disabled
Guid	Globally Unique Identifier of deduplication cluster
HBIfName	Heartbeat interface name (such as "eth0")
ReplicationOption	Indicates if replication is enabled or disabled
ReplicaCount	Number of replica target deduplication cluster instances assigned to this cluster (this cluster is the primary)
ReplicatorCount	Number of instances where this cluster is a target replica
SirNodeCount	Number of deduplication server nodes

VTL

Displays information about associated VTL servers.

Object	Description
falcSirCCVtlHostname	Hostname of the VTL server
falcSirCCVtlIP	IP address of VTL server

SIRReplication

Displays source and target replica information.

Object	Description
Type	Indicates if this is a primary (replicator) or target (replica) server
Name	Replication server name
Guid	Globally Unique Identifier of replication server
Protocol	Protocol of replication server (unknown, iSCSI, fibrechannel, tcp)
SirNodeCount	Number of deduplication nodes on the replication server
NodeIPCount	Number of IP addresses for the replication server
NodePortCount	Number of ports for the replication server

SIRReplicationSirNodeIP

Displays deduplication node IP information.

Object	Description
SirName	Deduplication node name of the replication server
Address	Replica node IP address
RepType	Indicates if this is a primary (replicator) or target (replica) server

SIRReplicationSirNodePort

Displays deduplication node port information.

Object	Description
SirName	Deduplication node name of the replication server
Port	Port information of replication server
RepType	Indicates if this is a primary (replicator) or target (replica) server

falcSirCCReclamationInfo

Displays reclamation information.

Object	Description
Option	Indicates if reclamation is enabled or disabled
Interval	Indicates time interval for reclamation, in seconds

SIRNode

Displays information about the deduplication nodes in each deduplication cluster.

Object	Description
Name	Deduplication server name
IPAddress	Deduplication server IP address
Type	Deduplication server type (active, standby)
PowerFactor	Number of CPUs x 2. This number determines the maximum number of deduplication and resolver processes that can run on each deduplication node.

SIRNodeFailover

Displays failover information for a deduplication node.

Object	Description
Name	Name of the failover deduplication server
HeartbeatIP	Heartbeat IP address for failover
HeartbeatNetmask	Heartbeat netmask for failover
HeartbeatGateway	Heartbeat gateway for failover
IpmiIP	IPMI IP address of failover
IpmiNetmask	IPMI netmask of failover
IpmiGateway	IPMI gateway of failover
IPCount	Number of IP adresses for the deduplication failover node

SIRNodeFailoverIP

Displays failover IP address information for a deduplication node.

Object	Description
SIRNodeName	Name of the deduplication server
Name	Heartbeat interface
RealIP	Heartbeat IP address
Netmask	Heartbeat netmask
VirtualIP	Virtual IP address

SIRNodeFibreChannel

Displays Fibre Channel information for a deduplication node.

Object	Description
NodeName	Deduplication node name
WWPN	World Wide Port Name

SIRHashIndexStorage

Displays information about the deduplication node hash index list.

Object	Description
SirName	Deduplication server name
Size	Size of hash index

SIRHashFolderStorage

Displays information about the deduplication node hash folder list.

Object	Description
SirName	Deduplication server name
Size	Size of hash folder

SIRHashDataStorage

Displays information about the deduplication node hash data list.

Object	Description
SirName	Deduplication server name
Size	Size of hash data

falcSirCluster - falcSirClusterStats

Displays cluster statistics about folders, repository usage, and deduplication results

Folder

Displays folder information for the deduplication cluster.

Object	Description
Barcode	Barcode of folder
BackupApp	Backup application: ArcServe, Atempo, Bacula, Commvault, DataProtector, IBMiSeries, Legato, MicrosoftTapeFormat, NDMP, Netbackup, NetbackupFalconstorOpenStorage, Netvault, OracleSecureBackup, Syncsort, TSM, Virbak, Unknown
DedupTime	Time of deduplication
DataType	Parser used: arcserve, mtf, netbackup, netbackup2, tsm, legato, generic, legato2, tsm2, commvault, dprotect, ndmp, arcserve2, osb, atempo, netvault, generic2, generic3, ost, IBMi, syncsort, virbak, bacula
MediaType	Media type
Source	Source VTL server
Details	Detailed description

falcSirCSSSirStatsSummaryInfo

Displays repository deduplication statistics (same as the information displayed in the *Deduplication Statistics Since <date/time>* section of the *Repository Dashboard Summary* in the console). .

Object	Description
StartTime	Time deduplication started
ResetTime	Reset
ElapsedTimeSinceStart	Time elapsed since start time
ElapsedTimeSinceReset	Time elapsed since last reset start time
DataWritten	Data written, in MB
DataStored	Data stored, in MB
RedunElimRatio	Redundancy elimination ratio: (data scanned) / (data stored)

falcSirCSSDedupeResults - DedupeStatsHour

Displays repository deduplication statistics on an hourly basis.

Object	Description
DataWritten	Data written, in MB
DataStored	Data stored, in MB
StartTime	Starting time of deduplication

falcSirCSSDedupeResults - DedupeStatsDaily

Displays repository deduplication statistics on a daily basis.

Object	Description
DataWritten	Data written, in MB
DataStored	Data stored, in MB
StartTime	Starting time of deduplication

falcSirCSSDedupeResults - falcSirCSSDedupeResultsInfo

Displays repository deduplication statistics (same as the information displayed in the *Deduplication Results* section of the *Repository Dashboard Summary* in the console).

Object	Description
Written	Data written, in MB
Stored	Data stored, in MB
RedundElimRatio	Redundancy elimination ratio: (data scanned) / (data stored)

falcSirCSSRepositoryUsage - falcSirCSSRepoObjCapacityInfo

Displays index cache capacity information (same as the information displayed in the *Repository usage* section of the *Repository Dashboard Summary* in the console).

Object	Description
Threshold	Index cache capacity threshold percentage
RetainedByReclamation	Percentage of repository retained by reclamation
UsedSinceReclamation	Percentage of repository used since reclamation
Free	Percentage of repository that is available

falcSirCSSRepositoryUsage - falcSirCSSRepoIndexDiskCapacityInfo

Displays index disk information (same as the information displayed in the *Repository usage* section of the *Repository Dashboard Summary* in the console).

Object	Description
Threshold	Index disk capacity threshold percentage
RetainedByPruning	Percentage of index disk retained by reclamation
UsedSincePruning	Percentage of index disk used since reclamation
Free	Percentage of index disk that is available
falcSirCSSIndexCapacity	Total capacity of index disk, in MB
falcSirCSSIndexSpaceUsed	Total index disk space used, in MB

falcSirCSSRepositoryUsage - falcSirCSSRepoDataDiskCapacityInfo

Displays data disk information (same as the information displayed in the *Repository usage* section of the *Repository Dashboard Summary* in the console).

Object	Description
Threshold	Data disk capacity threshold percentage
RetainedByReclamation	Percentage of data disk retained by reclamation
UsedSinceReclamation	Percentage of data disk used since reclamation
Free	Percentage of data disk that is available
Capacity	Total capacity of data disk, in MB
SpaceUsed	Total data disk space used, in MB

falcSirCSSRepositoryUsage - falcSirCSSRepoFolderDiskCapacityInfo

Displays folder disk information (same as the information displayed in the *Repository usage* section of the *Repository Dashboard Summary* in the console).

Object	Description
Threshold	Folder disk capacity threshold percentage
RetainedByReclamation	Percentage of folder disk retained by reclamation
UsedSinceReclamation	Percentage of folder disk used since reclamation
Free	Percentage of folder disk that is available
Capacity	Total capacity of folder disk, in MB
SpaceUsed	Total folder disk space used, in MB

Common MIBs

These MIBs monitor both VTL and deduplication servers.

ServiceEntry

Displays the current state of each VTL server process and module.

Object	Description
Name	Process/module name
Type	Type: Process or Module
CurrentState	Current state: service-up/service-down

falcServerInfo

Displays the current system status.

Object	Description
falcGeneralSystemStatus	Current status

falcEvents

The falcEvents MIB displays the same errors, warnings, informational messages, and attention required information displayed in the console.

The falcEvents MIB has the following tables:

Table	Description
ErrorEvent	Displays ID, date, and description for all error events logged in the Event Log
WarningEvent	Displays ID, date, and description for all warning events logged in the Event Log
InfoEvent	Displays ID, date, and description for all informational events logged in the Event Log
Attention	Displays date, category, and description for all attention required events logged on the *Attention Required* tab.
CriticalEvent	Displays ID, date, and description for all critical error events logged in the Event Log

CommonMIBs-Traps

This MIB contains warning, error, and critical messages generated by VTL servers that are not included in VirtualTapeLibraryMIB-Traps because they are too new to have been compiled into VirtualTapeLibraryMIB-Traps. Trap messages include a probable cause and a suggested action. A list of the most common error and warning messages are listed in the ‘Error codes’ section of the Troubleshooting chapter.

Troubleshooting

This section contains general troubleshooting information and a list of error codes generated by VTL and SIR servers.

Product registration

If you are unable to complete offline activation successfully, try the following solutions:

1. In order to prevent the possibility of unsuccessful email delivery to the FalconStor activation server, disable Delivery Status Notification (DSN) before you send the activation request email to `Activate.Keycode@falconstor.com`.

2. If you do not receive a reply to your offline activation email from the FalconStor activation server within one hour after sending it, check your email encoding and change it to UNICODE (UTF-8) if set otherwise, then send the email again.

3. If the reply email indicates that the license is successfully registered but the signature file is not attached, you may have set the name of the license information file improperly; you cannot use a single digit before the suffix in the file name. Change the registration file name to a valid alphanumeric string and then try to register again. If the issue persists, contact Technical Support.

General console operations

The VTL console is unable to connect to a VTL server

There are several operations that occur when the console connects to the server. A dialog indicates the current step. If there is a failure, the word *Failed* appears at the end of the step. Determining the current phase of connection can help you pinpoint the problem. It is also possible that the server is busy. Wait for a while and retry. At what step did the connection fail?

- **Connecting to the VTL server** - If the IP address of the server has recently changed, delete the server from the Console and re-add it. If you entered a server name, try entering its IP address instead. If this does not help or if the IP address has not changed, ping the target machine.

  If ping does not reply, ping other machines in the same subnet. If there is still no response, there is a network problem. Run a network command or utility to show the status of the network.
- **Verifying user name and password** - Check the user name and the password. You may use the root password or any other administrator or read-only user that you have created with VTL previously. Make sure the

user name and password exist on the server by opening a local session. The password is case-sensitive. Make sure the *Caps Lock* key is not pressed on the keyboard.

From the machine where the VTL console is installed open an SSH session to the VTL server. Log on to the server with the same user name and password. If the connection between the two machines is fine, the console should be able to connect to the server unless some important server module is not running, such as the communication module. To see the status of all modules, at the machine where VTL server is running, go to the system console and type:

```
vtl status.
```

If a module has stopped, restart it with the command:

```
vtl restart <module name>
```

Afterwards, go back to the console and retry connecting to the server.

- **Retrieving the server configuration** - If there is something wrong with the configuration, an error message may appear. Contact technical support.
- **Checking the VTL license** - Contact technical support.
- **Expanding the VTL server node** - This may be due to high memory usage. Check the memory consumption on the machine. If it is very high, stop all unnecessary processes. If the problem persists or if the memory consumption is normal, contact technical support.

Requested operations cannot be performed from the console

Sometimes the VTL server is very busy with operations that cause high CPU utilization (such as expanding tapes or data *compression).*

You can check the Event Log or syslog (/var/log/messages) for messages that show you the current activity of the system.

If you see messages such as *Server Busy* or *RPC Timeout*, you should wait awhile and retry your action after the current operation finishes.

If the problem persists or the server is not really busy, contact technical support.

Console operations are very slow

Check console machine memory usage

On the machine where you are using the VTL console, use the appropriate system utility (such as Task Manager) to show the memory usage of all running processes. If the memory usage is unusual, stop all unnecessary processes from running or provide more memory.

Check server activity

Sometimes the VTL server is very busy performing heavy processing. You can check the Event Log or syslog (/var/log/messages) for excessive pending SCSI commands on a single SCSI queue that may delay update requests coming from the console. Also, try starting a second instance of the console. If the second console cannot establish connections, that means the server is busy with previous RPC operations.

If this is the case, you should wait awhile and retry your action after the current processing finishes.

If the problem persists or the server is not really busy, contact technical support.

Physical resources

The VTL console does not show physical storage devices correctly

There are several steps to try when physical storage devices have been connected/ assigned to the VTL server yet they are not showing in the VTL console.

Rescan devices

Perform a rescan from the VTL console (right-click the *Physical Resources* object and select *Rescan*). Make sure that the *Discover New Devices* option is selected. Specify a *LUN Range* that you reasonably expect will include the LUN.

Check system log messages

Check the Event Log or syslog (/var/log/messages) for error messages that may correspond to the rescan operation and report failures on SCSI devices. It may be that even though the devices were discovered, they were not accessible due to errors.

Check device type

For external **SCSI devices**, check the following:

- Make sure the system is powered on. Perform a power cycle to make sure.
- Physically make sure all the cable connectors are securely plugged in.
- Verify SCSI termination. This can be quite involved. If you are not sure, you may have to contact the manufacturer of the devices and have their representatives assist with the troubleshooting.

Once the above conditions are verified, determine the SCSI HBA and the proper driver for it. This can normally be accomplished by going to the website of the HBA manufacturer. From the server console, make sure the correct driver for the HBA is loaded properly. If not sure, unload and load the driver again. While doing that, look into the syslog to see if any error messages have been logged corresponding to the action of loading the driver. Under some circumstances, the system may need to be power cycled (not just rebooted) to properly load the drive.

Some **Fibre Channel devices** use VSA (Volume Set Addressing) mode. This addressing method is used primarily for addressing virtual buses, targets, and LUNs. If this is the case, make sure to enable VSA on the VTL initiator driver and use persistent binding. Otherwise, VTL cannot manage the storage.

An HBA port is missing after rebooting and restarting VTL

Be sure to use the default QLogic HBA modules if the QLogic port is direct-connected to the storage.

Loop mode is required for the storage. The default QLogic driver uses "Loop preferred, then Point-to-Point".

Logical resources

Virtual tapes are displayed as "offline" in the console

If a physical resource that was used to create the virtual tape is missing, the tape's status will be offline (missing segment).

From the VTL console determine which physical resources comprise this virtual drive. To do this, highlight the tape in the tree and check the *Layout* tab or look under the *Storage Devices* object. For each physical device, check that:

- It is turned on
- It still exists (has not been removed)
- It is in a normal state and does not show any failure
- There is no failure at the connection level. Check FC connectivity to VTL to make sure that each physical resource is accessible. Refer to 'Fibre Channel connectivity issues' for more information.

Disks are displayed as "offline" in the console

All storage devices must be powered on before the appliance is started. If storage is not available when the appliances are up, reboot all appliances for the system to come up properly.

If you see a disk with a red dot indicating that it is offline, check the underlying physical device status, device connectivity, and switches.

In a VTL-SIR solution, the VTL servers associated with SIR servers access physical resources from SIR for successful deduplication. The SIR data disks are presented to the VTL servers and the SIR tape drives on the VTL servers are assigned to the SIR servers.

During a power outage/recycle of servers, the VTL and SIR servers lose access to each other's physical devices. To ensure that the tape drives are online on the SIR servers and the SIR data disks are online on the VTL servers, rescanning physical resources on the VTL and SIR servers is advised after power outage/recycling the environment.

In the console, if you see an SIR data disk with a red dot indicating that it is offline, right-click *Physical Resources* and select *Rescan*. Make sure that the *Discover New Devices* option is selected.

Tape expansion does not work

Highlight the tape in the console and check that the *Total Size* field shows the correct size of the expanded tape device.

If the console shows the correct size of the expanded virtual tape, the expansion has succeeded but the client machine is having trouble seeing the new size.

Make sure the client machine has been refreshed to see the updated status of its drives. You need to run the utility corresponding to your operating system to rescan the device and discover its new size.

Incorrect size - check Event Log

If the console does not show the correct size of the expanded virtual tape, the expansion was probably not successful. Check the Event Log to look for any error messages regarding the expansion. Errors may appear if:

- There is not enough physical disk space for the expansion. Add more physical storage or change the size of expansion.
- The physical partition is invalid. Check the storage device.
- An IO error occurred.
- An RPC timeout occurred when the expand command was issued. Try the following operation to see if the server is busy:
  - On the VTL server, run the command `top` or `ps –x`
  - Find and stop any unnecessary processes. If you find that the server is too busy, wait to see if the problem persists.

If it is possible to correct the problem, try to do so and then expand the virtual tape again.

Client cannot see tape library/drive as provisioned by VTL

Check device discovery by operating system

Check if the client's operating system sees the device or if it is the backup software that does not see the tape library or drive. Depending on the OS, the new device is indicated in the different ways:

- **Windows** - Tape libraries appear under *Medium Changers* and tape drives under *Tape drives*. Usually the tape drive is indicated as *\tape<index>*.
- **Linux** - The tape library is usually indicated by `/dev/sg<index>` (the *sg* module should be loaded) and the tape drive by `/dev/st/<index>`, `/dev/nst/<index>`, and `/dev/sg/<index>` (The *st* module should be loaded).
- **Solaris** - The tape library is usually indicated by `/dev/sg<index>` (the *sg* module should be loaded) and the tape drive by `/dev/rmt/<index>` (the *st* module should be loaded).
- **AIX** - The tape device is usually indicated by `/dev/rmt<index>` (for LTO1/LTO2) or `/dev/mt<index>` (for DLT/SDLT).

Operating system does not see device

If the operating system does not see the device, you need to troubleshoot virtual device discovery. To do this, in the console, select the virtual device. Check the device status. If the device status is *offline,* that is the problem as clients cannot see an offline device. Refer to the 'Virtual tapes are displayed as "offline" in the console' section for more information.

If the device status is *online*, check the client configuration.

- **Check client assignment** - From the console, right-click the specific client. If you do not see virtual devices on the *Resources* tab, assign them to that client. To share a device between several clients the mode should be *Read/Write non-exclusive*, otherwise device attachment fails.

- **Check VSA addressing** - Some hosts use VSA (Volume Set Addressing) mode. This addressing method is used primarily for addressing virtual buses, targets, and LUNs. If this is the case, make sure to enable VSA on the VTL target driver. Otherwise some clients cannot detect more than eight LUNs on VTL virtual devices.
- **Check FC connectivity** - Refer to 'Fibre Channel connectivity issues' for more information.

Operating system sees device

If the operating system sees the device but the **backup software does not see the device at all**, you need to check the drivers for the backup software. Make sure the driver used corresponds to the nature of the library and also the tape drive. Some backup products recommend using specific versions of drivers. Refer to the backup software manual for such settings or any necessary upgrade. Also, make sure that multiple backup software is not installed on the same backup server as they may conflict with each other.

If the operating system sees the device but the **backup software does not see the device in the expected place**, you need to check serialization. VTL libraries support serialization. Serialization is the conversion of the content of an object into a sequential stream. It identifies the owner of each component, such as robot, slots, and tape drives. If the device appears in the backup software, but it is not attached to the expected component, it may be related to the serialization. Refer to your backup software manual for any patch or upgrade related to serialization on the backup software.

Client sees the tape library/drive but cannot access it

Check device access by OS

Check if the client's operating system can access the device or if it is the backup software that cannot access the tape library or drive.

Depending on the OS you can use a raw device utility. Most of these tools work with tape drives; they are not capable of moving tapes into the drives. Even if some can move tapes, you need to know the exact address of the tape and the drive.

We recommend that you use the console to put a tape in a drive before running these tools. Also, stop the backup software before you use these utilities:

- **Windows** - For IBM Ultrium devices you can use `ntutil`, a command line tool that can check the tape device.
- **Unix systems** - You can use the `mt` or `tar` commands to access the tape device, for example: `mt -f /dev/rmt/0 status`

OS cannot access device

If the operating system *cannot access* the device, you need to troubleshoot virtual device access.

- Go to the storage to verify that it is not in error or in an abnormal state. The assigned devices have to be in read/write mode.
- Check the Event Log or syslog (/var/log/messages) for message indicating IO errors. Such messages usually begin with `log_scsi_error`.

- Check client driver - Go to the client machine and check the adapter driver version. It should be certified for use with VTL.

OS can access device

If the operating system *can access* the device, you need to troubleshoot the backup software. Verify that you have the correct drivers.

Client can no longer access a virtual device (tape library, drive, or SIR resource)

This can have different causes:

- Client machines may lose device access if you switch between a Multi-ID HBA and a single-ID HBA. If this occurs, you should reboot the client machine.
- If either the VTL or the SIR server is shut down for a long period, the devices offered to the clients will time out or be set to *offline*. If this occurs, you will need to perform a rescan from the host machine to regain access.
- ACSLS library users - If you did not select the *Firewall Support* option during configuration, the *portmap* process needs to be running. Otherwise, the system will fail to assign or retrieve the library's status after restarting VTL services or rebooting. To enable *portmap*, you will have to run the following command: chkconfig --`add portmap`

VIT tape is marked “Full”

If you see a VIT marked as “*full*”, check the log to see if there was enough disk space available during the backup but before the deduplication process started.

If there was not enough space, the tape is marked as “*full*” and this status is preserved after deduplication. If this occurs, you must use a different tape for backups.

Export to direct link tape fails

Writing data to a direct link tape will fail if the media types are not the same between your physical and virtual tapes. When you create a cache for your physical tapes, you must make sure that your physical tapes use the same media type as your virtual tapes.

Fibre Channel connectivity issues

This section provides more detail about FC connectivity issues that cause a client to be unable to see virtual tape libraries/drives. This assumes the devices are properly created and assigned to clients.

- **Check ports** to verify that the QLogic BIOS can see the target port of the DSI server to confirm that there is not a problem in the physical environment (HBAs, connections, zoning, etc.). If the target HBA port in the server is properly connected to the initiator HBA of Windows, the QLogic BIOS should see "FALCON IPSTOR DISK ..." for each target HBA connected after it scans devices.
- **Check WWPN** to verify that the client is associated with the proper WWPN. From the console, right-click the client and select *Properties*. Record initiator and target WWPNs. Highlight the *Physical Resources* object and locate the HBA that matches the recorded target HBA WWPN. Highlight the *SNS table* tab for that HBA and look for the WWPN that matches the recorded initiator WWPN. If the WWPN is not correct, unassign the client and assign it again using the appropriate mapping type. If multiple HBAs exist, either from the client host or from the VTL target, look up all entries from all target SNS tables.
- **Check switch mode and speed settings** to verify they are set properly. For example, a QLogic switch should have the port mode set to *F-Port* for point-to-point or *Fabric* for arbitrated loop and tuning should be set properly (i.e., normal, MIN-I, etc.). Certain switches can only support point-to-point.
- **Check switches** to verify that the mode and speed settings are set properly. For example, a QLogic switch should have the port mode set to *F-Port* for point-to-point or *Fabric* for arbitrated loop and tuning should be set properly (i.e., normal, MIN-I, etc.). Certain switches can only support point-to-point. Also, verify the switch status is "healthy", is using approved/tested firmware, and can see the client WWPN.
- **Check switch zoning parameters** to verify that the client HBA is in the same zone as the server target HBA. When using a Multi-ID HBA with dual mode, clients will need to be zoned to the *alias* port. If they are zoned to the *base* port, clients will not see any devices for you to assign.
- **Check FC connections** from the server target port and from the client to the switch.
- **Check all cables** to verify that they are connected properly and the lights are green. Try physically disconnecting and reconnecting the cable connectors, even if the light is green. After that, go back to the console and refresh the SNS. If possible, replace cables to check that they are functional.
- **Check HBA card port speed** in the BIOS and on the switch port.
- **Check the HBA driver** to verify that it is loaded on the client and its version is certified for use with VTL. Also, verify the HBA BIOS version is certified. There are times where it becomes necessary to restart an HBA driver (i.e., downstream I/O errors due to connectivity or a storage system problem).

However, the problem may be at the physical level and restarting the driver may not be enough to clear the underlying problem. For this reason, we recommend power cycling the server instead of just restarting the driver.

Fibre Channel connectivity and Solaris clients

Try the following if a Solaris client cannot detect a Fibre Channel device.

1. Check the /kernel/drv/st.conf file on the Solaris host.

   ```
   name="st" class="scsi" target=3 lun=0;
   name="st" class="scsi" target=3 lun=1;
   name="st" class="scsi" target=3 lun=2;
   ```

   For a target, the entry for each LUN has to comply with the above format. If you are unsure of the target, add LUN1 for several targets, reload the `st` driver and run `devfsadm -I st`

   Check if device files are created under the device file directory, /dev/rmt:

   ```
   #ls -l
   lrwxrwxrwx 43 Apr 16 01:50 0 -> ../../devices/pci@1f,0/pci@1/scsi@1/st@1,1
   ```

   The last two digits are the target and LUN.

   Using this information, add an entry for each LUN´s target in st.conf.

   Reload the `st` driver and run `devfsadm -I`

2. Reload the `st` driver.

   Run `modinfo` | `grep st` to find the number associated with the `st` driver, which is the first number in the entry.

   Run `modunload –I [`*`number from above command`*`]`

   For 32-bit Solaris, run `modload /kernel/drv/st`

   For 64-bit Solaris, run `modload /kernel/drv/sparcv9/st`

3. Check `st` instances.

   Solaris can only use 2,048 `st` instances. If there are too many entries in the /etc/path_to_inst file, delete some unused ones and run `devfsadm` again.

4. Check the device file.

   When device files are created in the /dev/rmt directory, confirm that it is for the assigned devices.

   Unmount all tapes in drives from the console. From Solaris, run `mt –f [`*`device file`*`] status` and you should see a message that the device is offline or has no tape loaded. Run `ls –l /dev/rmt/[`*`device file`*`]` to find out which LUN it is.

   From the console, mount the tape to the drive you want to test. The drive can be determined from the LUN. From Solaris, run `mt –f [`*`device file`*`] status`. If the device file is associated with the assigned drive, you should be able to get information about the drive and tape. The message content depends on the drive type.

To create the device file for media changer, first make sure the proper driver is installed. The configuration depends on the driver.

5. Reboot.

   When all settings have been verified but the VTL device is still not detected, reboot the system.

NDMP

NDMP backup jobs are failing

If you are using NDMP, you must define the hostname of the VTL server in the /etc/hosts file in the format "IPAddress Hostname".
For example: `10.7.2.41      Server41`

The VTL server NDMP daemon fails to start

Because some backup applications use NDMP, if you are running backup software on the VTL server, it should be started after VTL has started and should be stopped before VTL is stopped. Otherwise, the NDMP service that is loaded by the backup software may interfere with VTL's NDMP service.

Replication

There are several aspects to replication: replication configuration, replication process, replica resource promotion, replication configuration removal. If a problem occurs and you get a message on the *Attention Required* tab, check the Event Log or syslog (/var/log/messages). Look for error messages relating to replication. Some common problems are described below.

Replication configuration

Virtual tapes

Replication configuration fails with a “`Failed to add replication target`” error. This can occur if the replica server has a device assigned to it from the primary server. You will need to remove the device assignment before you can create your replication configuration.

Deduplicated tapes

Replication configuration fails between two VTL-S servers with a message saying that the target server’s WWPN is already assigned to client *x*. This can occur if the target server is assigned as a client of a virtual library on the primary server. You will need to remove the client assignment before you can create your replication configuration.

Replication process

Replication fails

When replication finishes successfully, you will see “`replica_fin 10000342 0`”. If you see a number other than zero, there was a problem. In the message below, replication was manually stopped from the primary server or it was stopped because the primary tape was moved into a drive.

```
Jan 13 15:52:54 VTL89-115 kernel: IOCORE1 [iocore|29557]
OnSANREPRequest, SANREP_STOP_REPLICATION
Jan 13 15:52:54 VTL89-115 ipstorcomm
[mgtpipe_exec.c:pipe_thread:3109][29861]:
Rcv'd mgtpipe cmd: 'replica_fin 10000342 1'
```

In the following message, you see `Failed to get virtual tape`. This indicates that the replica tape has been deleted but a request to the tape was delayed and is still trying to get its information. In this case, no action is required.

```
Jan 13 16:18:43 VTL89-115 ipstorcomm
[SANConsoleRPC_proc.c:sanconsolerpcgetvdevinfo2_1_svc:16499][29861]:
Failed to get virtual tape 10000219
```

Replication when a tape is corrupted

Replication appears to be successful, but you get a message in the Event Log similar to the following: "Encountered metadata inconsistency on Virtual Tape VID #. Write protecting tape".

This can be caused due to corruption on the virtual tape. Replication will proceed as long as there are sectors available, even if a tape is corrupted.

Replica resource promotion

You must have a valid replica resource in order to promote it. For example, if a problem occurred (such as a transmission problem or the replica resource failing) during the first and only replication, the replicated data would be compromised and therefore could not be promoted to a primary virtual tape.

Replication configuration removal

You try to remove a replication configuration for a tape but it fails. Search for a message that contains `[crres.c:ReplicationRemoval:6469]`. A return code of `ret=-2146631164` means that the tape that has no data on it (size is zero).

Import/Export

Import/Export does not work as expected

Check tape capacity mismatch

When you Import/export data between a physical tape device and a virtual device, you must make sure the tape devices are of the same type and the same capacity. If they do not have the same capacity, an end-of-media-hit condition occurs and import/export fails.

If data compression is used, make sure the *actual* capacity matches, not just the *compressed* capacity. Import/export will fail when the destination does not have enough space to hold uncompressed data coming from source.

Check job status

Highlight the *Import/Export Queue* and search for a job related to this operation. If the job is in progress, wait until it is completed. If the job is not there and the import/export operation is not done, look at console Event Log to see if there are any job failure messages.

Check barcodes of virtual tapes

When you import data from a physical tape, make sure the virtual tapes have different barcodes. Otherwise, the import operation fails. Use the *Inventory* feature in the console to get the updated bar codes and status from the physical library.

Check physical tape library and device status

Make sure the physical tape library does not show any abnormal situation. For example, the tape drives may require cleaning or tapes may need to be moved to the proper location.

Check element address on the physical library

When you import data, make sure the assignment of drive in VTL follows the element address of the drives in the physical library. Assign the tape drive in the order of their element address.

For example, an import job cannot be executed and the physical library has two DLT 8000 drives with the following configuration:

```
Library SCSI ID: 1
Drive at element address 1200: SCSI ID 10
Drive at element address 1201: SCSI ID 09
```

```
VTL Resources:
ABC-00003
DLT8000-0008: SCSI ID 09
DLT8000-0009: SCSI ID 10
```

In this case you need to unassign tape drives, select first DLT8000-0009, assign it, then select DLT8000-0008, and assign it to the physical library ABC-00003.

Check system log for errors

Check the Event Log or syslog (/var/log/messages). Look for error messages relating to the physical tape library or drive. If you find error messages but cannot find the cause, contact technical support.

NAS resources/shares

Prevent NAS resources from running out of space

Email Alerts allows you to monitor file system usage with the nasusagechk.sh trigger. If the current disk space on any file system larger than the size specified (default is 1024 GB) is lower than the specified threshold (default is 200 GB), an email alert is sent.

To set/change Email Alerts properties, right-click your server and select *Options* --> *Enable Email Alerts*.

Client cannot access NAS resources or shares

CIFS or NFS client

If clients are unable to connect to any NAS shares, try to ping the client machine from VTL and vice versa.

Also, run `ifconfig` or `ipconfig` and see if there are any dropped packets or transmission errors for the NICs.

CIFS client (Domain/Active Directory mode)

There are several things to try when a CIFS client can ping VTL but cannot connect to and access NAS resources.

- For Windows clients, try to access the VTL server using the IP address of VTL instead of the hostname (i.e., \\13.0.0.44 instead of \\myserver).
- If you can connect using the IP address, try to ping the server hostname. If you do not get any response when pinging the server name, the customer’s DNS may have a problem.
- Ping the domain controller to see if you get a response. If no response, VTL is not communicating with the domain controller. If you try to ping using the short name, the ping should return the long name.
- Check whether the VTL server can get user and group information from the domain controller.
  - Run: `getent passwd <user-name>` or `getent group <group-name>` where <`user-name`> is a specific user and <`group-name`> is the name of an existing group on the domain. Look at the results that are returned.
  - Run: `wbinfo --config=$ISHOME/etc/$HOSTNAME/smb.conf -u` and `wbinfo --config=$ISHOME/etc/$HOSTNAME/smb.conf –g` and check results related to users and groups.

CIFS client

Try the following when a CIFS client cannot access a NAS share:

- Check if there is any time difference between the domain controller and the VTL server. The difference should not be greater than three minutes.
- We recommend that you set up NTP and have the domain controller and the VTL server sync to the same NTP server. If NTP is already set up, use the `ntpdate` and `ntpd` commands to update the time.

```
service ntpd stop;
ntpdate ntp-myNTPserver;
service ntpd start
```

- The domain controller can be used as an NTP server in case the other is not available.
- VTL will get tickets from the domain controller. Run the command `klist`, and see if there is any ticket information returned. If not, it means that there is a problem with communication with the domain controller. You can restart the NAS MGTD module to refresh the connection and get new tickets.

If a Windows client is connected to VTL through a firewall and cannot connect to its share, make sure the following ports are not blocked: UDP 137, UDP 138, and TCP 139.

NFS client

Issue: A client must manually reconnect to a NAS share after a server reboot.

Solution: `/etc/fstab` must be modified in order to automatically mount the NFS share at every reboot of the client system in the following way:

```
<hostname or IP address>:/nas/NAS-00006/fds/test /mnt/nfs nfs auto,
user
```

Issue: An NFS client cannot mount NAS resources.

Solution: Verify the NFS client’s IP address and netmask and ping the server `showmount -e <VTL IP address>`.

Issue: A Linux or Unix client can ping VTL but cannot map NAS shares.

Solution:

- Check the console and see if the share is mounted. At the command line, use the command `# mount` to check and see if the share is mounted.
- Check the mount path on the server (located in `/nas/<resource>/fds/<share>`).
- Un-assign and re-assign the client.

Issue: A Linux or Unix client can ping VTL but cannot connect to and access NAS shares.

Solution:

- Check the mount path (the location where the NAS share is mounted (in the command to mount the file system on a NFS client):
  `mount vtl-hostname:/nas/nasresourcename/fds/foldername /mnt/share1`
  `/mnt/share1` would be the location for the share.
- Run `Showmount –e <VTL server name>` and see if the shares list returns.

Issue: A Linux or Unix client is connected to VTL through a firewall and cannot connect to its share.

Solution: Use the `rpcinfo –p` command to verify that the following ports are open:

- 111 - TCP/UDP

- 2049 - TCP/UDP

Issue: An AIX client cannot access an NFS share.

Solution: The client operating system requires connecting to a port less that 1024.

VTL includes an insecure checkbox if the client's OS does not use a reserved port for NFS (an Internet port that is less than IPPORT_RESERVED -- 1024). AIX is an example of an OS that needs to be flagged as insecure. Right-click the NFS client and select *Assign Share*. Check the *insecure* checkbox.

NFS client requires remapping user IDs to nobody

Squash can be used to map user IDs to *nobody*. This is set in NFS sharing.

root_squash	all_squash	Action
-	-	No UIDs mapped.
X	-	UID=0 (root user) is remapped to *nfsnobody:nasgrp* (default).
X	X	All UIDs are mapped to *nfsnobody:nasgrp*.

Domain controller authentication issue

1. Verify that the VTL and Active Directory server clock times are the same.
2. Make sure the Active Directory server is running normally.
3. Verify that the VTL NTP setting points to the Active Directory server.

VTL server cannot get some users/groups from the domain controller

Below is the command to clean the *winbind* cache. This cache must be purged when the VTL server cannot get some users or groups from the domain controller.

Method 1: This method is safe to run on a production server and will not affect client connections. This method does not require a restart of modules or downtime:

```
echo erase | ./tdbtool/ $ISHOME/var/tmp/$HOSTNAME/winbindd_idmap.tdb
> /dev/null
```

Method 2: This method will require downtime and will need to restart modules.

1. Stop CIFS modules by running `vtl stop NAS SMBD.`
2. Go to the directory: `cd $ISHOME/var/tmp/$HOSTNAME/`
3. Remove everything in that specific directory only.
4. Restart CIFS modules `vtl start NAS SMBD.`

Problem migrating data with ACLs an extended attributes

A copied file is not accessible or there is a difference in the rights between a copied file and the original.

A simple copy does not copy extended file attributes. Instead, the copied file inherits the file access rights of the owner of the session who did the copy. An additional operation is requested to also copy ACLs and extended attributes as well as data if the targeted file system is managing such attributes and ACLs.

1. Create a new file system with a size no less than the original one. In the following example, the old file system is mounted at /nas/NAS-00001 and the new file system is mounted at /nas/NAS-00002.

2. Log in as root on the VTL server, change the current path to the mount point of the old file system, and copy all files from the old file system mount point to the new file system mount point.

   ```
   cd /nas/NAS-00001
   ```

   ```
   cp -rf * /nas/NAS-00002
   ```

3. Copy the ACLs.

   ```
   getfacl -R . | (cd /nas/NAS-00002; setfacl --restore=- )
   ```

4. Copy the extended attributes.

   ```
   getfattr -R . | (cd /nas/NAS-00002; setfattr --restore=- )
   ```

5. Recreate Windows and NFS share information. You can recreate Windows and NFS shares on the new NAS resource from the console if there are not many of them.

   Otherwise, you may need to manually edit: $ISHOME`/etc/$HOSTNAME/nas.conf/nas.conf` and $ISHOME`/etc/$HOSTNAME/nas.conf/ipstorsmb.conf` when VTL is not running. Replace the old file system mount point with the new one, when applicable. (It is important to do this under the supervision of Technical Support.)

If a local file copy is not required, the following steps can be used instead of those above. For example, if it is certain that the original file system does not have any corruption; it may be preferable to copy the files instead of copying the raw data as mirroring does.

1. Mirror the file system with the new storage.

2. After synchronization, swap the mirror.

3. Promote the mirror to become a copy.

Truncated folder names

Folder names are truncated to 239 characters in the console and at the Linux prompt.

This is a limitation of the ext4 file system. Names cannot be longer than 239 characters and must be considered when creating file trees in VTL and when defining a naming convention in order to manage backup file organization.

Permission denied to folders

Folder names are using non-US accentuated characters and are not accessible from the host (i.e.: mes-données-sauvegardées).

Use US Code Page 437 characters when creating shares. Non-US Code Page 437 characters like accentuated characters or German ß are translated and authentication may fail as shown below:

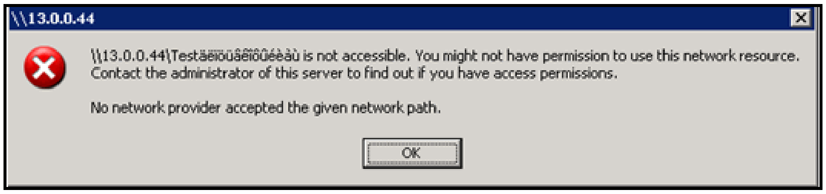

\\13.0.0.44

\\13.0.0.44\Testäëïöüâêîôûéèàù is not accessible. You might not have permission to use this network resource. Contact the administrator of this server to find out if you have access permissions.

No network provider accepted the given network path.

The following SMB.conf snip indicates that such accentuated characters are translated according to code page 437, translating the folder name. Do not modify the SMB.conf or even the folder name to manually attempt to fix it.

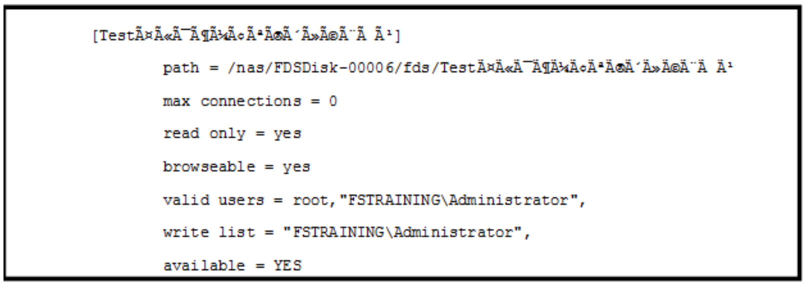

```
[TestÃ¤Ã«Ã¯Ã¶Ã¼Ã¢ÃªÃ®Ã´Ã»Ã©Ã¨Ã Ã¹]
        path = /nas/FDSDisk-00006/fds/TestÃ¤Ã«Ã¯Ã¶Ã¼Ã¢ÃªÃ®Ã´Ã»Ã©Ã¨Ã Ã¹
        max connections = 0
        read only = yes
        browseable = yes
        valid users = root,"FSTRAINING\Administrator",
        write list = "FSTRAINING\Administrator",
        available = YES
```

File access issues

The client can access shares, but cannot create, copy, or open files or cannot access a folder.

- Check file permissions on the client side.
  For Windows clients, right-click the file and check the permissions and see if they are read-only.
  For Linux clients, run `ll` and check the permissions of the file.
- In the console, click the share, and go to the *Clients* tab. Make sure the user or group that is assigned to the share has read/write privileges.

Change mount/unmount option

The mount option does not fit with the usage of the file system. i.e. forcing the file system in read-only mode.

1. Right-click the specific NAS resource and select *Unmount*.
2. Right-click *NAS Resources* and select *Properties*.

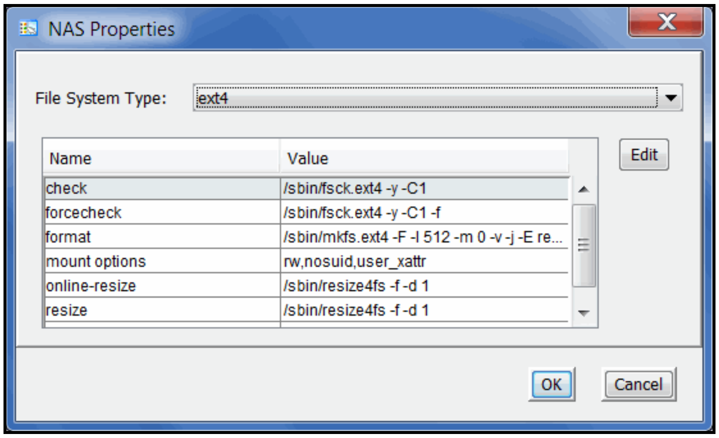

File System Type: ext4

Name	Value
check	/sbin/fsck.ext4 -y -C1
forcecheck	/sbin/fsck.ext4 -y -C1 -f
format	/sbin/mkfs.ext4 -F -I 512 -m 0 -v -j -E re...
mount options	rw,nosuid,user_xattr
online-resize	/sbin/resize4fs -f -d 1
resize	/sbin/resize4fs -f -d 1

3. Highlight the *mount options* in the table and click *Edit*.
4. Append the extra parameter `errors=remount-ro` at the end. Make sure this is separated with a comma (,).

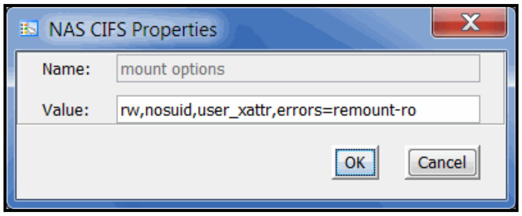

Name: mount options

Value: rw,nosuid,user_xattr,errors=remount-ro

5. Remount the NAS resource by right-clicking the specific NAS resource and select *Mount*.

System event messages

Information from the /var/log/messages file on the VTL server can be viewed via the Event Log in the console. A maximum of 10,000 records will be displayed in the Event Log.

For troubleshooting purposes, /var/log/messages keeps track of the last 20 MB of system events. When the file reaches 20 MB, it is renamed to `messages.n`, where *n* is a sequential number between 1 and 30. When it is renamed, it is also compressed to save space.

Take an X-ray of your system for technical support

Taking an X-ray of your system is useful for your technical support team to help solve system problems. Each X-ray contains technical information about your server, such as server messages and a snapshot of your server's current configuration and environment. You should not create an X-ray unless you are requested to do so by your technical support representative.

To create an X-ray file:

1. In the console, right-click your VTL server and select *X-Ray*.

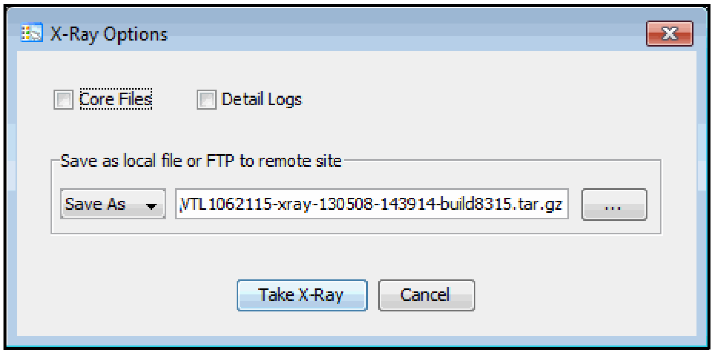

X-Ray Options

□ Core Files □ Detail Logs

Save as local file or FTP to remote site

Save As: VTL1062115-xray-130508-143914-build8315.tar.gz

2. Based on the discussion with your Technical Support representative, select the options you want to include.

   The system logs are always included in X-rays. The *Detail Logs* option allows for the collection of some additional log files that are not necessary for a standard X-ray and are only used for deeper troubleshooting. Including details logs can result in very large X-rays.

3. Set the X-ray file name and select whether you want to save the file locally or FTP it to a remote site.
4. If you are using FTP, click the *Setting* button and enter FTP information.

   Target Directory - The directory on the FTP server where the file will be stored. The directory name you enter here (such as `vtlxray`) is a directory on the FTP server (i.e., `ftp\vtlxray`). Do not enter an absolute path like `c:\vtlxray.`

   Username/Password - The user that the system will log in as. You must create this user on the FTP server with read/write access to the *Target Directory*. If the FTP server uses Active Directory, enter the name as `username@domain` instead of `domain\username`.

5. Click the *Take X-Ray* button.

Error codes

This section contains error messages generated by VTL/VTL-S and SIR servers.

Number	Type	Text	Probable Cause	Suggested Action
1016	C	Primary device ID %1 failed. The server is switching to the mirror device.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1017	E	The mirror of device ID %1 failed.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1022	E	Replication for virtual tape ID %1 failed; %2.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.
1023	E	Connection to physical device %1 failed; the path was switched to %2.	An adapter or cable might have a problem.	Check for a loose or damaged cable on the affected drive.
1030	E	Replication failed to start because a replication job was already in progress for virtual tape ID %1.	Only one replication job is allowed at a time for a tape.	Try again later. If replication was triggered by a schedule, adjust the schedule in the policy to avoid duplicate sessions.
1031	E	Replication failed to start because the replication control map was missing for virtual tape ID %1.	The configuration may not be valid.	Rescan devices to refresh the configuration.
1032	E	Replication failed to start because access to the replication control map for virtual tape ID %1 failed.	The virtual device may be offline or may have missing segments.	Check the underlying physical device and rescan devices.
1034	W	Replication failed for virtual tape ID %1 due to network transport error %2.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.

Number	Type	Text	Probable Cause	Suggested Action
1035	E	Replication failed for virtual tape ID %1 because the primary disk failed with error %2.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1038	E	Replication failed for virtual tape ID %1 because the local server could not allocate memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
1039	E	Replication could not proceed for virtual tape ID %1 because the replica device failed with error %2.	The replica reports the error specified in the message.	Check the device on the replica server and take necessary actions based on the error.
1046	E	Replica rescan failed for virtual tape ID %1 because the device had error %2.	The primary device reports the error specified in the message.	Check the device and take necessary actions based on the error.
1047	E	Replica rescan failed for virtual tape ID %1 because the replica device had error %2.	The replica device reports the error specified in the message.	Check the device and take necessary actions based on the error.
1048	E	Replica rescan failed for virtual tape ID %1 due to network transport error %2.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.
1049	E	Replica rescan could not proceed because the replication control map was missing for virtual tape ID %1.	The configuration may not be valid.	Rescan devices to refresh the configuration.
1050	E	Replica rescan could not proceed because access to the replication control map for virtual tape ID %1 failed.	The virtual device may be offline or may have missing segments.	Check the underlying physical device and rescan devices.
1052	E	Replica rescan failed for virtual tape ID %1; the replica status is %2.	The replication configuration may not be valid.	Check the configuration on the replica server. Check system logs on both servers for additional information.

Number	Type	Text	Probable Cause	Suggested Action
1053	E	Replica rescan could not proceed because a replication job was already in progress for virtual tape ID %1.	Only one replication job is allowed at a time for a device.	Try again later. If replication was triggered by a schedule, adjust the schedule in the policy to avoid duplicate sessions.
1055	E	Replication failed for virtual tape ID %1; the replica status is %2.	The replication configuration may not be valid.	Check the configuration on the replica server. Check system logs on both servers for additional information.
1056	E	Exchange of the replication control map between servers failed for virtual tape ID %1; error: %2.	This may be due to a connectivity issue.	Check connectivity between the primary and replica. Check system logs on both servers for additional information.
1059	E	Replication failed for virtual tape ID %1; error: %2.	Replication reports the error specified in the message.	Check the replica device and system logs; take necessary actions based on the error.
1060	E	Replica rescan failed for virtual tape ID %1; error: %2.	Replication reports the error specified in the message.	Check the replica device and system logs; take necessary actions based on the error.
1061	W	Storage path with ACSL %1 failed; alternate path %2 will be used.	This may be due to a connectivity issue.	Check the path connectivity between the server and the physical device.
1067	E	Replication could not proceed because connection to replica server %1 failed.	Either the network connection is down or the replica server is down.	Check the state of the replica server. Determine and correct either the network problem or server problem.
1069	E	Replication could not proceed because virtual tape ID %1 no longer has a replica.	The replica device may have been deleted or promoted while the primary server was down.	Reconfigure replication and create a new replica or use the replica that had been promoted.
1071	E	Replication could not proceed because remote device ID %1 does not exist or is not a replica.	The replica device may have been deleted.	Reconfigure replication.
1073	E	Replication could not proceed because the configuration file could not be opened.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.

Number	Type	Text	Probable Cause	Suggested Action
1074	E	Replication could not proceed because memory allocation failed.	The system may have been busy and did not have enough memory.	Check the memory usage of different processes and stop unnecessary processes to free up memory. You may need to restart server modules.
1075	E	Replication could not proceed because unexpected error %1 occurred.	Replication reports the error specified in the message.	Check system logs on both servers and take necessary actions based on the error.
1082	W	Replication for virtual tape ID %1 was cancelled.	This was most probably triggered by a user.	If this was not triggered by a user, check the system log to identify any related errors and take necessary actions based on the error.
1084	W	A SCSI command completed after recovering from an error. The device may have some reliability issues.	This may be due to a temporary error on the physical device.	Check the system log for additional information. Contact the hardware manufacturer for a diagnostic procedure.
1087	W	Replication could not proceed because virtual tape ID %1 is not currently available.	The tape is loaded in a drive and is in use.	Wait for the process to complete before trying again.
1088	E	Replication could not proceed because ID %1 could not located for the virtual tape.	The tape ID is missing at the kernel level.	Contact Technical Support.
1099	E	Replication could not proceed because virtual tape replica ID %1 for virtual tape ID %2 could not be expanded because the maximum licensed capacity on the replica server was reached.	All storage capacity licenses have been used.	Obtain additional license key codes.
1201	W	Kernel memory is low. Add more memory to the system if possible. Restart the server if possible.	There are too many processes for system resources.	Add more memory to the system; restart the server, if possible.

Number	Type	Text	Probable Cause	Suggested Action
1203	E	Storage path failed to be trespassed to %1.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1204	E	Storage path %1 failed to be added in the group.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1206	E	Storage path %1 failed to be activated.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1207	E	A critical storage path failure was detected. Path %1 will be removed.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1208	W	Storage path %1 does not belong to the active path group.	There was an attempt to access the storage via a path that is not part of an active group.	Use only active paths.
1210	W	There is no valid path available for device %1.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1211	W	There is no valid path group available.	Downstream storage path group is not correctly configured.	Check path group configuration on the storage device.
1212	W	There is no active path group for device GUID %1.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1214	E	Storage path %1 is not available.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1215	W	CLARiiON storage path is trespassing.	A downstream storage path failed or the path was manually trespassed.	If the path was not manually trespassed, check the physical device status, device connectivity and switches, and the storage log.

Number	Type	Text	Probable Cause	Suggested Action
1216	W	T300 storage path is trespassing.	A downstream storage path failed or the path was manually trespassed.	If the path was not manually trespassed, check the physical device status, device connectivity and switches, and the storage log.
1217	W	HSG80 storage path is trespassing.	A downstream storage path failed or the path was manually trespassed.	If the path was not manually trespassed, check the physical device status, device connectivity and switches, and the storage log.
1218	W	MSA1000 storage path is trespassing.	A downstream storage path failed or the path was manually trespassed.	If the path was not manually trespassed, check the physical device status, device connectivity and switches, and the storage log.
7001	E	Patch %1 failed; environment profile is missing in /etc.	Unexpected loss of environment variables defined in /etc/.is.sh occurred on the server.	Check server package installation.
7002	E	Patch %1 failed; it applies only to build %2.	The server is running a different build than the one for which the patch is made.	Get the patch, if any, for your build number or apply the patch on another server that has the expected build number.
7003	E	Patch %1 failed; you must be the root user to apply the patch.	The user account running the patch is not the root user.	Run the patch with the root account.
7004	W	Patch %1 installation failed; it has already been applied.	You tried to apply the same patch again.	No action is needed.
7005	E	Patch %1 installation failed; prerequisite patch %2 has not been applied.	A previous patch is required but has not been applied.	Apply the required patch before applying this one.
7006	E	Patch %1 installation failed because it could not copy new binaries.	Server modules are not in a consistent state or an unexpected error occurred on the binary file name or path in the patch.	Check the patch log file in / usr/local/<product>-archive.

Number	Type	Text	Probable Cause	Suggested Action
7008	W	Patch %1 rollback failed; there is no original file to restore.	You tried to roll back a patch that has not been installed or has already been rolled back.	No action is needed.
7009	E	Patch %1 rollback failed because it could not copy back previous binaries.	Unexpected error on the binary file name or path in the patch.	Check the patch log file in /usr/local/<product>-archive.
7010	E	Patch %1 failed; the file %2 has patch level %3, higher than this patch. You must first roll back %4.	A patch with a higher level patch has been applied that conflicts with this patch.	Roll back the higher level patch, apply this patch, and then reapply the higher level patch.
7011	E	Patch %1 failed; it applies only to kernel %2.	You tried to apply the patch applied to a server that is not running the expected OS kernel.	Apply the patch on a server that has the expected kernel.
7012	E	Patch %1 failed; the available free space is %2 bytes; you need at least %3 bytes to apply the patch.	Patch applied to a server running low on the disk used for server home directory.	Add more storage.
10001	E	User ID %1 has insufficient privileges.	Server modules are not running with root privileges.	Log in to the server with the root account before starting server modules.
10002	W	The server environment is not set properly.	Expected environment variables in /etc/.is.sh are missing.	Restore the file from an X-ray or a backup in /usr/local/<product>-archive.
10003	E	Initialization of configuration %1 failed.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
10004	E	SCSI device information failed to be retrieved.	The device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.

Number	Type	Text	Probable Cause	Suggested Action
10006	E	A write operation to configuration %1 failed.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
10054	E	Server FSID update failed.	This may be due to insufficient system memory or an unhealthy file system.	Check the server memory and file system status. You may need to restart server modules.
10059	E	The server configuration update for FC storage persistent encountered an error.	There is a conflict with the ACSL.	Use a different ACSL for binding.
10100	E	New SCSI devices failed to be scanned.	This may be due to unreliable storage connectivity, hardware failure, or system resources are running low.	Check the storage devices and the connectivity status. Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart the server machine.
10101	E	Update of configuration %1 failed.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
10102	E	SCSI devices failed to be added.	This may be due to unreliable storage connectivity, hardware failure, or system resources are running low.	Check the storage devices and the connectivity status. Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart the server machine.

Number	Type	Text	Probable Cause	Suggested Action
10204	E	Configuration file %1 could not be parsed.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. You may need to restart server modules. If the problem persists, contact Technical Support.
10207	E	Adapter %1 could not be added because there is not enough memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
10209	E	Physical device %1 could not be added because there is not enough memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
10210	W	Physical device %1 was marked as offline because its GUID, %2, did not match SCSI GUID %3.	This may be due to an old device being imported without proper initialization, invalid configuration, or corrupted device header.	Check the physical storage. Replace the device if it is not reliable. Fix any detected issues and rescan devices to refresh the configuration.
10212	W	Physical device %1 was marked as offline because the SCSI status indicated it was offline; GUID: %2.	The device may have been removed, turned off, or is not functioning properly.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
10213	W	Physical device %1 was marked as offline because it did not respond correctly to an inquiry; GUID: %2.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
10214	W	Physical device %1 was marked as offline because its GUID, %2, does not match a valid FSID.	The GUID recorded on the device header does not match the unique ID, called the FSID, which is based on the external properties of the physical device. This may be due to device changes while the server was down.	Make sure devices are not changed directly by 3rd-party applications. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
10215	W	Physical device %1 was marked as offline because its storage capacity has changed; GUID: %2.	The physical device geometry, including the number of sectors, is different from the original record.	Rescan devices to refresh the configuration.
10240	E	SCSI path %1 is missing.	This may be due to a disconnected storage cable, a re-zoned Fibre Channel switch, or a failed storage port.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
10241	E	Physical adapter %1 could not be located in /proc/scsi/.	The adapter driver may not be loaded.	Run 'lsmod' to check loaded drivers and try to load the driver if needed.
10242	E	Physical adapter number %1 is duplicated in /proc/scsi/.	The OS has assigned the same number to two different adapters probably due to continuously loading and unloading drivers.	Do not repeatedly load and unload the Fibre Channel drivers and the server modules individually. You may need to reboot the server.
10244	E	Device %1 LUN in FSID %2 does not match the actual LUN.	The FSID was generated using the physical device LUN but the LUN assignment may have changed on the storage controller.	Do not change the LUN after the device FSID has been generated. Revert back to the original LUN.
10245	E	FSID %1 does not match device ACSL %2, GUID %3.	The FSID was generated using the physical device LUN but the device SCSI path may have changed.	You may need to rescan devices.
10246	E	FSID generation for the device with ACSL %1 failed.	The physical device does not provide reliable data in the SCSI inquiry pages in order to generate a unique ID.	Prepare this device to become a virtual device and not an SED.
10247	E	GUID for the device with ACSL %1 is blank; FSID validation failed.	The device header may have accidentally been erased by a system command such as fdisk or format.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
10250	W	SCSI alias path %1 was removed because it did not have the same virtualization category as physical device %2.	The hardware configuration may have changed or there is a device or connectivity failure.	No action is necessary if this is due to a configuration change. Otherwise, check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
10251	W	SCSI alias path %1 was removed because it did not have the same GUID as physical device %2.	The hardware configuration may have changed or there is a device or connectivity failure.	No action is necessary if this is due to a configuration change. Otherwise, check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
11000	E	A socket failed to be created.	This may be due to insufficient system resources.	Check your network environment; you may need to restart the system.
11001	E	Setting the socket to re-use address failed.	This may be due to a network configuration error.	Check your network environment; you may need to restart the system.
11002	E	Binding the socket to port %1 failed.	Another process may be using the same port number.	Identify the process using the port and stop it.
11003	E	TCP service failed to be created.	This may be due to insufficient system resources.	Restart the server modules.
11004	E	TCP service failed to be registered; program %1 version %2.	This may be due to a network configuration error.	Restart the server modules.
11006	E	The server communication module failed to start.	This is due to the communication port being unavailable or a network error.	Restart the communication module.

Number	Type	Text	Probable Cause	Suggested Action
11007	W	There is not enough disk space available to successfully complete this operation and maintain the integrity of the configuration file. There is currently %1 MB of disk space available. The server requires %2 MB of disk space to continue.	The available space on the disk holding the configuration file is not enough.	Increase disk space.
11030	E	The Auto-Save option failed to set up a cron job.	The system command to add a cron job returned with an error.	Check the configuration in /etc/crontab and the status of the cron daemon.
11031	E	The Auto-Save option failed to create cron job script %1.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
11032	E	The Auto-Save option failed to connect to FTP server %1 on port %2.	This may be due to an invalid FTP address or port, or the FTP service is not running on the server.	Check that FTP service is enabled on the server; check network connectivity to the FTP site by manually running an FTP session.
11033	E	The Auto-Save option failed to log in to the FTP server with user %1.	The indicated user name is not valid.	Get a valid user name that can connect to the FTP server.
11034	E	The Auto-Save option failed because directory %1 does not exist.	The expected directory to back up the configuration files on the FTP site is missing.	Create the directory on the FTP site.
11035	E	The Auto-Save option failed to copy %1 to the FTP server.	The FTP user account may not have the write access to the directory on the FTP site.	Check the FTP user account and make sure it has valid access rights.
11036	E	The Auto-Save option failed to delete previous file %1 from the FTP server.	The FTP user account may not have the proper access rights for the directory on the FTP site.	Check the FTP user account and make sure it has valid access rights.

Number	Type	Text	Probable Cause	Suggested Action
11101	E	Client %1 failed to be added.	This error is most likely due to a system configuration error, or system resources running low.	Check OS resources running provided utilities such as 'top'.
11104	W	There are too many client connections.	The number of simultaneous client connections exceeded the limit that the current system memory can handle.	This is an unlikely condition as long as the recommended memory is available for the server.
11107	E	Client %1 encountered an illegal access error.	The client host attempted to connect to a session that is no longer available.	Restart the connection and ensure there is no security breach.
11109	E	Client %1 failed to open file %2.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
11112	E	Client %1 failed to parse configuration file %2.	The configuration file is corrupted, or manually tempered to the degree that is no longer recognizable by VTL. If corruption is the cause, then it is most likely due to a system drive hardware error.	If there is a valid configuration file saved, it can be restored to the system. Make sure to use reliable storage devices for critical system information.
11113	E	Client %1 failed to restart the authentication module.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11114	E	Client %1 encountered a memory allocation failure.	System resources are running low. This may be due to too little memory installed for the system, or some runaway process that is consuming too much of the memory.	Run 'top' to check the process that is using the most memory. If physical memory is below the server recommendation, install more memory on the system.
11170	E	Virtualization of %1 failed because its size in the configuration file was different from the one on the disk header. Run a rescan and try again.	Attempting to virtualize a LUN that has a different capacity than what was previously seen.	Rescan for new devices and try again.

Number	Type	Text	Probable Cause	Suggested Action
11201	E	There are too many console connections.	The number of RPC connections to the server via the console, CLI, or other components exceeded the limit that the system can handle.	Close some connections.
11202	E	User %1 had an illegal access error.	An attempt was made to connect to a session that is no longer available.	Restart the connection and ensure there is no security breach.
11203	E	User %1 failed to rescan SCSI devices.	This may be due to unreliable storage connectivity, hardware failure, or system resources are running low.	Check the storage devices and the connectivity status. Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart the server machine.
11204	E	An operation by %1 failed to check SCSI devices.	This may be due to unreliable storage connectivity, hardware failure, or system resources are running low.	Check the storage devices and the connectivity status. Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart the server machine.
11205	E	An operation by %1 failed to get information for file %2.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
11206	E	An operation by %1 resulted in a memory allocation failure.	The system does not have enough memory.	Check the memory usage of different processes and stop unnecessary processes to free up memory. You may need to restart server modules.
11207	E	An operation by %1 failed to open file %2.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
11208	E	An operation by %1 failed to read file %2.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.

Number	Type	Text	Probable Cause	Suggested Action
11209	E	User %1 has insufficient access privileges.	Access rights of the specified user are limited.	Retry the operation with a valid user account.
11211	E	An operation by %1 failed to save file %2.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
11212	E	An operation by %1 failed to create index file %2 for the Event Log.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
11213	E	An operation by %1 got an illegal time range (%2 - %3) for the Event Log.	The specified time range is not valid.	Retry with a valid range.
11214	E	An operation by %1 failed to get time range (%2 - %3) for the Event Log.	The Event Log does not cover the specified time range.	Retry with a valid range.
11215	E	An operation by %1 failed to open directory %2.	Either the directory is not available or the system is busy and does not have enough resources.	Make sure the directory exists and check the system status. You may need to restart server modules.
11216	E	An operation by %1 failed to fork a process due to insufficient system resources.	The system does not have enough memory.	Check the memory usage of different processes and stop unnecessary processes to free up memory. You may need to restart server modules.
11217	E	An operation by %1 failed to execute program %2.	Either the program is not available or the system is busy and does not have enough resources.	Make sure the program exists and check the system status. You may need to restart server modules.
11218	E	An operation by %1 failed to remove file %2.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.

Number	Type	Text	Probable Cause	Suggested Action
11219	E	User %1 failed to add device %2.	This may be due to insufficient system memory or disk space to update the device information or the underlying physical device may have a failure.	Check system resources, physical device status, device connectivity, switches, and storage log.
11220	E	User %1 failed to remove device %2.	The replica server could not be accessed in order to delete the associated replica device.	Make sure the replica server is accessible and retry.
11221	E	User %1 failed to add client %2 to virtual tape %3.	This may be due to insufficient system memory or disk space to update the device information.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
11222	E	User %1 failed to remove client %2 from virtual tape %3.	This may be due to insufficient system memory or disk space to update the device information.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
11231	E	User %1 failed to get the CPU status.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
11232	E	User %1 failed to get the memory status.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
11233	E	An operation by %1 failed to map the SCSI device name for %2 %3 %4 %5.	The device identified by its SCSI address (Adapter, Channel, SCSI ID, LUN) cannot be located due to a configuration change or storage failure.	Check storage and rescan devices to refresh the configuration.
11234	E	User %1 failed to test device %2.	The system command 'hdparm' to the device failed probably due to insufficient system resources or a storage device failure.	Run the command manually on the system and check storage devices.

Number	Type	Text	Probable Cause	Suggested Action
11237	E	An operation by %1 failed to get file %2.	The file is in use by another process.	The console automatically retries every 3 seconds. If this occurs repeatedly, close the console and retry.
11238	E	An operation by %1 failed to restart the authentication module.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11240	E	User %1 failed to start server modules.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11242	E	User %1 failed to stop server modules.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11244	E	User %1 failed to get the user list.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
11257	E	User %1 failed to add client %2.	This error is most likely due to system configuration error, or system resources running low.	Check OS resources using provided utilities such as 'top'.
11261	E	User %1 failed to get the client connection status for virtual tape %2.	Failed to inquire a SAN Client connection status due to system configuration error, storage hardware failure, or system resource access failure. This should rarely happen.	Check the system resource, such as memory, system disk space. Check the system log for specific reason of the failure.
11262	E	An operation by %1 failed to parse configuration file %2.	The configuration file is not readable by the server.	If there is a valid configuration file saved, it can be restored to the system. Make sure to use reliable storage devices for the critical system information.
11263	E	User %1 failed to restore configuration file %2.	Error encountered when writing the server configuration file to the system drive. This can only happen if the system drive ran out of space, is corrupted, or there has a hardware failure.	Check the system drive using OS-provided utilities. Free up space if necessary. Replace drive if not reliable.

Number	Type	Text	Probable Cause	Suggested Action
11265	E	An operation by %1 failed to restart the IO core module.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11266	E	User %1 failed to erase virtual tape %2 partition.	Storage hardware failure occurred.	Check the storage devices, e.g., power status; controller status, etc. Check the connectivity, e.g., cable connectors. With fibre channel switches, even the connection status light indicates the connection is good, it is still not a guarantee. Push the connector in to make sure. Check the specific storage device using OS-provided utilities such as 'hdparm'.
11278	E	User %1 failed to swap device ID %2 with its mirror.	Hardware problem with the repository mirror disk.	Check the repository mirror disk.
11280	E	User %1 failed to create secure communication credentials for server %2.	The remote server cannot be accessed to establish a secure communication channel.	Check that the specified server can be reached, the TCP ports required for communication are open, and both servers are using the same password security rules.
11289	E	User %1 failed to restart server modules.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11291	E	An operation by %1 failed to update metadata of physical device %2.	Storage hardware failure occurred.	Check the storage devices, e.g., power status; controller status, etc. Check the connectivity, e.g., cable connectors. With fibre channel switches, even the connection status light indicates the connection is good, it is still not a guarantee. Push the connector in to make sure. Check the specific storage device using OS-provided utilities such as 'hdparm'.

Number	Type	Text	Probable Cause	Suggested Action
11292	E	User %1 failed to swap IP address %2 with %3.	Storage hardware failure occurred.	Check the storage devices, e.g., power status; controller status, etc. Check the connectivity, e.g., cable connectors. With fibre channel switches, even the connection status light indicates the connection is good, it is still not a guarantee. Push the connector in to make sure. Check the specific storage device using OS-provided utilities such as 'hdparm'.
11295	E	An operation by %1 failed because the configuration file has an invalid format.	The configuration file is not readable by the server.	If there is a valid configuration file saved, it can be restored to the system. Make sure to use reliable storage devices for the critical system information.
11301	E	An invalid password for user %1 was used by a client with IP address %2.	An incorrect username/ password was provided to connect to this server	Provide the correct username/password.
11315	E	You do not have a license for the %1 protocol.	An attempt was made to enable a protocol such as Fibre Channel, iSCSI with no appropriate license keys.	Install appropriate license keys that include support for the needed protocol.
11316	E	You have exceeded the backup cache capacity allowed by your license; used space: %1 MB, requested space %2 MB, licensed capacity: %3 MB.	Physical capacity usage has exceeded the capacity allowed by the license.	Contact DSI to purchase additional capacity license.
11406	E	The failover configuration failed to be prepared; %1.	Suspend fail over initiated while packaging fail over configuration information.	No action required. (Operation will be retried automatically)
11407	E	The failover configuration failed to be extracted; %1.	Suspend fail over initiated while packaging fail over configuration information.	No action required. (Operation will be retried automatically)

Number	Type	Text	Probable Cause	Suggested Action
11500	E	Virtual tape %1 expansion failed because there was not enough disk space.	There is no more storage space available for expanding a virtual tape.	Add more storage.
11512	E	User %1 failed to failed to enable replication of virtual tape %2 to server %3; watermark: %4 MB, time: %5, interval: %6, watermark retry: %7, suspended: %8.	It can be connection error while the primary server is synchronizing the initial status of the virtual tape to the replica on the target server, or the virtual tape is loaded to the drive at the moment by backup software or the console on different machine.	Check if the network is working properly first and correct the problem first if it is. If the virtual tape is moved to the drive, reconfigure the replication later.
11514	E	User %1 failed to remove the replica of virtual tape %2 from server %3; watermark: %4 MB, time: %5, interval: %6, watermark retry: %7, suspended: %8.	Error encountered when writing the updated configuration to the VTL configuration file to the system drive. This can only happen if the system drive ran out of space, or is corrupted, or if there is hardware failure in the system drive.	Check to make sure there is enough free space on the system disk. (This should never happen since the VTL system has a mechanism to automatically prune the syslog from using up all system disk free space). If enough space is available on system disk, check the integrity of system disk file system using fsck utility.
11516	E	User %1 failed to create replica virtual tape %2.	Could not update the virtual tape partition information on disk.	Check the storage system and make sure that storage is working properly.
11518	E	User %1 failed to start replication of virtual tape %2.	The replication triggered manually by the user failed. It can be due to one of the following reasons. Network problems, Virtual tape is loaded in a drive, replication is in progress or the replica no longer exists.	If the replication is manually triggered by the user, check the replication status at the right panel of the replica before starting another replication. If the replication is triggered by the scheduler, adjust the schedule in the replication policy to avoid replicationg too often. Check if the network is working properly as well as the server activity. Remove replication setup from the virtual tape console if it no longer has a valid replica.

Number	Type	Text	Probable Cause	Suggested Action
11522	E	User %1 failed to promote virtual tape replica %2 to a virtual tape.	Failed to update the virtual tape partition information.	Check if the physical disk is working properly or the server is busy. Retry the operation when the physical disk is working properly or when the server is not busy.
11524	E	User %1 failed to take an X-ray.	When any server process cannot be started, it is most likely due to insufficient system resources, invalid state left by a server process that may not have been stopped properly, or due to an unexpected OS process failure that left the system in a bad state. This should happen very rarely. If frequent occurrences are encountered, there must be external factors that contributed to the behavior that must be investigated and removed before running the server.	If the system resources are low, run 'top' to check the process that is using the most memory. If the physical memory is below the server recommendation, install more memory on the system. If the OS is suspected to be in a bad state due to an unexpected failure in either hardware of software components, restart the server machine.
11534	E	An operation by %1 failed to reset the replication control map for virtual tape %2.	Storage hardware failure.	Check the storage devices for power status, controller status, etc. Check for proper connectivity. Fibre Channel Switch connection status lights do not guarantee a solid connection. Disconnect/Reconnect the Fibre Channel connector for verification. check the specific storage device using OS provided utility such as 'hdparm'.
11535	E	User %1 failed to update replication parameters for virtual tape %2; server %3, watermark: %4 MB, time: %5, interval: %6, watermark retry: %7, suspended: %8.	Error encountered when writing the updated configuration to the VTL configuration file to the system drive. This can only happen if the system drive ran out of space, or is corrupted, or if there is hardware failure in the system drive.	Check if the system drive is out of space or has any corruption.

Number	Type	Text	Probable Cause	Suggested Action
11537	E	User %1 failed to import physical device %2.	This may due a specific version of the VTL server limiting the support of the storage.	Check license agreement for the version of VTL server.
11539	E	User %1 failed to import physical device %2.	Storage hardware failure.	Check the storage devices for power status, controller status, etc. Check for proper connectivity. Fibre Channel Switch connection status lights do not guarantee a solid connection. Disconnect/Reconnect the Fibre Channel connector for verification. check the specific storage device using OS provided utility such as 'hdparm'.
11542	E	User %1 failed to remove virtual tape replica %2.	Error encountered when writing the updated configuration to the VTL configuration file to the system drive. This can only happen if the system drive ran out of space, or is corrupted, or if there is hardware failure in the system drive.	Check if the system drive is out of space or has any corruption.
11554	E	User %1 failed to set the self-checking interval to %2 sec in failover properties.	IP Network communication failure.	Check network connectivity.
11569	E	User %1 failed to set Fibre Channel WWPN %2 to %3 mode.	This is possibly due to the fibre channel driver being improperly loaded, or the wrong version of the driver is loaded. VTL FC target mode requires the VTL version of the driver to be used.	Run 'lsmod' to check the proper VTL driver is loaded. If it is, check to make sure it is the correct version. The correct revision should be located in the VTL/lib directory.
11578	E	User %1 failed to get Fibre Channel initiator information.	This is possibly due to the fibre channel driver being improperly loaded, or the wrong version of the driver is loaded. VTL FC target mode requires the VTL version of the driver to be used.	Run 'lsmod' to check the proper VTL driver is loaded. If it is, check to make sure it is the correct version. The correct revision should be located in the VTL/lib directory.

Number	Type	Text	Probable Cause	Suggested Action
11581	E	User %1 failed to set the NAS option to %2.	Failed to start the NAS processes. It is most likely due to insufficient system resources, invalid state left by last running server processes that may not have been stopped properly, or due to an unexpected OS processes failure that left the system in a bad state. This should happen very rarely. If frequent occurrences are encountered, there must be external factors that contributed to the behavior that must be investigated and removed before running the server.	If system resources are low, run 'top' to check the process that is using the most memory. If the physical memory is below the server recommendation, install more memory on the system. If the OS is suspected to be in a bad state due to unexpected failure in either hardware of software components, restart the server machine.
11632	E	User %1 failed to set failover options on the partner server; heartbeat interval: %2 sec, auto-recovery interval: %3 sec.	IP Network communication failure.	Check Network connectivity.
11633	E	User %1 failed to set failover options on the partner server; heartbeat interval: %2 sec, auto-recovery interval: disabled.	IP Network communication failure.	Check Network connectivity.
11648	E	The inquiry string on SCSI device %1 failed to be retrieved.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
11649	E	Conversion of the inquiry string on SCSI device %1 failed.	The device SCSI inquiry string contains invalid information.	Check the device configuration and status.
11650	E	The capacity of SCSI device %1 failed to be retrieved.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
11653	W	SCSI device %1 was ignored due to unsupported type %2.	The disk is not from a supported vendor.	Make sure you are using supported devices.
11655	E	SCSI device %1 capacity is not valid.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
11656	W	SCSI device %1 was ignored due to an unsupported Cabinet ID.	The Cabinet ID of the device is not supported.	Make sure you are using supported devices.
11657	W	SCSI device %1 was ignored due to a missing %2 vendor in the inquiry string.	The disk is not from a supported vendor.	Make sure you are using supported devices.
11741	E	An operation by %1 failed to create virtual library with %2 slots because only %3 slots were available.	Specified slot count for the virtual library exceeds the total slot count supported by the system.	Specify appropriate slot count.
11742	E	User %1 created only %2 out of %3 virtual drives due to a memory allocation failure.	System is out of memory, which prevents the creation of the specified number of virtual tape drives.	Increase system memory.
11744	E	An operation by %1 rolled back the configuration update for test mode promotion of %2 %3(s).	Failed to write the VTL configuration file.	Check /usr partition for free space. If the partition is out of space, delete unwanted files to create space.
11745	E	An operation by %1 rolled back the disk partition update for test mode promotion of %2 %3(s)\.	Failed to write update partition information to disks.	Check physical connectivity to the storage system/LUN.
11746	E	An operation by %1 rolled back test mode promotion of %2 %3(s).	Cannot add device to IOCore module.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
11782	E	Barcode [%1] of the source tape ID %2 already exists on target server %3. Auto-replication could not be configured.	A tape with the same barcode exists on the remote server.	Remove the tape with the same barcode from the remote server or change its barcode.
11788	E	Appliance hardware problem %1 was detected.	Detected an error on the VTL appliance.	Check the error message for information on the error to determine the solution.
11791	E	Virtual tape %1 failed to shrink to %2 MB; error: %3.	Virtual tape partition information couldn't be updated. This could happen in rare cases when the system is extremely busy and updates to disks take too long.	No user action required since it doesn't cause any problem for backup/restore. If this error happens, the tape will not be resized.
11793	W	Appliance hardware problem %1 was detected.	Detected a hardware problem with the appliance.	Check the error message and take appropriate hardware maintenance.
11796	E	The tape caching data migration trigger policy failed to be set.	Connectivity issues with the storage where VTL repository is located.	Check storage connectivity.
11797	E	The tape caching reclamation trigger policy failed to be set; error: %1.	Connectivity issues with the storage where VTL repository is located.	Check storage connectivity.
11906	E	File %1 failed to open.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
12002	E	Directory %1 failed to open.	Either the directory is not available or the system is busy and does not have enough resources.	Make sure the directory exists and check the system status. You may need to restart server modules.
12003	E	File %1 failed to open.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
12509	E	The inventory of physical library ID %1 failed.	Physical library inventory process failed.	Check the server log for more information on the cause of HW failure.
12511	E	A tape in library ID %1 failed to move from slot %2 to drive ID %3.	Physical tape move from slot to drive failed.	Check the server log for more information on the cause of HW failure.

Number	Type	Text	Probable Cause	Suggested Action
12514	E	Job %1 failed to restart; command: %2, return code: %3.	Import/Export tape job failed probably due to the tape not being present or the tape drive not being available.	Check the tape exists and the tape drive is available for job to continue, and then restart the job.
12521	E	A tape in library ID %1 failed to be moved from drive ID %2 to slot %3.	Physical library may have hardware issues.	Check the server log for more information on the cause of HW failure.
12525	E	Tape %1 failed to be reclaimed.	The virtual drive is busy with I/O and is not responsive to the upper layer calls.	Try again when the system is less busy, or determine the cause of the extensive I/O, and correct the situation if necessary.
12559	E	Jobs %1 failed to be resumed; error: %2, %3.	The jobs may already be running.	No action is required.
12564	E	Hardware or software compression failed to be set.	The physical device used by the database may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
12565	E	A tape in library ID %1 failed to be moved from import/export slot %2.	The library is in maintenance mode or the destination slot is not available.	Ensure that the library is operational and there is no tape in the slot.
12567	E	A tape in library ID %1 failed to be moved from slot %2 to an import/export slot.	The library is in maintenance mode or the import/export slot is not available.	Ensure that the library is operational and there is no tape in the slot.
12578	E	An import job from physical library ID %1 for tape %2 to virtual library ID %3 for tape %4 failed to be submitted.	This may be due to a memory allocation error, the tape does not have a valid barcode, or the maximum number of tapes in a job has reached.	Check the system log for more details and take necessary actions based on reported errors.
12579	E	An import job from physical drive ID %1 to virtual library ID %2 for tape %3 failed to be submitted.	This may be due to a memory allocation error, the tape does not have a valid barcode, or the maximum number of tapes in a job has reached.	Check the system log for more details and take necessary actions based on reported errors.

Number	Type	Text	Probable Cause	Suggested Action
12581	E	A tape stacking job failed to be submitted for physical tape %1 [%2]; error: %3.	The job reports the error specified in the message.	Take necessary actions based on the error.
12583	E	A tape stacking job failed to be submitted for physical library ID %1.	This may be due to a memory allocation error, the tape does not have a valid barcode, or the maximum number of tapes in a job has reached.	Check the system log for more details and take necessary actions based on reported errors.
12585	E	Importing stacked physical tapes to virtual tapes detected that the virtual and physical tapes did not have the same capacity.	The stacked tape import feature requires the same size for physical and virtual tapes.	Use tapes with the same capacity.
12586	E	Memory allocation failed for stacked tapes.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
12587	E	Memory allocation failed for virtual tapes.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
12588	E	A job to import stacked tapes from physical library ID %1 to virtual library ID %2 failed to be submitted; error: %3.	The job reports the error specified in the message.	Take necessary actions based on the error.
12590	E	Job %1 failed to be purged.	The physical device used by the database may have a failure or the job may have already been purged.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
12592	E	Tape %1 properties failed to be set.	The tape may be in use.	Retry later.

Number	Type	Text	Probable Cause	Suggested Action
12598	E	Physical tape %1 failed to be purged from the database.	The physical device used by the database may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
12600	E	%1 failed to assign %2 to client %3 (%4, %5).	The maximum number of Fibre Channel devices assigned to the client was reached.	Unassign devices that are no longer needed.
12601	E	A tape that was ejected failed to be loaded in virtual drive ID %1.	The eject command may not have completed.	Make sure the tape is completely ejected and try again.
12603	E	The ACSLS/LS physical drive ID %1 failed to be assigned to the server; error: %2.	This may be due to network or FC connection problems.	Check the ACSLS is running properly and is connected to both the server and the physical library. Check the FC connection between the server and the library.
12605	E	Physical tapes failed to be ejected; error: %1.	The server may not be connected to the physical library or the tape is no longer in the library.	Check the FC connection between the server and the library; if the tape is still loaded, try again or eject the tape manually.
12611	E	The ACSLS/LS library failed to be added as physical library ID %1.	This may be due to network connection problems or the maximum number of supported libraries has been reached.	Check the ACSLS is running properly and is connected to both the server and the physical library.
13101	E	Communication with server %1 failed; error: %2	This may be due to a communication error with the partner server.	Check the connection with the partner server.
13102	E	Execution of %1 failed.	The self-monitor module did not start properly.	Try to manually start the module.
13300	E	Authentication to server %1 failed; the failover module stopped.	The credentials used by failover servers are no longer valid.	Reset the root password of both hosts and reconfigure failover.
13301	E	Authentication to the local server failed; the failover module stopped.	The credentials used by failover servers are no longer valid.	Reset the root password of both hosts and reconfigure failover.

Number	Type	Text	Probable Cause	Suggested Action
13302	E	The configuration file of partner server %1 failed to be transferred to server %2; error: %3.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13303	E	Dynamic configuration of server %1 failed to be transferred to partner server %2; error: %3.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13304	E	File %1 failed to be renamed.	The file name already exists or the file system is inconsistent or read-only.	Check the file system.
13305	E	A write operation failed on file %1.	The file does not exist or there may be insufficient disk space, a system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
13306	E	File %1 failed to be opened.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
13307	E	Credential information of server %1 failed to be transferred to server %2; error: %3.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13308	E	Invalid failover configuration was detected. Failover will not occur.	The partner's configuration file is missing.	Check connectivity with the partner to make sure the configuration can be transferred.
13309	E	Server %1 is unable to communicate with server %2.	This may be due to a network error or a connectivity issue with the Configuration Repository.	Restart the network and check the Configuration Repository storage device.
13316	E	A virtual IP address failed to be added; error: %1.	This may be due to a network error.	Restart the network. If the problem persists, you may need to reboot.

Number	Type	Text	Probable Cause	Suggested Action
13317	E	A virtual IP address failed to be released; error: %1.	The server is still holding on to the IP address of the failed server.	Take necessary actions based on the error specified in the message.
13319	E	The failover module failed to stop. You may need to reboot the server.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to reboot.
13320	E	Update of configuration files failed on server %1; %2.	This may be due to a communication error with the partner server.	Check the connection and the filesystem on the partner server.
13700	E	Memory allocation failed; the self-monitoring module stopped.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory.
13701	E	Virtual IP address %1 failed to be released; the operation is being retried.	The server is still holding on to the IP address of the failed server.	No action is needed since this is temporary.
13702	E	Virtual IP address %1 failed to be retrieved; the operation is being retried.	This may be due to a network error.	Restart the network.
13703	E	The self-monitoring module failed to stop.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to reboot.
13710	W	The live trial license expired for %1. Contact DSI or its representative to purchase a license.	The live trial grace period has been exceeded.	Obtain a new license.
13711	W	The following options are not licensed: %1. Contact DSI or its representative to purchase a license.	The specified option is not licensed properly.	Obtain proper licenses.
13800	C	The following failure was detected on server %1: %2	The server reported the specified error. The partner server may take over depending on the current condition.	Based on the reported error, take appropriate actions.

Number	Type	Text	Probable Cause	Suggested Action
13807	C	A failure was detected on server %3. Server %1 will take over server %2.	A critical error occurred that triggered failover.	After takeover is complete, check the server log and fix any detected issue prior to failback.
13817	C	Server %1 could not fail back because it failed to update the partner server configuration.	This is related to other error conditions.	Check the system log to identify any related errors and take necessary actions based on the error.
13818	E	Access to the Configuration Repository failed.	The virtual device holding the Configuration Repository is no longer available.	The partner server will take over anyway. Shut down this server in order to avoid any conflicts.
13820	W	Server %1 heartbeat failed to be detected; %2.	This server did not receive a heartbeat from the partner server and is trying to contact other network entities to determine if this server is disconnected from the network or if the partner server is down.	Check related messages in the log for more details.
13821	E	The server could not reach other entities in the network. Takeover was not initiated because it is assumed the failure is on this server.	This server did not receive a heartbeat from the partner server and failed to contact any other network entities in the subnet.	Check network connectivity.
13822	C	Server %1 will not take over because storage devices could not be accessed.	This is due to a storage or connectivity issue that is affecting both servers.	Check the storage and connectivity for both servers and fix any detected issue.
13823	W	Server %1 failed to acknowledge the takeover request in time. Server %2 will forcefully take over.	The partner server is busy or not fully operational.	After takeover is complete, check the partner server and fix any detected issue prior to failback.
13827	E	The Configuration Repository update process with PID %1 failed to be stopped.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
13834	E	Copying files from the Configuration Repository failed.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13836	E	Getting configuration files from the Configuration Repository failed. Check and fix the repository disk.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13841	E	Takeover cannot complete because the partner server status is %1 instead of DOWN or READY.	The self-monitor module may not be running or may be busy on the partner server.	Check that the module is running.
13843	E	This server failed to get the original configuration file from the Configuration Repository of the partner server before failback.	The physical device used by the Configuration Repository may have a failure or there may be a connectivity issue with the partner server.	Check the physical device status, device connectivity and switches, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.
13845	W	The Configuration Repository disk failed. Server %1 is still in takeover mode.	The physical device used by the Configuration Repository may have a failure or there may be a connectivity issue with the partner server.	Check the physical device status, device connectivity and switches, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.
13849	W	The following heartbeat IP addresses are down:%1	This may be due to a network error.	Check connectivity between servers; check network parameters, including jumbo frame configuration, if applicable.
13850	E	Server %1 could not locate the Configuration Repository disk; error: %2.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13851	E	This server could not take over; reason: %1.	The server reports the error specified in the message.	Check the server and take necessary actions based on the error.

Number	Type	Text	Probable Cause	Suggested Action
13853	E	This server is unable to take over and notified the partner server to remain active.	The server detected a failure on its partner but a local failure prevented takeover.	Check the status of both servers.
13856	E	Server %1 failed to communicate with server %2 through IP %3.	This may be due to a network error.	Check connectivity between servers; check network parameters, including jumbo frame configuration, if applicable.
13859	E	This server will reboot because it detected a storage or network failure while it was taking over its partner.	The server cannot take over if it has any hardware issue.	Check network or storage connectivity.
13860	E	After server recovery, configuration files %1 and %2 failed to be merged.	The configuration files may be inconsistent; this may result in losing some configuration changes made during failover.	Check the configuration and re-apply your changes, if necessary.
13861	E	File %1 failed to be renamed to %2.	The file name already exists or the file system is inconsistent or read-only.	Check the file system.
13863	C	This server was forced to resume its operations; %1.	A user initiated a forceup command even though the server may not be fully functional.	Check the server status.
13864	C	This server is stopping its operations; reason: %1.	The partner server is requesting this server stop and be taken over.	Check the server status.
13878	E	This server has an invalid failover configuration; reason: %1.	The server reports the error specified in the message.	Check the server and take necessary actions based on the error.
13879	C	This server has detected a kernel module failure.	A kernel module on the partner server is not healthy.	Check the partner server; you may need to reboot it.
13882	E	The following error was detected on the Configuration Repository: %1.	The physical device used by the Configuration Repository may have a failure or there may be a connectivity issue with the partner server.	Check the physical device status, device connectivity and switches, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
13900	E	This server failed to take over its partner server due to %1.	This may be due to an issue with the power control device that is needed to turn off the partner or a problem retrieving information from the Configuration Repository.	Check the power control device on the partner server and connectivity with the Configuration Repository storage.
13901	E	Reading %1 from the Configuration Repository failed.	The physical device used by the Configuration Repository may have a failure or there may be a connectivity issue with the partner server.	Check the physical device status, device connectivity and switches, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.
13910	E	The server failover status could not be updated in the Configuration Repository; reason: %1.	The physical device used by the Configuration Repository may have a failure or there may be a connectivity issue with the partner server.	Check the physical device status, device connectivity and switches, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.
13921	C	The following failure was detected: The NAS module terminated abnormally.	A user may have stopped the module or the module had a failure. In a failover configuration, the partner server will try to take over this server.	If the module was not stopped intentionally, check the system log to get more information about the reason for the module failure.
13947	E	Server %1 cannot forcefully take over server %2 since an IPMI/HP iLO power control command failed.	The power control command ipmitool or hponcfg could not query or set the power status of the partner server.	Check the network connectivity with the partner server. Try to run the related command manually to check the power control module status.
14000	E	Virtual tape ID %1 is missing.	The device may have been deleted by another user.	Perform the operation on an existing device.
14001	E	A disk partition checksum mismatch was detected for virtual tape %1.	This is due to some configuration inconsistency issues.	Contact Technical Support.
14003	E	The configuration file failed to be parsed.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. You may need to restart server modules. If the problem persists, Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
14004	E	Virtual tape ID %1 failed to be renamed to %2.	The configuration file could not be updated probably due to insufficient memory, system disk failure, or an unhealthy file system.	Ensure there is enough memory and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
14005	E	Disk partition information could not be updated.	The device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
14006	E	Configuration %1 failed to be written.	The configuration file could not be updated probably due to insufficient memory, system disk failure, or an unhealthy file system.	Ensure there is enough memory and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
14008	E	Physical device %1 failed to be renamed to %2.	The configuration file could not be updated probably due to insufficient memory, system disk failure, or an unhealthy file system.	Ensure there is enough memory and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
14010	E	Client %1 is missing.	The client may have been deleted by another user.	Perform the operation on an existing client.
14011	E	Client %1 failed to be renamed to %2.	The configuration file could not be updated probably due to insufficient memory, system disk failure, or an unhealthy file system.	Ensure there is enough memory and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
14016	E	User %1 failed to renew caching for tape ID %2.	There may not be enough storage available.	Add storage, if applicable.

Number	Type	Text	Probable Cause	Suggested Action
14017	C	The server configuration failed to be saved to the Configuration Repository. Check the storage connectivity.	The Configuration Repository device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
14021	E	The physical device with ACSL %1 failed to be erased.	The device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
14023	E	The physical device with ACSL %1 failed to be claimed.	The device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
16001	E	An operation by %1 failed to convert the file system on %2.	The operation failed due to a system or mount error.	Contact Technical Support.
16100	W	Connection to SMTP server %1 on port %2 failed for email alerts; reason: %3.	The IP address or the port ID of the SMTP server may not be valid.	Check network connectivity, SMTP server IP address and port ID.
16101	W	Email alerts could not be sent via SMTP server %1 on port %2; reason: %3.	The IP address or the port ID of the SMTP server may not be valid.	Check network connectivity, SMTP server IP address and port ID.
18000	E	The Deduplication Repository disk %1 with GUID %2 is missing.	The server could not detect a data disk from another node.	Confirm services are running on the node that owns the missing device. Check the underlying physical storage, FC connectivity, and switch zoning.
18001	E	The repository module failed to start because the Deduplication Repository disk is missing. Check that each member of the cluster is operating properly.	The server could not connect to a physical device used by the deduplication repository.	Check the underlying physical storage, FC connectivity, and switch zoning. Ensure that each cluster node can connect to the storage.

Number	Type	Text	Probable Cause	Suggested Action
18050	E	The reclamation triggering module stopped because index pruning has failed.	The index disk may be in an invalid state.	Contact Technical Support.
18055	E	Deduplication index reclamation failed; reason: %1.	The operation failed due to the specified reason.	Take necessary actions based on the reported error.
18060	E	Deduplication Repository space reclamation did not complete; reason %1	The operation failed due to the specified reason.	Take necessary actions based on the reported error.
18066	E	Deduplication index pruning did not complete; reason: %1.	The operation failed due to the specified reason.	Take necessary actions based on the reported error.
18067	E	Deduplication index pruning did not start; reason: %1.	The operation failed due to the specified reason.	Take necessary actions based on the reported error.
18068	E	Space reclamation did not start; reason: %1.	The operation failed due to the specified reason.	Take necessary actions based on the reported error.
18069	E	Space reclamation will fail because folder %1 does not exist.	The folder referenced by the VIT is not present in the deduplication repository.	Contact Technical Support.
18070	E	A hardware failure was detected on the data disk. The repository module has been shut down.	For consistency and safety, the repository has been shut down.	Contact Technical Support immediately.
18071	C	A data consistency error was detected. The repository module has been shut down.	For consistency and safety, the repository has been shut down.	Contact Technical Support immediately.
18072	C	A hardware failure was detected on the metadata disk used for the index. The repository module has been shut down.	For consistency and safety, the repository has been shut down.	Contact Technical Support immediately.
18073	C	An invalid block was detected while reading the metadata disk for the index. The repository module has been shut down.	For consistency and safety, the repository has been shut down.	Contact Technical Supoort immediately.

Number	Type	Text	Probable Cause	Suggested Action
18075	W	Deduplication index pruning could not start because the Deduplication Repository was in read-only mode.	For consistency and safety, the deduplication repository is set to read-only due to an unexpected error, which is described in the system log.	Check messages in the system log to get more details and contact Technical Support.
18076	C	The deduplication repository has been set to read-only; reason: %1.	For consistency and safety, writes are not permitted to the deduplication repository due to the specified reason.	Contact Technical Support.
18077	C	A hardware failure was detected on the folder disk. Replication and deduplication processes were stopped.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
18087	W	The Deduplication Repository usage threshold of %1 was reached. %2 of repository capacity has been used.	The specified threshold has been reached.	Run reclamation, add storage, or change the threshold.
18090	E	Deduplication index reclamation failed because there is no folder.	You may have deleted, erased, or relabeled all tapes when you detected the SIR repository was full. At least one VIT tape is required for reclamation to proceed.	Contact Technical Support.
18101	E	The power control module status for node %1 at IP address %2 failed to be retrieved.	The power control module may not be configured properly or may not be functioning correctly.	Make sure a power control module such as IPMI or HP iLO is present on the server, is enabled, and is configured correctly.
18104	E	This standby server failed to take over node %1 because that node reported its hostname as %2, which is unknown to this server.	The hostname is not in the cluster configuration file probably because the server was not registered correctly.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
18105	E	This standby server failed to take over node %1 because it cannot retrieve required parameters in the cluster configuration.	The standby server needs cluster parameters in order to take over.	Contact Technical Support.
18107	E	This standby server failed to take over node %1 because it could not add virtual IP %2 of the failed node.	The heartbeat IP address may not have been properly set.	Contact Technical Support.
18108	E	This standby server failed to take over node %1 because it could not spoof the WWPN of the failed node.	The standby server cannot assume the FC identity of the failed node.	Contact Technical Support.
18109	E	This standby server failed to take over node %1 because it could not change its hostname to the name of the failed node.	The standby server cannot assume the identity of the failed node.	Contact Technical Support.
18110	E	This standby server failed to take over node %1 because it could not load physical resources of the failed node.	One or more physical devices may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
18111	E	This standby server failed to take over node %1 because it could not retrieve the failed node configuration from the Configuration Repository.	The standby server needs the failed node's configuration in order to assume its identity.	Contact Technical Support.
18112	E	This standby server failed to take over node %1 because it could not start up the deduplication module.	The standby server needs to start certain services to become active and take over the failed node.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
18113	E	This standby server failed to take over node %1; reason: %2.	The standby server reports the error specified in the message.	Take necessary actions based on the specified reason. If failover is desired, it should be performed manually.
18114	E	This standby server could not see physical device %1 [%2], used for the Deduplication Repository.	The standby server could not connect to a physical device used by the deduplication repository.	Check the underlying physical storage, FC connectivity, and switch zoning. Rescan devices on the standby server.
18116	E	This standby server detected physical device %1 [%2], used for the Deduplication Repository, is offline.	The standby server could not connect to a physical device used by the deduplication repository.	Check the underlying physical storage, FC connectivity, and switch zoning.
18119	E	This standby server failed to check the health of node %1 at IP address %2 because of an authentication error.	The remote server cannot be accessed via a secure communication channel.	Check that the server can be reached. You may need to restart services on this node.
18120	E	This server cannot be selected as a standby server in a failover configuration because it is running as an active node.	This server is running the deduplication module as an active node.	Select an appropriate server for the standby server. If the module has been started manually, restart services on this server and try again.
18133	W	This standby server is being forced to take over node %1, which may not be fully shut down.	The standby server cannot take over if it is unable to confirm that the failed node is shut down.	Make sure that the node that is to be taken over has actually been shut down, and if not, shut it down.
18135	E	User %1 failed to set a reclamation schedule for deduplication.	An I/O operation to the deduplication configuration file failed probably due to insufficient system resources.	Ensure there is enough memory and disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.

Number	Type	Text	Probable Cause	Suggested Action
19004	C	Storage usage has reached the threshold; %1% of total storage is being used; total storage capacity is %2; available storage space is %3.	Storage utilization has reached the limit specified by the user.	Check storage utilization and delete unused virtual tapes to free up space or add more storage.
19056	E	Remote copy failed; operation: %1; error: %2; server: %3, virtual tape ID: %4, target server: %5, replica tape ID: %6.	The operation reports the error specified in the message.	Check the log for related errors and take necessary actions.
19057	E	Remote copy of virtual tape failed because target server %1 could not be connected.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.
19058	E	Remote copy of virtual tape ID %1 to the target server failed because the replica tape was missing.	This is due to some configuration inconsistency issues.	Contact Technical Support.
19059	E	Remote copy of virtual tape ID %1 to the target server failed because the replica tape was invalid.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.
19060	E	Remote copy of the virtual tape to the target server failed because the configuration file could not be opened.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
19061	E	Remote copy of the virtual tape to the target server failed because there was not enough memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory.
19063	W	Remote copy of virtual tape ID %1 to the target server was cancelled.	This was most probably triggered by a user.	If this was not triggered by a user, check the system log to identify any related errors and take necessary actions based on the error.

Number	Type	Text	Probable Cause	Suggested Action
19064	E	Remote copy of virtual tape ID %1 to the target server failed because the tape could not be loaded.	This may be due to insufficient system resources or storage issues.	Check the physical device status, device connectivity and switches, and the storage log.
19065	W	Remote copy of virtual tape ID %1 to the target server failed because the tape was in a drive.	The virtual tape cannot be in a tape drive when remote copy is in progress.	Run remote copy when the virtual tape is moved back to the vault or a slot.
19066	E	Remote copy of virtual tape ID %1 to the target server failed because the tape could not be located by the tape emulation module.	This is due to some configuration inconsistency issues.	Contact Technical Support.
19068	E	Replica ID %1 of virtual tape ID %2 failed to be promoted because the target server was busy or there was a communication failure; error: %3.	This may be due to too many pending requests or a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable, and try again.
19073	E	Remote copy failed because target server %1 was busy.	This may be due to too many pending requests.	Try again later.
19074	E	Replication failed because replica server %1 was busy.	This may be due to too many pending requests.	Try again later.
19075	E	Remote copy of virtual tape ID %1 from %2 to %3 failed due to a version mismatch.	The target server does not have the required version for replication.	Upgrade the target server or select a different target.
19076	E	Replication of virtual tape ID %1 from %2 to %3 failed due to a version mismatch.	The replica server does not have the required version for replication.	Upgrade the target server, select a different target, or stop replication.
19077	E	Remote copy failed because target server %1 did not have enough space.	The target server is running low on disk space.	Check the disk usage and try to reclaim disk space or add more storage.

Number	Type	Text	Probable Cause	Suggested Action
19078	E	Replication failed because replica server %1 did not have enough space.	The replica server is running low on disk space.	Check the disk usage and try to reclaim disk space or add more storage.
19200	E	User %1 failed to get the encryption keys.	This may be due to insufficient system resources.	Ensure there is enough memory and the system disk is healthy. Check the event log messages for more information.
19201	E	User %1 failed to get the encryption key.	This may be due to insufficient system resources.	Ensure there is enough memory and the system disk is healthy. Check the event log messages for more information.
19202	E	User %1 failed to create encryption key %2.	A key with the same name already exists; it may have been created via another console.	Retry the operation using a different name.
19204	E	User %1 failed to delete encryption key %2.	The key may have been deleted or renamed via another console.	Wait for the console to refresh and retry, if necessary.
19206	E	User %1 failed to update information for encryption key %2.	The key may have been deleted or renamed via another console.	Retry the operation.
19208	E	User %1 failed to export encryption keys to package %2.	This may be due to insufficient system resources.	Ensure there is enough disk space and the system disk is healthy and the file system is not read-only. Check the event log messages for more information.
19210	E	User %1 failed to get encryption key package information.	The format or contents of the key package are not valid.	Use a valid key package.
19211	E	User %1 failed to import encryption keys from a package.	The format or contents of the key package are not valid.	Use a valid key package.
19381	E	The Deduplication Repository could not be reconfigured using the Flexible method.	The configuration file could not be updated for the repository storage allocation mode.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.

Number	Type	Text	Probable Cause	Suggested Action
31003	E	File %1 failed to be opened.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
31004	E	Adding user %1 to the NAS server failed.	The user name already exists in the /etc/passwd system file, nasgrp is missing from /etc/group, or the password file cannot be updated.	Make sure the user name does not already exist. If nasgrp does not exist, add it with the 'groupadd nasgrp' command. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31005	E	Memory allocation failed.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory.
31011	E	%1 failed to unmount.	The NAS resource may still be in use or system resources are insufficient to run the system command to unmount the file system.	Check the system status and retry later. If the problem persists, you may need to reboot.
31012	E	The predefined group for Samba users, nasgrp, is missing.	The predefined NAS group does not exist in /etc/group probably due to a configuration error.	Add the group with the 'groupadd nasgrp' command.
31013	E	%1 mount failed.	This may be due to an invalid mount point, obsolete NAS configuration, or insufficient system resources.	Ensure there is enough disk space and the system disk is healthy and the file system is not read-only. Check the event log messages for more information. If necessary, try to mount the resource manually to see the actual error.
31017	E	A write operation failed on file %1.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
31020	E	File %1 failed to be renamed to file %2.	The file name already exists or the file system is inconsistent or read-only.	Check the file system.

Number	Type	Text	Probable Cause	Suggested Action
31028	E	File %1 failed to be locked.	The file is still in use by another process.	Retry later. If the error persists, you may need to restart the server modules.
31029	E	File %1 failed to be created.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
31030	E	Directory %1 failed to be created.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
31031	E	Directory %1 failed to be removed.	Another process may be accessing the directory or the system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
31032	E	Execution of program %1 failed.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to reboot.
31035	E	Group %1 failed to be added to the NAS server.	All the group IDs in the reserved range have been used.	Add a new GID range for NAS Windows clients.
31036	E	User %1 failed to be deleted from the NAS server.	The user is currently logged in and cannot be deleted.	Wait until the user logs off or stop running processes that belong to this user account.
31039	E	File %1 failed to be renamed to file %2.	The file name already exists or the file system is inconsistent or read-only.	Check the file system.
31041	E	A NAS user cannot be created because the maximum number of reserved UIDs was reached.	All the user IDs in the reserved range have been used.	Add a new UID range for NAS Windows clients.

Number	Type	Text	Probable Cause	Suggested Action
31042	E	A NAS group cannot be created because the maximum number of reserved GIDs was reached.	All the group IDs in the reserved range have been used.	Add a new GID range for NAS Windows clients.
31045	E	File /etc/passwd failed to be updated.	This may be due to insufficient system resources or an unhealthy file system.	Check the server memory and file system status. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31047	E	The NAS management module is not running.	A user may have stopped the module or the module had a failure. In a failover configuration, the partner server will try to take over this server.	If the module was not stopped intentionally, check the system log to get more information about the reason for the module failure.
31048	E	NAS resource %1 failed to mount.	This may be due to an invalid mount point, obsolete NAS configuration, or insufficient system resources.	Ensure there is enough disk space and the system disk is healthy and the file system is not read-only. Check the event log messages for more information. If necessary, try to mount the resource manually to see the actual error.
31050	E	NAS resource ID %1 failed to unload during failover.	The device information could not be retrieved from the configuration file probably due to a storage error.	Check the underlying physical storage.
31051	E	NAS resource ID %1 failed to load during failover.	The device information could not be retrieved from the configuration file probably due to a storage error.	Check the underlying physical storage.
31054	E	The server hostname failed to be retrieved.	This is due to a network configuration issue.	Check that the host name returned by the 'uname -a' command corresponds to definitions in the DNS database or the /etc/hosts file.
31056	W	The mount operation was delayed because the file system on NAS resource ID %1 is in an inconsistent state.	This may be due to an ungraceful shutdown or the file system was mounted many times.	If the issue persists, run 'fsck' to check the file system.

Number	Type	Text	Probable Cause	Suggested Action
31058	W	Some NAS resources could not be unmounted.	The resources may still be in use.	You may need to restart the machine.
31060	W	Some NAS resources could not be mounted.	This may be due to insufficient system resources or storage issues.	Check the physical device status, device connectivity and switches, and the storage log.
31061	E	NAS process %1 failed.	This may be due to insufficient system resources.	Check the event log messages for more information and take necessary actions.
31062	E	A read operation failed on file %1.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. You may need to run 'fsck' to check the file system.
31063	E	A write operation failed on file %1.	This may be due to an unhealthy file system or a storage issue.	Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Check the underlying physical storage.
31064	E	An invalid file system type of %1 was retrieved from the configuration file.	The configuration file may be corrupted.	Contact Technical Support.
31066	E	The configuration file cannot be read.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. You may need to run 'fsck' to check the file system.
31067	E	Dynamic configuration file %1 failed to be parsed.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31068	E	Device information is missing from dynamic configuration file %1.	This is due to some inconsistency issues.	Contact Technical Support.
31069	E	The file system format operation on NAS resource %1 failed.	The file system superblock could not be erased probably due to an unhealthy file system or a storage issue.	Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Check the underlying physical storage.
31071	E	The status of file %1 failed to be retrieved.	The 'stat' system command failed to run for the file.	Run the command manually to display errors.

Number	Type	Text	Probable Cause	Suggested Action
31072	E	The CIFS native configuration failed to be updated.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31073	E	The NFS native configuration failed to be updated.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31074	E	Configuration file %1 failed to be parsed.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31075	E	NAS resource %1 failed to unmount during failback.	The resource may still be in use.	You may need to restart the machine.
31076	C	Due to a storage failure, the partner server had to reboot to resume NAS services.	There was a storage error preventing recovery without a reboot.	No action is required.
31078	E	NAS resources failed to load during failback when processing %1.	The device information could not be retrieved from the configuration file probably due to a storage error.	Check the underlying physical storage.
31079	E	On NAS resource %1, parameter %2 is missing from a file system command.	The file system command requires the parameter, which is not specified.	Specify the parameter in the command associated with the file system.
31086	E	Replication throttling for VITs failed to be set to %1 KB/s.	The configuration file could not be updated probably due to a disk failure or an unhealthy file system.	Ensure the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.

Number	Type	Text	Probable Cause	Suggested Action
32087	E	File system properties failed to be set on directory %1 and its subdirectories.	This may be due to a disk failure or an unhealthy file system.	Ensure the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
32090	E	The description of user %1 failed to be changed.	The user may not exist or the NAS user configuration file is missing.	Check the validity of the user name and make sure NAS users and groups are still available.
32091	E	Group %1 failed to be deleted.	The group may not exist or the NAS user configuration file is missing.	Check the validity of the group name and make sure NAS users and groups are still available.
32092	E	The description of group %1 failed to be changed.	The group may not exist or the NAS user configuration file is missing.	Check the validity of the group name and make sure NAS users and groups are still available.
32093	E	Members of group %1 failed to be changed.	The group or the group members may not exist.	Check the validity of the group name and members.
32094	E	Connection to Domain Controller %1 failed.	This may be due to an invalid user/password or a network error.	Validate the domain account and check connectivity.
40003	C	The virtual tape library database on device ID %1 is inconsistent.	The database consistency check failed, probably due to a storage error.	Check the underlying physical storage.
40004	E	The following device is not supported: [%1][%2][%3].	The device could not be found in the list of supported devices.	Contact Technical Support.
40007	E	Unloading a tape from virtual drive ID %1 failed; error code: %2.	The virtual drive is busy and is not responsive.	Take necessary actions based on the error stated in the message.
40009	E	The Move Medium command failed on virtual library ID %1 from source element address %2 to %3.	The tape library may be unavailable, the source element may be missing, or the destination may be occupied.	Check the library status, make sure the source element is available, and the destination element is not occupied.
40018	E	The tape in physical library ID %1, in drive ID %2, with barcode [%3], failed to be loaded.	There is a hardware problem with the physical tape drive.	Check with the tape drive vendor.

Number	Type	Text	Probable Cause	Suggested Action
40020	W	Job %1 was cancelled.	The physical library may have been reset or the related tape was requested by a backup software.	No action is required.
40029	E	There is not enough memory to complete the operation.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
40032	W	Physical tape [%1] is not available to start an auto archive job. The job will be started when the tape becomes available.	The physical tape may have been removed from the slot.	Insert a tape with an appropriate barcode in the physical library and perform an inventory operation.
40036	E	The connection to tape database %1 failed.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
40039	E	The status of physical library ID %1 could not be reported because the 'Read Element Status' SCSI command failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40040	E	Initialization of device %1 failed; error code: %2.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
40043	E	The 'Move Medium' command for physical library ID %1 failed; source address: %2, destination address: %3, error code: %4.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.

Number	Type	Text	Probable Cause	Suggested Action
40044	E	The 'Unload' command for physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40046	E	Writing to physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40047	E	Writing a filemark to physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40048	E	The 'Mode Sense' command for physical drive ID %1 failed; page code: %2, error code: %3.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40049	E	The 'Mode Select' command for physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40050	E	The 'Rewind' command for physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.

Number	Type	Text	Probable Cause	Suggested Action
40051	E	The 'Inquiry' command for physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40067	E	Physical drive ID %1 failed to be added to tape database %2.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
40074	E	Export job %1 failed; error code: %2, source tape: [%3], destination tape: [%4].	The job reports the error specified in the message.	Check the event log messages prior to this message for more information and take necessary actions.
40076	E	Import job %1 failed; error code: %2, source tape: [%3], destination tape: [%4].	The job reports the error specified in the message.	Check the event log messages prior to this message for more information and take necessary actions.
40077	E	Import job %1 failed due to a duplicate virtual tape barcode; destination tape: [%2].	A virtual tape with the same barcode already exists in the virtual library.	Delete the existing tape or use a different barcode.
40078	E	Import job %1 for a standalone drive failed due to a duplicate virtual tape barcode; destination tape: [%2].	A virtual tape with the same barcode already exists in the virtual library.	Delete the existing tape or use a different barcode.
40084	W	Tape [%1] is blank and could not be exported.	Import/export jobs do not allow blank tapes.	Make sure that cleaning tapes are not used for import/export jobs.
40089	E	A process failed to be created to communicate with an ACSLS/LS server because of insufficient system resources; error: %1.	System resources are running low.	Check the system status, log, and the error stated in the message. You may need to reboot the server.

Number	Type	Text	Probable Cause	Suggested Action
40090	E	The ACSLS/LS communication module failed to be executed; error: %1.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '$ASCSLSHOME/log/ acsss_event.log' for details.
40091	E	%1 failed to establish ACSLS/LS communication; error: %2.	This may be due to a network error or the ACSLS server is not responsive.	Check the ACSLS server status and the connectivity between servers. Take necessary actions based on the error stated in the message.
40092	E	The ACSLS/LS server could not resolve the VTL server name; DNS or the /etc/hosts file returned IP address %1 for %2.	This is due to a network configuration issue.	Check that the VTL host name returned by the 'uname -a' command corresponds to definitions in the DNS database or the / etc/hosts file.
40093	E	The connection to the ACSLS/LS server using IP address %1 and ACS %2 failed.	This may be due to a network error or the IP returned by host name resolution is not valid.	Check the ACSLS server status and the connectivity between servers. Check that either a DNS server is running or the '/etc/hosts' file contains the correct IP address.
40094	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 (%3) ACS %4.	The ACSLS server did not provide the requested information in time.	Check the ACSLS server's event log '$ASCSLSHOME/ log/acsss_event.log' for details and take necessary actions.
40095	E	The ACSLS/LS server failed to mount %1 on physical drive ID %2; error from %3 (%4): %5.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '$ASCSLSHOME/log/ acsss_event.log' for details.
40096	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 (%3) when attempting to mount %4 on drive ID %5.	The ACSLS server did not provide the requested information in time.	Check the ACSLS server's event log '$ASCSLSHOME/ log/acsss_event.log' for details and take necessary actions.

Number	Type	Text	Probable Cause	Suggested Action
40097	E	The ACSLS/LS server failed to dismount %1 from drive ID %2; error from %3 (%4): %5.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '$ASCSLSHOME/log/acsss_event.log' for details.
40098	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 (%3) when attempting to dismount %4 from drive ID %5.	The ACSLS server did not provide the requested information in time.	Check the ACSLS server's event log '$ASCSLSHOME/log/acsss_event.log' for details and take necessary actions.
40099	E	The ACSLS/LS server failed to retrieve drive information in ACS %1; error from %2 (%3): %4.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '$ASCSLSHOME/log/acsss_event.log' for details.
40100	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 (%3) when attempting to retrieve drive information in ACS %4.	The ACSLS server did not provide the requested information in time.	Check the ACSLS server's event log '$ASCSLSHOME/log/acsss_event.log' for details and take necessary actions.
40101	E	The ACSLS/LS server failed to retrieve volume information in ACS %1 and pool %2; error from %3 (%4): %5.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '$ASCSLSHOME/log/acsss_event.log' for details.
40102	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 when attempting to retrieve volume information in ACS %3 and pool %4.	The ACSLS server did not provide the requested information in time.	Check the ACSLS server's event log '$ASCSLSHOME/log/acsss_event.log' for details and take necessary actions.
40103	E	The ACSLS/LS server failed to retrieve LSM information in ACS %1; error from %2 (%3): %4.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '$ASCSLSHOME/log/acsss_event.log' for details.

Number	Type	Text	Probable Cause	Suggested Action
40104	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 (%3) when attempting to retrieve LSM information in ACS %4.	The ACSLS server did not provide the list of tape drives in time.	Check the ACSLS server's event log '$ASCSLSHOME/log/acsss_event.log' for details and take necessary actions.
40112	W	Physical tape [%1] is not available in physical library ID %2 to start the tape caching job. The job will be started when the tape becomes available.	The tape is missing from the physical library.	Make sure the tape is in the library.
40114	W	Export jobs are not allowed because tape [%1] in physical library [%2][%3] is used for tape caching.	A wrong physical tape was selected.	Retry with a valid tape.
40118	E	A block compressed by hardware compression failed to be decompressed using software decompression; error code: %1.	The compressed data may not be valid or the compression software had a failure.	You may need to use a hardware compression adapter.
40120	W	Tape [%1] had no data so no export job was submitted for tape stacking.	The tape is empty.	Retry with a valid tape.
40123	E	The tape in library ID %1, drive ID %2, barcode [%3], failed to load because it is a cleaning tape.	A wrong slot was used for the tape operation.	Retry with a valid tape and slot.
40124	E	A write command to the Configuration Repository failed.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
40130	W	The cache could not be renewed for tape ID %1. Data will be redirected to physical tape [%2].	The space reserved for tapes may not be sufficient.	Increase the space available for tapes.
40132	E	The tape shredding job failed for tape [%1]; error code: %2.	A write operation has failed probably due to a storage device issue.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues, rescan devices to refresh the configuration, and try again.
40136	E	A tape position could not be set by the SCSI SPACE command on physical drive ID %1; error code: %2.	The physical tape or drive may have a failure.	Fix any detected issue and try again.
40137	E	An import/export job failed to be added to the queue because the maximum of %1 jobs was reached; job ID: %2, physical tape: [%3].	The maximum number of jobs has been reached.	Delete some jobs.
40141	E	Tape stacking scan job %1 failed; error code: %2, physical tape: [%3].	The job reports the error specified in the message.	Take necessary actions based on the error stated in the message.
40142	E	Tape stacking import job %1 failed; error code: %2, physical tape: [%3].	The job reports the error specified in the message.	Take necessary actions based on the error stated in the message.
40143	E	Tape stacking export job %1 failed; error code: %2, virtual tape: [%3], physical tape: [%4].	The job reports the error specified in the message.	Take necessary actions based on the error stated in the message.
40148	E	Tape stacking export job %1 failed because there is not enough memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.

Number	Type	Text	Probable Cause	Suggested Action
40152	E	Tape stacking import job %1 failed because there is not enough memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
40153	E	A stacked tape could not be imported because slot %1 in virtual library ID %2 is occupied.	Another virtual tape is in the import destination slot.	Choose another slot.
40163	E	The ACSLS/LS server failed to find a compatible drive to mount tape [%1]; error from %2 (%3): %4.	The ACSLS/LS server reports the error specified in the message.	Take necessary actions based on the error stated in the message.
40164	E	After %1 second(s), %2 server (%3) did not respond to the mount request for physical tape [%4]; the operation timed out.	The ACSLS/LS server may have been busy.	Check the server status and try again.
40165	W	The door was opened on physical library ID %1 %2 %3, which can cause tape operations to stop.	A user opened the library door.	Close the library door.
40167	W	Encryption key %1 does not exist. Decryption and writing operations were disabled for tape ID %2.	The key may have been deleted.	Recreate the key.
40168	W	There is no license for encryption. Decryption and writing operations were disabled for encrypted tape ID %1.	There is no valid license for this option.	Obtain a valid license keycode.
40172	W	The tape shredding job was cancelled for tape [%1], ID %2.	A user requested to stop the operation.	No action is required.

Number	Type	Text	Probable Cause	Suggested Action
40173	E	Physical drive ID %1 [%2][%3] became offline due to a hardware issue; error: %4.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
40175	W	Virtual tape ID %1 [%2] became read-only due to metadata inconsistencies. %3	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
40178	E	User %1 failed to run patch %2.	Execution of a patch could not be launched.	Check to see if there are further errors reported by the patch that describe the reason.
40183	W	Tape duplication job %1 could not create copy job %2 because there were no more physical tapes with the same barcode [%3] in any library.	Either tapes are missing from the library or they have a different barcode.	Try again using physical tapes with the correct barcodes in the library.
40184	W	Tape duplication job %1 running copy job %2 detected that tape [%3] would not be able to be deleted at the end of export because all related duplication jobs did not complete successfully.	To export tape data to a physical tape, multiple jobs were created but one of them did not complete.	Normally, the tape will be moved during the next export job. If this error persists, check related messages and take necessary actions.
40185	W	Tape duplication job %1 running copy job %2 detected that tape [%3] could not be reclaimed because all physical copies were not created.	To export tape data to a physical tape, multiple jobs were created but some of them did not complete.	Wait until all jobs complete.
40186	W	Tape duplication job %1 detected that tape [%2] could not be reclaimed because a related tape copy job did not complete successfully.	To export tape data to a physical tape, multiple jobs were created but one of them did not complete.	Normally, the tape will be moved during the next export job. If this error persists, check related messages and take necessary actions.

Number	Type	Text	Probable Cause	Suggested Action
40187	W	Tape duplication job %1 detected that tape [%2] could not be moved to the vault because a related tape copy job did not complete successfully.	To export tape data to a physical tape, multiple jobs were created but one of them did not complete.	Normally, the tape will be moved during the next export job. If this error persists, check related messages and take necessary actions.
40188	W	Tape stacking import job %1 completed; virtual tape library ID: %2; source tape: [%3], throughput: %4 MB/ min. [%5] tapes were not imported due to duplicate barcodes.	Some tapes are using the same barcodes as other tapes in this library probably because they were previously used in another library.	Try again with different barcodes for virtual tapes to eliminate duplicates.
40189	W	Tape stacking scan job %1 detected that physical tape [%2] is not a stacked tape.	The physical tape header does not indicate that the tape is a stacked tape.	Make sure a stacked tape is used for this operation.
40191	E	The hardware compression card failed; error code: %1, complete code: %2, hardware complete code: %3, tape: [%4], ID [%5].	The compression card reported a failure.	Check the error details and take necessary actions to fix the hardware; you may need to contact the compression card vendor.
40193	E	A query to the %1 server with IP %2 and ACS %3 failed because ACS ID %4 was invalid.	The ACS ID specified in the ACSLS configuration is incorrect.	Use a correct ACS ID in the ACSLS configuration.
40194	E	The portmap/rpcbind daemon, required for ACSLS/LS servers, is not running. Start the daemon and then rescan physical devices.	The daemon may have been inadvertently stopped.	Check and make sure the daemon is running.
40195	W	An export job could not be submitted for tape [%1] because it is being deduplicated.	An attempt was made to manually export a tape while a deduplication operation was in progress.	Run export jobs after deduplication is completed.

Number	Type	Text	Probable Cause	Suggested Action
40197	W	The tape database consistency check detected non-critical errors. Check the system log for more information.	The underlying physical device may have some issues that compromise the integrity of the tape database.	Check the physical device status, device connectivity and switches, and the storage log.
40198	E	The tape database consistency check detected critical errors. Check the system log for more information.	The underlying physical device may have some issues that compromise the integrity of the tape database.	Check the physical device status, device connectivity and switches, and the storage log.
40199	E	The consistency check for compressed data detected an invalid data signature on tape [%1], ID %2.	The data integrity signature in compressed blocks does not match the original signature probably due to some storage issues.	Check the physical device status, device connectivity and switches, and the storage log.
40201	E	Tape ID %1 could not be created because the ID already exists.	This may be due to some inconsistency issues with the tape database.	Try to create the tape again; if the issue persists, contact Technical Support.
40202	W	Tape [%1], ID %2 could not be found.	This may be due to some inconsistency issues with the tape database.	Contact Technical Support.
40203	W	Tape [%1], ID %2 became read-only due to a maximum capacity mismatch.	The tape does not have the expected value for the maximum capacity.	Contact Technical Support.
40204	W	Tape [%1], ID %2 became read-only due to an allocated size mismatch.	The tape does not have the expected allocation size.	Contact Technical Support.
40205	W	Tape [%1], ID %2 became read-only due to a property mismatch for stub tape %3.	The stub tape does not have an expected property.	Contact Technical Support.
40207	W	Virtual tape ID %1 [%2] became read-only due to tape header inconsistencies.	This is due to some inconsistency issues.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
40208	E	ACSLS library ID %1 is offline; ACS: %2, pool: %3. Check the physical library and then perform an inventory to bring the library back online.	The ACSLS library might have some issues.	Check ACSLS logs and take necessary actions.
40209	E	ACSLS drive ID %1 (%2,%3,%4,%5) is offline. Check the physical library and then perform an inventory to bring the drive back online.	The ACSLS library might have some issues.	Check ACSLS logs and take necessary actions.
40210	E	Physical library ID %1 is disabled for maintenance. Enable the library when maintenance completes.	A user has disabled the library for maintenance purposes.	Wait until maintenance is complete.
40211	W	Metadata was rolled back on virtual tape %1 because of a write failure; details: %2.	The write error may be due to a storage issue.	Check the physical device status, device connectivity and switches, and the storage log.
40214	W	Physical tape [%1] was not available to start a tape caching job. The job will be started when the tape becomes available.	The current physical tape position got changed by third party application during tape stacking job running.	Make sure physical drives are not shared with other applications.
40216	E	Tape stacking job %1 did not complete because the current physical tape position (%2) did not match with the number of blocks (%3) transferred to physical tape [%4] in physical drive ID %5. The start position is %6; the expected current position is %7.	This is due to some inconsistency issues.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
40220	W	Turbo deduplication was stopped for tape [%1], ID %2, in virtual drive ID %3; this tape will be deduplicated using the post-processing method.	The tape scanner encountered an error.	Check the log files to identify the actual reason and take necessary actions.
40222	W	Inline deduplication was stopped for tape [%1], ID %2, in virtual drive ID %3.	The tape scanner encountered an error.	Check the log files to identify the actual reason and take necessary actions.
40223	W	Turbo deduplication was stopped for tape [%1], ID %2, in virtual drive ID %3.	The tape scanner encountered an error.	Check the log files to identify the actual reason and take necessary actions.
40224	W	Inline deduplication was not used for tape [%1], ID %2, in virtual drive ID %3; this tape will be deduplicated using the post-processing method.	The tape scanner does not support this tape format for inline deduplication or encountered an error.	Check the log files to identify the actual reason and take appropriate actions, if necessary.
40226	W	Tape [%1], ID %2, could not be loaded into virtual drive ID %3.	This is due to some other errors.	Check the system log to identify any related errors and take necessary actions.
40227	E	Tape [%1] could not be imported into virtual library ID %2 from stacked physical tape [%3] by job ID %4.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
40228	E	The tape import/export job queue could not be initialized on tape database %1.	This may be due to insufficient system resources.	Restart the server modules.
40404	E	Physical tape [%1] cannot be labeled as a stacked tape in physical drive ID %2.	The physical tape has some data and cannot be labeled without overwrite mode.	Use an empty tape or use overwrite mode if you do not want to keep existing data on this tape.
40405	W	Encryption is not activated, resulting in I/O failures for encrypted tape [%1], ID %2.	The encryption activation password was not entered after server startup.	Activate encryption in order to access encrypted virtual tapes.

Number	Type	Text	Probable Cause	Suggested Action
50000	E	iSCSI target name is missing in login session from initiator %1.	The iSCSI initiator may not be compatible.	Check the iSCSI initiator on the client side.
50002	E	iSCSI login to non-existent target %1 was requested by initiator %2	The iSCSI target does not exist any longer.	Check the iSCSI initiator on the client side and the iSCSI configuration on the server. Remove targets from the configuration if they do not exist.
50003	E	iSCSI CHAP authentication method was rejected in the login request to target %1 from initiator %2.	The CHAP settings are not valid.	Check the iSCSI CHAP secret settings on the server and the client sides.
50115	E	Deduplication cluster %1 could not detect Fibre Channel targets of %2.	This may be due to invalid Fibre Channel zoning or connectivity.	Check the Fibre Channel zoning and make sure the cluster nodes can see each other's target ports.
50204	E	A virtual drive with serial number %1 could not be found on deduplication cluster node %2.	This may be due to a previous redundant node configuration that has not been removed completely.	You need to remove the obsolete configuration. Contact Technical Support for help.
50206	E	Deduplication data disk %1 is not available.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
50501	E	Deduplication Policy %1: The policy failed to be saved; error: %2.	This may be due to insufficient system resources.	Restart the server modules.
50503	E	Deduplication Policy %1: The policy failed to be deleted; error: %2.	This may be due to insufficient system resources.	Restart the server modules.

Number	Type	Text	Probable Cause	Suggested Action
50514	E	Deduplication Policy %1: The scan process failed to start for tape [%2] with ID %3; reason: %4, server: %5, cluster: %6.	Either the deduplication server is not healthy or it cannot be connected.	Make sure services are up and running on both the local server and the deduplication server. Refer to the error specified in the message and check the log for additional infomation and take necessary actions. You may need to restart services on this server to release tapes for deduplication.
50516	E	Deduplication Policy %1: The scan process failed for tape [%2] with ID %3; error: %4, server: %5, cluster: %6.	Either the deduplication server is not healthy or it cannot be connected.	Make sure services are up and running on both the local server and the deduplication server. Refer to the error specified in the message and check the log for additional infomation and take necessary actions.
50538	W	Space reclamation failed to start on node %1; reason: %2.	Space reclamation encountered a network connection issue and reported the error specified in the message.	Check the log for related errors and take necessary actions.
50546	E	Space reclamation failed for tape %1 [%2].	The tape cannot be read due to an issue with the underlying physical device or the connection to the device.	Check the physical device status, device connectivity and switches, and the storage log.
50549	W	Deduplication Policy %1: The policy was not started because the repository on deduplication server %2 is still being loaded.	The deduplication repository is in the startup phase and is not ready yet.	Wait until the server startup phase is complete.
50550	E	Deduplication Policy %1: The policy was not started because deduplication server %2 could not be accessed.	This may be due to a network connectivity issue with the deduplication server or the server is not healthy.	Check that the deduplication server can be reached and all modules are up and running.

Number	Type	Text	Probable Cause	Suggested Action
50551	W	The deduplication policy was not started because the repository on deduplication server %1 is still being loaded.	The deduplication repository is in the startup phase and is not ready yet.	Wait until the server startup phase is complete.
50552	E	The deduplication policy was not started because deduplication server %1 could not be accessed.	This may be due to a network connectivity issue with the deduplication server or the server is not healthy.	Check that the deduplication server can be reached and all modules are up and running.
50553	W	Deduplication Policy %1: Tape [%2] scan process was cancelled by a backup client or a user.	A user requested to stop the operation or the tape was requested by backup software.	No action is required.
50554	E	Space reclamation failed because VIT %1 [%2] has an old format.	The VIT was created by a previous version.	Run the deduplication policy to convert the VIT to the new format.
50557	W	Deduplication Policy %1: Tape [%2] unique data replication was cancelled by a backup client or a user.	A user requested to stop replication or the tape was requested by backup software.	No action is required.
50561	W	Deduplication Policy %1: Tape [%2] unique data replication was cancelled because it was outside of the specified time range for the policy. Replication will be triggered during the next specified time range.	The replication duration was longer than the policy time range.	Extend the policy time range or wait for the next time range.
50563	W	Deduplication Policy %1: The policy stopped because there were no tapes in the policy.	The deduplication policy does not have any tapes.	Add tapes to the policy before running the policy.

Number	Type	Text	Probable Cause	Suggested Action
50568	E	The replica tape ID %1 could not be resolved to an LVIT due to the following reason: %2	The replication process for a VIT reports the error specified in the message.	Check the log for related errors and take necessary actions.
50569	W	Deduplication Policy %1: Processing will be stopped for tape [%2]; reason: %3	The deduplication process reports the error specified in the message.	Check the log for related errors and take necessary actions.
50570	W	Deduplication Policy %1: Processing is not allowed at this moment for tape [%2]; reason: %3. Post-processing deduplication will be performed for the tape.	The deduplication process reports the error specified in the message.	Check the log for related errors and take necessary actions.
50573	E	Deduplication Policy %1: Processing failed for tape [%2] in drive ID %3; reason: %4	The deduplication process reports the error specified in the message.	Check the log for related errors and take necessary actions.
50575	E	Deduplication Policy %1: Processing for tape [%2] with ID %3 failed %4 times. The job will not be retried.	Too many errors occurred during tape data deduplication.	Check the system log for related errors and take necessary actions.
50602	E	NAS replication failed due to a network connectivity error.	This may be due to a network connectivity issue.	Check connectivity between the primary and replica; check network parameters.
50603	E	A NAS resource being replicated is full.	The file system is full.	Remove some unnecessary files or expand the resource.
50609	E	The Deduplication Repository is full on the replica server.	The replica server reported that its repository was full.	Run reclamation on the replica server.
50612	E	A NAS operation failed because information could not be exchanged between the server and the console.	The system location that is used temporarily for information exchange cannot be accessed probably due to low system resources.	If this issue persists, you may need to restart the server machine.
50613	E	The NAS server encountered a memory allocation failure.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory.

Number	Type	Text	Probable Cause	Suggested Action
50615	E	NAS replication failed.	An unexpected error occurred during replication.	Wait for the next scheduled start time or run replication manually.
50617	E	Replication cannot proceed because the replica server is running an older version.	Replication is only allowed from an older version to a newer version.	Make sure you are using compatible versions on both servers.
50631	E	The NAS deduplication configuration failed to be saved; error: %1.	This may be due to insufficient disk space, an unhealthy file system, or an issue with the underlying physical device or device connectivity.	Ensure there is enough disk space and the disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
50653	E	NAS outgoing replication session ID %1 terminated.	Replication was stopped unexpectedly, probably due to services being restarted.	Wait for the next scheduled start time or run replication manually.
50662	E	NAS outgoing replication session failed for file %1.	The stub file could not be read. This may be due to an unhealthy file system or an issue with the underlying physical device or device connectivity.	Check the system log to determine the nature of the error. Fix any file system, hardware, or connectivity issues.
50669	E	NAS reclamation session ID %1 failed.	Due to previous file system errors, the reclamation process stopped.	Check that file systems are mounted and can be browsed. Check related errors and take necessary actions.
50675	E	NAS incoming replication stopped; user: %1, host: %2, session ID: %3.	The server did not receive data from the source probably due to a network issue or an error on the source server.	Check connectivity between the primary and replica; check network parameters and the source server.
50682	E	NAS incoming replication failed for file %1.	The server could not reach the source to acknowledge the replication probably due to a network issue.	Check connectivity between the source and replica; check network parameters.
50688	E	NAS integrity check detected a checksum mismatch on file %1; type: %2, recorded: %3, expected: %4.	The deduplicated file did not contain expected data.	Check the system log to see if there are any file system, storage, or connectivity issues. If there are no issues with your environment, Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
50693	E	NAS integrity check failed; start time: %1, end time: %2, files processed: %3, policy ID: %4.	A stub file could not be read because the related target file was deleted or was incomplete.	Run the integrity check again.
50695	E	NAS integrity check detected some problems; total files: %1, passed: %2, failed: %3, mismatched: %4, skipped: %5, ineligible: %6.	Some deduplicated files did not contain expected data.	Check the system log to see if there are any file system, storage, or connectivity issues. If there are no issues with your environment, Contact Technical Support.
50699	E	NAS integrity check failed to read file %1.	This may be due to an unhealthy file system or an issue with the underlying physical device or device connectivity.	Check the system log to determine the nature of the error. Fix any file system, hardware, or connectivity issues.
50702	E	Scan of tape [%1] in drive %2 on server %3 failed to start; reason: %4.	The deduplication server may have lost access to the VTL resources.	Check FC connectivity to ensure that the deduplication cluster can access tape drives and tapes.
50706	E	Scan of tape [%1] in drive %2 failed to create a folder resource; reason: %3.	The folder disk may not be accessible.	Make sure the storage that contains the folder disk is accessible by the deduplication cluster.
50707	E	Scan of tape [%1] in drive %2 failed to store data in the repository; reason: %3.	This may be due to a storage or a connectivity issue.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues.
50709	E	Scan of tape [%1] in drive %2 on server %3 failed; reason: %4.	This may be due to the tape drive being inaccessible, the deduplication server may be unavailable or busy due to heavy I/O, or a storage issue.	Take necessary actions based on the specified reason in the message, which may include checking FC connectivity to ensure that the deduplication cluster can access tape drives and repository storage.
50710	W	Scan of tape [%1] in drive %2 was cancelled.	Scanning was cancelled manually or the tape was requested by backup software.	No action is required.

Number	Type	Text	Probable Cause	Suggested Action
50714	E	Scan of tape [%1] in drive %2 failed to convert the tape to a VIT after running for %3.	Due to previous errors, the scan process stopped.	Check related errors and take necessary actions.
50719	E	Scan of tape [%1] in drive %2 on server %3 failed to generate the index information.	The deduplication server may have lost access to the VTL resources used for the index information or there is not enough storage on the VTL server.	Check FC connectivity to ensure that the deduplication cluster can access tape drives and tapes; check that there is enough storage on the VTL server and add more, if necessary.
50723	E	Scan of tape [%1] in drive %2 failed to validate data integrity; reason: %3.	The scanner process reports the error specified in the message.	Check the log for related errors and take necessary actions.
50724	E	Scan of tape [%1] detected an inconsistency in the metadata to be used for deduplication; the source tape did not get deduplicated.	This may be due to data inconsistencies.	Contact Technical Support.
50802	E	The resolving process failed for VIT %1 using source drive %2 and destination drive %3 on VTL server %4; reason: %5.	The resolver process reports the error specified in the message.	Check the log for related errors and take necessary actions.
50803	E	Replication failed to process a data block from VIT %1 using source drive %2 and destination drive %3; reason: %4.	A block of data could not be replicated to the target due to the specified reason in the message.	Take necessary actions based on the specified error.
50805	E	Replication of VIT %1 failed; source drive: %2, destination drive: %3, VTL server: %4, reason: %5, time elapsed: %6, total data: %7, transferred data: %8.	Unique data replication reports the error specified in the message.	Check the log for related errors and take necessary actions.

Number	Type	Text	Probable Cause	Suggested Action
50806	W	Replication of VIT %1 was cancelled; source drive: %2, destination drive: %3.	Replication was cancelled manually or the tape was requested by backup software.	No action is required.
50902	E	The deduplication client module initialization failed; error: %1.	The deduplication client module reports the error specified in the message; this is preventing the scanner or resolver process from starting.	Take necessary actions based on the error.
50906	E	The deduplication client module failed to store a data block for folder %1; reason: %2.	Data failed to be written to the deduplication repository.	Take necessary actions based on the specified error.
50907	E	The deduplication client module failed to reach the deduplication server; error: %1.	The deduplication server may not be running or there may be a connectivity issue.	Take necessary actions based on the specified error.
50908	E	The deduplication client module failed to initialize the deduplication server; error: %1.	The deduplication server is not configured or it has already been started.	Check configuration and take necessary actions based on the specified error.
50909	E	The deduplication client module failed to get the deduplication server configuration; error: %1.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50910	E	The deduplication client module failed to set the deduplication server configuration; error: %1.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50911	E	The deduplication client module failed to get hashed data; error: %1.	This may be due to data inconsistencies.	Contact Technical Support.
50912	E	The deduplication client module failed to get the location of hashed data; error: %1.	This may be due to data inconsistencies.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
50913	E	The deduplication client module failed to create a folder; error: %1.	Folder disks may not be large enough or cannot be accessed.	If this is due to insufficient capacity, run reclamation and add folder disks if necessary. Otherwise, check storage connectivity and the status of the folder disks.
50915	E	The deduplication client module failed to close folder %1; reason: %2.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50916	W	The deduplication client module failed to discard folder %1; reason: %2.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50917	E	The deduplication client module failed to flush data for folder %1; reason: %2.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50919	W	The deduplication client module failed to initialize; error %1.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50920	E	Data could not be written to the deduplication index; reason: %1.	The index disk may be full or a SCSI write error occurred.	If the index disk is full, run pruning. Otherwise, check storage connectivity and the status of the index disk.
50950	W	The SSD index cache is being rebuilt from the index drive due to an ungraceful shutdown.	This is due to a server failure, an ungraceful shutdown, or failover to the standby server.	Wait for the index cache to complete loading; this may take a while.
50952	E	The deduplication server initialization failed; reason: %1.	The deduplication client module could not start successfully due to the specified reason.	Ignore this messages if it occurs when the system is initially configured. Otherwise, take necessary actions based on the error.
50956	E	Deduplication data disk %1 is full.	There is not enough space on the data disk for backup.	Run space reclaimation to free disk space or add more storage.
50958	E	The deduplication index cache is full.	There is not enough memory or SSD disk space.	Run space reclaimation to free disk space or add more storage.

Number	Type	Text	Probable Cause	Suggested Action
50963	C	The deduplication index disk is too small; it is below the minimum required of %1 MB. Increase the size by at least %2 MB.	This is due to an improper configuration.	Add index disks as recommended.
50964	C	The deduplication index disk free space is %1 percent, which is critically low.	The capacity of index disks is not large enough even after pruning.	Add index disks.
50965	C	The deduplication folder disk free space is %1 percent, which is critically low.	The capacity of folder disks is not large enough even after reclamation.	Add folder disks.
50966	C	The deduplication data disk free space is %1 percent, which is critically low.	The capacity of data disks is not large enough even after reclamation.	Add data disks.
50967	C	The deduplication index cache is critically low; only %1 percent is free.	The amount of memory allocated for the repository index cache is not enough and is still low after reclamation.	Check the repository sizing guidelines and add required memory to the server.
50968	E	Encryption is not activated resulting in deduplication failure for the encrypted repository.	No password was entered for encryption activation.	Activate encryption using a valid password.
51007	W	Report %1 failed to be sent via email because mail server %2 could not be connected.	This may be due to a network error.	Check connectivity to the mail server.
51008	W	Report %1 failed to be sent by mail server %2.	This may be due to a network error or limitations on the mail server.	Check connectivity to the mail server and review the mail server log to identify the issue.
51009	W	There is no SMTP server provided to email reports.	The SMTP mail server is unavailable.	Check the SMTP server information is valid and the server is up and running.

Number	Type	Text	Probable Cause	Suggested Action
51205	E	Encryption failed to be activated for virtual tapes or the deduplication repository.	An invalid password was entered for encryption activation.	Activate encryption for virtual tapes or the deduplication repository using a valid password.
51701	E	The file system is now mounted in read-only mode due to a file system error.	This may be due to a storage issue.	Check the system log for additional information and take necessary actions based on the error.
51801	E	The NAS option failed to be enabled; error: %1.	The system may have been busy and did not have enough resources.	Check the system status. Retry later. If this error persists, you may need to restart server modules.
51803	E	The NAS option failed to be disabled; error: %1.	The system may have been busy and did not have enough resources.	Check the system status. Retry later. If this error persists, you may need to restart server modules.
52009	W	The file system on device ID %1 could not be detected; the reclamation session will be aborted.	The NAS file system is not mounted; reclamation requires all file systems to be mounted.	Mount the file system or restart NAS services.
52093	E	NAS Fast Copy process with ID %1 stopped; reason: %2.	The process reports the error specified in the message; this may be due to a storage issue.	Check the system log for additional information and take necessary actions based on the error.
52102	W	In network bonding configuration %1, slave interface %2 is down.	This may be due to an issue with the network port or adapter.	Check the network port or adapter related to the device and fix any problems.
52104	W	In network bonding configuration %1, slave interface %2 was removed.	This was most probably triggered by a user.	If this was not triggered by a user, check the system log to identify any related errors and take necessary actions based on the error.
52106	E	Network bonding device %1 is down.	This may be due to an issue with the network port, adapter, or connectivity.	Check the network port or adapter related to the device and fix any hardware or connectivity problems.

Number	Type	Text	Probable Cause	Suggested Action
53000	W	User account %1 used an incorrect password too many times when attempting to query SNMP information. The SNMP module has been shut down temporarily and will automatically restart in %2 minutes.	For security purposes, invalid passwords result in module interruption.	Try again later.
60003	E	The OST server module failed to allocate memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
60004	E	The OST server module failed to write to image ID %1; error: %2.	This may be due to insufficient disk space, an unhealthy file system, or an issue with the underlying physical device or device connectivity.	Ensure there is enough disk space and the disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
60005	E	The OST server module failed to read from image ID %1; error: %2.	This may be due to an unhealthy file system or an issue with the underlying physical device or device connectivity.	Check the system log to determine the nature of the error. Fix any file system, hardware, or connectivity issues.
61002	E	User %1 failed to enable the OST option.	The system may have been busy and did not have enough resources.	Check the system status. Retry later. If this error persists, you may need to restart server modules.
61004	E	User %1 failed to disable the OST option.	The system may have been busy and did not have enough resources.	Check the system status. Retry later. If this error persists, you may need to restart server modules.
62000	W	Replication will be retried because connection to replica server %1 failed.	Either the network connection is down or the replica server is down.	Check the state of the replica server. Determine and correct either the network problem or server problem.

Number	Type	Text	Probable Cause	Suggested Action
62001	W	Replication will be retried because the configuration file could not be opened.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
62002	W	Replication will be retried because memory allocation failed.	The system may have been busy and did not have enough memory.	Check the memory usage of different processes and stop unnecessary processes to free up memory. You may need to restart server modules.
62003	W	Replication will be retried because virtual tape ID %1 no longer has a replica.	The replica device may have been deleted or promoted while the primary server was down.	Reconfigure replication and create a new replica or use the replica that had been promoted.
62004	W	Replication will be retried because remote device ID %1 does not exist or is not a replica.	The replica device may have been deleted.	Reconfigure replication.
62005	W	Replication will be retried because virtual tape ID %1 is not currently available.	The tape is loaded in a drive and is in use.	Wait for the server to retry.
62006	W	Replication will be retried because Replication could not proceed because ID %1 could not located for the virtual tape.	The tape ID is missing at the kernel level.	Contact Technical Support.
62007	W	Replication will be retried because replica server %1 was busy.	This may be due to too many pending requests on the replica server.	Wait for the server to retry or run replication manually.
62008	W	Replication will be retried because replica server %1 did not have enough space.	The replica server is running low on disk space.	Check the disk usage and try to reclaim disk space or add more storage.
62009	W	Replication will be retried because unexpected error %1 occurred.	Replication reports the error specified in the message.	Check system logs on both servers and take necessary actions based on the error.

Number	Type	Text	Probable Cause	Suggested Action
62010	W	Replication will be retried because virtual tape replica ID %1 for virtual tape ID %2 could not be expanded because the maximum licensed capacity on the replica server was reached.	All storage capacity licenses have been used.	Obtain additional license key codes.
62012	W	Replication will be retried for virtual tape ID %1 because the replica status is %2.	The replication configuration may not be valid.	Check the configuration on the replica server. Check system logs on both servers for additional information.
62013	W	Replication will be retried for virtual tape ID %1 because it failed with error %2.	Replication reports the error specified in the message.	Check the replica device and system logs; take necessary actions based on the error.
62014	W	Replication will be retried for virtual tape ID %1 due to network transport error %2.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.
62015	W	Replication will be retried for virtual tape ID %1 because the primary disk failed with error %2.	The primary device reports the error specified in the message.	Check the device and take necessary actions based on the error.
62016	W	Replication will be retried for virtual tape ID %1 because the replica device failed with error %2.	The replica reports the error specified in the message.	Check the device on the replica server and take necessary actions based on the error.
62017	W	Replication will be retried for virtual tape ID %1 because the exchange of the replication control map between servers failed; error: %2.	This may be due to a connectivity issue.	Check connectivity between the primary and replica. Check system logs on both servers for additional information.
62018	W	Replication will be retried for virtual tape ID %1 because it failed; %2.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.

Number	Type	Text	Probable Cause	Suggested Action
62084	E	User %1 failed to rename client %2 to %3 due to a memory allocation error.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory.

For any error not listed in these tables, contact DSI Technical Support.

Appendix

This appendix contains information about:

- Ports used by VTL/SIR
- LUN migration
- Enabling FIPS
- Setting LD_LIBRARY_PATH

Port usage

VTL and SIR servers use the ports listed in the following tables for incoming requests. Network firewalls should allow access through these ports for successful communications. In order to maintain a high level of security, you should disable all unnecessary ports. The ports are not used unless the associated option is enabled in VTL/SIR. For DSI appliances, the ports marked ● are enabled by default.

Port	Protocol	Open On	Used By	Description
20	TCP/UDP	VTL and SIR server	FTP client	Standard FTP port used for file data transfer.
21	TCP/UDP	VTL and SIR server	FTP client	Standard FTP port used for sending commands.
22 ●	TCP	VTL and SIR server	Host client	Standard Secure Shell (SSH) port used for remote connection to the server.
25	TCP/UDP	VTL and SIR server	Mail server	Standard SMTP port used for Email Alerts.
80 ●	HTTP	VTL and SIR server	FalconStor web license server	Standard Internet port for online registration of license key codes and also used by the `dwlpatch.pl` utility for the automatic download of server patches. License registration material is sent back using HTTP protocol, where a local random port number is used on the server, just like a typical Web-based page. The firewall does not block the reply if the 'established bit' is set to let established traffic in.
80 ●	HTTP	VTL and SIR server	Console machine	Standard HTTP port used to access the DSI Management Console.

Port	Protocol	Open On	Used By	Description
81 ●	HTTP	VTL and SIR server	Console machine	Standard HTTP port used to access the DSI Management Console via Web Start.
111	TCP/UDP	VTL server	NFS client	NFS and ACSLS* port used for RPC program number mapper. The port is assigned via the SUNRPC protocol. The ports vary, so it is not feasible or convenient to keep checking them and reprogramming a firewall. Most firewalls have a setting to "Enable NFS", upon which they will change the settings if the ports change.
123	UDP	VTL and SIR server	NTP server	Standard Network Time Protocol (NTP) transport layer used to access external time servers.
137	UDP	VTL server	CIFS client	ipstornmbd NETBIOS Name Service used for CIFS protocol.
138	UDP	VTL server	CIFS client	ipstornmbd NETBIOS Datagram Service used for CIFS protocol.
139	TCP	VTL server	CIFS client	ipstorsmbd NETBIOS Session Service used for CIFS protocol.
161	TCP/UDP	VTL and SIR server	SNMP server	Standard Simple Network Management Protocol (SNMP) port used to query VTL/ SIR MIBs.
199	UDP	VTL and SIR server	SNMP client	Standard SNMP multiplexing (SMUX) protocol port used to query Dell OpenManage system MIBs.
445	TCP	VTL server	CIFS client	Microsoft Directory Services for Windows shares using Active Directory.
623 ●	UDP	VTL and SIR server	Failover partner server	IPMI power control port used to power off the failed server in a failover configuration.
705	UDP	VTL and SIR server	SNMP client	Standard SNMP AgentIX port used to query agents such as Fujitsu ServerView.
1311	HTTPS	VTL and SIR server	Dell OpenManage server	Management port used for hardware configuration of Dell servers.
2049	TCP/UDP	VTL server	NFS client	nfsd NFS server (NFS).
3260	TCP	VTL and SIR server	iSCSI client	Communication port between iSCSI clients and the server.
5001	TCP	VTL and SIR server	Clients and replica server	isttcp port used to test network connection and measure bandwidth performance.

Port	Protocol	Open On	Used By	Description
10000	TCP	VTL server	Symantec NetBackup server	Communication port between Symantec NetBackup and VTL for NDMP backup/ restore and for the FalconStor OpenStorage Option (FSOST) Copy to Tape feature.
10161	TCP/UDP	VTL and SIR server	SNMP	SNMP communication port.
11576 ●	TCP	VTL and SIR server	CLI, FSOST, and the DSI Management Console	Secure RPC communication port between DSI Management Console and the server and also between FSOST and VTL server.
11577 ●	TCP/UDP	VTL server	Replication servers	Communication port between servers for data replication. This port is only open while replication is being performed.
11579	TCP/UDP	VTL and SIR server	Replication servers	Replication authentication.
11580 ●	TCP	VTL server	Failover servers	Communication port used between a pair of failover servers. It is not required for a standalone server.
11582 ●	TCP	VTL and SIR server	CLI	Communication port for used to send CLI commands to the server.
11583	TCP	VTL and SIR server	DSI Management Console	Communication port used to send report requests (report schedules, global replication report, statistics log, and configuration updates) to the configuration management module on the server.
11584	TCP	VTL server	Replication servers	Communication port between replication servers for data replication of deduplicated data.
11676	TCP	VTL server	VTL and SIR server	Communication port for deduplication management and repository health check.
11676 ●	TCP	SIR server	SIR cluster nodes	Communication port between SIR cluster nodes for deduplication management and repository health check.
11780	TCP	SIR server	Replication servers	Communication between replication servers if the replica uses SIR software prior to version 7.5 (backward compatibility).
11781 ●	TCP	SIR server	Replication servers	Unencryption port for replication.
11782 ●	TCP	SIR server	Replication servers	Encryption port for replication.

Port	Protocol	Open On	Used By	Description
18651	TCP	VTL server	Replication servers	Communication port between servers for non-encrypted replication.
18652	TCP	VTL server	Replication servers	Communication port between servers for encrypted replication.
18720	TCP	VTL server	FSOST client	OST-related operations
18750	TCP	VTL server	FSOST	Communication port between FSOST and the VTL server.
20491	TCP/UDP	VTL server	NFS client	NFS interface-related operations.
20492	TCP/UDP	VTL server	NFS client	NFS interface-related operations.
20493	TCP/UDP	VTL server	NFS client	NFS interface-related operations.
46373	TCP	VTL server	FSOST and NDMP	Used to transfer an FSOST image to physical tape.

* PortMapper requires dynamic ports to be open. This requires the ACSLS to be in the same VLAN with ACSLS server.

Notes:

- Although you may temporarily open some ports during initial setup of the VTL server, such as the telnet port (23) and FTP ports (20 and 21), you should shut them down after you have done your work.
- Make sure there are no blocked ports and loopback device access is open.

LUN migration

VTL offers a command line tool to migrate data on a LUN to one or more other LUNs. With this tool, you can specify the target LUN(s) or let the system auto-select the target LUN(s). With its built-in restart capability, incomplete/failed migration jobs can be restarted from where they left off.

After data is moved to the target LUN, the source LUN will be unassigned.

Requirements

The following requirements must be met before running a LUN migration job:

- The source LUN must have data on it. You cannot migrate an empty LUN. You will get an error if you try and the job will fail.
- You need to move all of the tapes on the LUN to be migrated to the virtual vault before you run a migration job or your job will fail.
- The target LUN must have enough space or you will get an error and the job will fail. Because data is copied over sector by sector, you must have enough space for each sector on the target LUN. For example, the source LUN is 10 GB and there are two target LUNs that are 5 GB each. Even though the total of the two target LUNs is 10 GB, the job will fail if the first sector on the source LUN is 6 GB (since each target LUN is only 5 GB and neither is large enough to accept the 10 GB segment).
- Database LUNs (LUNs with VTL database segments on it) can be selected as the source but the database segment will not be moved. Database LUNs cannot be selected as the target. You will get an error if you try and the job will fail. If you run automatic migration when there are only DB LUNs, you will get a *createmigratedummytape fail* error and the job will fail.
- **It is very important** that there is no I/O occurring on the source or target LUN while LUN migration is in progress.
- **It is very important** to turn off space reclamation before LUN migration. To do this, right-click the VTL server object and select *Options* --> *Deduplication* --> *Reclamation* --> *Reclamation Properties.* Be sure to turn it back on after the migration is complete.

Manual migration

Manual migration lets you specify the target LUN(s). To run a manual migration job, go to $ISHOME/bin and run the following command.

```
lunmigration.sh <source LUN ACSL> <target LUN ACSL>
```

The format for specifying the ACSL is `a:c:s:l`

For multiple target LUNs, you must separate each target LUN with a comma:

```
lunmigration.sh <source LUN ACSL> <target LUN ACSL 1>, ..., <target LUN
ACSL n>
```

For example: `lunmigration.sh 1:0:0:6 1:0:0:7,1:0:0:8`

Automatic migration

Automatic migration lets the system auto-select the target LUN(s). To run an automatic migration job, go to $ISHOME/bin and run:

```
lunmigration.sh <source LUN ACSL> AUTO
```

`AUTO` must be in uppercase.

For example: `lunmigration.sh 1:0:0:6 AUTO`

Restart a job

To restart an incomplete/failed migration job from where it left off, run the following before re-running the migration job:

```
lmclean.sh
```

Enable FIPS

VTL and SIR can be configured to be Federal Information Processing Standard (FIPS) compliant. FIPS is a Federal government standard for security and interoperability and is useful for government organizations that have FIPS certification as a product requirement.

FIPS certification is supported for VTL and SIR servers:

- Running Red Hat/Oracle Enterprise Linux 5.7, 5.10, and 6.5.
- Installed with DSI's VTL USB key without any modifications. It is not supported if the operating system was modified after installation.

The following procedure can be performed at any time, before or after VTL and SIR have been configured.

1. Disable PRELINKING in the /etc/sysconfig/prelink file.

   ```
   vi /etc/sysconfig/prelink
   ```

   Change the line "PRELINKING=yes" to "PRELINKING=no"

2. Undo prelink for all binaries in the system.

   ```
   prelink -u -a
   ```

3. Regenerate the initrd.

   ```
   mkinitrd --with-fips -f /boot/initrd-$(uname -r).fips.img $(uname
   -r)
   ```

4. Add a new boot menu entry in the boot loader configuration file.

   ```
   vi /boot/grub/menu.lst
   ```

   - Create a new boot entry by copying the current boot entry.
   - Append `fips=1` in the kernel line.
   - Edit the initrd line to point to /initrd-2.6.18-274.18.1.el5.fips.img.

   The new boot menu entry should look like the following:

   ```
   title FalconStor VTL-S 2.6.32-504.1.3.el6.x86_64(FIPS)
   root (hd0,0)
   kernel /vmlinuz-2.6.32-504.1.3.el6.x86_64 ro root=LABEL=/1 rd_NO_LUKS
   rd_NO_LVM LANG=en_US.UTF-8 rd_NO_MD SYSFONT=latarcyrheb-sun16
   crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM notsc divider=10
   fips=1
   initrd /initrd-2.6.32-504.1.3.el6.x86_64.fips.img
   ```

   Edit the default boot setting to the new entry you just created.

5. Reboot the server and confirm that the server is booting with the new FIPS entry.

6. Verify FIPS mode is enabled after bootup.

   ```
   cat /proc/sys/crypto/fips_enabled
   ```

   This should return a “1”.

Set LD_LIBRARY_PATH

LD_LIBRARY_PATH is an environment variable containing a list of directories used when searching for libraries. The default value of LD_LIBRARY_PATH can be removed in order to control the loading of shared libraries during runtime.

To do this:

1. Open `/etc/profile`.
2. Replace the following:

```
if [ -f /etc/.is.sh ]    # ISENV
then                     # ISENV
    . /etc/.is.sh        # ISENV
fi                       # ISENV
```

with

```
if [ -f /etc/.isnold.sh ]    # ISENV
then                         # ISENV
    . /etc/.isnold.sh        # ISENV
fi                           # ISENV
```

3. Log out of the server and then log in again in order for the change to take effect.

Index