

FALCONSTOR®

VIRTUAL TAPE LIBRARY

USER GUIDE

FalconStor® Virtual Tape Library User Guide

Version 10.10

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*)”. Added or changed content will begin with “(Added/Changed *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*.

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

Copyright © 2003-2022 FalconStor Software. All Rights Reserved.

FalconStor® is a registered trademark of FalconStor Software, Inc. in the United States and other countries.

Linux® is a registered trademark of Linus Torvalds.

Windows® is a registered trademark of Microsoft Corporation.

All other brand and product names are trademarks or registered trademarks of their respective owners.

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.

Contents

Introduction

VTL components	11
Terminology	12
Additional resources	12

Getting Started

Run the FalconStor Management Console	13
Connect to your server	14

VTL Tape Configuration

Prepare a VTL server via the configuration wizard	15
Enter license keys	15
Set up network	15
Set hostname	18
Prepare devices	19
Enable Configuration Repository	19
Create Virtual Tape Library database	19
Enable Virtual Tape Encryption	20
Create virtual libraries	21
Assign virtual library to clients	21
Enable Deduplication	22
Create Deduplication Policy	22
Add tape backup clients	23
Rename a client	23
Configure client access to the VTL server	24
Discover virtual tape libraries from your backup application server	25
Use your operating system to scan for hardware changes	25
Create and run backup jobs	27
Confirm successful backups	28

Console

Console user interface	29
Console modes	30
Understanding the objects in the tree	31
Backup server object	31
Multi-Node Group object	32
Activities object	32
Status object	33
Virtual Tape Library System object	34
Virtual Tape Libraries	34
Virtual Tape Drives	34

Virtual Vault	34
Replica Resources	35
Deduplication Policies	35
Reports object	35
Repositories object	36
Clients object	36
Physical Resources object	36
Group Reports object	36
Set console options	38
Perform system maintenance	40
VLAN tagging	41
Network configuration	42
Set hostname	42
Set date and time	42
Restart VTL	43
Restart network	43
Reboot	43
Halt	43
Add and register licenses	44
Offline registration	44
Monitor space usage	46
Manage user accounts	47
Strong passwords	47
Add or modify users	48
Change password	48
Unlock an account	49
Event Log	50
Attention Required tab	52
Monitor performance	53
System performance	53
Object performance	53
Server properties	54
Activity Database Maintenance settings	54
SNMP Maintenance settings	54
Performance settings	56
Auto Save Config settings	56
Storage Monitoring settings	57
Location settings	57
Apply software patch updates	58
Server patches	58
Console patches	59
Mirror repository disks to protect your configuration	60
Check mirroring status	61
Replace a failed disk	61
Fix a minor disk failure	62
Replace a disk that is part of an active mirror configuration	62
Swap the primary disk with the mirrored copy	62
Replace a failed physical disk without rebooting your server	62
Remove a mirror configuration	63

Manually save/restore the Virtual Tape Library configuration	64
Information and requirements	64
After restoring your configuration	64
Save your configuration	64
Restore configuration	65

Multi-Node Groups

Create a group	67
Add servers to a group	67
Remove a server from a group	69

Physical Resources

Physical Resources object	70
Physical resources icons	70
Prepare physical storage devices	71
Rescan physical devices	73
Storage Devices object	74
Filter storage devices	75
Change the LUN reservation	75
Test physical device throughput	76
Set autopathing	77
Storage Pools object	80
Storage pool access rights	80
Create a storage pool	80
Update storage pool settings	81
Migrate existing virtual tape libraries to storage pools	81

Virtual Tape Libraries, Tape Drives, and Tapes

Virtual tape library, virtual tape drive, and virtual tape icons	84
Create virtual tape libraries	86
Create standalone virtual tape drives	91
Create virtual tapes	92
How virtual tapes are allocated	95
Locate and display virtual tapes in the Console	96
Search by barcode	96
Display virtual tapes	96
Sort all tapes	97
Filter the display of tapes	97
Assign virtual tape libraries and drives to clients	99
Client multipathing	99
Assign a library to a client	100
Assign a client to a virtual tape library or drive	101
Unassign virtual tape libraries, drives, and iSCSI targets from clients	102
Set virtual tape library system properties	102
Use virtual tape drive compression	103

Enable/disable compression	103
Change firmware of a virtual library or drive	104
Shred a virtual tape	104

Deduplication

How tape deduplication works	106
Deduplication methods - at a glance	107
Tape deduplication policies	109
Create tape deduplication policies	109
Modify a tape deduplication policy	119
Add/remove tapes from a tape deduplication policy	121
Manually run a tape deduplication policy	121
Suspend a tape deduplication policy	121
Delete a tape deduplication policy	121
Suspend replication on a target in a tape deduplication policy	122
Resume replication on a target in a tape deduplication policy	122
Manage active tape deduplication policies	123
Monitor deduplication and view statistics	124
Repository statistics	124
Tape information	127
Deduplication Policies object	127
Individual tape deduplication policies	128
Deduplication Job Queue	134
Virtual index tape status	135
Reclaim disk space	136
Space reclamation	136
Index pruning	137
Reclamation requirements	137
Reclamation thresholds	137
Run reclamation	138
Expand deduplication repository	139

Encryption

Deduplication repository encryption	140
Virtual tape encryption	141
Enable encryption	142
Activate encryption for a server	143
Change the encryption activation password for a server	143
Manage encryption keys	144
Create a key	144
Change a key name or password	145
Delete a key	146
Export a key	146
Import a key	147

Data Replication

Replication of virtual tapes without deduplication	150
Replication requirements for virtual tapes	151
Configure replication for virtual tapes	151
Set replication throttling for virtual tapes	157
Check replication status for virtual tapes	157
Promote a virtual tape replica resource	158
Promote a virtual tape replica resource without breaking a replication configuration	158
Change your virtual tape replication configuration options	159
Suspend/resume virtual tape replication schedule	159
Start/stop replication of a virtual tape	159
Remove a virtual tape replication configuration	159
Replication of tapes with deduplication	160
Replication requirements for deduplicated tapes	162
Configure replication for deduplicated tapes	163
Overview of steps to configure replication for deduplicated tapes	163
Add a target deduplication replication server	163
Edit a replication target	165
Set replication throttling for deduplicated tapes	166
Check replication status for deduplicated tapes	167
Access data on a replicated VIT	167
Stop replication of a VIT	167
Remove replication for deduplicated tapes	167
Auto Replication	168
Remote Copy	169

Object Storage

Deduplication data repository on object storage	171
Requirements	171
Configuration	172
Expand deduplication data repository capacity	172
Tape migration to object storage	173
Migration and recovery jobs	173
Configuration	175
Migrate virtual tape data to object storage	176
Convert a virtual tape to a stub tape	176
Recover data from object storage	177
Manage migrated/stub tapes	177
Object storage accounts	179
Add an object storage account	179
Manage object storage accounts	184

iSCSI Configuration

iSCSI users	185
Windows configuration	186

Requirements	186
Enable iSCSI	186
Prepare client initiators	186
Add an iSCSI client	187
Create targets for the iSCSI client to log onto	188
Assign a virtual tape library to the iSCSI target	189
Log the client onto the target	189
Linux configuration	190
Requirements	190
Enable iSCSI	190
Prepare client initiators	190
Add an iSCSI client	191
Create targets for the iSCSI client to log onto	191
Assign a virtual tape library to the iSCSI target	191
Log the client onto the target	191
IBM i configuration	192
Requirements	192
Enable iSCSI	192
Prepare client initiators	192
Add an iSCSI client	192
Create iSCSI target	192
Assign a virtual tape library to the iSCSI target	195
Log the client onto the target	195

IBM i System Configuration

Overview	196
Before you begin	197
Set up the tape library	197
Import cartridges	199
Export cartridges	199

Server Maintenance

Start/stop server modules	200
Important notes about stopping a VTL server	200
Server modules	201

Reports

Report types	203
Create a report	204
Create a one-time report	204
Schedule a report	206
View a report schedule	207
Manage job	207
Create a group report	208
View a report	209

Manage reports	210
Set report properties	210
Export data from a report	212
Email a report	212
Refresh report display	212
Print a report	212
Delete a report	212
Information reports	213
Deduplication - Policy Status	213
Deduplication - Tape Activity	214
Deduplication Replication Status	217
Deduplication Repository - Reclamation	220
Import/Export Jobs	221
Object Storage Migration Jobs	222
Replication Status	224
For virtual tapes	224
For virtual tape replicas	224
Virtual Library and Drive Assignment	225
Virtual Library Information	226
Virtual Tape Activity	227
Virtual Tape Information	229
Overall Summary View	229
Deduplication View	229
Replica Resources View	230
Vault View	230
Detailed Tape View	231
Migration View	232
Usage reports	233
Deduplication - Tape Usage	233
Deduplication Repository - Memory and Space Usage	234
Disk Space Usage History	235
LUNs	237
Allocation reports	239
Disk Space Allocation for Virtual Tapes in Libraries	239
Status report	239
History report	241
Physical Resource Allocation	242
Configuration reports	243
Physical Resources Configuration	243
Storage Pools Configuration	244
Performance report	245
VTL Performance	245

Email Alerts

Configure Email Alerts	248
Email format	254
Modify Email Alerts properties	254
Limit repetitive emails	254

Customize the email for a specific trigger	254
Script/program trigger information	256
Add a new script	256

Command Line

Usage	258
Common arguments	259
Commands	260
Server login/logout	260
Server info	261
Server licensing	262
Physical devices	264
Virtual devices	274
Clients	278
Users	283
Virtual libraries and drives	286
Virtual tapes	293
Deduplication	301
Replication	307
Data encryption	315
Import/Export	318
Object Storage	321
Mirroring	330
Alarm policies	333
Support utilities	335
Reports	337

SNMP Integration

VirtualTapeLibraryMIB	355
falcVtlMonitorMIB - falcVtlMonCapacity	355
falcVtlCapCacheGeneralInfo	355
LibCacheUsage	355
PolicyCacheUsage	356
falcVtlMonitorMIB - falcVtlMonPerformance	356
falcVtlPerfOneDayIntervalDataInfo	356
falcVtlMonPerformanceInfo	356
falcVtlHistoryMIB	357
Activity	357
DashStatistics	357
TapeHistory	358
falcVtlServer	360
Processor	360
NetInterface	360
FailoverInfo	360
falcVtlServerOptionsInfo	360
falcVtlServerInfo	361

falcVtlLibrarySystem	362
VirtualLibrary / falcVtlVirtualLibsInfo	362
VirtualDrive / falcVtlVirtualDrivesInfo	363
VirtualTape / falcVtlVirtualTapesInfo	363
VtlJob / falcVtlJobQueueInfo	365
ReplicaResource	365
ReplicaPhyAllocationLayout	366
ReplicationPolicy / falcVtlReplicaResourcesInfo	366
DeduplicationPolicy / falcVtlDeduplicationPoliciesInfo	367
falcVtlPhysicalResources	368
Storage HBAs	368
StorageDevices	368
StoragePools	369
falcVtlPhysicalResourcesInfo	369
falcVtlSanClients	370
SANClient	370
falcVtlSanClientsInfo	370
Deduplication Repository MIBs	371
falcSirMonitorMIB	371
falcSirMonSwapMemoryInfo	371
falcSirMonCapacityInfo	371
falcSirDedupeRatioRangesInfo	372
falcSirServerPerformanceInfo	372
falcSirMonNodePerformanceInfo	373
falcSirCluster - falcSirClusterConf	374
VTL	374
SIRReplication	374
SIRReplicationSirNodeIP	374
SIRReplicationSirNodePort	375
falcSirCCReclamationInfo	375
SIRNode	375
SIRHashIndexStorage	375
SIRHashFolderStorage	376
SIRHashDataStorage	376
falcSirCluster - falcSirClusterStats	377
Folder	377
falcSirCSSSirStatsSummaryInfo	377
falcSirCSSDedupeResults - DedupeStatsHour	378
falcSirCSSDedupeResults - DedupeStatsDaily	378
falcSirCSSDedupeResults - falcSirCSSDedupeResultsInfo	378
falcSirCSSRepositoryUsage - falcSirCSSRepoObjCapacityInfo	378
falcSirCSSRepositoryUsage - falcSirCSSRepoIndexDiskCapacityInfo	379
falcSirCSSRepositoryUsage - falcSirCSSRepoDataDiskCapacityInfo	379
falcSirCSSRepositoryUsage - falcSirCSSRepoFolderDiskCapacityInfo	379
Common MIBs	380
ServiceEntry	380
falcServerInfo	380
falcEvents	380
CommonMIBs-Traps	381

Troubleshooting

Product registration	382
General console operations	382
Physical resources	385
Logical resources	386
Replication	389
Replication process	389
Deduplication	391
System event messages	392
Take an X-ray of your system for technical support	392
Error codes	394

Best Practices

Deduplication repository sizing	443
Index/folder disk sizing	443
CPU cores	443
Deduplication system sizing	444
Backup cache sizing	444
Memory sizing	445

Appendix

Port usage	446
IP address and netmask update	448
Storage LUN migration	450
FIPS security	452
Shared library environment variable	453
Linux auditing	454
Block device support	455
Configuration	455

Index.	456
-----------------------	------------

Introduction

FalconStor Virtual Tape Library (VTL) is an optimized backup and deduplication solution that offers high-speed backup/restore, global data deduplication, enterprise-wide replication, and tape integration, without requiring changes to the existing environment.

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.

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.

For additional data protection, the data on virtual tapes can be migrated to object storage for long-term data archiving.

VTL components

There are several components:

- VTL servers - Appliance providing VTL and deduplication functionality.
- VTL and deduplication storage - Physical devices used for tapes, database, and deduplicated data.
- FalconStor Management Console - The graphical administration tool, installed on a separate workstation, lets you configure and manage VTL, deduplication, and VTL servers.
- Clients - The backup application servers that use VTL via iSCSI or Fibre Channel.

Terminology

Clients	Backup application servers that are assigned to virtual tape libraries and drives for backup.
Deduplication	Deduplication is a process that frees considerable space by passing only single instances of unique data to the repository. The original virtual tape is replaced with a virtual index tape (VIT) that contains pointers to the data in the repository.
Replication	Replication is a process that copies data from one server to a target server for disaster recovery purposes.
Storage pool	A group of one or more physical devices.
Storage web link	External web links to manage storage adapters on a server.
Vault	The vault is comparable to the I/E slots in a physical tape library and is a storage area for tapes that are not inside a virtual tape library.
Virtual tape library	Virtual tape libraries emulate physical tape libraries and are used for backup by third-party tape or disk backup software, database backup utilities, and archiving applications.
WORM	A write-once-read-many (WORM) property for tapes in a library with IBM or HP drives that support ULTRIUM5 media type and above. A WORM tape cannot be overwritten by backup software. It allows non-rewritable and non-erasable data to be written, providing extra data security by prohibiting accidental data erasure. Since tapes are written once, they cannot be altered or overwritten by some virus/ ransomware/other malicious software. A WORM tape can only be imported to a WORM tape. Importing a WORM tape requires a virtual WORM tape. WORM tapes are not supported with Veritas NetBackup WENCR media (WORM media on which NetBackup encrypts data).

Additional resources

You can download software builds, patches, and other documentation related to your FalconStor product from the FalconStor support community sites at support.falconstor.com (account required).

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.

VTL supports a broad range of server hardware, tape devices, and backup software. The Certification Matrix at www.falconstor.com identifies all certified hardware and software.

If you need technical support, create a support ticket on a FalconStor support community site.

Getting Started

Once you configure your virtual appliance(s), you should launch the console and configure them using the configuration wizard.

Run the FalconStor Management Console

The FalconStor Management Console is the graphical administration tool that enables you to manage your VTL 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.

To install the FalconStor Management Console software, go to a FalconStor support community site.

You can install the management console onto any Windows machine. Note that you must be a Power User or Administrator.

Afterward, launch the console from your desktop or by selecting *Start --> Programs --> FalconStor*. Select the version of the VTL console that corresponds to your installation and then click *VTL Console*.

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.

Note that depending upon the settings that were selected when your system was installed, you may not be allowed to have multiple concurrent console sessions with the same user login.


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 FalconStor appliances, the default hostname has the format *FSxxxxx*, 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 FalconStor, 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. Refer to [“VTL Tape Configuration”](#) for more information.

VTL Tape Configuration

This section describes how to configure a VTL server for tape backups.

Prepare a VTL server via the configuration wizard

Note: If you are using VTL in a Fibre Channel environment, refer to the [“Fibre Channel Configuration”](#) 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.

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

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.

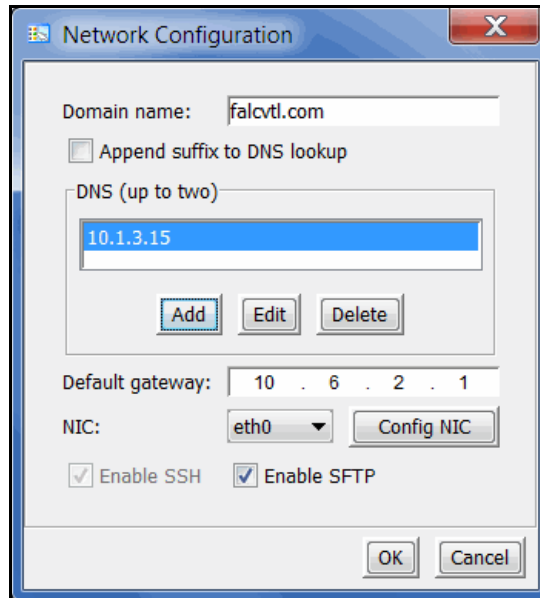
Note: After completing the configuration wizard, if you need to add license keys, you can right-click your server and select *License*.

Set up network

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

Note: We recommend that you finalize all IP addresses before replication and deduplication, are configured. If you need to change an IP address afterward, refer to [‘IP address and netmask update’](#).

1. Enter information about your network configuration.



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).



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

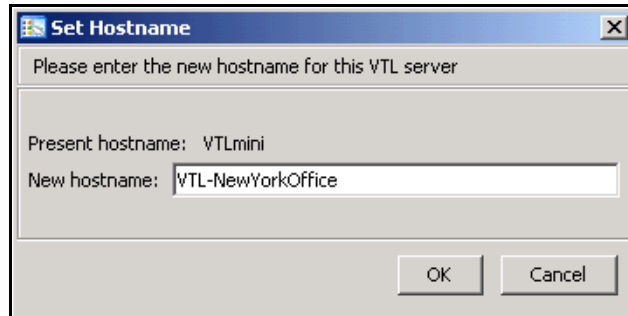
MTU - Maximum transfer unit (MTU) of each IP packet. If your network supports it, set this value to 9000 for jumbo frames. Check network settings for speed and Maximum Transmission Unit (MTU) values on network devices and routers. Most network devices have an MTU value of 1500 bytes, but tunnel interfaces, used for example on IBM power systems, have an MTU of 1476 bytes since they use some bytes for IP headers.

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.



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

(Required for Fibre Channel tape backups) 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

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

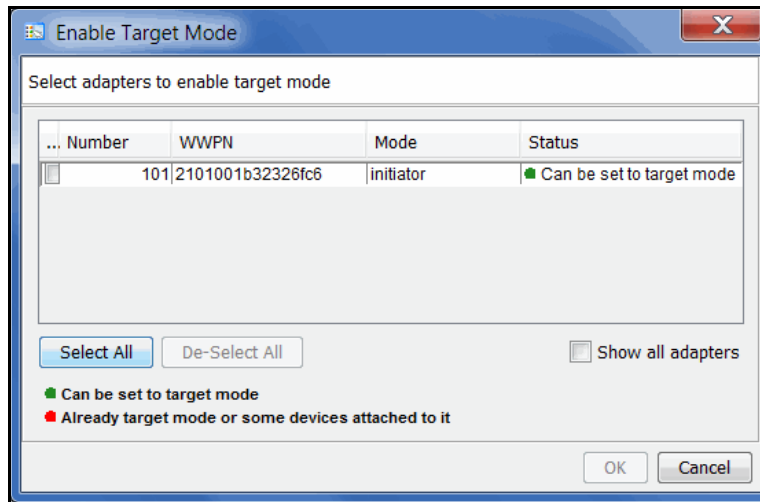
If you haven't already done so, you will need to switch any initiator zoned with a backup application server to target mode so that the backup application server can see the VTL server.

You will 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).

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



(Multi-ID HBAs) Click *Ok*.



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, 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 - Includes deduplication index, folder, and data disks, as well as the associated mirror devices.
- Tapes - Used only for tape storage.

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. 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

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.

3. 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 server and select *Options --> Enable Virtual Tape Encryption*.

Create virtual libraries

Select the tape library that you are emulating.

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.

Refer to '[Create virtual tape libraries](#)' for detailed information about creating virtual tape libraries.

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 FC tape backups) This step allows you to select the clients (backup application 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

If you added backup application servers, do the following:

1. Select a backup application 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

Data deduplication eliminates redundant data for tapes, minimizing replication time and storage requirements.

You will create a deduplication repository (index, folder, and data disks). The deduplication repository *data* storage can be configured on either object storage or SCSI devices while the *index* and *folder* disks will be configured on SCSI devices.

Create Deduplication Policy

If you 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*.

Add tape backup clients

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

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.

Special characters such as \, /, *, ?, ", <, >, |, %, \$, or spaces are not supported.

When replication is configured, a client with the name of the replica is created. You cannot add a client with the same name as that of a replica.
3. If you started the wizard from the configuration wizard, select the protocol being used by the client.

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 client 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 clients

In order for Fibre Channel backup application 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 virtual tape libraries from your backup application server”](#) for more information.

Additional information about steps 1-3 can be found in the [“Fibre Channel Configuration”](#) chapter.

iSCSI clients

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 virtual tape libraries from your backup application server’](#) for more information.

Additional information about steps 1-5 can be found in the [‘iSCSI Configuration’](#) chapter.

Discover virtual tape libraries from your backup application server

To enable your backup application 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.

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.

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 FalconStor 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.

Console

The FalconStor Management Console allows you to manage your VTL, servers.

Console user interface

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

The screenshot displays the FalconStor Virtual Tape Library Console in Configuration Mode. The interface includes a menu bar (File, Edit, View, Tools, Help), a toolbar, and the FalconStor logo. A tree view on the left shows the hierarchy: Servers > OBD190 > System > Network Interface. The main pane shows a configuration table for the selected server.



Name	Value
Server Name	OBD190
Login Machine Name	10.8.25.190
Login User Name	root
Processor 1 - 16	Intel(R) Xeon(R) CPU E5520 @ 2.27GHz 2262 MHz
Network Interface	eth0 - mtu 1500 inet 10.8.25.190 mac 00:26:b9:59:bb:61
Network Interface	eth1 - mtu 1500 inet 1.0.0.85 mac 00:26:b9:59:bb:63
Network Interface	eth2 - mtu 1500 inet 1.0.106.1 mac 00:26:b9:59:bb:65
Network Interface	eth3 - mtu 1500 inet 1.0.47.1 mac 00:26:b9:59:bb:67
Network Interface	eth4 - mtu 1500 inet 1.0.0.78 mac 00:1b:21:a2:38:48
Network Interface	eth5 - mtu 1500 inet 1.0.29.2 mac 00:1b:21:a2:38:49
Admin Mode	Read/Write
Server Status	Online
System Up Time	2 days 13 hours 30 minutes 25 seconds
VTL Up Time	2 days 13 hours 14 minutes 57 seconds
Fibre Channel WWPN	21-00-00-0d-77-97-7e-64 [target]
Fibre Channel WWPN	21-01-00-1b-32-b7-7e-64 [initiator]
Fibre Channel WWPN	21-00-00-0d-77-31-cb-48 [target]
Fibre Channel WWPN	21-00-00-0d-77-31-cb-49 [target]
Virtual Tape Encryption	Disabled
Deduplication Repository Encryption	Disabled

Disk Space Usage

Total Disk Capacity:	2.66 TB	VTL Reserved:	(976.55 GB) 35.92%
Total Space Reserved:	2.66 TB	SIR Reserved:	(1.35 TB) 50.98%
Total Space Unconfigured:	0.00 KB	Others Reserved:	(356.32 GB) 13.11%
		Unconfigured:	(0.00 KB) 0.00%

Refresh

12/17/2018 10:11:51 [OBD190] Logged in Server:OBD190 12:10 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 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



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 email alerts, manage software licenses, perform system maintenance, 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*, *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, 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 replication errors and import/export job status.
- *SIR Replication* - This tab only appears if replication has been configured. The tab provides the names of *primary* (source) and *replica* (target) pairs.

Multi-Node Group object



If you have configured multi-node groups, the group object contains the 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, 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 a server. Click an object to display related information in the right-hand pane.

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.



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



Tape Import/Export Queue - Lists import, migration to object storage, and recovery from object storage 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.

Status object



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

- *Dashboard Summary* - This object can have up to three tabs.
 - *VTL Space Usage* - Displays capacity information for VTL storage.
 - *VTL Performance* - Displays performance statistics for VTL.
 - *Deduplication Repository* - Displays 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 tape properties for the library (Auto Replication/Auto Migration to Object Storage, Tape Capacity on Demand, maximum tape capacity, barcode/slot information)
- Assign clients
- Auto load/unload tapes
- Change firmware
- 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 (barcode, tape capacity on demand, write protection, and Auto Replication/Auto Migration to Object Storage)
- 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 a virtual tape 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 (comparable to the I/E slots in a physical tape library). 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, migrated to

object storage, converted to a stub tape, or moved to a virtual library or standalone drive. Stub tapes in the vault can be reconstructed from object storage. 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.

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.

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.

Reports object



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

- VTL performance
- Physical resources - allocation and configuration
- Disk space - allocation and usage
- Storage pool configuration
- LUN usage
- Fibre Channel adapters configuration
- Status for deduplication and deduplication replication
- Virtual tape/library information
- Virtual library and drive assignment
- Import job status
- Object storage migration job status
- Deduplication policy status, tape activity and tape usage
- Deduplication repository usage, performance, reclamation

Repositories object



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

- Database - The database stores information about libraries, clients, and replication setup.
- Configuration Repository - The configuration repository stores information about the deduplication configuration.
- Index/folder disks - These resources store pointers to unique deduplicated data and information related to deduplication sessions.
- Deduplication data disks - 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*. For each iSCSI client, you can create/assign/unassign 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.

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.

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. Hard disks are used for creating virtual tape libraries/drives, virtual tapes, deduplication repository, and configuration repository.

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

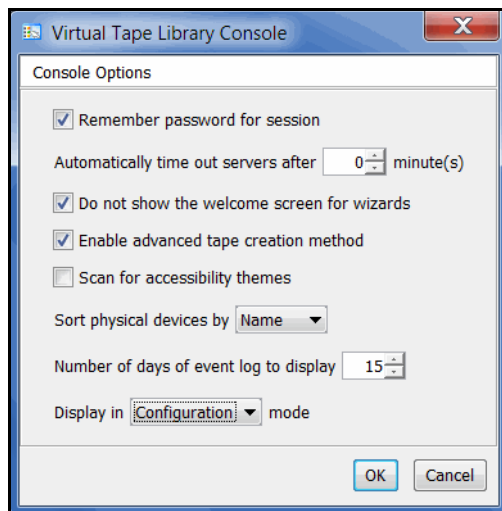


If you have a multi-node group configured, you will see a *Group Reports* object. This object provides reports that can be generated for all servers in a group. This includes standard reports that are generated on each server in the group and contain data specific to that server. You can also run a consolidated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

Set console options

To set options for the console:

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



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 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.


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*.



2. Click *Add* and add a VLAN.



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 - Maximum transfer unit (MTU) of each IP packet. If your network supports it, set this value to 9000 for jumbo frames. Check network settings for speed and Maximum Transmission Unit (MTU) values on network devices and routers.

Most network devices have an MTU value of 1500 bytes, but tunnel interfaces, used for example on IBM Power Systems, have an MTU of 1476 bytes since they use some bytes for IP headers.

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:

- We recommend that you finalize all IP addresses before replication and deduplication are configured. If you need to change an IP address afterward, refer to [‘IP address and netmask update’](#).
- You cannot change the network configuration of a server that is in a multi-node group.

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 replication configured for deduplicated tapes
- The server has deduplication enabled

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.

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 bring the server to its lowest state (no services running). The power will remain on.

Add and register licenses

To license Virtual Tape Library and its options, make sure you have obtained your keycode(s) from FalconStor 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:

The screenshot shows a Windows-style dialog box titled "Register License" with a close button (X) in the top right corner. The dialog has a tab labeled "Offline Registration". Inside the dialog, the text reads: "Follow the steps to finish offline registration:" followed by four numbered steps:

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: ... Save
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: ... Send


At the bottom of the dialog, there are three buttons: "Back", "Finish", and "Cancel".

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 *Save*.




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

3. Copy the saved file to a computer with an Internet connection and email it to the 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.

4. When you receive a reply, save the attached signature file (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 name of the file.
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.



7. Click the *Finish* button.

Notes:

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

Monitor space usage


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

The top section displays information about VTL resources.



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.

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
- Encryption keys - for virtual tape encryption

With the *Strong Passwords* option, passwords must contain a minimum of 14 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.

Any time a password is changed, at least five characters must be changed between the old and new password.

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.

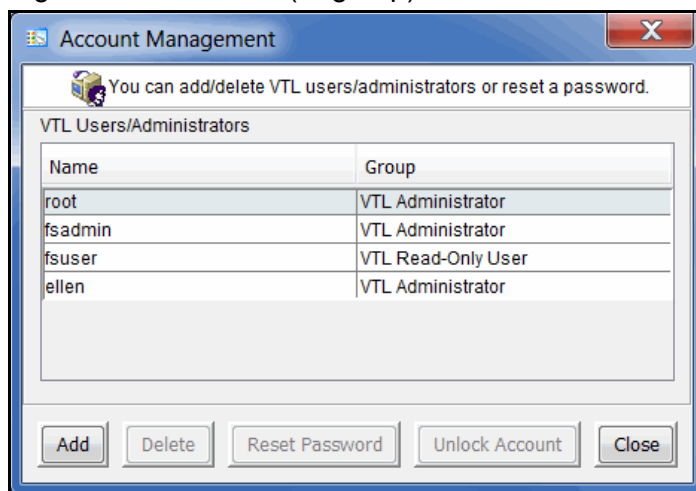
The *Strong Passwords* option includes a locking mechanism for user accounts that repeatedly fail to log in because of an incorrect password. The default number of times the wrong password can be entered is three. If a user is locked out, a system administrator can unlock the account or the user can 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.

Add or modify users

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



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 FalconStor 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	<p>I: This is an informational message. No action is required.</p> <p>W: This is a warning message that states that something occurred that may require maintenance or corrective action, although the system is still operational.</p> <p>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.</p> <p>C: These are critical errors that stop the system from operating properly.</p>
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:

- Hardware appliance errors
- Replication errors
- Import job status
- Migration to object storage and reconstruction from object storage job status

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

Select *Dashboard Summary --> VTL Performance* to display information about VTL performance.

Each section displays read and write throughput for all related resources. Performance statistics are acquired from all adapters 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.

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.

To hide a read or write performance chart, click the appropriate checkbox.

Server properties

To set properties for a specific server or group, right-click the server/group and select *Properties*.

Activity Database Maintenance settings

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.

SNMP Maintenance settings

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:



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:

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.

Performance settings

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: $(2000 / 1024) * 8 = 15.6\text{Mb/s}$

Depending upon the settings that were selected when your system was installed, you may see an option to enable a preferred management IP address.

If this server is a replication target server, you can use two different IP addresses to isolate server management from replication traffic of non-deduplicated virtual tapes. Each IP address must be on its own subnet.

Once enabled, enter a preferred management IP address for the console. The IP address that is used for replication traffic is specified when you configure replication for a virtual tape.

Auto Save Config settings

Depending upon the settings that were selected when your system was installed, the *Auto Save Config* tab allows you to 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.

Storage Monitoring settings

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.

Location settings

Depending upon the settings that were selected when your system was installed, the *Location settings* tab allows you to 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.

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 To apply a patch:

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 FalconStor support community sites. 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.

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.
4. Select the patch and click *OK*.
5. Check the Event Log to confirm that the patch uninstalled successfully.

Console patches

You need an account with administrator privileges to install the 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.

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.

Note: Before creating a mirror for the tape database, all tape activity should be reduced to limit the chance that expansion might occur on the virtualized device created 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 (under the *Repositories* object).

- *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 (under the *SCSI Devices* object).

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.

Manually save/restore the Virtual Tape Library configuration

VTL includes a console option and a command line utility (*vtlrecover*) that enables you to protect your databases and recover your 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

- All appliance hardware, including FC HBAs and network adapters, storage, and connectivity must be intact and functioning properly.
- Patches are not backed up or restored. Patches need to be saved and restored prior to running the restore process.
- Your database mirroring configuration will not be restored. If you had database mirroring configured before recovery, you will need to reconfigure it after recovery.
- The tape import queue and scheduled reporting will not be saved and cannot be restored.

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 save your configuration after you make any major Virtual Tape Library configuration changes.

- Console
1. Highlight the server in the tree.
 2. Select *Tools* menu --> *Save Configuration*.
 3. Specify a filename and location outside the appliance for the saved configuration file.
The output file includes all configuration information needed for recovery.

- Command line
1. Run the following command: `$ISHOME/bin/vtlrecover save archive.tar`
The output file includes all configuration information needed for recovery.
 2. Copy the output file to a safe remote location outside the appliance.

Restore configuration

Restoring a configuration is for disaster recovery purposes and should not be used in day-to-day operations. Changes made since the configuration was last saved will not be included in this restored configuration.

Use the following procedure:

1. If the Linux operating system is lost, reinstall Linux using an approved procedure.

For a FalconStor appliance, use the FalconStor 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. Restore the configuration:

From the console, highlight the server in the tree, select *Tools* menu --> *Restore Configuration* and locate the saved configuration file.

From the command line, run the following command:

```
$ISHOME/bin/vtlrecover restore archive.tar.bz2
```

NOTE: Depending on how many tapes are present on a 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.

Multi-Node Groups

If you have multiple VTL 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 servers. 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.

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, 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*.



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 servers per 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.
- 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, and storage monitoring triggers, 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.

Your console will now look similar to the following:

This is a group.

These servers are members of this group.

Server	Status	Members reference
OBD190	Logged In	OBD190 10.8.25.190 VTL01 10.8.25.130
VTL01	Logged In	OBD190 10.8.25.190 VTL01 10.8.25.130

12/17/2018 13:36:00 [OBD190] Device 2081 (ID: ADIC-Scalar 100-02081) tape library properties was changed. Server:OBD190 12:10 PM

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.

Physical Resources

Physical resources include:

- Storage HBAs - Fibre Channel and SCSI physical adapters.
- Storage devices - Physical disks directly accessing an existing local disk or partition.
- Storage pools - Groups of one or more virtualized physical devices.

There is one object for each type of physical resource. Each object displays the resources that have been created on the selected server.

Physical Resources object





The right pane of the *Physical Resources* object displays a tab for each type of physical resource.

From the *Physical Resources* object, you can prepare and rescan physical devices.

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

Physical resources icons

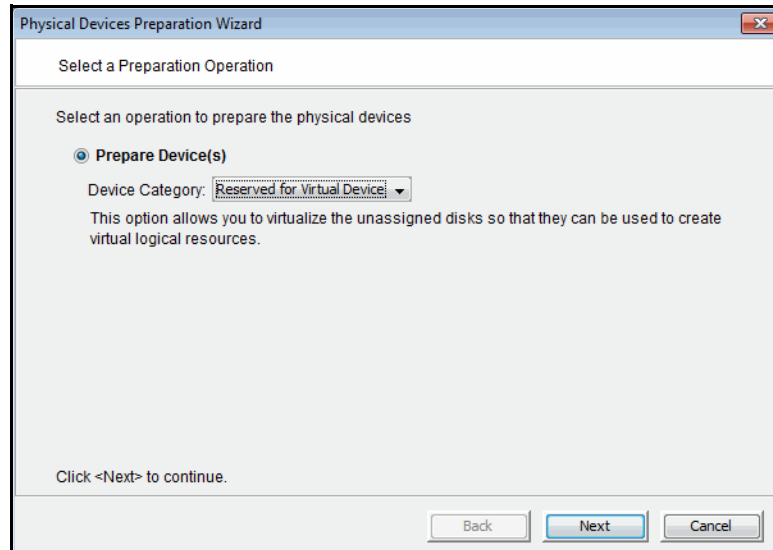
The following table describes the icons that are used for physical resources:

Icon	Description
	The  icon on a disk indicates that this disk has been virtualized.
	The  icon on a disk indicates that this is shared storage and is being used by another server. The <i>Owner</i> field lists the other server.

Prepare physical storage devices

This procedure virtualizes physical disks so that you can create logical resources for storage devices. It can also be used to unassign physical disks.

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



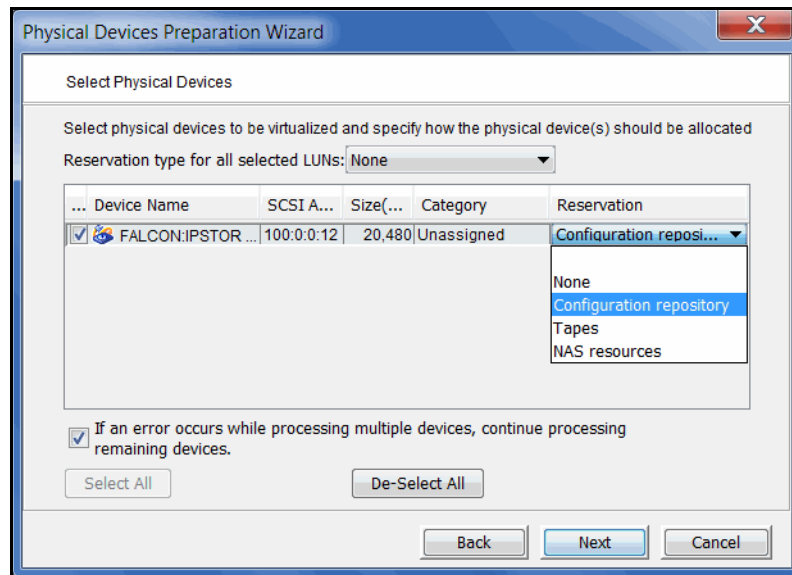
The Device Category *Reserved for Virtual Device* is selected by default. This is the action you want if you want to virtualize physical disks.

The Device Category *Unassigned* is used to remove “ownership” of devices that have been virtualized for your server. 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.

To determine how a device is being used, expand the *Physical Resources* object until you see the device and then select it. On the *Layout* tab, review the *Type* and *Used by Virtual Device* columns to identify the entity that “owns” the device.

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

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



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*.

- 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 - Includes deduplication index, folder, and data disks, as well as the associated mirror devices.
 - Tapes - Used only for tape storage.
- Confirm that all selections are correct and click *Finish*.
- Type *Yes* at the warning message and then click *OK*.

Rescan physical 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*.



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.

Storage HBAs object

Storage HBAs are the Fibre Channel and SCSI physical adapters used by your server.

When you select a Fibre Channel HBA from the list displayed in the right pane, you will see detailed information in the lower area of the pane. Refer to [“Verify your hardware configuration”](#) for more information.

You can do the following from the *Storage HBAs* object:

- Rescan adapters. You can also select an HBA from the list displayed in the right pane to rescan the specific adapter. Refer to [“Rescan physical devices”](#) for more information.
- Rename adapter information.
- Refresh the SNS table for a Fibre Channel HBA.
- Enable target mode for a Fibre Channel HBA. Refer to [“Set QLogic ports to target mode”](#) for more information.

Storage Devices object

Storage devices include hard disks, tape drives, and tape libraries. Hard disks are used for creating virtual tape libraries/drives, virtual tapes, deduplication repository, and configuration repository.

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.

To see existing devices, expand the *Storage Devices* object and then select *Fibre Channel Devices* or *SCSI Devices*. The icon displayed next to each device describes its purpose or status (refer to “[Physical resources icons](#)”).

When you select a device in the right pane, details are displayed in the lower area of the 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](#)’).

You can do the following from the *Storage Devices* object:

- [Prepare physical storage devices](#). You can also virtualize/unassign physical devices from the *Fibre Channel Devices* or *SCSI Devices* object. To prepare a single physical disk, right-click the device and select *Enlist* (or *Discharge*). Refer to “[Prepare physical storage devices](#)” for more information.
- Change the LUN reservation for physical devices. Refer to “[Change the LUN reservation](#)” for more information.
- Test the performance of physical devices. Refer to “[Test physical device throughput](#)” for more information.
- Set autopathing. Refer to “[Set autopathing](#)” for more information.
- Save/restore the system preferred path. Refer to “[Set autopathing](#)” for more information.
- Rename a device.
- View disk properties.

Filter storage devices

All devices are displayed by default. To filter the list, click *Filter* to display the *Filter Options* dialog.

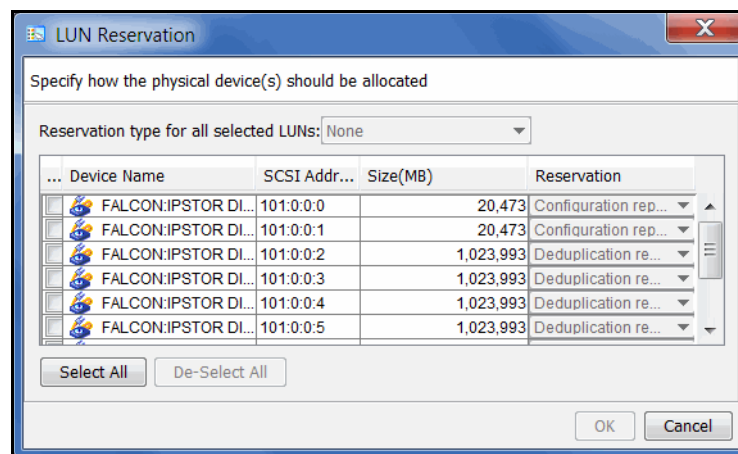
1. Select the checkbox(es) for the types of devices you want to see in the list.
2. Click *Search* to update the list.
3. To change filter selections, de-select one or more selections or click *Clear* to de-select current choices, then click *Search*.

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

Change the 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* and select *LUN Reservation* to display a list of all virtualized devices.



You can also display this dialog from the *Fibre Channel Devices* or *SCSI Devices* object to display only those devices. You can also right-click on a specific device in the right pane to change the LUN reservation for that device.

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 check the performance of your physical devices to see:

- 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.
To test multiple devices, right-click the *Storage Devices* object.
2. Select *Test* from the menu.

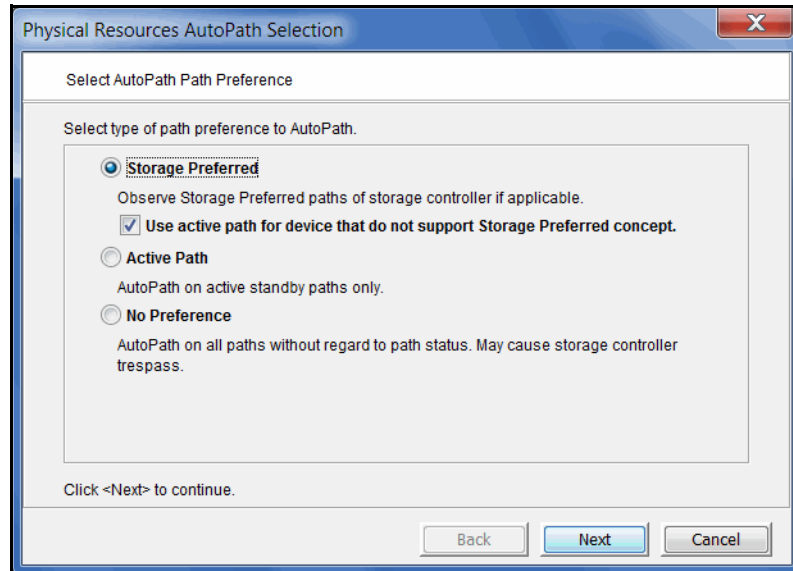
The system will test each device and then display the throughput results for different disk spindle locations (beginning, middle, and end) on a new *Throughput* tab. If you tested multiple devices from the *Storage Devices* object, aggregate results will be shown in a dialog and the *Throughput* tab 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* and select *Autopath*.
2. Select the autopath method you want to use.



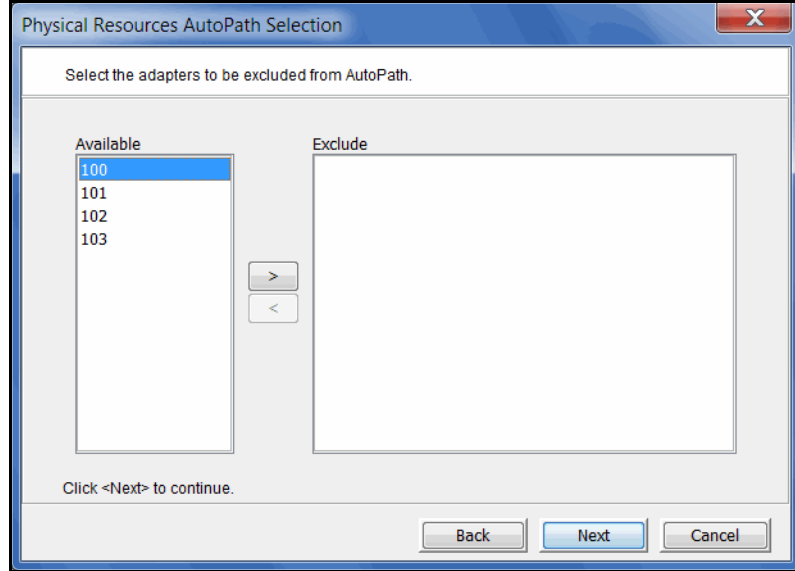
Storage Preferred - (Default) Detects and follows the preferred paths for each LUN as they are set by the storage controller on supported systems and uses them as preferred paths by VTL. If there is no preferred storage path, select the sub-option and VTL will use the Active Path.

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.

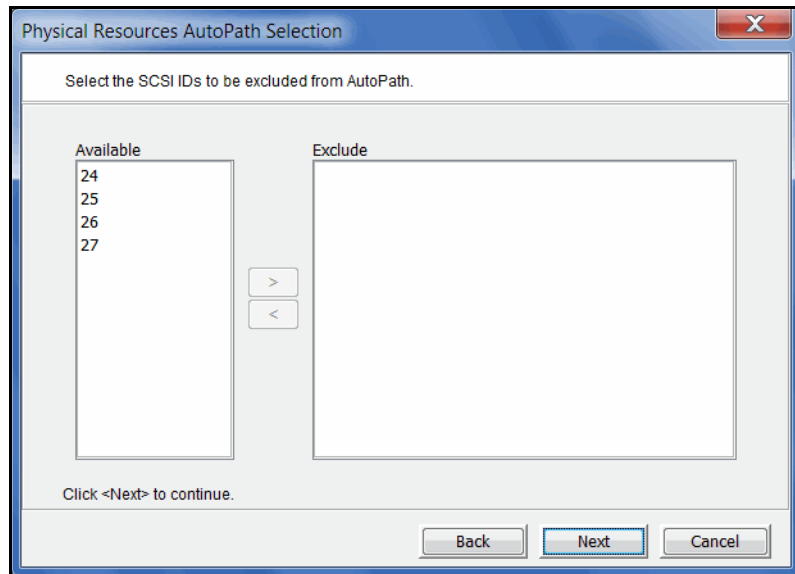
No Preference - Do not use the other methods. VTL uses all available independent paths. This can cause paths to trespass.

3. Select any adapters that should be excluded.

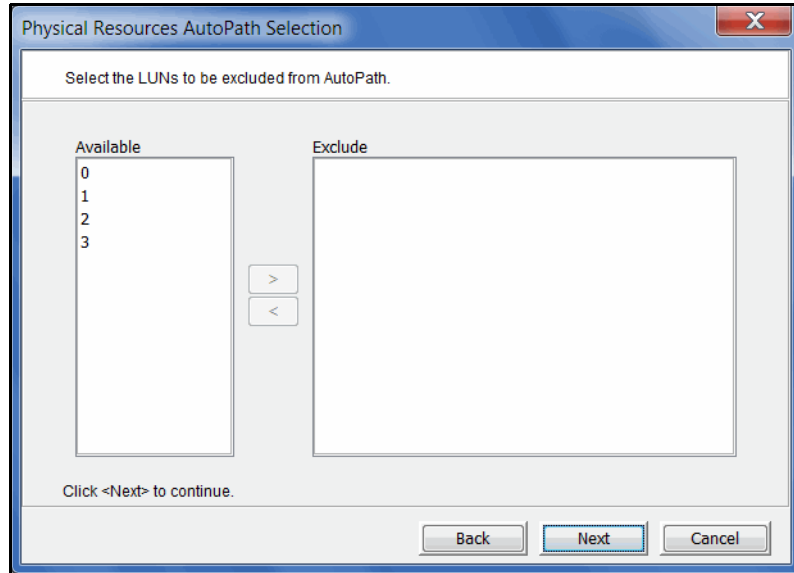


You may need to do this if you have a specific path that you want to maintain. In that case, you would exclude the adapter, SCSI ID, and LUN for that path.

4. Select any SCSI IDs that should be excluded.



5. Select any LUNs that should be excluded.



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* 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* 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.

Storage Pools object

A storage pool is a container that can consist of zero, one, or more physical devices. These physical devices can use various interface protocols (such as SCSI or Fibre Channel).

Any physical device that has been reserved for tapes can be added to a storage pool. Unprepared devices can be added as well.

Each storage pool can be assigned to one or more VTL administrators.

Storage pools work with all automatic allocation mechanisms in VTL. Tape Capacity on Demand automatically allocates storage space from a specific pool when storage is needed.

As your storage needs grow, you can easily extend your storage capacity by adding more devices to a pool and then creating more logical resources or allocating more space to your existing resources. The additional space is immediately and seamlessly available.

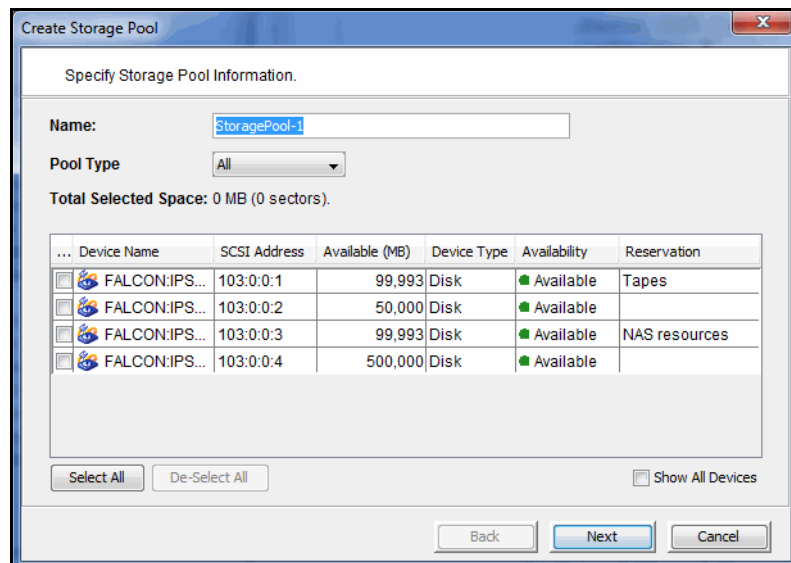
Storage pool access rights

Only root users can create, delete, and manage (add/remove devices) storage pools. When a storage pool is created, the root user can assign administrators to access it.

VTL administrators assigned to a storage pool can create resources, such as virtual tapes and replica resources, from the pool.

Create a storage pool

1. Expand *Physical Resources* --> *Storage Device*, right-click on *Storage Pools* and select *New*.



All eligible devices are listed.

2. Enter a storage pool name (maximum 64 characters).

3. Specify *Tapes* as the pool type.

Eligible devices for the pool type are listed.

4. Select the physical devices to include.

If a device is unassigned, you will be prompted to prepare it.

You can also create an empty storage pool, one without any devices, for future use.

Note: If you select all of the disks that are reserved for a specific type to a pool, it is important to use this storage pool when creating applicable resources. For example, if all physical devices reserved for tapes are added to a *Tapes* storage pool, all tapes should be created using this storage pool. Otherwise, a thin tape (a tape with Tape Capacity on Demand enabled) created outside a storage pool will have no devices on which to expand when expansion is needed.

5. Select the VTL administrators that can create resources from this storage pool.
6. Verify all information and then click *Finish* to create the storage pool.

Update storage pool settings

The root user can do the following to an existing storage pool:

- Rename the pool
- Add/delete physical devices
- Change user access
- Delete the pool

To update storage pool settings, right-click on a storage pool and select the appropriate option.

Migrate existing virtual tape libraries to storage pools

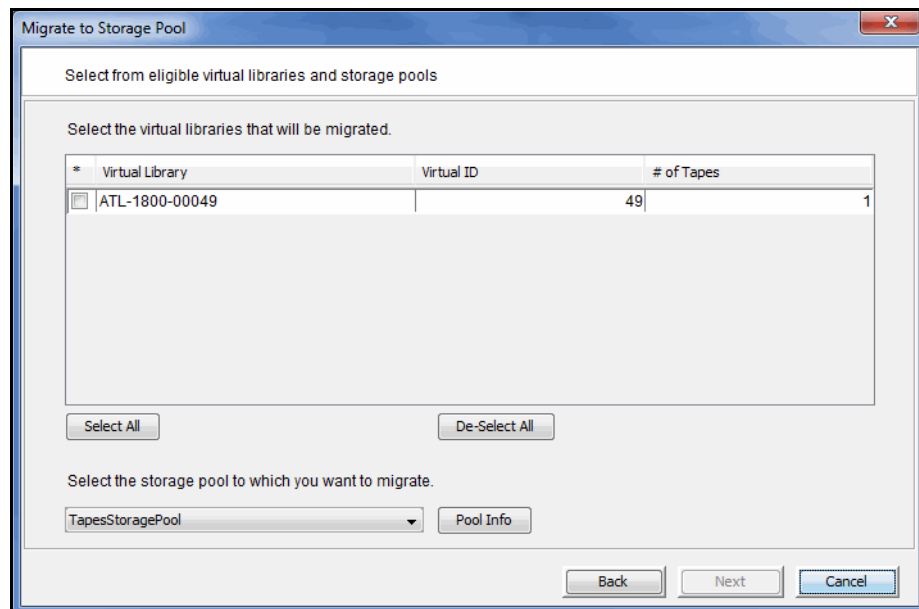
If you have existing virtual tape libraries with virtual tapes that are not in a storage pool, you can migrate the libraries and their virtual tapes to an existing storage pool. This is useful if you have upgraded from an earlier VTL version or if you started defining storage pools after your system was configured.

The migration operation will scan all tapes from the selected virtual libraries and assign the disks used by those tapes to the selected storage pool. This will cause unpooled storage capacity to decrease, which might be needed for operations such as tape creation and expansion for any virtual tape not using a storage pool.

Notes:

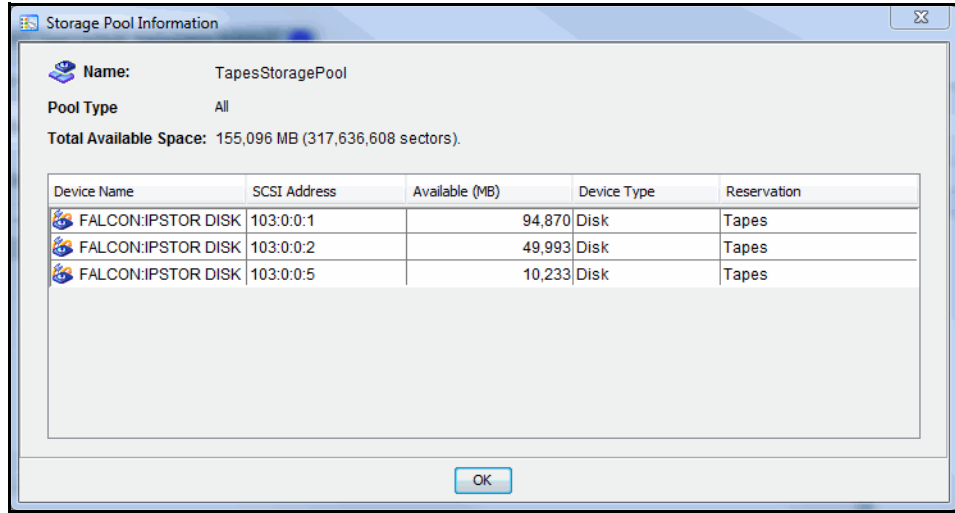
- There cannot be any virtual tapes in drives during migration.
- There should be no IO activity at the time of migration.

1. Right-click the *Virtual Tape Libraries* object and select *Migrate to Storage Pool*.
2. Select the libraries to be migrated and the storage pool to which you want to migrate.
3. Select one or more virtual tape libraries and select a storage pool.



Only one storage pool can be selected for the target.

You can click the *Pool Info* button to see details about the selected pool. Here you will see the physical devices ins the pool and the available space.









4. Confirm all information and click *Finish*.



Virtual Tape Libraries, Tape Drives, and Tapes

Virtual tape libraries, tape drives, and tapes are all managed under the *Virtual Tape Library System* object.

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

The following table describes the icons that are used for virtual tape libraries, virtual tape drives, and virtual tapes:

Icon	Description
	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.
	The C icon on a virtual tape drive indicates that the drive has compression enabled.
	The E icon on a virtual tape indicates that the tape is encrypted. The icon appears in red if encryption is not activated.
	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 <i>General</i> tab for current information about the tape.
	Used with Tape Migration to Object Storage, the O icon indicates that this is a stub tape.
	The W icon on a virtual tape indicates that the tape is a write once, read many (WORM) tape.

Icon	Description
	<p>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)</p> <ul style="list-style-type: none"> • 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.
	<p>An R icon that is divided diagonally will appear for a tape in a deduplication policy with any type of Advanced Replication enabled.</p> <p>The upper left section represents replication to Replication Target 1; the lower-right section represents replication to Replication Target 2.</p> <p>Green, yellow, and red indicators apply as described above.</p>

Create virtual tape libraries

You can create virtual tape libraries that emulate physical tape libraries. If you have a preconfigured VTL appliance with a default virtual tape library that does not meet your needs, 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*.

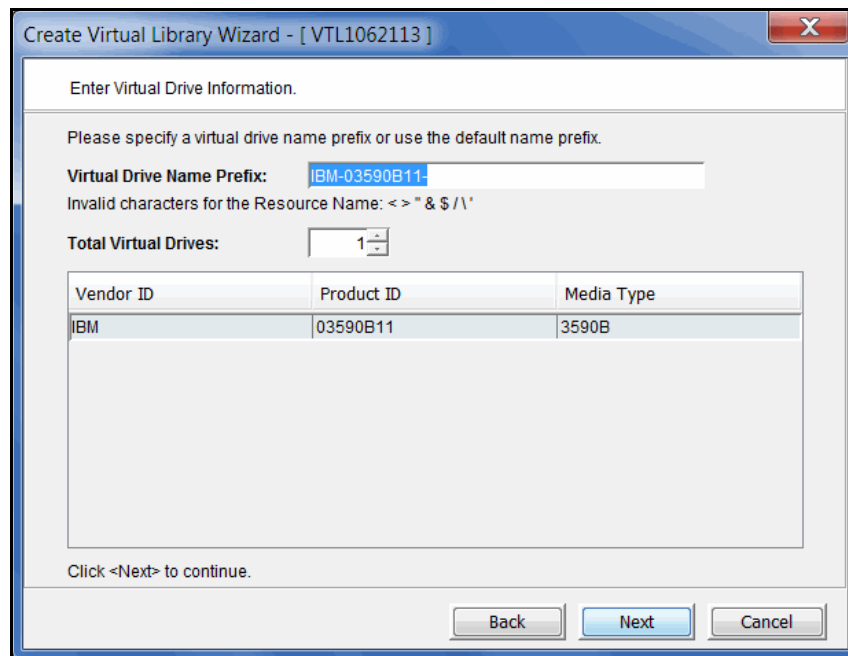
Note: If you have recently added additional storage to your VTL system, before you can use it to create a virtual tape library, you must reserve it for virtual use. To do this: Right-click *Physical Resources* and select *Prepare Devices*. Set hard drives to *Reserved for Virtual Device*.

1. Select the physical tape library that you are emulating.

Special characters such as \, /, *, ?, ", <, >, |, %, \$, or spaces are not supported.

For IBM i clients, the virtual tape library type must be FalconStor FALCON TS3500L32 (03584L32) or FALCON TS3500L32 (03584L32) and the media type must be ULTRIUM3 (LT03) or newer.

2. Enter information about the tape drives in your library.



Enter Virtual Drive Information.

Please specify a virtual drive name prefix or use the default name prefix.

Virtual Drive Name Prefix:

Invalid characters for the Resource Name: < > * & \$ / \'

Total Virtual Drives:

Vendor ID	Product ID	Media Type
IBM	03590B11	3590B

Click <Next> to continue.

Back Next Cancel

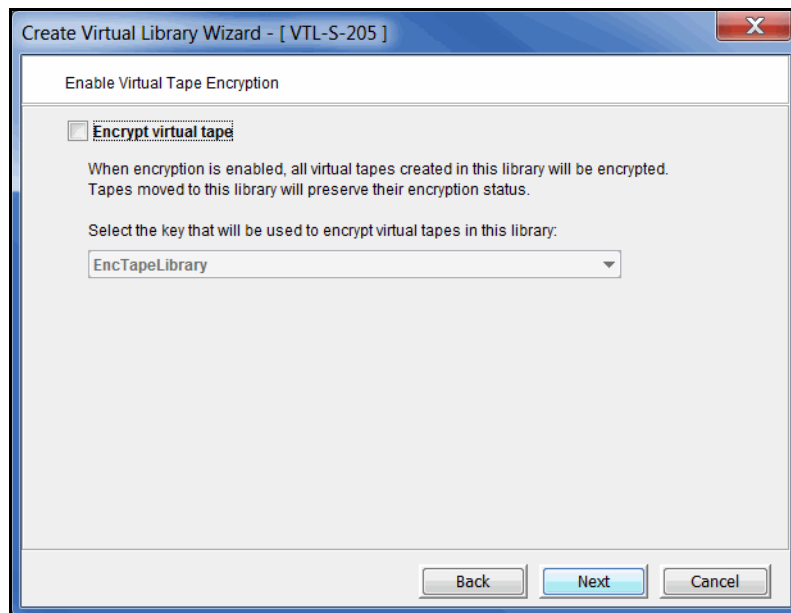
Special characters such as \, /, *, ?, ", <, >, |, %, \$, or spaces are not supported.

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)*.

3. Indicate if you want to use encryption for this virtual tape library.



When encryption is enabled, each new tape that is created in the library is encrypted with the selected key. Each encrypted tape always retains its key, even if it is moved to another library.

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

If encryption is ever disabled for this library, tapes created afterward will not be encrypted. Therefore, each library can have both encrypted and unencrypted tapes. 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.

4. Determine if you want to use Auto Replication or Automatic Migration for this virtual library.

Auto Replication and Auto Migration to Object Storage are mutually exclusive.

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. Refer to [“Auto Replication”](#) for more information.

Auto Migration exports virtual tapes to object storage when they are ejected by backup software. If selected, determine whether you want the virtual tape moved (default) or copied after the data is migrated to object storage.

- When *Move* is selected, the source virtual tape will be converted to a stub tape and all related disk space on the VTL cache will be freed once migration to object storage completes successfully. Specify the number of days (up to 90, where 0 means immediate) to keep virtual tapes before they get converted to stub tapes.
- When *Copy* is selected, the source tape will stay in the vault after migration is complete and will not be converted to a stub tape. Virtual tapes can still be manually converted to stub tapes.

In order to use *Auto Migration*, you must have a configured object storage account. Refer to [“Tape migration to object storage”](#) for more information.

5. Enter barcode information for the virtual library.

The screenshot shows a dialog box titled "Create Virtual Library Wizard - [H21-38]". The main heading is "Enter Virtual Library Information." Below this, it says "Please enter information for the virtual library." The fields are as follows:

- Name:** IBM-TS3500L32-00075
- Barcode Starts:** 004B00
- Barcode Ends:** 004BZZ
- Barcode Suffix:** Add media type suffix to tape barcode when new tapes are created.
- Number of Slots:** 253
- Auto Loader:**
- Import|Export Slots:** 10

At the bottom, there is a prompt: "Click <Next> to continue." and three buttons: "Back", "Next", and "Cancel".

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.

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

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

Auto Loader - Set the auto-loader for those libraries that support the feature.

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.

6. Enter the guidelines for expanding virtual tape capacity.

The screenshot shows a Windows-style dialog box titled "Create Virtual Library Wizard - [VTL1062113]". The main area is titled "Enter virtual tape properties" and contains the text "Please enter virtual tape properties." Below this, there is a checked checkbox for "Tape Capacity On Demand". Underneath, there are three spinners: "Initial Tape Size" set to 5 GB, "Incremental Size" set to 11 GB, and "Maximum Capacity" set to 680 GB. The "Media Type" is set to "ULTRIUM4". At the bottom, there are "Back", "Next", and "Cancel" buttons, and a prompt "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.

7. Verify all information and then click *Finish* to create the virtual tape library.

You will be prompted to create virtual tapes. Answer *Yes* to continue. Refer to the following section for more information about creating virtual tapes.

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 and specify how many drives you are creating.
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.

4. Verify all information and then click *Finish* to create the virtual tape drive.

To create a tape for this standalone virtual drive, right-click the drive and select *New Tape*. Refer to [“Create virtual tapes”](#) for more information.

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 the virtual tape to be a write once, read many (WORM) tape.

A WORM tape cannot be overwritten by backup software.

Only virtual tapes in a library with IBM or HP drives that support ULTRIUM5 media type and above can be configured as WORM tapes.

WORM tapes are not supported with Veritas NetBackup WENCR media (WORM media on which NetBackup encrypts data).

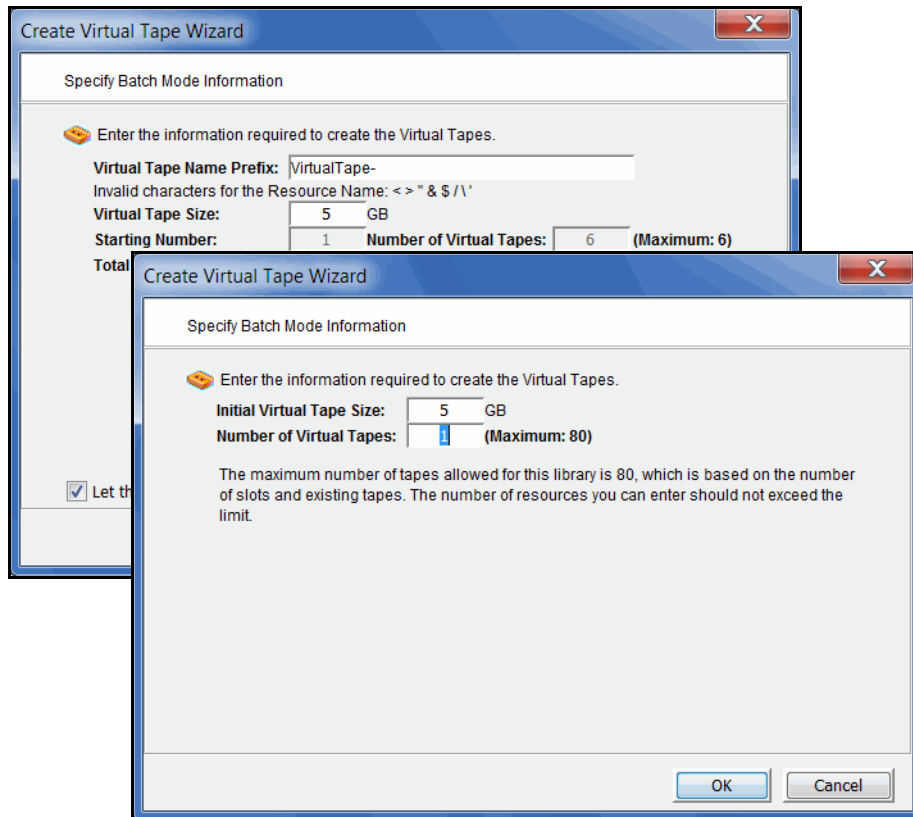
3. (*Custom* virtual tape creation mode) Select the storage pool or physical device(s) from which to create this virtual tape.

The list includes only physical devices that were previously reserved for virtual tapes and storage pools with the *Tapes* or *All* device category.

Storage space is allocated from the local server even if this server is part of a multi-node group.

- 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.

Special characters such as \, /, *, ?, ", <, >, |, %, \$, or spaces are not supported.

If *Tape Capacity on Demand* is enabled for the library, you can specify the tape size; if it is disabled for the library, the tape size is fixed to the maximum capacity.

- 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.

6. 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 FalconStor 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*.



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 Replication/Auto Migration to Object Storage).

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. 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. 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.

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.
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 clients

You can assign a virtual tape library or drive to the target of a backup application server listed in the VTL console under the *Clients* object. The backup application server can then access the assigned virtual tape library/drive(s).

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

There are three ways to assign a library or drive to a client (backup application 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 application 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 application server to assign it to.

Client multipathing

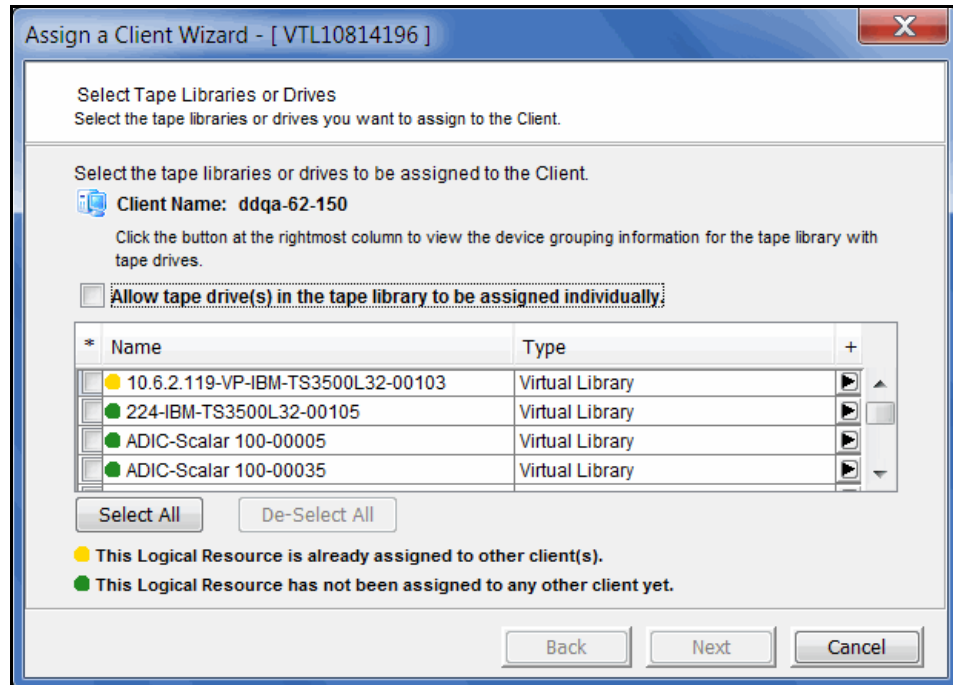
In a multi-path configuration, there can be more than one physical path between a host and an assigned tape library or drive. This requires that the host have multiple initiator ports. For each path, you need to create a client associated with each initiator port and assign the same tape library/drive to those clients representing storage paths.

Ensure that you have a tape device driver that supports multipathing on your host. For example, for a Windows host accessing an IBM library, refer to IBM's [website](#).

Assign a library to a client

If you started from the configuration wizard or a client object, follow these steps to continue:

1. Select a virtual tape library or drive.



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 application 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 application server's operating system to discover the VTL server.

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.

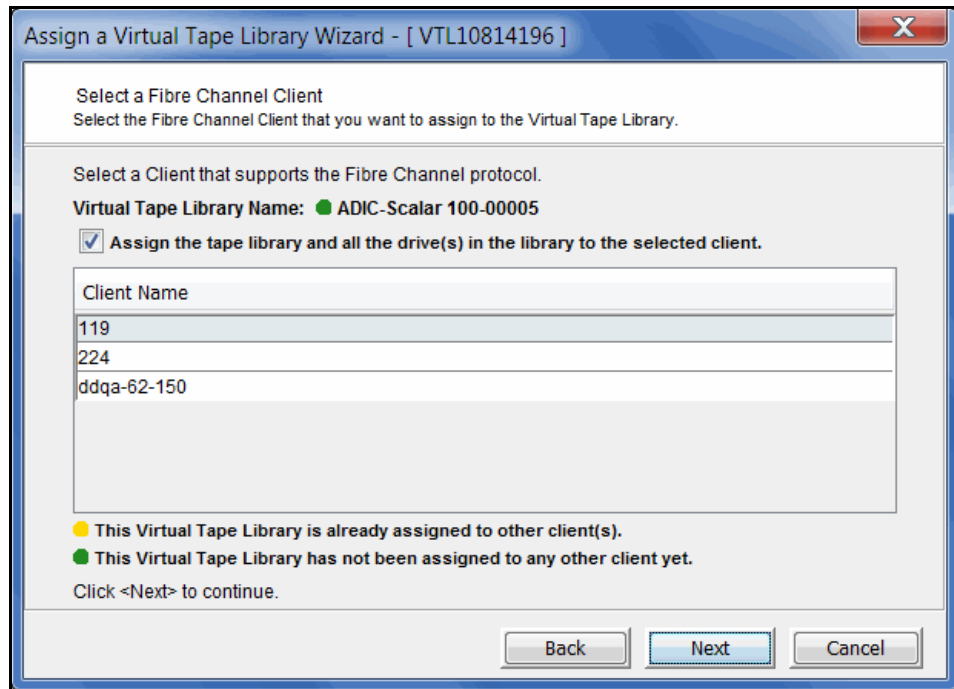
For Windows, *Control Panel --> Computer Management --> Device Manager -->* right-click the device in the right pane --> *Scan for hardware changes*.

4. Use your backup software to discover the library.

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

Assign a client to a virtual tape library or drive

1. Right-click a virtual tape library or drive and select *Assign*.
2. Select the appropriate protocol for the backup application server to which you want to assign the library or drive.
3. Select a backup application server.



4. Click *Next* and then click *Finish* when you are done.

Unassign virtual tape libraries, drives, and iSCSI targets from clients

You can unassign a virtual tape library, drive, or iSCSI target from a client under the *Clients* object.

- FC client - Right-click the client and select *Unassign*. You can also select the *Resources* tab, right-click a library or drive and select *Unassign*.
- iSCSI client - Highlight the client, select the *Resources* tab, right-click a target, library, or drive and select *Unassign*.

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*.
2. Select the options you want to use.

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.



Change firmware of a virtual library or drive

You can change the firmware of a virtual library. 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.
- The maximum number of concurrent tape shredding jobs is 32.

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.
Once a WORM tape is shredded, it will automatically be deleted.
4. Type *YES* to confirm and click *OK*.

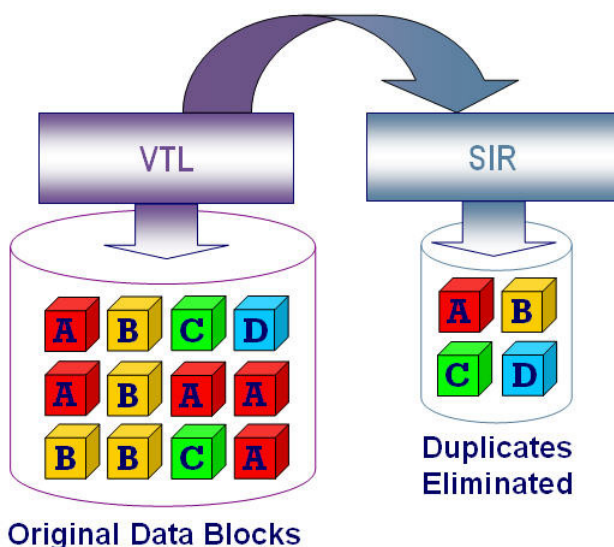
You can view the status by highlighting the virtual tape in the vault. The status bar displays the progress.

If you want to cancel the shredding process, right-click the tape or the *Virtual Vault* object and select *Tape Shredding --> Cancel*.

Deduplication

FalconStor's data deduplication solution eliminates redundant data without impacting your established backup window. FalconStor data deduplication offers as much as a 30:1 reduction of backup data, minimizing replication time and storage requirements.

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.

The Virtual Tape Library common console allows you to view real-time deduplication activity, as well as historical statistics, using one efficient interface.

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.

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, FalconStor offers the following deduplication methods: Turbo deduplication, Inline deduplication, and Post-processing.

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.

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.

The following table compares the different deduplication methods:

	Inline Deduplication	Turbo Deduplication	Post-Processing (Without Inline or Turbo)
Backup Landing Storage and Performance	<ul style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> 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.

Tape deduplication policies

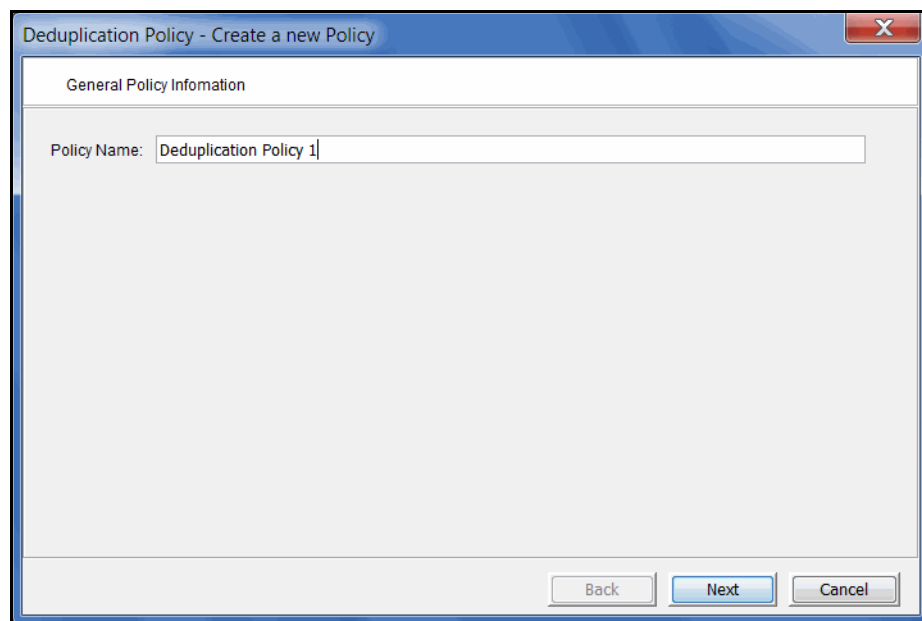
Create tape deduplication policies

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. You can have a maximum of 64 deduplication policies.

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: We recommend that you finalize all IP addresses before deduplication is configured. If you need to change an IP address afterward, refer to '[IP address and netmask update](#)'.

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.



Use standard characters. Unicode characters are not permitted.

3. Select *Localcluster* as your deduplication cluster.

4. Specify when deduplication should occur.

No Schedule (Manual) - Deduplication must be manually initiated.

Inline Deduplication - Data deduplication occurs while backup is in progress. Indicate if the deduplication method should switch to post-processing if inline deduplication is not possible at the start of processing (or at the first write from the backup application)

If the *Switch to post-processing* option is not selected and inline is not possible, the current backup job will fail. If the *Switch to post-processing* option is selected, 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. Regardless of whether or not the *Switch to post-processing* option is selected, 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.

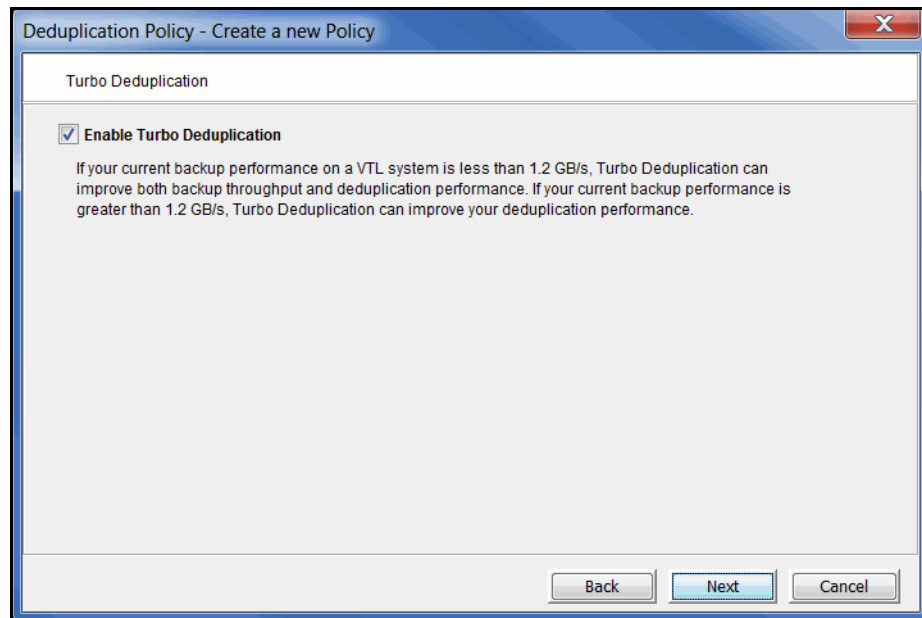
Scheduled Data Deduplication - Deduplication will occur based on the schedule specified.

- *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.

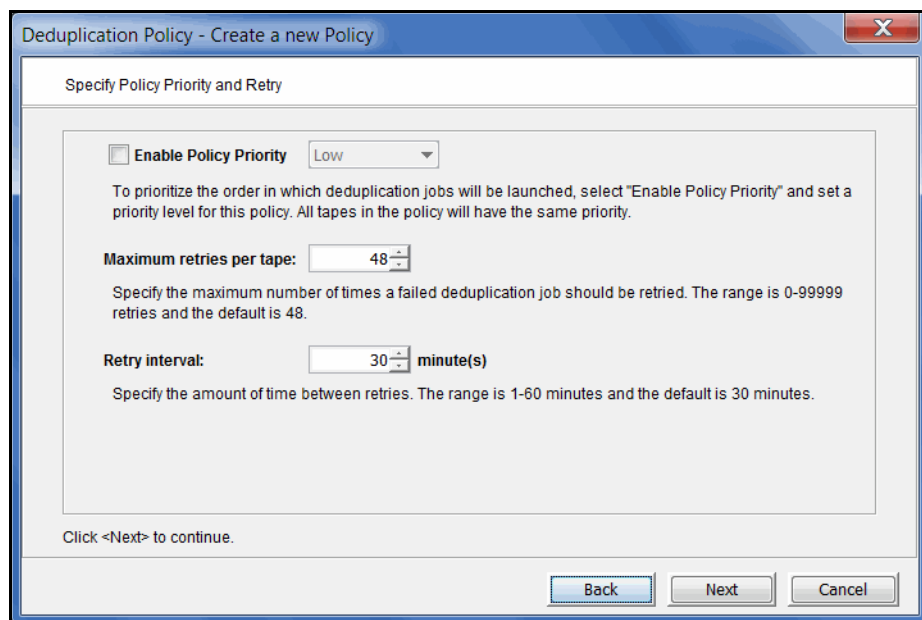
- *Minimum New Data* - Specify the minimum amount of new data that must have been backed up in order for deduplication to occur.

- *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*.



With Turbo deduplication, VTL will pre-process tapes during backup. Refer to '[Deduplication methods - at a glance](#)' for more information about Turbo deduplication.

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.



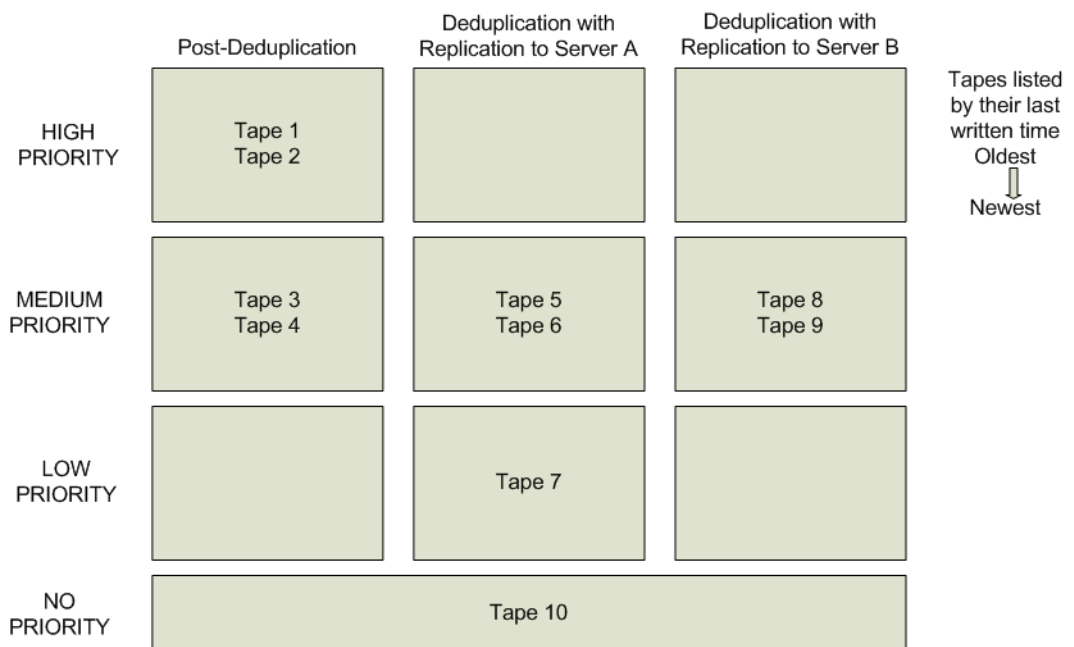
Setting a priority is valid for policies configured for Turbo deduplication, Post-processing deduplication, and Inline deduplication with replication. It also becomes valid for Inline deduplication policies if the job fails and the deduplication becomes a post-processing job.

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:



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.

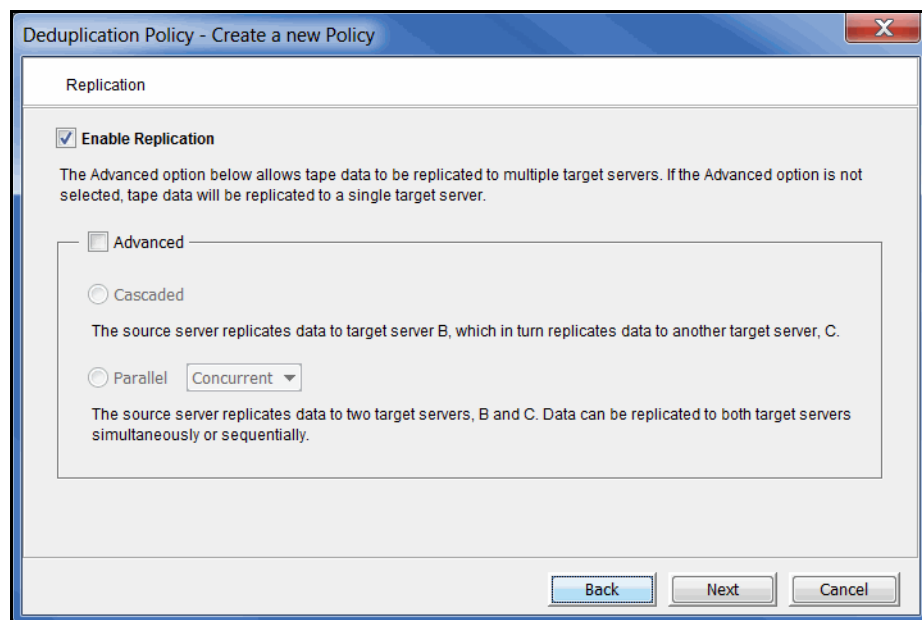
If a deduplication, VIT replication, or resolver process fails, the tape will be returned to the job queue and the entire job will be retried up to the maximum number of times specified in the *Maximum retries per tape* field. The range is 0-99999 retries and the default is 48.

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*.



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 2, where the name of the policy is the server name followed by the policy name.

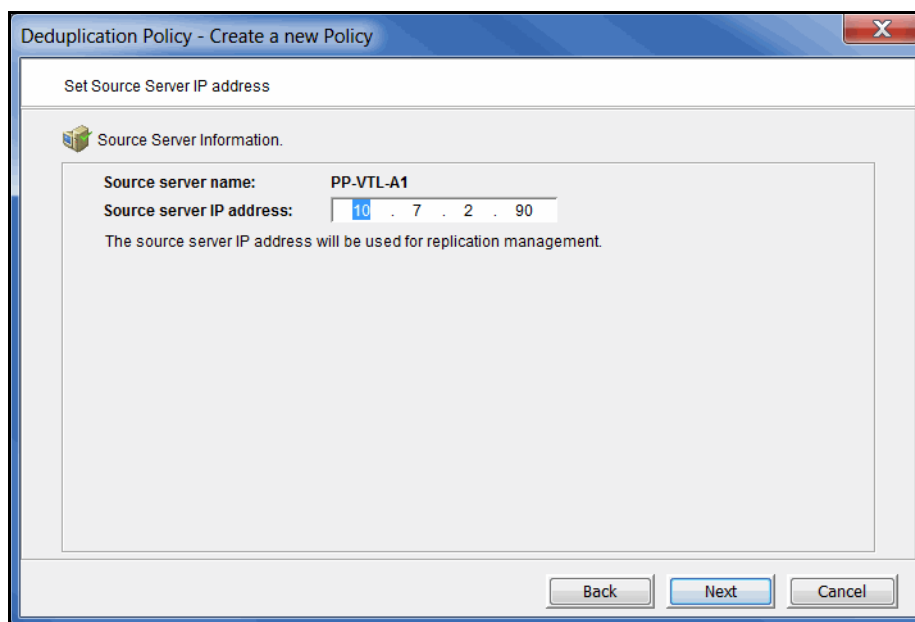
- In *Parallel* mode, source server A replicates data to two target servers, 1 and 2. You can choose either *Concurrent* (replication to both target servers at the same time), or *Serial* (replication first to Replication Target 1 and then to Replication Target 2).

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.



Deduplication Policy - Create a new Policy

Set Source Server IP address

Source Server Information.

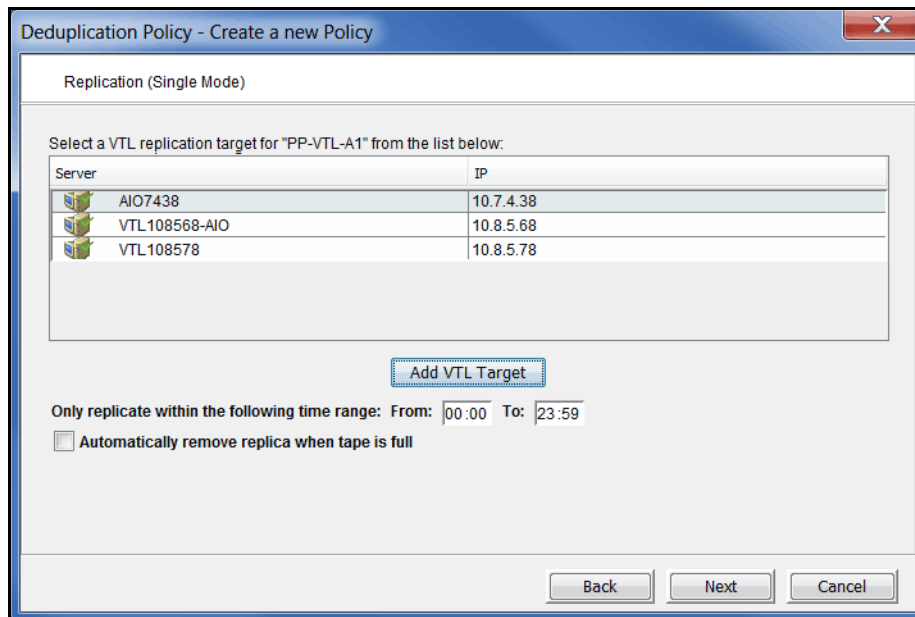
Source server name: PP-VTL-A1

Source server IP address: 10 . 7 . 2 . 90

The source server IP address will be used for replication management.

Back Next Cancel

9. If you selected *Single* mode, select a target server.



The dialog lists backup servers that have already been configured as replication targets for the source server.

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.

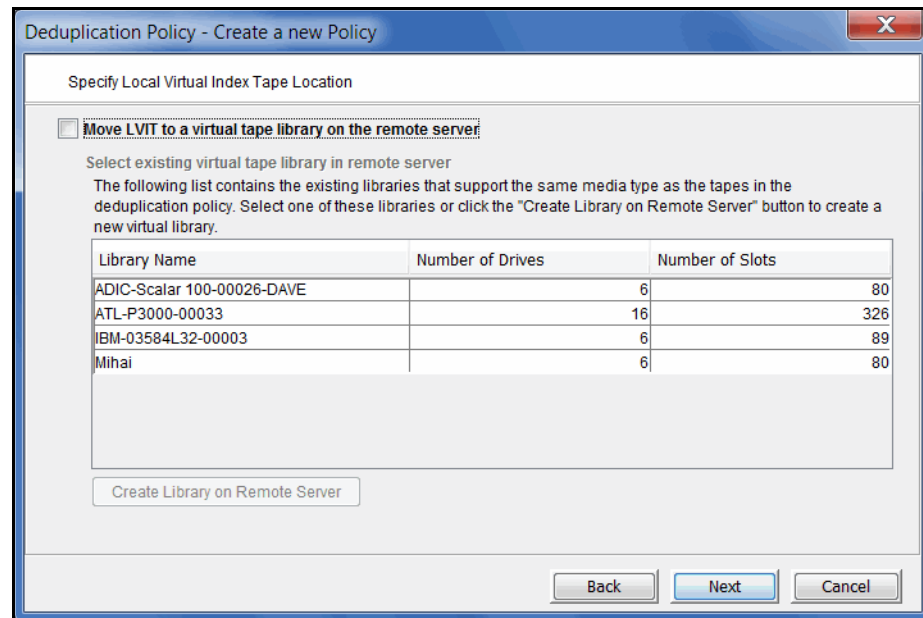
Once you select a replication target for a policy, you cannot edit the policy to change the target.

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. Select the storage pool or physical device(s) from which to create the replica resource.

11. 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.



You can create a library for this purpose on the target server.

12. If you selected *Cascaded* mode, choose two replication targets and an LVIT location for each target in the next four dialogs.
- 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.
 - In the next dialog, select a replication target for Replication Target 1. This will be displayed as *Replication Target 2* in the console.
 - Next, select an LVIT location for Replication Target 1. The same guidelines and options apply as for *Single* mode.
 - Next, select an LVIT location for Replication Target 2.

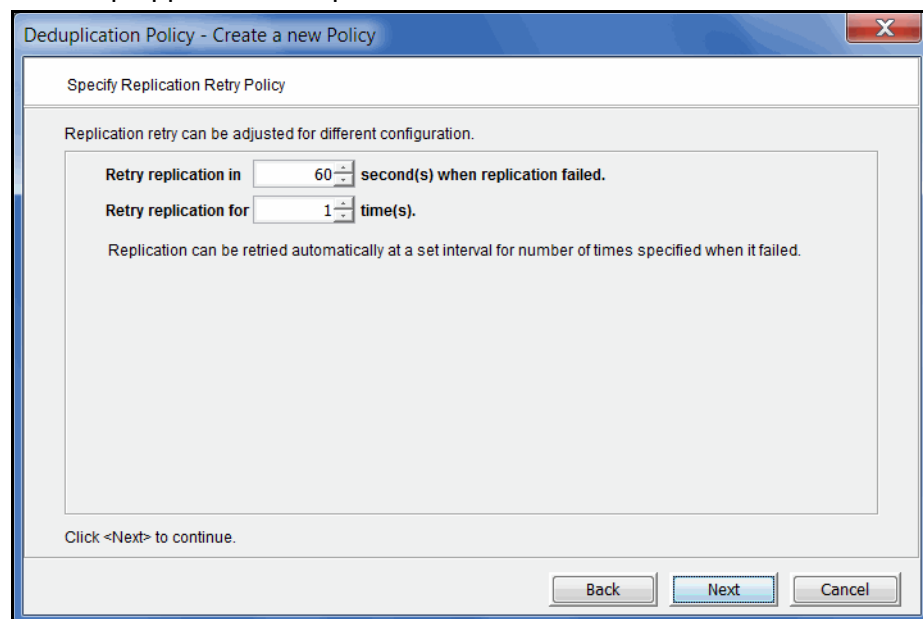
13. If you selected *Parallel* mode, choose two replication targets and an LVIT location for each target server in the next four dialogs.

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.

- 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.
14. If replication is enabled, indicate how often the system should retry if a scheduled replication attempt fails.

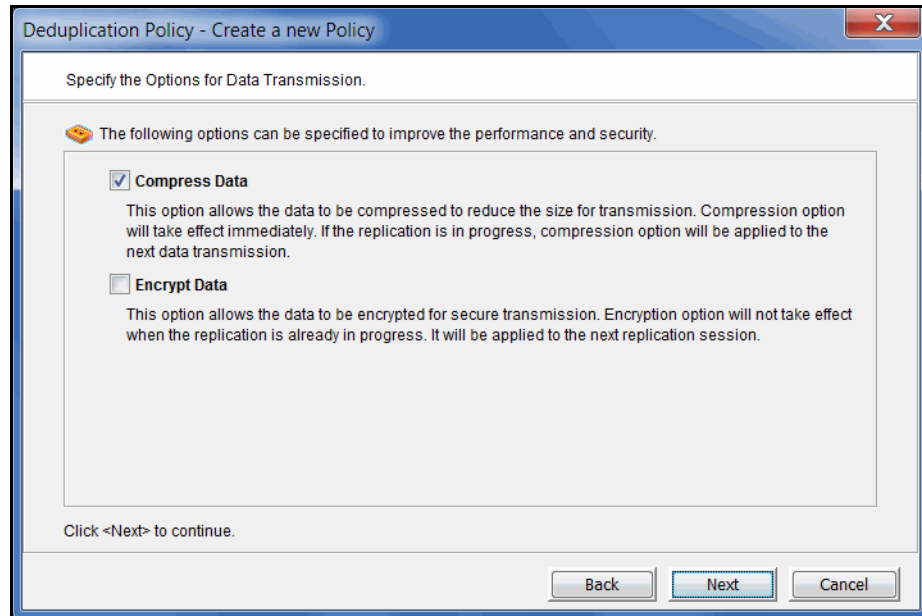
This step applies to all replication modes.



The replication retry policy is specific to index replication failures and will only retry index replication-specific functions.

15. If replication is enabled, indicate if you want to use *Compression* and/or *Encryption*.

This step applies to all replication modes.



The *Compression* option provides enhanced throughput during replication by compressing the data stream.

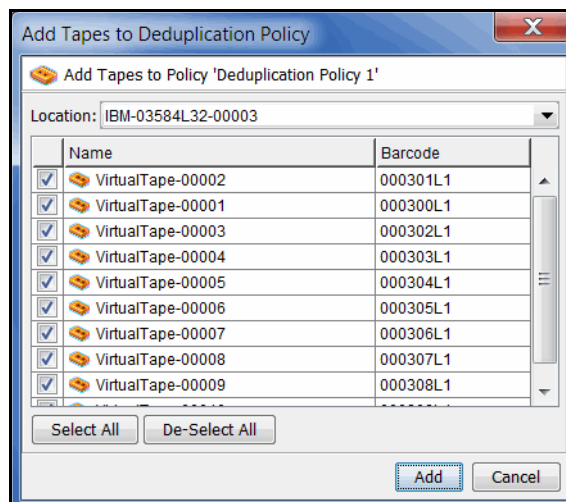
The *Encryption* option secures data transmission over the network during replication. Initial key distribution is accomplished using the authenticated Diffie-Hellman exchange protocol. Subsequent session keys are derived from the master shared secret, making it very secure.

Compression/encryption for transmission over a network should not be set if the source tapes are already encrypted.

16. Click *Finish* to finalize the policy.

Once the policy is created, you will be prompted to add tapes to the policy.

17. Select the virtual tape(s) that you want to include in this policy.



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.
- 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 deduplication and replication properties. You cannot change the replication target for a policy and you cannot change the deduplication policy name.

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*.

Suspend a tape deduplication policy

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

To suspend a policy, right-click the policy and select *Suspend Policy*.

To suspend multiple policies, right-click the *Deduplication Policies* object and select *Suspend Policy*.

Delete a tape deduplication policy

To completely remove a policy, right-click the policy and select *Delete*.

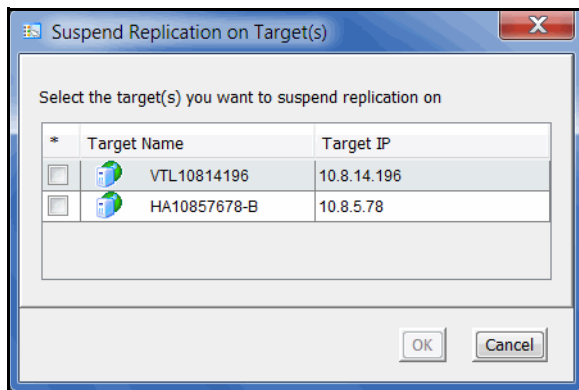
If you want to delete multiple policies, right-click the *Deduplication Policies* object and select *Delete*.

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*.

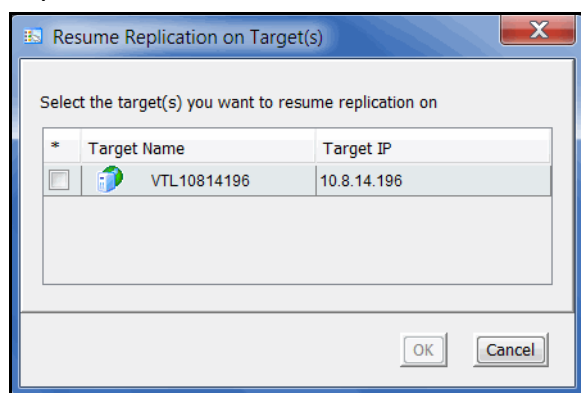


The dialog lists all replication targets. If the policy includes cascaded replication, only Replication Target 1 will be listed, whether you suspend replication from the source server or from Replication Target 1.

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*.



Select the checkbox next to each target for which you want to resume replication and click *OK*.

Manage active tape deduplication policies

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.

You can also monitor the Deduplication Job Queue from the command line.

Type `sirnodeqcli` at the command line 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 FalconStor 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
```

Monitor deduplication and view statistics

Repository statistics

To view repository statistics, go to the *Status* object --> *Dashboard Summary* --> *Deduplication Repository* tab.

Deduplication Repository usage

This section graphically displays the current state of deduplication storage. 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.

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.

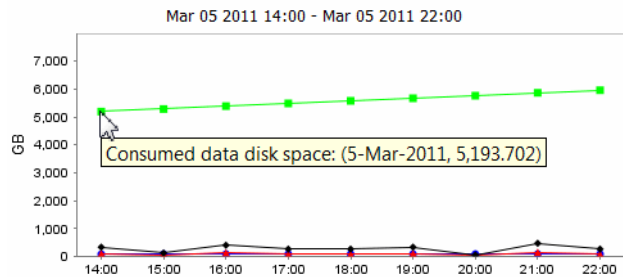
These statistics show deduplication activity over time. Viewing data in this way allows you to calculate the redundancy elimination ratio for any period of time.

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.



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

The management console offers the following tape deduplication-related information:

- 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 and schedule information (such as status, history, and next run time).

Individual tape deduplication policies

When you highlight a policy in the tree, you can view information about that policy.

<i>General Info</i> tab	The <i>General Info</i> tab shows how many tapes are included in this policy, schedule, and replication information. If <i>Advanced Replication</i> is enabled, information is provided for Replication Target 1 and Replication Target 2.
<i>Tapes</i> tab	The <i>Tapes</i> 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. <i>Maximum Capacity</i> - Maximum uncompressed storage capacity of the tape. This is determined when the tape was created. <i>Written</i> - 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. <i>New</i> - The amount of data (before compression) that has not yet been deduplicated, including newly appended data to a tape. <i>In SIR</i> - 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. <i>Unique</i> - The actual physical storage used to store tape data. This includes the effect of deduplication compression. <i>Dedupe ratio</i> - The ratio between the scanned data and unique data moved to the deduplication repository. A ratio of >10000:1 indicates that extremely little data has changed and the amount of unique data is very small or zero. <i>Last run Dedupe</i> - The last time the tape was deduplicated. <i>Last run Replicated</i> - The last time the tape was replicated. <i>Next run</i> - 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.

<i>Active Policies</i> tab	The <i>Active Policies</i> tab lists information about currently running policies and replication jobs. The data is automatically refreshed. All values are rounded.
<i>Tape History</i> tab	The <i>Tape History</i> tab lists all of the deduplication and replication jobs that have run and provides statistics for each. All values are rounded.
<i>Run History</i> tab	The <i>Run History</i> 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.
<i>Event Log</i> tab	The <i>Event Log</i> tab displays informational events and errors pertaining to this policy.

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.

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.

Reclaim disk space

During the deduplication process, only single instances of unique data are passed to the deduplication repository. The original virtual tape is replaced with a VIT 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).

When a VIT 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)

Notes:

- 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.
- If you re-label all deduplicated tapes to overwrite data, space reclamation will not free up space occupied by those labels until you perform additional backup to write new data to a tape.

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 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
- The 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, expand the *Status* object, select the *Dashboard Summary* object, and select the *Deduplication Repository* 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

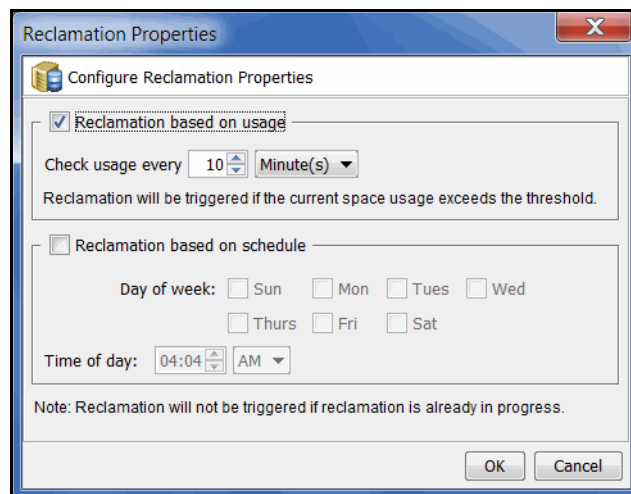
Reclamation can be run in the following ways:

- 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 server and select *Deduplication --> Reclamation --> Reclamation Properties*.



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 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.

Expand deduplication repository

Refer to [“Expand deduplication data repository capacity”](#).

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

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

Encryption requires that an activation password be created. Data in the repository will not be accessible unless the activation password has been entered.

Notes:

- 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.
- Once the deduplication repository is created, encryption cannot be enabled or disabled 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.

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 FalconStor Management Console or via the command line interface. Encrypted virtual tapes will not be accessible for backup or restore unless the activation password is 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.
- When encryption is enabled, you can only replicate to a server with virtual tape encryption enabled.

Enable encryption

Virtual tape encryption can be enabled as follows:

1. Run `unlockdataencryptionoption` from the command line of your virtual tape server.
Refer to ['Enable virtual tape or deduplication repository encryption'](#) for more information.
2. Right-click your VTL server and select *Options --> Enable Virtual Tape Encryption*.



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; previously stored data is not encrypted. 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  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.

Manage encryption keys

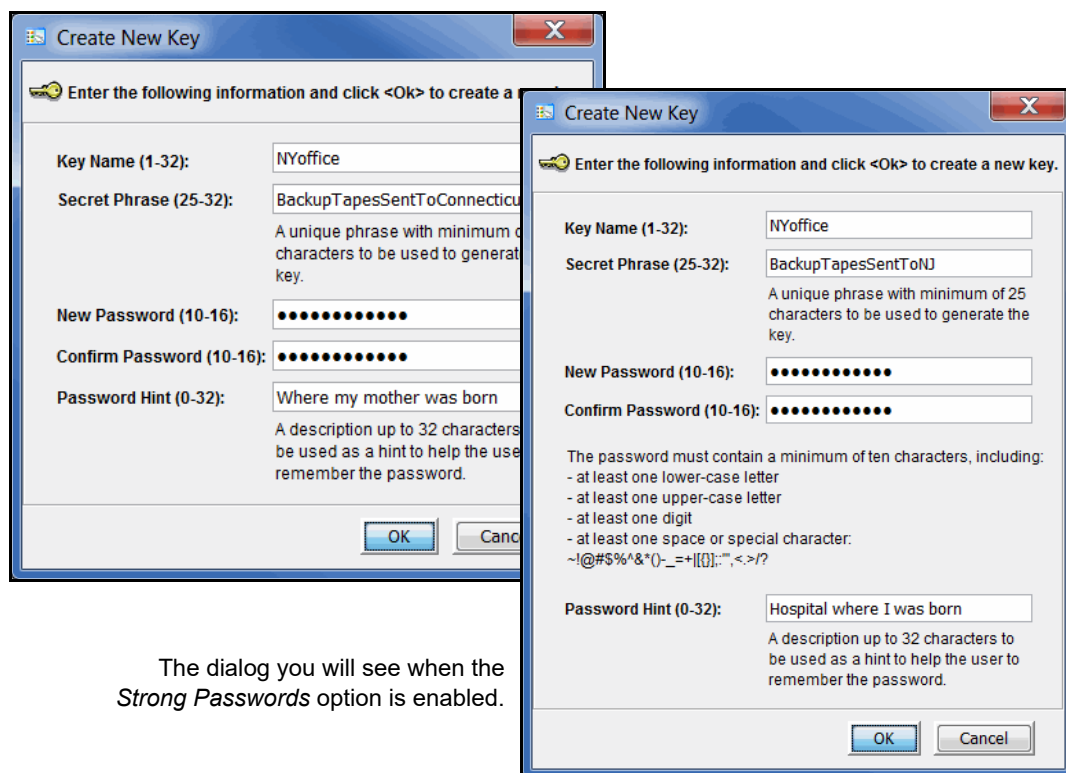
Keys can be created and managed for virtual tape encryption. Deduplication repository encryption uses an internal key that is created when encryption is enabled; it is not visible in *Key Management*.

Create a key

Each key 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*.



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 and all imported keys from the 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.

Data Replication

Replication protects data by maintaining a copy of the data on the same VTL server or on another VTL server.

You can select one of the following methods for replicating data.

Method	Description	Limitations
Tape replication without deduplication	Replicates <i>changed</i> data from a primary virtual tape to the same server or another server at prescribed intervals, based on user defined policies.	Tape replication is mutually exclusive with Auto Replication and Auto Migration to Object Storage.
Tape replication with deduplication	Replicates <i>changed</i> data from a primary VIT to another server at prescribed intervals, based on user-defined policies within a deduplication policy. Cascaded and parallel replication are only available for deduplicated tapes to replicate data to an additional target server.	Replication is not applicable for Auto Migration to Object Storage stub tapes.
Tape 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 and Auto Migration to Object Storage are mutually exclusive.
Tape Remote Copy	Replicates the contents of a single tape <i>on demand</i> .	Remote copy is mutually exclusive with Auto Replication and Auto Migration to Object Storage.

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 without deduplication


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 software does not have access to the replica resource on the target server.

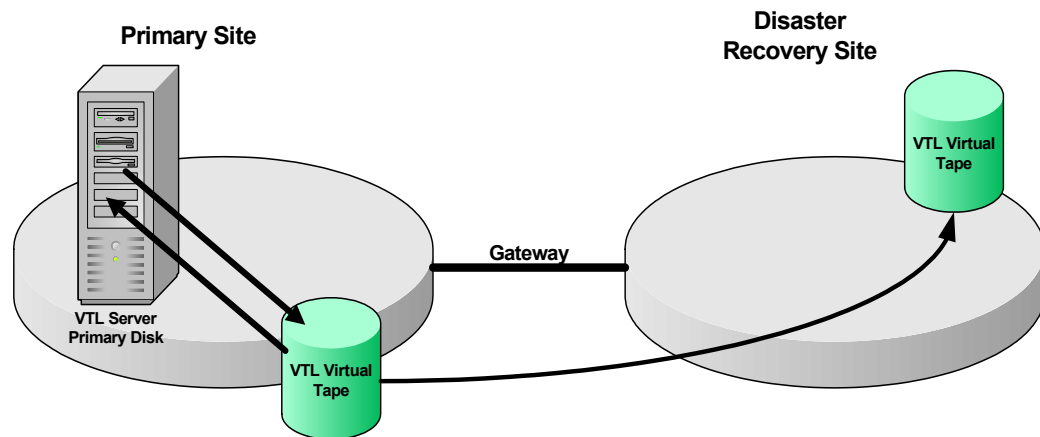
If a disaster occurs and the replica is needed, the administrator can *promote* the replica to become the primary virtual tape so that backup servers can access it.

Replication from the source server to the target server occurs over IP using the TCP protocol. The target server can be a remote server or be the same server as the source server for local 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 of non-deduplicated tape data allows fast, data synchronization of storage volumes on the same VTL server to maintain a local copy of virtual tape data. The replica device can be located on a SAN network and not be necessarily near the source device.



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: We recommend that you finalize all IP addresses before replication is configured. If you need to change an IP address afterward, refer to [‘IP address and netmask update’](#).

1. Right-click one or more virtual tapes in a virtual tape library or in the virtual vault and select *Replication --> Add*.

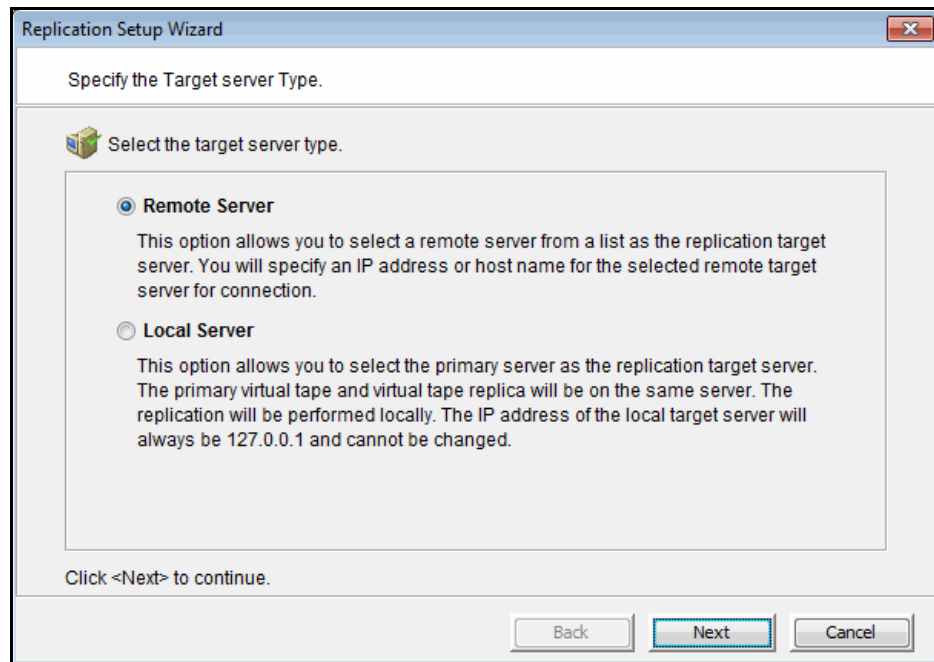
You can also right-click the virtual tape library and select *Replication --> Add*.

You cannot configure replication for tapes that are loaded in tape drives.

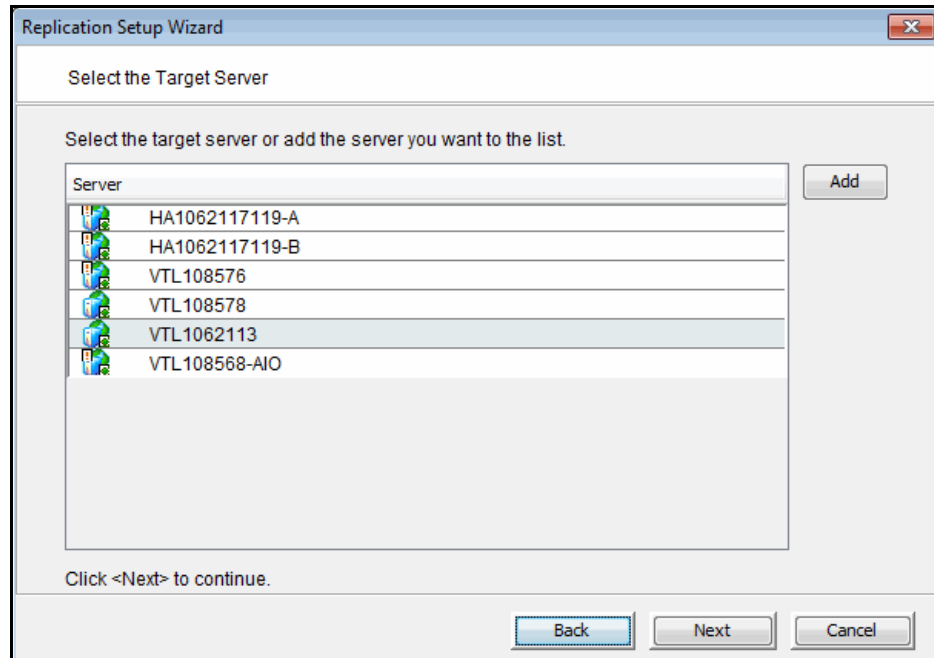
Each virtual tape can only have one replica resource.

Note: If you get a message that Replication cannot be enabled because *Auto Replication* is enabled, you must first disable *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 *Replication/Migration* tab.

2. Indicate whether you want to use remote replication or local replication.

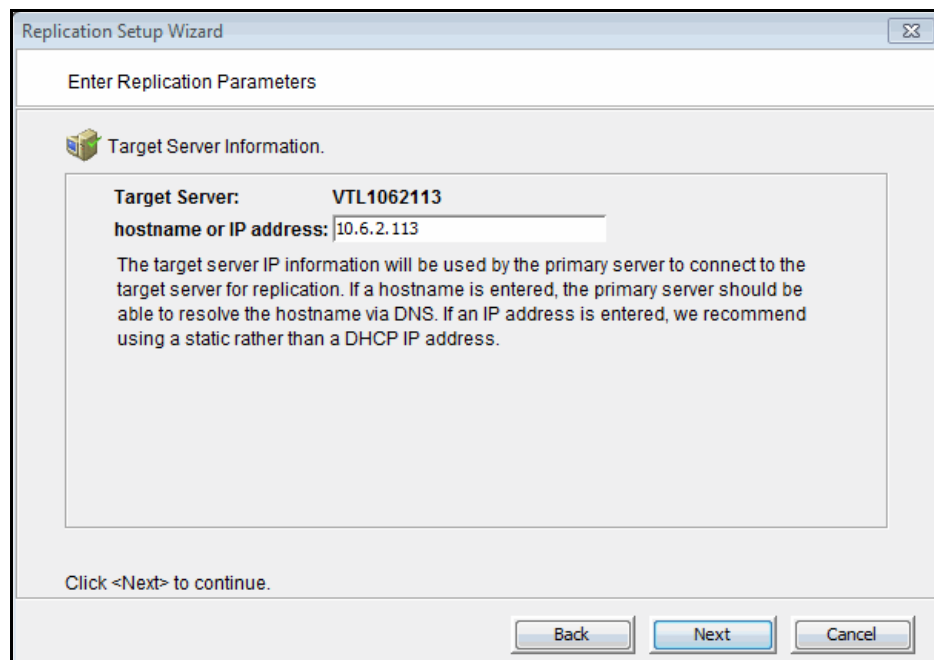


- (Remote Replication only) Select the server that will contain the replica.

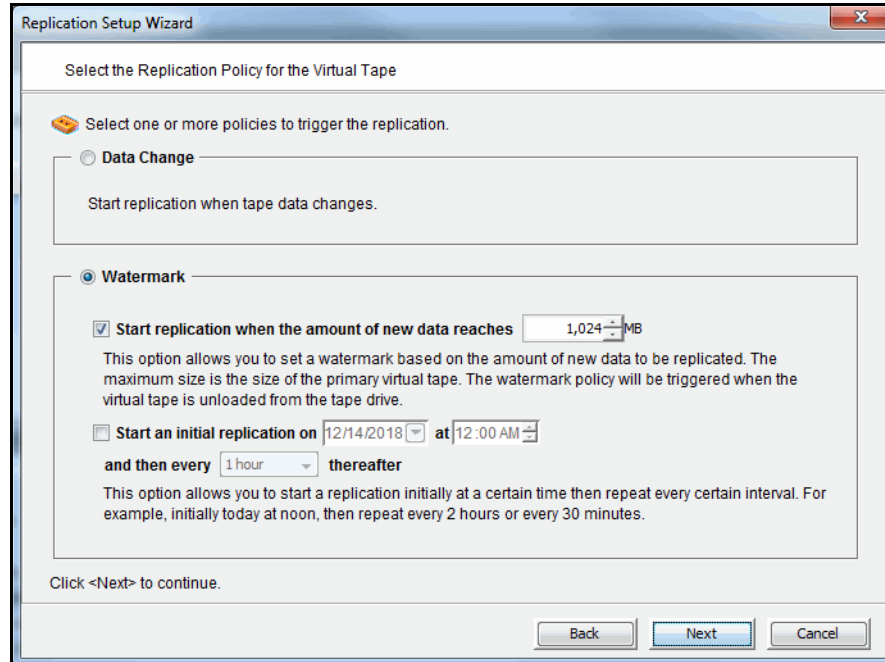


If the server you want does not appear on the list, click the *Add* button.

- For remote replication, confirm/enter the target server's IP address.



5. Configure how often, and under what circumstances, replication should occur.



You must select at least one type of policy (*Data Change* or *Watermark*); you can select multiple *Watermark* policies.

Data Change - Trigger replication whenever data changes on a virtual tape.

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.

6. Indicate what to do if a replication attempt fails.

The screenshot shows a dialog box titled "Replication Setup Wizard" with the subtitle "Specify Replication Timeout and Retry Policy". It contains the following text and controls:

- Icon: A yellow speech bubble icon.
- Text: "Replication timeout and retry can be adjusted for different configuration."
- Form fields:
 - "Timeout replication in" with a spinner box set to "60" and the unit "seconds".
 - "Retry replication in" with a spinner box set to "60" and the unit "seconds when replication failed".
 - "Retry replication for" with a spinner box set to "1" and the unit "times".
- Text: "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."
- Text: "Click <Next> to continue."
- Buttons: "Back", "Next", and "Cancel".

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.

7. (Remote Replication only) Indicate if you want to use *Compression* and/or *Encryption*.

The screenshot shows a dialog box titled "Replication Setup Wizard" with the subtitle "Specify the Options for Data Transmission". It contains the following text and controls:

- Icon: A yellow speech bubble icon.
- Text: "The following options can be specified to improve the performance and security."
- Form fields:
 - 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.
- Text: "Click <Next> to continue."
- Buttons: "Back", "Next", and "Cancel".

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.

8. Select the storage pool or physical device(s) to be used for the replica resource.
The list includes only physical devices that were previously reserved for virtual tapes and storage pools with the *Tapes* device category.
9. (Local Replication only) Enter a name for the replica resource.

Replication Setup Wizard

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.

Back Next Cancel

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.

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 source tape.

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 application server, rescan devices or restart the backup application server to see the promoted virtual tape.

The promoted virtual tape has the same properties as the primary tape; if the primary tape was a WORM tape, the promoted replica will be a WORM 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.

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*.
You cannot remove replication if a tape is loaded in a tape drive.
2. Determine if you want to promote or delete the replica.
3. If deleting, confirm that you want to remove the replica.

Replication of tapes with deduplication

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.

Unique data replication from the source server to the target server occurs via TCP and an IP connection between VTL servers associated with the deduplication repositories is required.

If advanced replication (cascaded or parallel) is configured, data can be replicated to an additional remote location.

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.

- 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 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.
- *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.
- On the target server, you must prepare physical resources for deduplication use and enable deduplication.
- Network firewalls should allow access through TCP ports 11782 (encrypted data replication) and 11781 (unencrypted data replication).

Configure replication for deduplicated tapes

Note: We recommend that you finalize all IP addresses before replication is configured. If you need to change an IP address afterward, refer to ['IP address and netmask update'](#).

Overview of steps to configure replication for deduplicated tapes

You must do the following to configure replication with deduplication:

1. Add a target deduplication replication server to enable replication between the source and target 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.
2. Create a deduplication policy and enable replication.

When you enable replication, you will have to add a VTL target server. Refer to ['Create tape deduplication policies'](#) for more information.

Add a target deduplication replication server

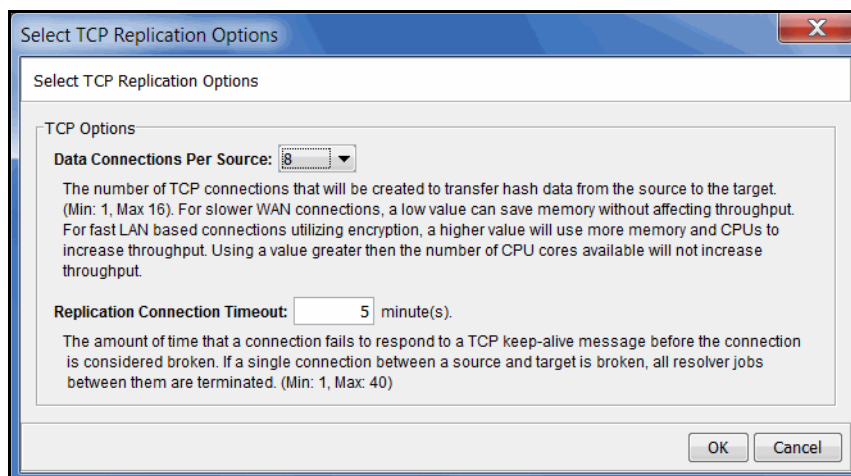
If you are using deduplication, before you can configure replication for tapes in a deduplication policy, you must enable replication between the source and target servers.

To do this:

1. Right-click the primary server and select *Deduplication --> Replication --> Add Target*.
2. Select the target server or click *Add Target* if the server is not listed.
3. You can choose to use TCP 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.

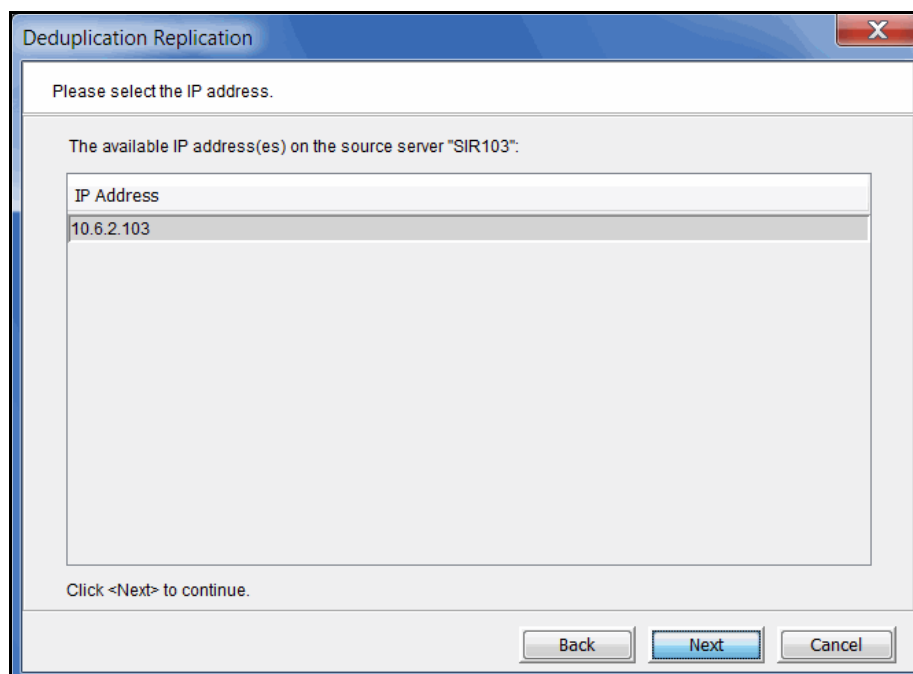


The suggested defaults should be sufficient for most configurations.

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. Select the IP address to use for replication on the source server.



Note: If you are using network address translation (NAT), contact Technical Support before continuing.

- Select the IP address to use for replication on the target server.



- Confirm the information and click *Finish*.

The target 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

You can change the security level and/or advanced options from the console. To do this:

- Right-click the server and select *Deduplication --> Replication --> Edit target*. All targets will be listed.
- Select the target.
If you are not connected to that server, you will have to enter login information for the target server.
- Change the security level or click *Advanced* to modify advanced options recommended by your network administrator.
- Confirm the information and click *Finish*.

The target server is reconfigured. The revised 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. If throttling is not used, replication will use the maximum bandwidth that is available.

1. Right-click a server 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 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.

<i>Active Policies</i> tab (source server)	The <i>Active Policies</i> 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 <i>Reports</i> object) provides a centralized view for displaying replication status for all deduplication policies. Refer to ‘Deduplication Replication Status’ for more information.
<i>Unique Replication Queue</i> tab (target server)	The <i>Unique Replication Queue</i> tab (under <i>Activities</i> 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.
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. 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 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.

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:

- *Auto Replication* and *Auto Migration to Object Storage* are mutually exclusive.
- *Auto Replication* and *Remote Copy* are mutually exclusive.
- Do not enable auto replication for libraries or tapes for which you will be defining a deduplication policy. *Auto Replication* 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

Remote Copy replicates the full contents of a single tape to a local or remote server.


Notes:

- Remote copy is mutually exclusive with Auto Replication and Auto Migration to Object Storage.
- Remote Copy is not applicable for tapes in a deduplication policy or tapes with replication.
- The virtual tape cannot be in a drive.
- 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.



You can select a tape library, standalone tape drive, or the virtual vault.

Notes:

- You can only copy a tape to a tape library if there is no virtual tape in the library with the same barcode. The virtual vault can have duplicate barcodes.
- 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.

Object Storage

Object storage is an architecture to manage data as objects, providing scalability, durability, and security in the cloud.

FalconStor VTL offers the ability to host the deduplication data repository on object storage in the cloud as well as archive virtual tape data to object storage in the cloud.

Deduplication data repository on object storage

During deduplication, unique data is copied to the deduplication repository, which includes deduplication data, index, and folder storage.

Object storage in the cloud can be used to host the VTL deduplication data repository while the index and folder disks remain on SCSI or FC devices.

Benefits of this solution include reduced storage costs, reduced network load, and increased scalability.

Amazon Web Services (AWS) and Hitachi Content Platform (HCP) authentication methods are supported. With AWS authentication (Generic S3), HCP follows the Amazon S3 method to authenticate users; an S3 access key and secret key are required. With HCP authentication, HCP uses its own standard method (via tenant user name and password).

Requirements

The following are the requirements for hosting the deduplication data repository on object storage in the cloud:

- You must have an HCP 8.1 or above account.
- Hitachi API for AWS S3 (or HS3) must be enabled if HS3 is used for data access or if Generic S3 is used.
- The HCP node's domain name must be resolvable via a DNS server or the `/etc/hosts` file and reachable by the VTL server.
- DNS zone configuration must be used for IO load balancing among HCP nodes.
- You must have a stable network connection to the repository in the cloud. If the connection has problems that cause the deduplication process to time out, inline deduplication will fail and will not switch to post-deduplication.
- The namespace (bucket) that will be used for the deduplication repository must be empty.
- The namespace must be assigned to the tenant user.
- The HCP option *Optimized for Cloud Protocols* is not needed.

Configuration

Configuring the deduplication data repository on object storage in the cloud is done when you enable deduplication for a VTL server.

Select *Object Storage* as the downstream storage to use for the deduplication data repository.

Expand deduplication data repository capacity

Right-click the server, select *Deduplication --> Add/Expand Deduplication Data Repository*.

You can enter a new capacity in the dialog.

Note that you must have enough available capacity with your storage provider; the system cannot check whether there is sufficient free space.

Tape migration to object storage

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

FalconStor's Tape Migration to Object Storage can be used to archive virtual tape data to the cloud.

In VTL, a policy can be set for a virtual tape library to automatically migrate tape data to object storage as soon as a backup is complete and a tape is ejected to the virtual vault. In this way, migration is transparent to the backup software.

Migration can occur in two modes, *Move* or *Copy*. For space optimization, *Move* mode frees up all disk space used by the virtual tape once data is moved to object storage. In this mode, the source virtual tape is converted to a stub tape after migration completes successfully. When there is a need to recover data from object storage, virtual tapes are recovered from stub tapes so that they can be accessed by the backup software.

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

Auto Migration to Object Storage, Auto Replication, and replication without deduplication are mutually exclusive.

Replication of data on a deduplicated virtual index tape (VIT) can occur before it becomes a stub tape. If you want to replicate data on VITs, you can use *Copy* mode and replicate before converting VITs to stub tapes, or if you use *Move* mode, set a grace period that is large enough to give you time to replicate the data before VITs becomes stub tapes.

Migration and recovery jobs

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

All object migration and recovery jobs are listed in the Tape Import/Export Queue (under the *Activities* object). While a job is running, you can see its status there. Up to 10 migration and 5 recovery jobs can run at the same time.

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

Should a job fail, it will be retried based on the retry settings of the Tape Import/Export Queue, as long as the tape remains in the vault. When a migration job is retried, it continues from the failed point, not from the beginning of a tape. Migration

and recovery jobs share the same job properties with import/export jobs (number of retries and retry interval for failed jobs). To change the retry settings, right-click the *Tape Import/Export Queue* object and select *Properties*.

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

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

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

If the data is not compressed by VTL, the migration process will compress the data in flight and at rest.

Configuration

A VTL administrator must do the following to configure tape migration to object storage:

1. Create at least one object storage account with your object storage provider.
To isolate data between different users, you can create separate accounts.

2. Add your object storage account to your VTL server.
Refer to [“Add an object storage account”](#) for more information.

3. Enable Auto Migration for a virtual tape library.

This is typically done when you create a virtual tape library. Refer to [“Create virtual tape libraries”](#) for more information. You can also enable it for an existing library by right-clicking the virtual library and selecting *Properties*.

Each library can use a different object storage account.

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

If you enable Auto Migration on a virtual tape library with existing tapes, those tapes will inherit the library’s Auto Migration properties when they are ejected to the virtual vault. Once inherited, you will not be able to change the object storage account for those tapes.

Migrate virtual tape data to object storage

Migration is automatically triggered when a virtual tape is ejected to the virtual vault. If you enabled Auto Migration on a VTL system with existing tapes that can be migrated, you can manually trigger migration to object storage. To do this, right-click a virtual tape in the virtual vault and select *Migrate to Object Storage*.


You can watch the status of the migration job in the Tape Import/Export Job Queue.

Convert a virtual tape to a stub tape

If object migration is configured in *Copy* mode, the source tape is not automatically converted to a stub tape after migration.

If object migration is configured in *Move* mode, a grace period may have been set, so even though migration is complete, the source tape may not yet be converted to a stub tape.

To manually convert a virtual tape to a stub tape, right-click a virtual tape in the virtual vault that has completed migration and select *Convert to Stub Tape*.

Afterward, the virtual tape will have a blue  icon to indicate that it is a stub tape.

Recover data from object storage

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

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

You must make sure there is enough backup cache capacity available for the recovered data before you begin the recovery process. Reconstructed tapes are not deduplicated even if the original tape was in a deduplication policy. The recovered tape will be write-protected after recovery. To reuse the tape, right-click the tape and select *Property* to change the write protection property.

1. In the virtual vault, right-click a stub tape and select *Recover Virtual Tape from Object Storage*.
2. Select the virtual library to which you want the recovered tape to reside.
3. Select the virtual tape library slot and recovery mode.

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

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

4. If your AWS S3 storage class is *Glacier*, select your AWS retrieval method.

If your AWS retrieval method is *Bulk* or *Standard*, the retrieval job may take hours to complete and you will only see progress in the console when the data download begins.

If the data download fails, it will be retried and will resume from the point of failure, not from the beginning of the tape.

You can watch the status of the recovery job in the Tape Import/Export Job Queue. If the recovery job fails, it will be retried based on the retry settings of the Tape Import/Export Queue. It can also be manually restarted from the Tape Import/Export Job Queue.

Manage migrated/stub tapes

Migration is triggered when a virtual tape is ejected to an I/E slot by backup software, which is the virtual vault in VTL. After migration, the source tape (*Copy* mode) or stub tape (*Move* mode) remain in the virtual vault.

Stub tapes cannot be moved out of the virtual vault. A virtual tape that is being migrated cannot be moved out of the vault; you must cancel the job before moving the tape out of the vault.

When you click on a virtual tape, the *Performance* tab displays the performance of any active migration (or recovery) job.

If you delete a virtual tape in the virtual vault that has been migrated, the associated objects for the tape will be purged at the object storage vendor.

If you delete a stub tape, its data at the object storage vendor will be deleted.

Object storage accounts

Object storage accounts are created with your object storage provider. VTL uses these accounts to host the deduplication data repository on object storage in the cloud as well as archive virtual tape data to object storage in the cloud.

Add an object storage account

1. Right-click a VTL server and select *Object Storage Account*.
2. Select a provider from the drop-down box and click *Add*.

Microsoft Azure

Account Name - Specify a unique name for this account.

Azure Account Name - Specify a user account.

Account Key - Specify the account secret key used for authentication.

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

Endpoint Suffix - Optionally, specify an endpoint suffix.

Blob Container - Specify the blob container in the storage account.

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

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

Amazon Web Services (AWS) S3

Account Name - Specify a unique name for this account.

IAM Key ID - Specify the AWS Identity and Access Management (IAM) key ID (16 to 128 bytes).

Secret Key - Specify the secret access key of the key ID that is used to access object storage.

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

Region - Specify the region of Amazon AWS S3 service. The default is *US East (N. Virginia)*.

Storage Class - Specify the storage class. The default is *Glacier*.

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

Dual Stack - Indicate if you have a dual stack network where IPv4 and IPv6 protocols are available for object storage access. With dual stack, object storage access points can be resolved to IPv6 addresses in addition to IPv4 addresses.

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

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

Hitachi Content Platform (HCP)

The screenshot shows a dialog box titled "Add Object Storage Account" with the following fields and options:

- Account Name : [Text Input]
- Namespace : [Text Input]
- Domain Name: [Text Input]
- Tenant Name: [Text Input]
- User Name: [Text Input]
- Password: [Text Input]
- Protocol : https http
- Data Access Method : REST HS3
- End-to-End Encryption Reserved for SIR
- Comments : [Text Input]

Buttons: OK, Cancel

Account Name - Specify a unique name for this account.

Namespace - Specify the name of an existing namespace assigned to the tenant. Objects created by tape migration will be kept in this namespace. Objects in one namespace are not visible in any other namespace.

Domain Name - Specify the domain name you use to access the provider.

Tenant Name - Specify the tenant for the provider.

User Name - Specify the username for the provider.

Password - Specify the password for the user.

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

Data Access Method - HCP namespace contents can be accessed via *REST* or *HS3*. *REST* (default) is a standard HTTP protocol. *HS3* is an HTTP-based API that is compatible with Amazon S3. In order to use *HS3*, *HS3* must be enabled from the HCP management console. In addition, the *Optimized for cloud protocols only* option must be selected from the HCP management console when *HS3* is selected.

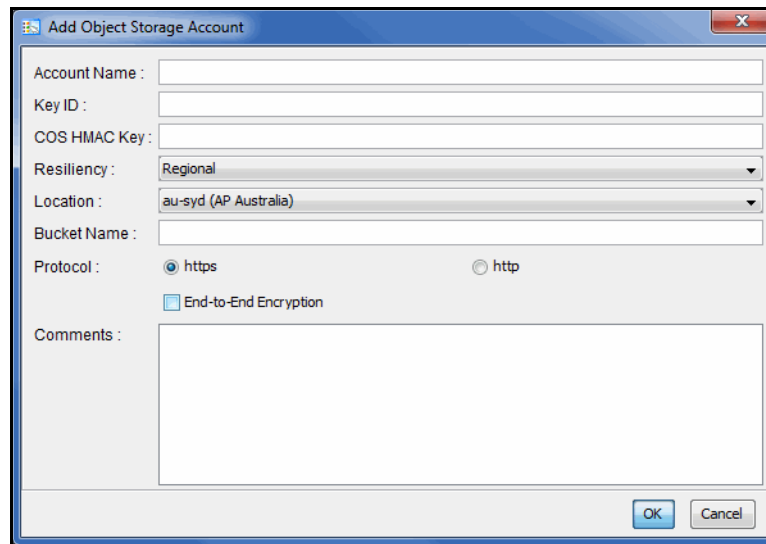
End-to-End Encryption - Enable or disable end-to-end encryption. When enabled, all data is encrypted before saving it to object storage; data is always encrypted in-flight and at-rest regardless of the protocol (*http/https*).

Reserved for SIR - Specify if this account will be used to host the deduplication data repository on object storage instead of being used to archive virtual tape data to object storage. You cannot use the same account for both.

Either *End-to-End Encryption* or *Reserved for SIR* can be selected, but not both.

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

IBM Cloud Object Storage (COS)



The screenshot shows a dialog box titled "Add Object Storage Account". It contains the following fields and controls:

- Account Name : [text input]
- Key ID : [text input]
- COS HMAC Key : [text input]
- Resiliency : [dropdown menu, selected: Regional]
- Location : [dropdown menu, selected: au-syd (AP Australia)]
- Bucket Name : [text input]
- Protocol : [radio buttons, selected: https, unselected: http]
- End-to-End Encryption : [checkbox, unselected]
- Comments : [text area]
- OK [button]
- Cancel [button]

Account Name - Specify a unique name for this account.

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

COS HMAC Key - Specify the secret access key (password) for authentication, up to 4,096 characters.

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

Location - Specify a location, based on *Resiliency*.

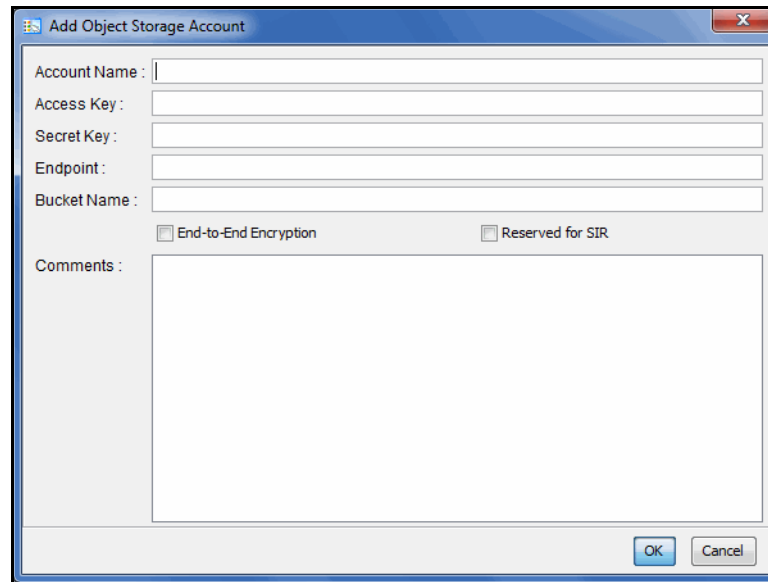
Bucket Name - Specify the name of an existing bucket.

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

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

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

Generic S3



The screenshot shows a dialog box titled "Add Object Storage Account". It features a standard Windows-style title bar with a close button. The main area contains the following elements from top to bottom: five text input fields labeled "Account Name", "Access Key", "Secret Key", "Endpoint", and "Bucket Name"; two checkboxes labeled "End-to-End Encryption" and "Reserved for SIR"; a large text area labeled "Comments"; and finally, "OK" and "Cancel" buttons at the bottom right.

Create a Generic S3 object storage account if your storage provider is Wasabi.

Account Name - Specify a unique name for this account.

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

Secret Key - Specify the secret key for the access key ID.

Endpoint - Specify the full URI/URL path to the object storage.

Bucket Name - Specify the name of an existing bucket.

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

Reserved for SIR - Specify if this account will be used to host the deduplication data repository on object storage instead of being used to archive virtual tape data to object storage. You cannot use the same account for both.

Either *End-to-End Encryption* or *Reserved for SIR* can be selected, but not both.

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

3. Click *OK* when done.

Manage object storage accounts

VTL administrators can add, edit, and delete object storage accounts.

For AWS, the region and bucket cannot be changed. For security reasons, it is a common practice to change the account secret key periodically.

Notes:

- If the deduplication data repository is on object storage, the object storage provider account cannot be deleted.
- Changing an HCP user name/password or Generic S3 access/security keys can result in an IO disruption. Be sure to suspend all active IO on VTL before changing credentials and restart VTL services afterward.

Deleting an account deletes all data stored at the object storage. You may want to move objects belonging to the soon-to-be-deleted account to a new location (in a new account) before deleting the account. In addition, you must do the following before deleting an account:

- Remove the migration configuration from virtual tape libraries using that account.
- Recover stub tapes to virtual tapes.
- Cancel all active migration and recovery jobs in the Tape Import/Export Queue.
- Delete all queued and on-hold migration and recovery jobs in the Tape Import/Export Queue.

iSCSI Configuration

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.

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.

iSCSI target mode is supported for the following client platforms:

- [Windows](#)
- [Linux](#)
- [IBM i](#)

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.
- iSCSI software initiator installed on each Windows 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.msp>
- For improved performance, if your network supports it, turn on jumbo frames and set the maximum transfer unit (MTU) of each IP packet to 9000.

Enable iSCSI

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*.

Prepare client initiators

Before an iSCSI client can be served by a backup appliance, the two entities need to mutually recognize each other.

You will need to register each iSCSI client as an initiator to your VTL server. This enables the VTL server to see the initiator.

To do this, launch the iSCSI initiator on the client machine and identify your VTL server as the target server. Then, enter the IP address or name (if resolvable) of your VTL server and use the default port (3260).

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.

Add an iSCSI client

1. In the VTL console, expand the *Clients* object, right-click *iSCSI Clients* and select *Add*.
2. Enter the client name.
3. 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.

4. 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.

5. 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 Clients* object.
3. Right-click an *iSCSI Client* object and select *Create Target*.

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.

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*.

Assign a virtual tape library to the iSCSI target

1. Right-click an *iSCSI Client* object and select *Assign*.
2. 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.
3. On the next screen, change the LUN for the resource if you need to resolve a conflict.
4. 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.

Linux configuration

Requirements

- A VTL server with an Ethernet adapter installed.
- iSCSI software initiator installed on each Linux backup application server. iSCSI initiator software/hardware is available from many sources.
- For improved performance, if your network supports it, turn on jumbo frames and set the MTU of each IP packet to 9000.

Enable iSCSI

Refer to ['Enable iSCSI'](#) for more information.

Prepare client initiators

Before an iSCSI client can be served by a backup appliance, the two entities need to mutually recognize each other.

You will need to register each iSCSI client as an initiator to your VTL server. This enables the VTL server to see the initiator.

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
```

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

Refer to [‘Assign a virtual tape library to the iSCSI target’](#).

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 i configuration

IBM i is an operating system that runs on IBM's Power Systems and Pure Systems.

Requirements

- A VTL server with an Ethernet adapter installed.
- iSCSI software initiator installed on each IBM i backup application server.
- For improved performance, if your network supports it, turn on jumbo frames and set the MTU of each IP packet to 9000.

Enable iSCSI

Refer to ['Enable iSCSI'](#) for more information.

Prepare client initiators

For IBM i version 7.3 or higher, use the *IBM Navigator for i* GUI to configure iSCSI on an IBM i client.

Add an iSCSI client

From the FalconStor console, add the iSCSI client. You will need to enter the client iSCSI initiator iqn. Refer to ['Add an iSCSI client'](#).

Create iSCSI target

For IBM i version 7.3 or higher, use the *IBM Navigator for i* GUI to create an iSCSI target. Select system Services → iSCSI Tab → Action → Create iSCSI Target → Enter target iqn, target IP address or hostname, and CHAP parameters, if necessary.

For older IBM i versions, use SQL commands as described below. For command details, refer to <https://www.ibm.com/support/pages/ibm-i-removable-media-support-iscsi-support-iscsi-vtl#:~:text=IBM%20i%20partitions%20running%20IBM,and%20restore%20operations%20over%20Ethernet>

You must have *IOSYSCFG special authority. The SQL schema name is QSYS2.

List iSCSI connections

Retrieve iSCSI configuration information:

```
SELECT * FROM Qsys2.Iscsi_Info;
```

The ISCSI_INFO view returns iSCSI configuration information for the system. One row is returned for each iSCSI target that is configured.

If no iSCSI targets are configured, one default row is returned where the INITIATOR_NAME column contains the default initiator name generated by the system.

The following describes the columns in the view:

- TARGET_NAME: The iSCSI target name
- TARGET_HOST_NAME: The TCP/IP host name of the iSCSI target
- TARGET_PORT: The TCP/IP target port being used for iSCSI connection
- INITIATOR_NAME: The iSCSI initiator name associated with the iSCSI target
- INITIATOR_CHAP_NAME: The Challenge-Handshake Authentication Protocol (CHAP) name used by the IBM i iSCSI initiator to authenticate with the iSCSI target; null if CHAP authentication is not in use
- RESOURCE_NAME: Resource Name of the iSCSI Virtual I/O Processor (IOP)
- TARGET_STATUS: Current status of the iSCSI target

Create iSCSI connections

The ADD_ISCSI_TARGET procedure configures an iSCSI target as the VTL server. A maximum of 16 iSCSI targets can be configured on the system.

The following describes the columns in the view:

- TARGET_NAME: The iSCSI target name.
- TARGET_HOST_NAME: The TCP/IP host name or IP address of the iSCSI target
- TARGET_PORT: An integer value (1-65535) specifying the remote TCP/IP port that the remote iSCSI target is listening on; default 3260
- INITIATOR_NAME: The initiator name for local system. You can supply an initiator name or use the default initiator name. Specify NULL to use the system-generated initiator name. The system generates a default initiator name using the Universal Unique Identifier (UUID) for the partition. The UUID is available when the IBM i is running on POWER8 or later hardware. If the IBM i is running on POWER7 hardware or earlier, an initiator name must be specified. To view the default initiator name generated by the system before configuring an iSCSI target, use the iSCSI_info view. The IBM i uses the iSCSI Qualified Name (IQN) format for its initiator name. The following guidelines should be used when constructing an initiator name. An IQN type name consists of the following components:
 - The string "iqn." to distinguish the name as an iqn type name.
 - A date code in the format `yyyy-mm`. This date must be a date during which the naming authority owned the domain name used in this format and should be the first month in which the domain name was owned by this naming authority at 00:01 GMT of the first day of the month.
 - A dot (.).
 - The reverse domain name of the naming authority creating this iSCSI name, for example "com.ibm".
 - Optional: Colon (:) followed by product and/or system specific

information. The IBM i default name uses `ibmi.-i`. The UUID is a 32-character hexadecimal identifier for the partition. The initiator index is a zero-based index. Since one initiator is supported on IBM i, the initiator index is always '0'.

The following is an example of a default initiator name generated for IBM i:

```
iqn.1924-02.com.ibm:ibmi.4520920efdc3454db06b96a56d912aa5-i0
```

- **INITIATOR_CHAP_NAME:** The initiator CHAP name to be used by the iSCSI target to authenticate the IBM i initiator.
- **INITIATOR_CHAP_SECRET:** The initiator CHAP secret to be used by the iSCSI target to authenticate the IBM i initiator. If specified, the initiator CHAP secret must be at least 12 characters in length.

Examples

```
CALL Qsys2.Add_Iscsi_Target (
  Target_Name => 'FalconStor_iSCSI_iqn_name',
  Target_Host_Name => 'VTL Name or VTL IP',
  Target_Port => 3260,
  Initiator_Chap_Name => '',
  Initiator_Chap_Secret => '');
```

```
CALL QSYS2.ADD_ISCSI_TARGET(
  TARGET_NAME=>'iqn.2000-03.com.swvtl:vtl.vtltest.test-47',
  TARGET_HOST_NAME=>'vtltest.ibm.com',
  INITIATOR_CHAP_NAME=>'username',
  INITIATOR_CHAP_SECRET=>'ChapSecretPW');
```

```
CALL QSYS2.ADD_ISCSI_TARGET(
  TARGET_NAME=>'iqn.2000-03.com.swvtl:vtl.vtltest.test-47',
  TARGET_HOST_NAME=>'vtltest.aaa.demo.aaa.com',
  INITIATOR_NAME=>'iqn.1924-
02.com.ibm:ibmi.45200920efdc3454db06b96a56d912aa5-i0',
  TARGET_PORT=>3260,
  INITIATOR_CHAP_NAME=>'username',
  INITIATOR_CHAP_SECRET=>'ChapSecretPW');
```

Update iSCSI target

The `CHANGE_ISCSI_TARGET` procedure changes the configuration of an existing iSCSI target. For example:

```
CALL Qsys2.Change_Iscsi_Target (
  Target_Host_Name => 'VTL Name or VTL IP',,
  Target_Name => 'FalconStor_iSCSI_iqn_name',
  Initiator_Chap_Name => '',
  Initiator_Chap_Secret => '');
```

Remove iSCSI target

The REMOVE_ISCSI_TARGET procedure permanently ends use of an iSCSI target.

```
CALL Qsys2.Remove_Iscsi_Target(  
  Target_Name => 'Falcon_Stor_iSCSI_ign_name',  
  Target_Host_Name => 'VTL Name or VTL IP');
```

Log the client onto the target

After changes to the iSCSI configuration, an Initial Program Load (IPL) of the iSCSI Virtual I/O Processor (IOP) must be done for the changes to take effect.

If the VTL server has been power cycled, the iSCSI controller also needs to be IPLed in order to recover the connection.

Use the Hardware Service Manager function in the System Service Tools to IPL the iSCSI Virtual IOP.

1. Run the IBM i command line, STRSST.
2. Select option 1 (Start a Service tool).
3. Select option 7 (Hardware service manager).
4. Select option 2 (Logical Hardware Resources).
5. Select option 1 (System bus resources).
6. Scroll through the bus resources and look for one with a Type-Model or 298A-001. Select option 6 for that resource.
7. Select option 4 to (IPL I/O Processor).
8. Exit hardware service manager and SST.
9. Wait several seconds until you see a new TAPMLBxx device and new tapes device based on libraries defined on the Falconstor VTL server.

Note: IBM i system value QAUTOCFG should be set to '1'.

Assign a virtual tape library to the iSCSI target

Refer to ['Assign a virtual tape library to the iSCSI target'](#).

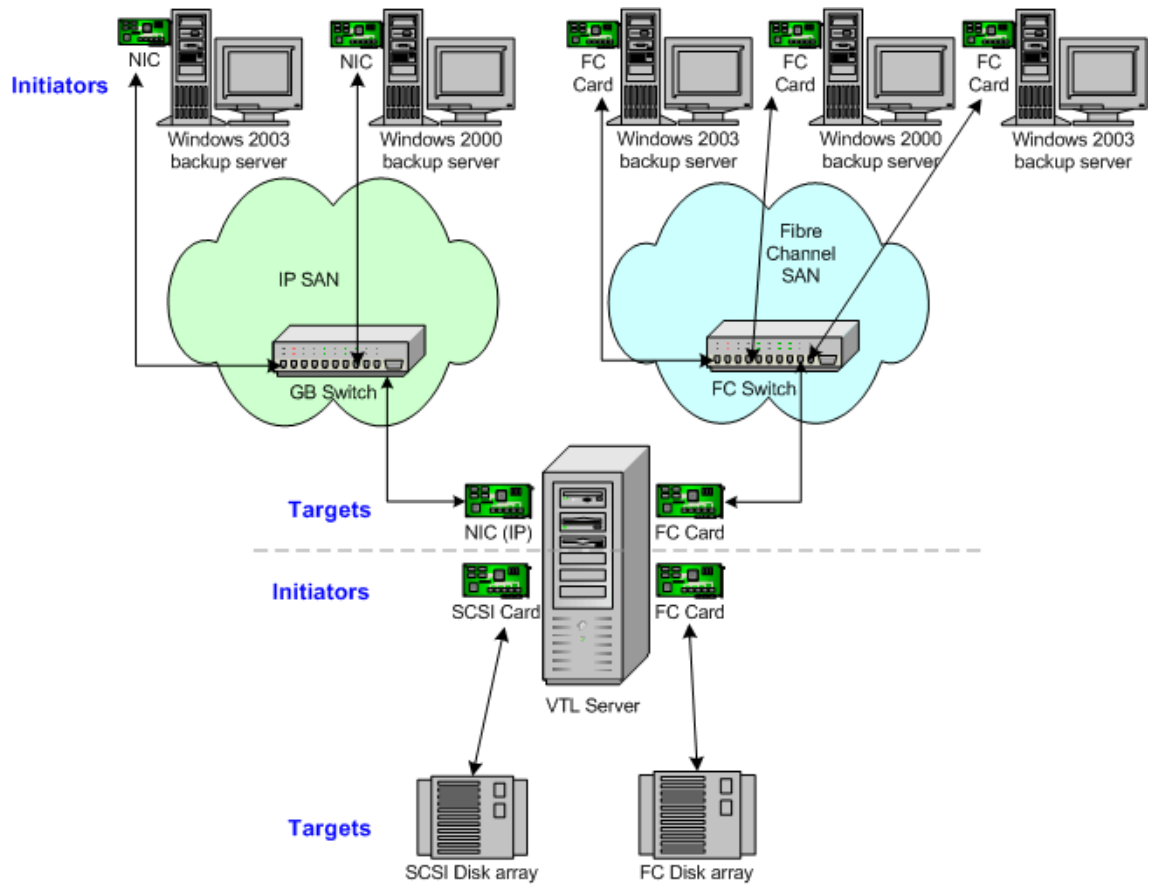
Log the client onto the target

Refer to the documentation provided by the vendor of your iSCSI initiator.

Fibre Channel Configuration

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.

Note: Fibre Channel switches can use Registered State Change Notification (RSCN) suppression to manage registered state change notifications.

By default, RSCN suppression is disabled on all ports. RSCN suppressed off/ disabled means that device changes on the port generate an RSCN to all other end devices that are zoned with this one. Do not change the RSCN default settings, otherwise, any device changes on a port will not generate an RSCN to any other end device, which can result in FC connectivity issues.

Configure Fibre Channel hardware on server

VTL supports the use of QLogic HBAs for the VTL server.

Ports

Your VTL appliance will be equipped with several Fibre Channel ports. Some of these ports will interface with storage arrays while other 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*.

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.
- 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.
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.
General zoning requirements	<p>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.</p> <p>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.</p>

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 StorSafe appliance is rebooted (or StorSafe 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, storage array, or a host client 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.

Configure Fibre Channel hardware on clients

Fibre Channel connection to upstream clients can be via a direct connection or an FC switch.

For a direct connection, make sure topology for each HBA port in `fshba.conf` on the server matches the topology of the associated HBA port on the storage, such as the data rate. The setting in `fshba.conf` is called `connection_option`. Both *Point-to-Point* and *Loop* topologies are supported; we recommend *Point-to-Point* topology. IBM i clients do not support direct connections using FC 16 Gb or higher HBAs; use a FC switch or iSCSI instead.

For all clients, *except* Solaris SPARC, using a Switched Fabric topology, we recommend *Point-to-Point* topology for each HBA that logs into the switch.

We recommend hard coding the link speed of the HBA to be in line with the switch speed.

Configure Fibre Channel on storage

Fibre Channel connection to downstream storage can be a direct connection or can be via a FC switch. For a direct connection, make sure the topology for each HBA port in `fshba.conf` on the server matches to the topology of the associated HBA port on the storage. The setting in `fshba.conf` is called `connection_option`.

Load balance paths for downstream storage

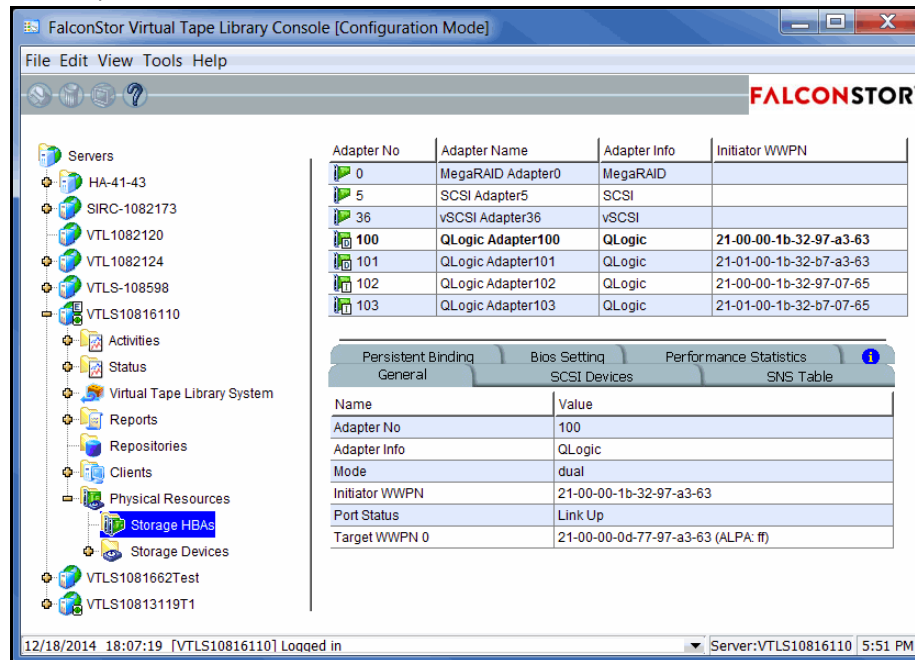
For optimal performance, if you have more than one path available for your storage LUNs, you can set VTL to evenly distribute I/O between all storage LUNs. To do this:

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.

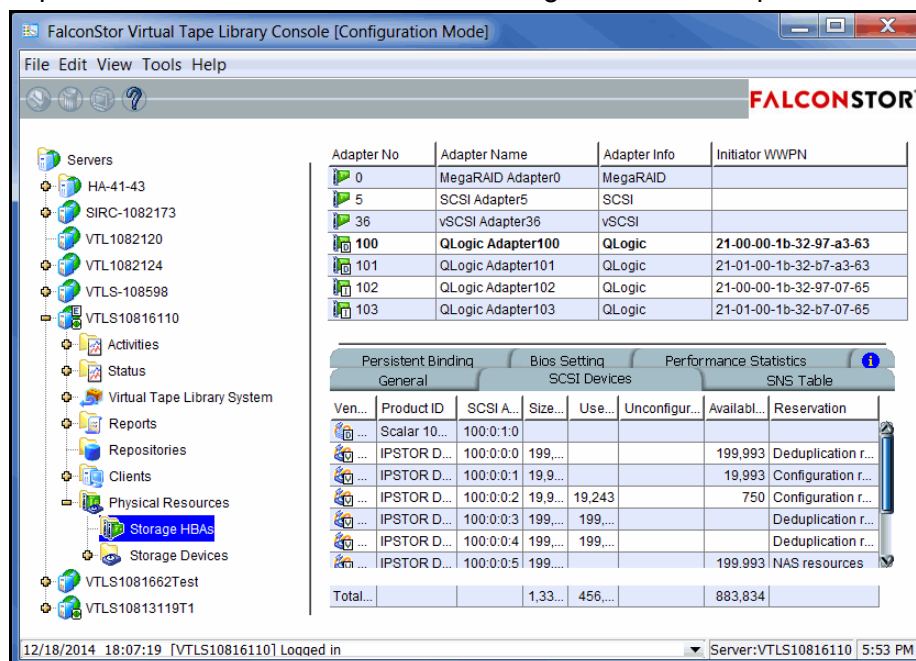
Verify your hardware configuration

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.

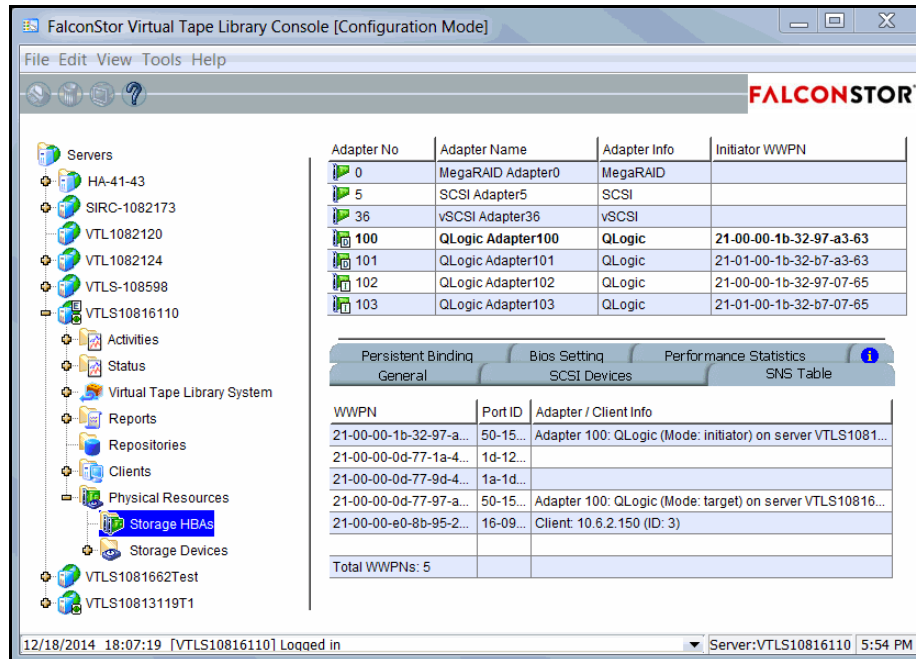


SCSI Devices tab The *SCSI Devices* tab lists the SCSI storage devices attached to this adapter. If you expect to see a device that is not listed, right-click the adapter and select *Rescan*.



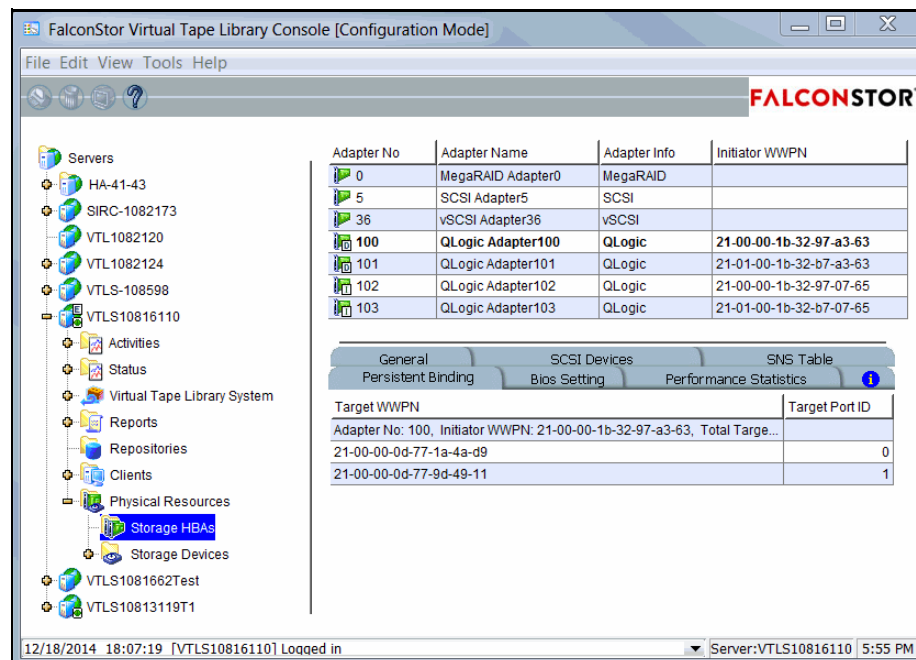
SNS Table tab

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*.

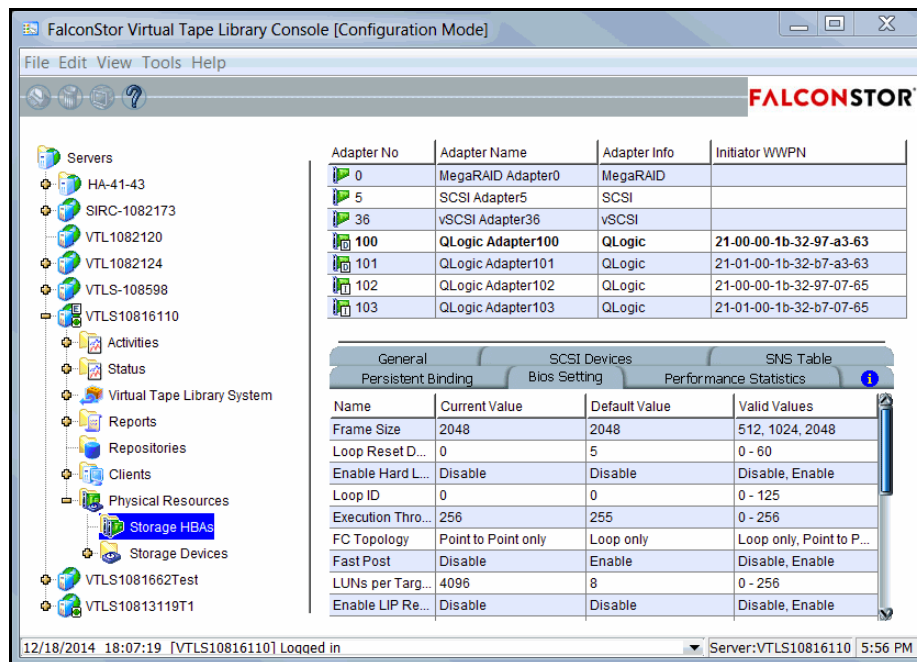


Persistent Binding tab

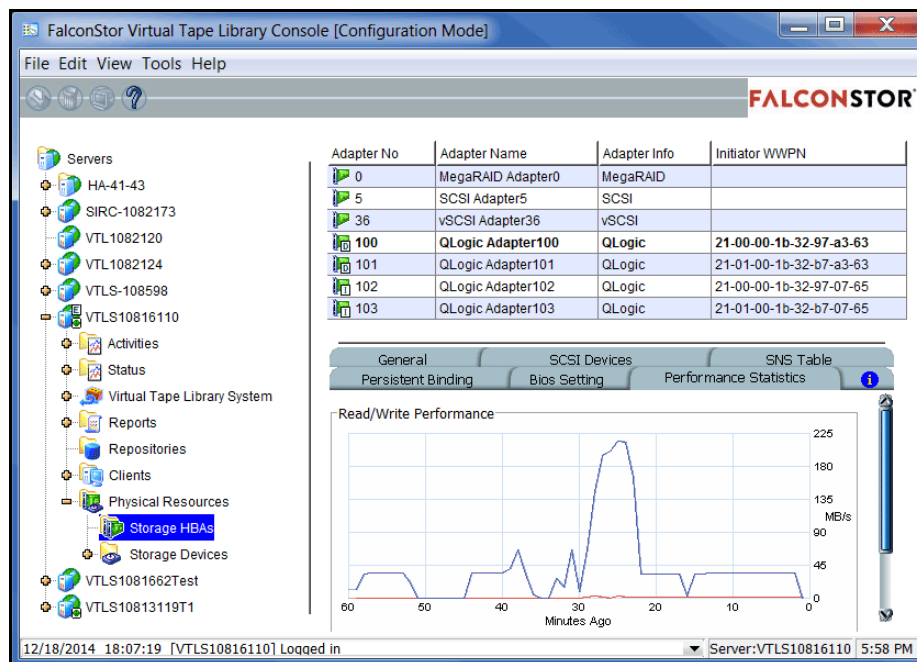
(Initiator ports only) The *Persistent Binding* tab lists all of the target ports to which this adapter is bound.



Bios Setting tab The *Bios Setting* tab lists all of the HBA settings for this adapter so that you can confirm what is set.



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.



Note that SIR replication is not included in the performance statistics.

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.

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.

The screenshot shows the FalconStor Virtual Tape Library Console in Configuration Mode. The left sidebar displays a tree view of the system hierarchy, including Servers, Physical Resources, and Storage HBAs. The main area shows a table of adapters and a configuration panel for the selected QLogic Adapter100.

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

Persistent Binding		Bios Setting		Performance Statistics	
General		SCSI Devices		SNS Table	
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

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.

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]

Associate World Wide Port Names with clients

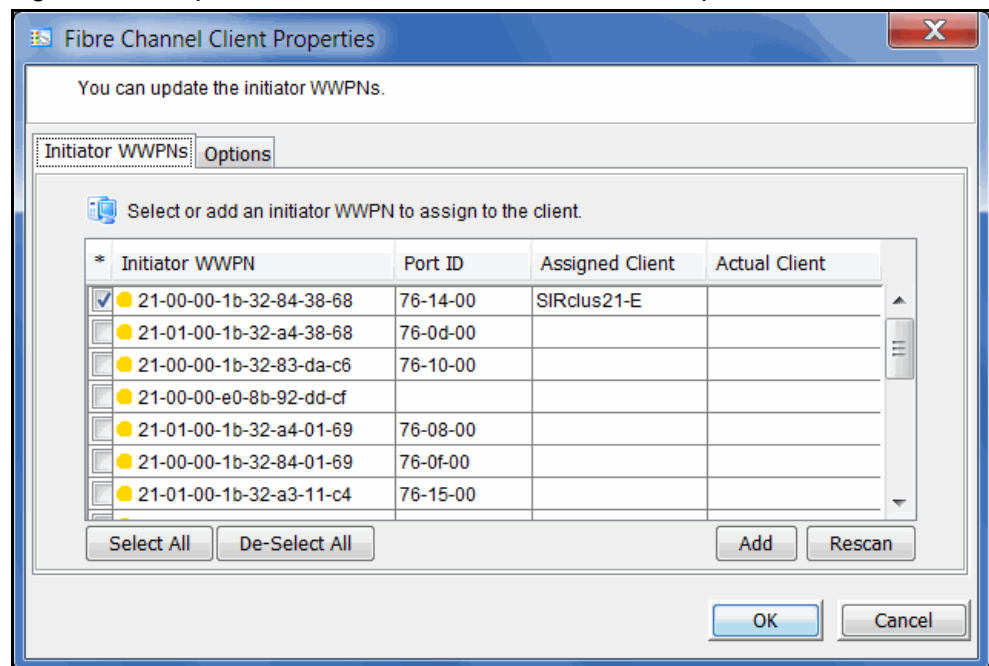
Similar to an IP address, the WWPN uniquely identifies a port in a Fibre Channel environment. Unlike an IP address, the WWPN is vendor assigned and is hardcoded and embedded.

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*.



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.




IBM i System Configuration

Overview

IBM i System 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.

You can connect an IBM i system to your VTL appliance, which emulates IBM libraries as FALCON TS3500L32 (03584L32) or FALCON TS3500L32 (03584L32) with media types ULTRIUM3 (LT03) or newer.



Before you begin

Before you can use Virtual Tape Library for IBM i systems, your environment must meet the following criteria:

- You must be using a supported version of IBM Backup Recovery & Media Services. Refer to the FalconStor Certification Matrix for a list of supported versions.
- There must be an iSCSI connection between the IBM i system host and the VTL appliance. For more information about iSCSI configuration, refer to ["iSCSI Configuration"](#).
- IBM i systems and the VTL appliance must use certified FC HBAs or iSCSI initiators. Refer to the VTL release notes for a list that are currently certified.

Set up the tape library

With Virtual Tape Library for IBM i systems, you use the procedures described earlier in this guide to create a virtual tape library and assign it to an IBM i system host.

Note: For IBM i clients, the virtual tape library type must be FalconStor FALCON TS3500L32 (03584L32) or FALCON TS3500L32 (03584L32) and the media type must be ULTRIUM3 (LT03) or newer.

Once you have created a virtual tape library and assigned it to the host, you should perform the following tasks to ensure that the IBM i system can see and properly work with the library. (For more information about working with an IBM i system, refer to your IBM i system documentation.)

1. At the IBM 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

Cartridges that have been removed from the tape library inventory are referred to as *exported*.

To export a cartridge and move it to the vault, 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.

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.

Server modules

When you run a server command, you will see a list of modules. The list will vary, depending upon which options you are using. A description of each module is listed below.

Module	Description
Authentication Module	Manages authentication requests between replica servers.
Block Device Module	Represents a generic block-to-SCSI driver that provides the SCSI interface for VTL to access non-SCSI block devices.
Base Module	Provides basic memory management and SCSI device management to IO modules.
CLI Proxy Module	Facilitates communication between the CLI utility and the VTL server.
Communication Module	Handles console-to-server communication and manages overall system configuration information.
Compression Module	Provides software compression for replication and deduplication.
Deduplication Module	Provides the repository server.
Email Alerts Module	Sends alerts via email to indicate alarming situations.
Event Module	Provides message logging interface to the system log.
FC Initiator Module	Represents the QLogic Fibre Channel initiator module, which provides interaction between the VTL server and the FC storage.
FC Target Module	Provides Fibre Channel target functionality.
FUSE Module	Represents 'File system in Userspace (FUSE)' module to allow creation of file systems in the user space without editing the Linux kernel code.
IO Core Module	Provides core IO services.
iSCSI Target Module	Provides iSCSI target functionality via a network adapter.
iSCSI Daemon	Represents a user daemon, which handles the login process to VTL iSCSI targets from iSCSI initiators.
Logger Module	Provides the information logging function for VTL reports.
Memory Map Module	Allows mapping of physical memory from kernel space to user space.
Node Manager	Facilitates communication among VTL servers.
Path Manager Module	Manages I/O traffic over multiple paths to multi-path storage devices.

Module	Description
Reclamation Triggering Module	Monitors deduplication usage to trigger reclamation at pre-set thresholds or when scheduled.
SNMPD Module	Interacts with SNMP management software to reply to MIB browsing queries and to send SNMP traps for abnormal situations.
Statistics Database Daemon	Provides deduplication statistics.
TCP Stream Replication Source Module	Provides outgoing deduplication unique data replication support.
TCP Stream Replication Target Module	Provides incoming deduplication unique data replication support.
TLE Core Module	Provides the tape drive/library emulation.
TLE Upcall Kernel Module	Represents the kernel module, handling interactions between kernel mode and user mode components, which manage tape operations.
TLE Upcall User Module	Represents the user module, handling interactions between kernel mode and user mode components, which manage tape operations.
Transport Module	Handles I/O transport for replication.
Upcall Module	Handles interactions between kernel mode and user mode components.
Virtual Device FSVSHOST Module	Provides device access to the repository.
Virtual SCSI Device Driver	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 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 FalconStor 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.

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](#)
- [Object Storage Migration Jobs](#)
- [Replication Status](#)
- [Virtual Library and Drive Assignment](#)
- [Virtual Library Information](#)
- [Virtual Tape Activity](#)
- [Virtual Tape Information](#)

Usage

- [Deduplication - Tape Usage](#)
- [Deduplication Repository - Memory and Space Usage](#)
- [Disk Space Usage History](#)

- [LUNs](#)
- Allocation
 - [Disk Space Allocation for Virtual Tapes in Libraries](#)
 - [Physical Resource Allocation](#)
- Configuration
 - [Fibre Channel Adapters Configuration](#)
 - [Physical Resources Configuration](#)
 - [Storage Pools Configuration](#)
- Performance
 - [VTL Performance](#)

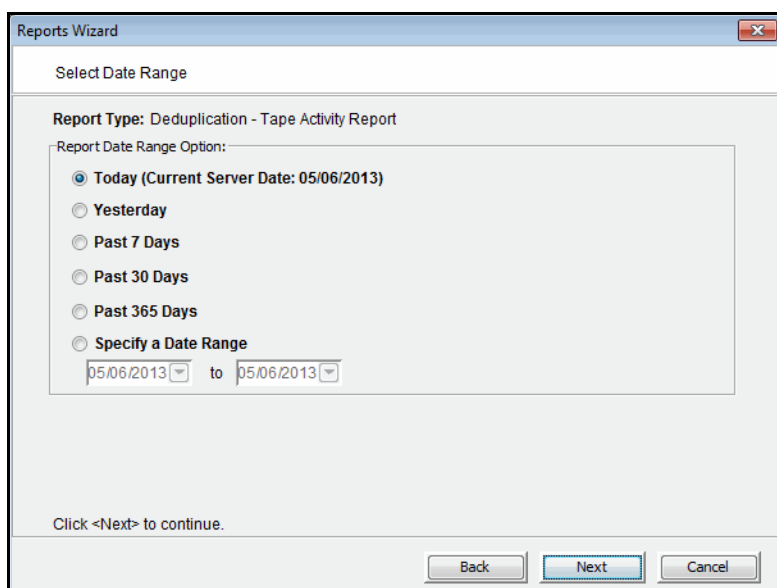
Create a report

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.



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 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, job type, status, 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).

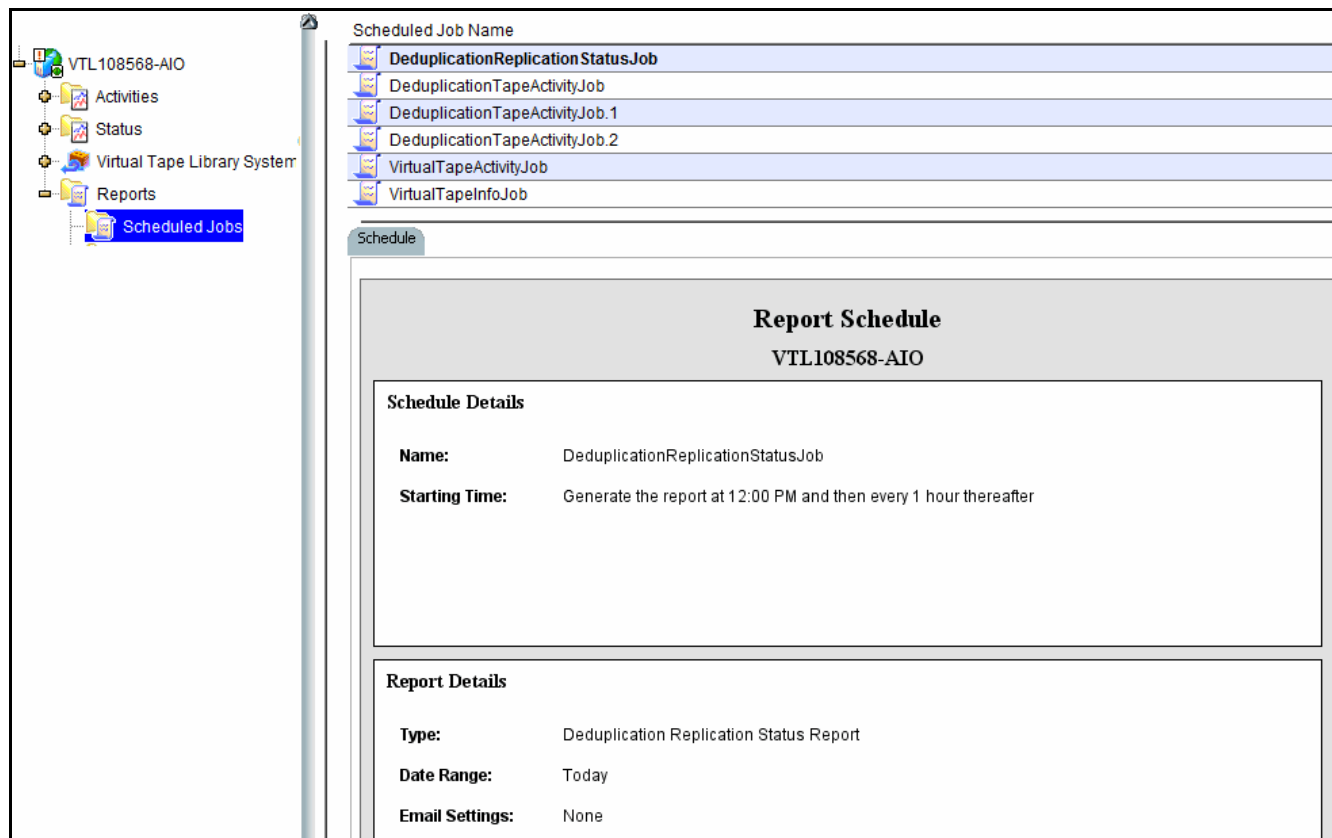
4. Depending upon which report you select, additional windows appear to allow you to filter the information it will include.

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.



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.

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.

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.

- On the *Retention* tab, specify how long generated reports should be retained.



- On the *Other* tab, define how a day is defined in your environment: from midnight to midnight, or from noon to noon.



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.

Information reports

Deduplication - Policy Status

This 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 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*.

Deduplication - Tape Activity

This 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.

The screenshot shows the 'Reports Wizard' dialog box with the following configuration:

- Report Type:** Deduplication - Tape Activity Report
- Tape/Job Options:**
 - Tape Barcode Range: From 'First' to 'Last'
 - Job Status: 'Complete'
 - Group Together by Job Status
 - Show as Tape List
 - Display Latest Jobs Only
 - Show the Last Completed Job in the Policy
 - Show the Jobs in the Last Policy Run
 - Sort by End Time
 - Include Active Tapes

Buttons at the bottom: Back, Next, Cancel.

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	Deduplication - Tape Activity Report		05/06/2013 21:28
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
<i>Include Policies (ID): All</i>			
<i>Filter: show active tapes</i>			

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	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	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	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
<i>Note: Includes failed jobs</i>			


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.



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.

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

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.

Deduplication Repository - Reclamation

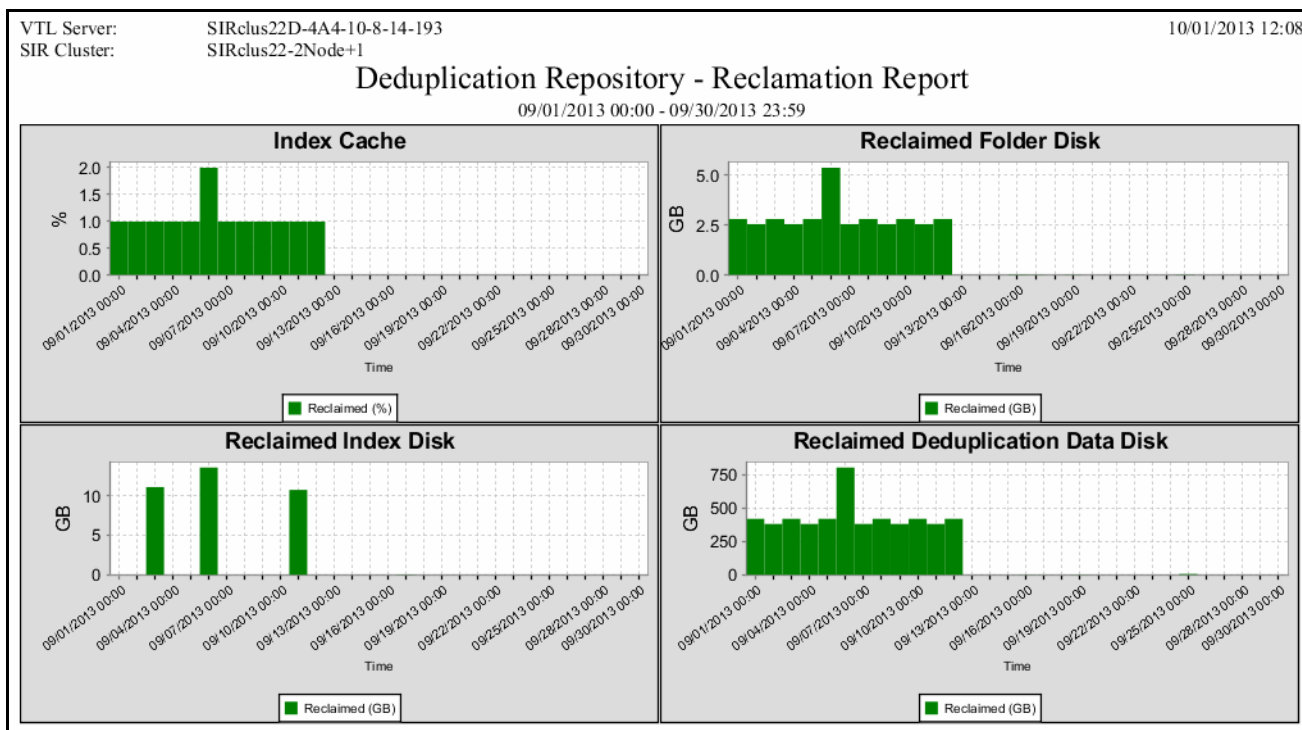
The Repository Dashboard Summary tab displays information about repository reclamation. This status report is an alternate way to view the reclamation information on these tabs.

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.

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.



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.

Import/Export Jobs

This report lists all import/export jobs that were placed in the queue during the specified period of time, regardless of job status. Options in the wizard include job type and status.

You must select at least one job type and one job status.

- | | |
|------------|--|
| Job type | <ul style="list-style-type: none">• Export to Standalone Drive - For these jobs, you can include results for jobs that used <i>Copy Mode</i>, <i>Move Mode</i>, or both.• Import from Standalone Drive - For these jobs, you can include results for jobs that used <i>Copy Mode</i>, <i>Recycle Mode</i>, or both. |
| Job status | For all selected job types, the report will include information on jobs with the selected types of status. |

The summary page displays the number of jobs found for all job types and all types of job status.

The detail pages display results based on the job type(s) and types of job status 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, job throughput, and an indicator of whether tapes were encrypted.

Object Storage Migration Jobs

This report lists all object storage migration and restoration jobs that were run during the specified period of time. Options in the wizard include job type and status.

You must select at least one job type and one job status.

The summary page displays the number of jobs found for each job type and job status.

OBD190		Migration Job Report					01/03/2019 14:37
		12/04/2018 00:00 - 01/02/2019 23:59					
	Running	Failed	Completed	Cancelled	On Hold		
Migration	0	8	44	17	0		
Recovery	0	0	7	6	0		

The detail pages display results based on the job type(s) and types of job status you selected. Jobs are ordered by job ID. For each job, the report lists the job type, barcodes of source/destination tapes; locations of source/destination tapes, job status, start/end time of job, amount of data transferred, job throughput, and an indicator of whether tapes were encrypted.

OBD190		Migration Job Report						01/03/2019 14:37	
		12/04/2018 00:00 - 01/02/2019 23:59							
Job Id	Type	From Tape / To Tape	From Location / To Location	Status	Start Time / End Time	Transfer (MB)	Throughput (MB/sec)	Encrypted	
85	Migration	081E000L	VLIB: 2078 Slot: 0	Cancelled	12/10/2018 16:27 12/10/2018 16:31	290	1	No	
86	Migration	081E000L	VLIB: 2078 Slot: 0	Cancelled	12/10/2018 17:21 12/10/2018 17:24	492	2	No	
87	Migration	081E000L	VLIB: 2078 Slot: 0	Cancelled	12/10/2018 17:33 12/10/2018 17:36	0	0	No	
88	Migration	081E000L	VLIB: 2078 Slot: 0	Failed	12/10/2018 17:40 12/10/2018 17:41	1,394	19	No	
89	Migration	081E000L	VLIB: 2078 Slot: 0	Failed	12/10/2018 18:32 12/10/2018 19:02	534	0	No	
90	Migration	081E000L	VLIB: 2078 Slot: 0	Failed	12/10/2018 21:30 12/10/2018 21:36	625	2	No	
91	Migration	081E000L	VLIB: 2078 Slot: 0	Failed	12/10/2018 21:59 12/10/2018 21:59	491	22	No	
92	Migration	081E000L	VLIB: 2078 Slot: 0	Failed	12/10/2018 22:17 12/10/2018 22:19	5,446	54	No	
93	Migration	081E000L	VLIB: 2078 Slot: 0	Completed	12/10/2018 22:23 12/10/2018 22:24	10,000	140	No	
94	Migration	081E000M	VLIB: 2078 Slot: 0	Completed	12/11/2018 12:48 12/11/2018 12:48	10	5	No	
95	Migration	081E000M	VLIB: 2078 Slot: 0	Completed	12/11/2018 12:54 12/11/2018 12:54	1,000	83	No	
96	Migration	081E000M	VLIB: 2078 Slot: 0	Cancelled	12/11/2018 14:58 12/11/2018 14:58	4,259	137	No	
97	Migration	081E000M	VLIB: 2078 Slot: 0	Cancelled	12/11/2018 15:02 12/11/2018 15:02	695	38	No	

Type: All
Status: All

5 / 9

Replication Status

This 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 includes 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		Replication Status Report						05/06/2013 23:12
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 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, 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, and assigned 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 lists the name of all devices on the system and for each, displays the type of device (library or drive), ID, vendor, product ID and serial number.

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 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.



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 report shows activity for all virtual tapes within the specified period of time for three types of operations: Backup, Migration to Object Storage, and Recovery from Object Storage.

You can filter report data by specifying barcode information: tapes within a specified barcode 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 of 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.

There is no compression ratio for object migration jobs. For object reconstruction jobs, the compression ratio is only meaningful when either VTL or the backup software compressed the data because data from object storage will be compressed by VTL during recovery.

Typical displays include:

- For backup operations:

H12-202		12/27/2018 22:51				
Virtual Tape Activity Report						
12/20/2018 00:00 - 12/26/2018 23:59						
<i>Operation Type: Backup</i>						
Start Time	End Time	Barcode	Data Compression Ratio	Duration (H:M:S)	Performance (MB/s)	
12/26/2018 18:06	12/26/2018 18:06	02B7003P	3.37:1	0:00:02	64.00	
12/26/2018 18:04	12/26/2018 18:05	02B7003P	1.00:1	0:01:19	51.85	
12/25/2018 15:19	12/25/2018 15:21	004F0007	1.00:1	0:02:20	42.76	
12/25/2018 15:17	12/25/2018 15:19	00420004	1.02:1	0:01:42	40.23	
12/25/2018 10:45	12/25/2018 10:45	02B7003O	3.37:1	0:00:02	64.00	
12/25/2018 10:35	12/25/2018 10:35	02B7003N	3.37:1	0:00:04	128.00	
12/24/2018 16:24	12/24/2018 16:27	004F0006	1.00:1	0:02:09	46.38	
12/24/2018 12:22	12/24/2018 12:25	004F0004	1.00:1	0:03:11	52.23	
12/24/2018 11:43	12/24/2018 11:45	004F0004	1.00:1	0:02:07	47.11	
12/21/2018 16:09	12/21/2018 16:09	02B7003M	3.37:1	0:00:03	170.67	
12/26/2018 10:03	12/26/2018 10:07	78650001	N/A	0:04:04	192.45	
12/25/2018 17:20	12/25/2018 17:40	78650000	N/A	0:19:57	19.69	
12/25/2018 14:53	12/25/2018 14:57	00420004	N/A	0:03:04	22.28	
12/25/2018 14:53	12/25/2018 14:57	004F0007	N/A	0:03:59	25.03	

- For object storage migration operations:

H12-202		12/27/2018 22:51				
Virtual Tape Activity Report						
12/20/2018 00:00 - 12/26/2018 23:59						
Operation Type: Migration						
Start Time	End Time	Barcode	Data Compression Ratio	Duration (H:M:S)	Performance (MB/s)	
12/26/2018 18:10	12/26/2018 19:02	02B7003P	N/A	0:51:54	1.33	
12/26/2018 18:07	12/26/2018 18:08	02B7003P	N/A	0:01:10	1.34	
12/25/2018 15:29	12/25/2018 15:42	004F0007	N/A	0:13:29	14.80	
12/25/2018 15:21	12/25/2018 15:33	00420004	N/A	0:12:49	10.47	
12/25/2018 15:07	12/25/2018 15:12	004F0007	N/A	0:05:49	17.14	
12/25/2018 14:58	12/25/2018 15:02	00420004	N/A	0:03:52	17.67	
12/25/2018 14:23	12/25/2018 14:28	004F0006	N/A	0:04:07	24.23	
12/25/2018 14:18	12/25/2018 14:35	02B70031	N/A	0:16:28	1.23	
12/25/2018 10:46	12/25/2018 10:47	02B7003O	N/A	0:00:31	1.23	
12/25/2018 10:35	12/25/2018 10:37	02B7003N	N/A	0:02:04	1.23	
12/24/2018 13:51	12/24/2018 14:02	004F0004	N/A	0:10:18	25.82	
12/24/2018 11:56	12/24/2018 12:01	004F0004	N/A	0:04:29	22.25	
12/21/2018 16:11	12/21/2018 16:12	02B7003M	N/A	0:00:59	2.58	

- For object storage recovery operations:

H12-202		12/27/2018 22:51				
Virtual Tape Activity Report						
12/20/2018 00:00 - 12/26/2018 23:59						
Operation Type: Recovery						
Start Time	End Time	Barcode	Data Compression Ratio	Duration (H:M:S)	Performance (MB/s)	
12/25/2018 15:59	12/25/2018 16:04	004F0007	1.00:1	0:05:34	35.84	
12/25/2018 15:59	12/25/2018 16:03	00420004	1.02:1	0:04:03	33.76	
12/25/2018 15:50	12/25/2018 15:55	004F0007	1.00:1	0:05:22	37.17	
12/25/2018 15:50	12/25/2018 15:53	00420004	1.02:1	0:03:43	36.78	
12/25/2018 15:14	12/25/2018 15:16	004F0007	1.00:1	0:02:28	40.43	
12/25/2018 15:09	12/25/2018 15:11	00420004	1.02:1	0:02:12	31.06	
12/25/2018 14:19	12/25/2018 14:27	004F0004	1.00:1	0:07:55	33.60	
12/24/2018 16:55	12/24/2018 17:07	004F0004	1.00:1	0:11:25	23.30	
12/24/2018 16:47	12/24/2018 16:54	004F0004	1.00:1	0:06:13	26.68	
12/24/2018 16:42	12/24/2018 16:46	004F0004	1.00:1	0:03:19	27.90	
12/24/2018 16:28	12/24/2018 16:38	004F0004	1.00:1	0:09:42	27.42	
12/24/2018 15:59	12/24/2018 16:12	004F0004	1.00:1	0:12:35	21.14	
12/24/2018 15:20	12/24/2018 15:35	004F0004	1.00:1	0:15:06	17.61	
12/24/2018 12:02	12/24/2018 12:06	004F0004	1.00:1	0:03:22	29.62	
12/24/2018 10:53	12/24/2018 10:54	02B7003L	3.37:1	0:01:16	53.89	
12/21/2018 14:28	12/21/2018 14:44	02B7003L	3.37:1	0:15:39	1.09	

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, Replica Resources View, Vault View, Detailed Tape View, and Migration 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; encryption status, WORM status, whether migration, deduplication, and replication are required; tape location (library/vault); and the deduplication policy (if any) to which the tape belongs.

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.

Virtual Tape Information Report											
05/07/2013 00:00 - 05/07/2013 22:06											05/07/2013 22:06
Deduplication View											
Deduplication Policy	Barcode	Needs Dedupe	Needs Replication	Written (GB)	Last Successful Deduplication		Last Successful Replication		VIT Size (GB)	Ratio	
					Started	Finished	Data Processed (GB)	Unique Data (GB)			
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	0.000	- 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	0.000	- 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	0.000	- 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

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, and whether the tape is a FVIT or LVIT.

Virtual Tape Information Report							
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
00840QL6	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

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 migration, deduplication, or replication is required, the parent library, and the deduplication policy (if any) to which it belongs.

Virtual Tape Information Report									
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	
0124000ORIG	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

Detailed Tape View This view presents information for tapes in a specific library, including bar code, tape ID, tape location (library/vault), whether or not the tape is full, WORM status, current allocation, maximum capacity, used size, amount of data written, and number of used segments.

Virtual Tape Information Report										
03/11/2019 00:00 - 03/11/2019 23:21										
Tape Detail View										
Barcode	Tape ID	Location	Tape Full	WORM Tape	Current Allocation (GB)	Maximum Capacity (GB)	Used Size (GB)	Written (GB)	Used Segments	
004A0009	10000537	HP-ESL G3-00074	Yes	Yes	5.000	5.000	4.828	4.828	1	
004A000A	10000538	HP-ESL G3-00074	No	Yes	165.000	5,100.000	122.815	122.811	3	
004A000B	10000539	HP-ESL G3-00074	No	Yes	1.000	5,100.000	0.699	58.523	1	
004A000C	10000540	HP-ESL G3-00074	No	Yes	5.000	5,100.000	0.001	0.001	1	
00760000	10000361	HP-MSL2024-00118	No	No	5.000	5.000	2.295	2.295	1	
00760001	10000362	HP-MSL2024-00118	No	No	5.000	5.000	3.061	3.061	1	
004F00L7	10000311	IBM-ULT3583-TL-00079	No	No	5.000	5,100.000	4.729	4.844	2	
004F01L7	10000312	IBM-ULT3583-TL-00079	No	No	1.000	5,100.000	0.191	14.539	1	
004F02L7	10000421	IBM-ULT3583-TL-00079	No	Yes	1.000	5,100.000	0.093	7.922	1	
004F03L7	10000422	IBM-ULT3583-TL-00079	No	Yes	0.000	5,100.000	0.000	0.000	0	
004F04L7	10000423	IBM-ULT3583-TL-00079	Yes	Yes	5.000	5.000	4.828	4.828	1	
004F05L7	10000570	IBM-ULT3583-TL-00079	No	No	165.000	5,100.000	115.568	115.563	2	
004F07L7	10000571	IBM-ULT3583-TL-00079	No	No	5.000	5,100.000	4.828	4.827	1	
004F08L7	10000541	IBM-ULT3583-TL-00079	No	Yes	5.000	5,100.000	0.016	0.015	1	
004F09L7	10000542	IBM-ULT3583-TL-00079	No	Yes	0.000	5,100.000	0.000	0.000	0	
004F0AL7	10000543	IBM-ULT3583-TL-00079	No	Yes	0.000	5,100.000	0.000	0.000	0	
004F0BL7	10000544	IBM-ULT3583-TL-00079	No	Yes	0.000	5,100.000	0.000	0.000	0	
00390000	10000013	STK-9730-00057	Yes	No	30.000	30.000	29.828	29.828	7	
00390001	10000014	STK-9730-00057	Yes	No	30.000	30.000	29.828	29.828	6	
00390002	10000015	STK-9730-00057	Yes	No	30.000	30.000	29.828	29.828	7	
00390003	10000016	STK-9730-00057	Yes	No	30.000	30.000	29.828	29.828	5	

Filters
Barcode: All
Included libraries(id): All
Included policies(id): All
Views: Overall Summary, Detailed Tape View ,

Migration View This view presents information for migrated tapes, including bar codes, data size, whether or not the tape is full, whether or not the tape is a stub tape, migration time, object storage, and object location.

H12-202		Virtual Tape Information Report					12/27/2018 05:21
12/27/2018 00:00 - 12/27/2018 05:21							
Migration View							
Barcode	Data size (GB)	Tape Full	Is stub	Migration time	Object Storage	Object Location	
00420004	8.011	No	No	12/25/2018 16:03	AWS S3	http://172.22.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.00420004-10003989	
004F0007	11.689	No	No	12/25/2018 16:04	AWS S3	http://172.22.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.004F0007-100039..	
02B70031	4.000	No	Yes	12/25/2018 14:35	AWS S3	us-east-1/595b5327-d535-4ede-8fe8-adde0e2f0131/02B70031-10003956	
02B70032	4.000	No	No	12/18/2018 17:14	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B70032-10003..	
02B70033	4.000	No	No	12/18/2018 17:30	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B70033-10003..	
02B70034	4.000	No	No	12/18/2018 23:22	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B70034-10003..	
02B70035	4.000	No	No	12/18/2018 18:07	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B70035-10003..	
02B70036	4.000	No	No	12/18/2018 16:51	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B70036-10003..	
02B70037	4.000	No	No	12/18/2018 21:43	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B70037-10003..	
02B70038	4.000	No	No	12/18/2018 21:09	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B70038-10003..	
02B70039	4.000	No	No	12/18/2018 17:59	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B70039-10003..	
02B7003A	4.000	No	No	12/18/2018 18:01	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B7003A-10003..	
02B7003B	4.000	No	No	12/18/2018 17:54	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B7003B-10003..	
02B7003C	4.000	No	No	12/18/2018 17:14	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B7003C-10003..	
02B7003D	4.000	No	No	12/18/2018 16:56	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B7003D-10003..	
02B7003E	4.000	No	No	12/18/2018 17:27	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B7003E-10003..	
02B7003F	4.000	No	No	12/18/2018 18:28	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B7003F-10003..	
02B7003G	4.000	No	No	12/18/2018 17:48	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B7003G-10003..	
02B7003H	4.000	No	No	12/18/2018 19:29	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B7003H-10003..	
02B7003I	4.000	No	No	12/18/2018 17:51	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B7003I-100039..	
02B7003J	4.000	No	No	12/18/2018 18:32	AWS S3	http://192.168.12.13:8000/595b5327-d535-4ede-8fe8-adde0e2f0131.02B7003J-10003..	

Filters
Barcode: All
Included libraries(id): All
Included policies(id): All
Views: Object Storage Migration View

Usage reports

Deduplication - Tape Usage

This 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 Veritas NetBackup or IBM Tivoli Storage Manager backup applications are being used, the report includes information about compressed data.
- **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							
Tape Name	Barcode	VIT	Capacity (MB)	Written (MB)	New (MB)	In SIR (MB)	Unique (MB)	Dedupe Ratio	CB	Client 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 *Deduplication Repository* tab displays 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.

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).

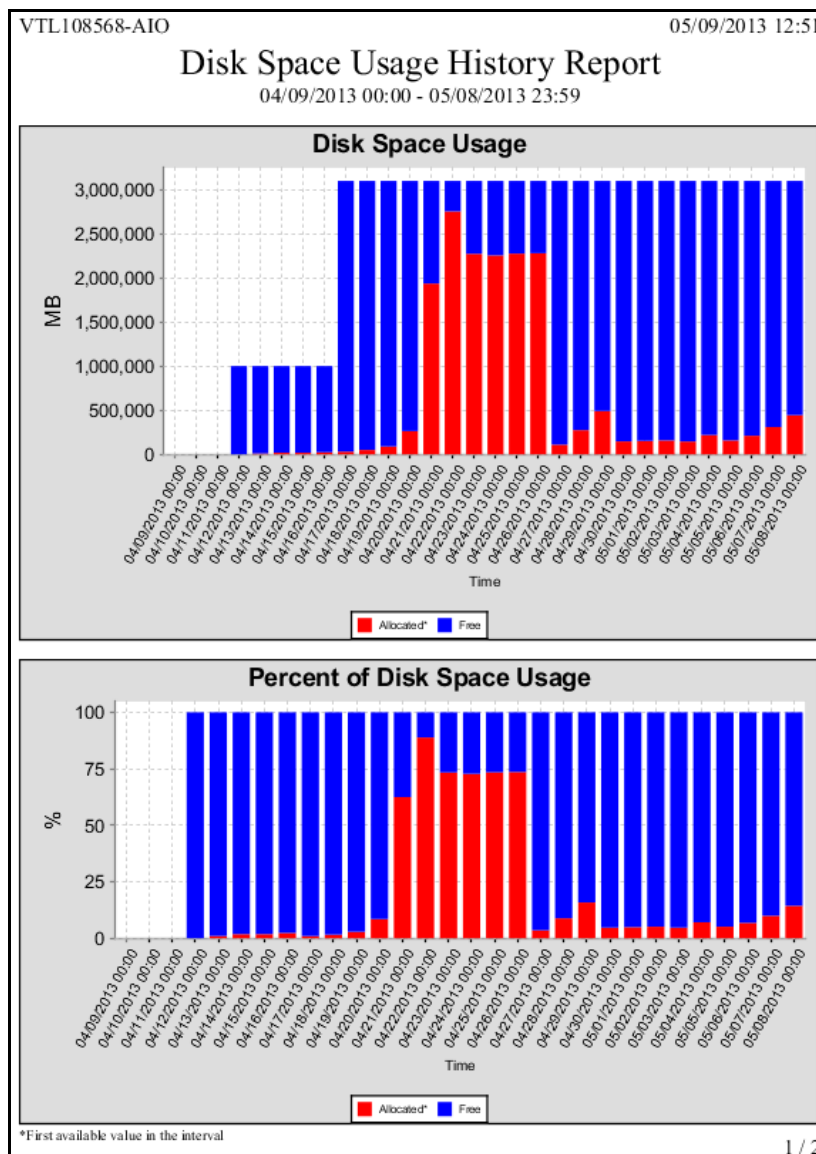
Note: If data cannot be collected for an interval, the report will derive the value based on the next good value.

In the bar charts, each bar represents the total amount of repository space or index cache used as of the beginning of the interval.



Disk Space Usage History

This 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.

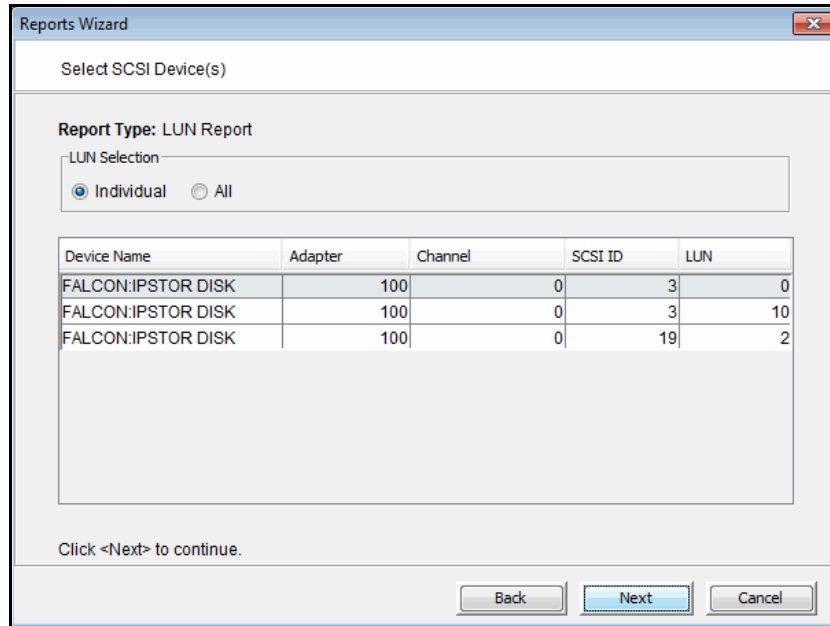


The results table includes a row of data for each interval in the graph: the start and stop time, the total capacity for all disks and the total allocated/free capacity expressed as a percentage of total capacity.

VTL108568-AIO		05/09/2013 12:51				
Disk Space Usage History Report						
04/09/2013 00:00 - 05/08/2013 23:59						
Start Time	Stop Time	Capacity*	Allocated*		Free	
		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 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).



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 (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 100:0:1 (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

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.



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

SCSI Address	Vendor ID	Product ID	Capacity	Allocated		Free			
			GB	GB	%	GB	%	Allocated	Free
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)		
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		

Total Disk Allocation
 System Physical Disk Capacity 3,024.54 GB

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.



Physical Resource Allocation

This 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. The report also includes the number of virtual tapes allocated on the device.

HA1062117119-A		Physical Resources Allocation Report								05/09/2013 10:56	
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 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 (dual, initiator, target) and port status 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		
<u>WWPN</u>	<u>Port ID</u>	<u>Switch Port</u>	<u>Adapter/Client Info</u>
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 resources, this 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, 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 or virtual tapes.

This report is available only as a one-time report.

Storage Pools Configuration

This report displays details about each storage pool, including details about each device in the storage pool. For each device, the name, SCSI address, total/used/available space is displayed.

This report is available only as a one-time report.


VTL107298		10/21/2016 15:01		
Storage Pools Configuration Report				
NASStoragePool	ID: 4			
Types	NAS			
Tag				
User ACL	ellen			
Devices	2			
Size	Total: 599,986 MB, Used: 0 MB, Available: 599,986 MB			
Device Name	SCSI Address	Total(MB)	Used(MB)	Available(MB)
FALCON:IPSTOR DISK	103:0:0:3	99,993	0	99,993
FALCON:IPSTOR DISK	103:0:0:4	499,993	0	499,993
TapesStoragePool	ID: 5			
Types	ALL			
Tag				
User ACL	fsadmin			
Devices	1			
Size	Total: 99,993 MB, Used: 0 MB, Available: 99,993 MB			
Device Name	SCSI Address	Total(MB)	Used(MB)	Available(MB)
FALCON:IPSTOR DISK	103:0:0:1	99,993	0	99,993

Performance report

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.



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			
<i>Device Type : VTL Server</i>			
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			
<i>Device Type : Adapter</i>			
<i>Device ID : 100</i>			
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			
<i>Device Type : LUN</i>			
<i>Device ID : 0:0:0:4</i>			
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			
<i>Device Type</i> : Client			
<i>Device ID</i> : 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					
<i>Device Type</i> : Virtual Tape Library					
<i>Device ID</i> : 27					
Start Time	End Time	Uncompressed Data (GB)		Compressed Data (GB)	
		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.

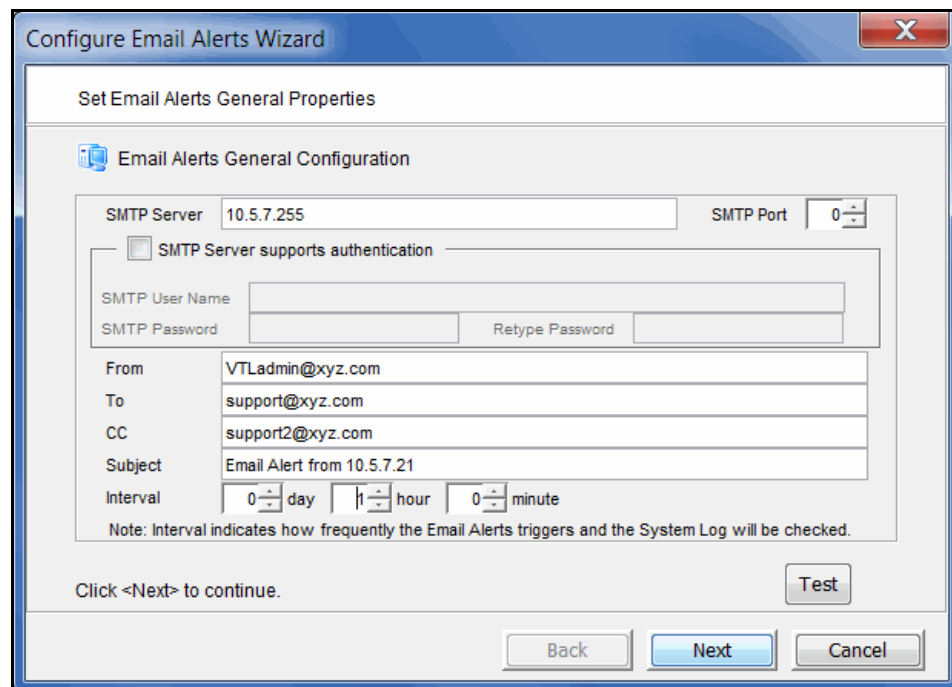
With *Email Alerts*, the performance and behavior of servers can be monitored so that system administrators are able to take corrective measures within the shortest amount of time, ensuring optimum service uptime and IT efficiency.

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.

1. In the console, right-click your server and select *Options --> Enable Email Alerts*.
2. Enter general information for your Email Alerts configuration.



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.

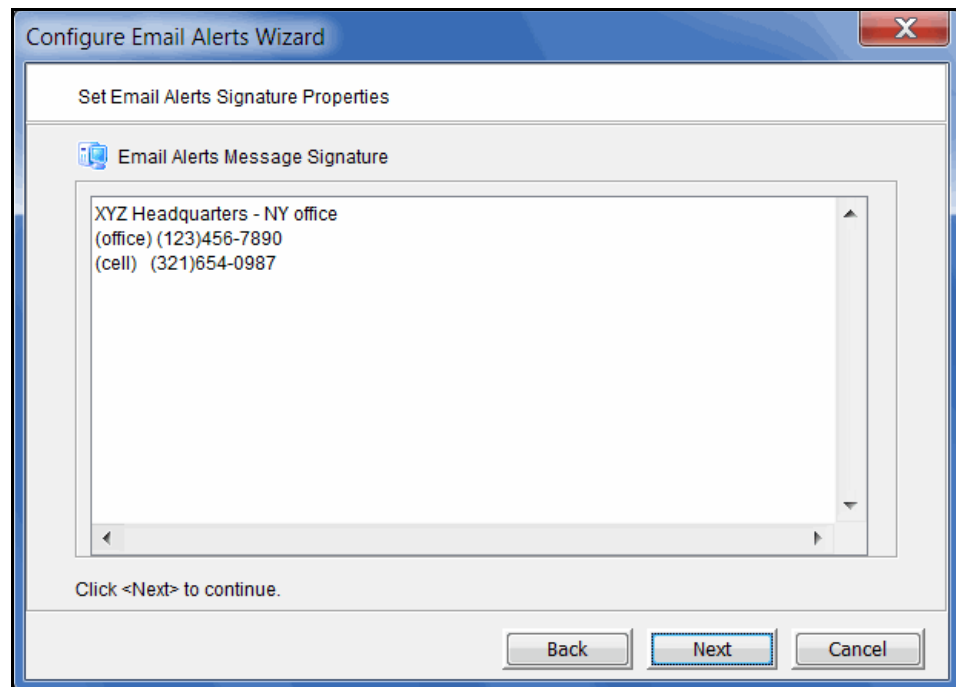
CC - Specify any other email accounts that should receive emails from Email Alerts.

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.

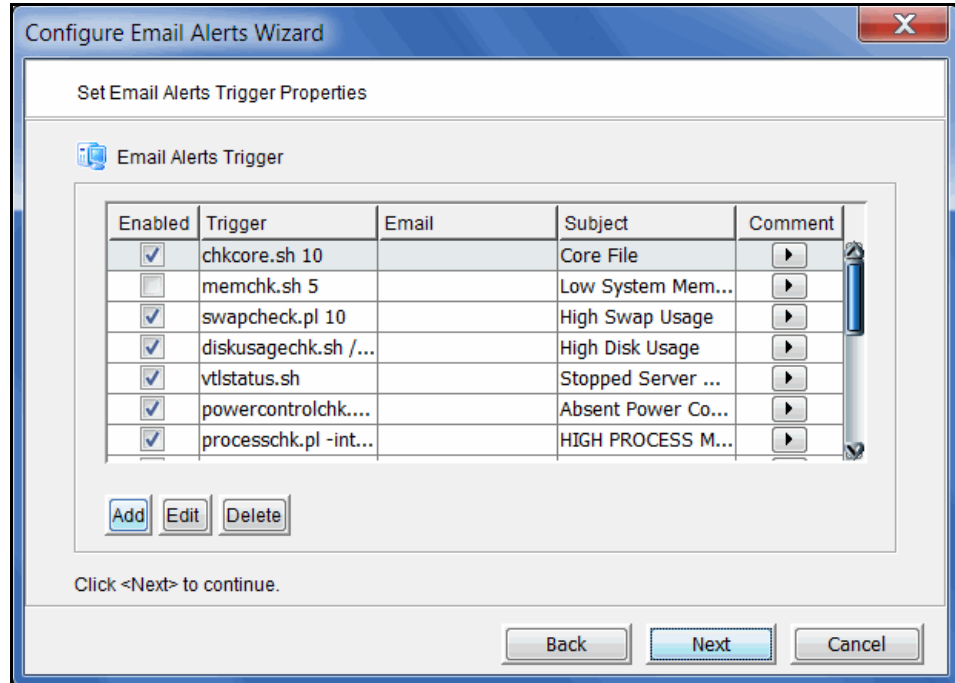
Interval - Specify how frequently the Email Alerts triggers should be checked.

Test - Click to test the configuration by sending a test email to the address defined in the *To* field.

3. Enter the contact information that should appear in each Email Alerts email.



4. Set the triggers that will cause Email Alerts to send an email.



Triggers are the scripts/programs that perform various types of error checking. By default, FalconStor 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.

serverstatus.sh (VTL status check) - Sends an alert if a server module has stopped.

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.

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.

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.

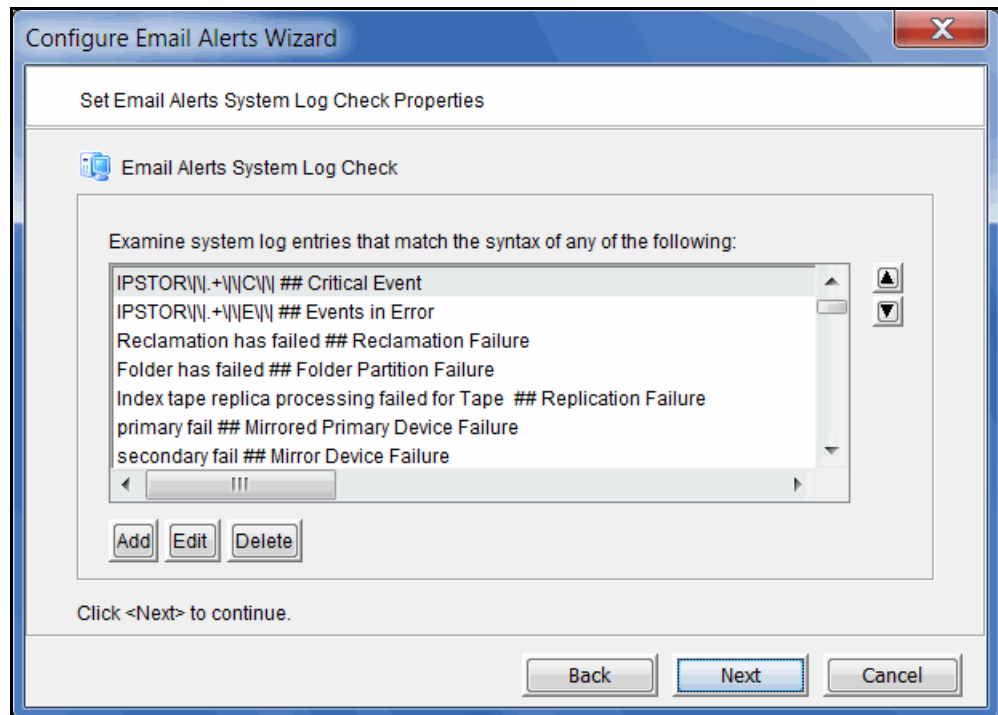
Some of the more common events checked in the system log are:

- Replication failure
- Storage path failure
- Mirror failure
- Mirror swap
- SCSI error
- Abandoned commands
- FC pending commands, busy FC, or FC loop down
- 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.

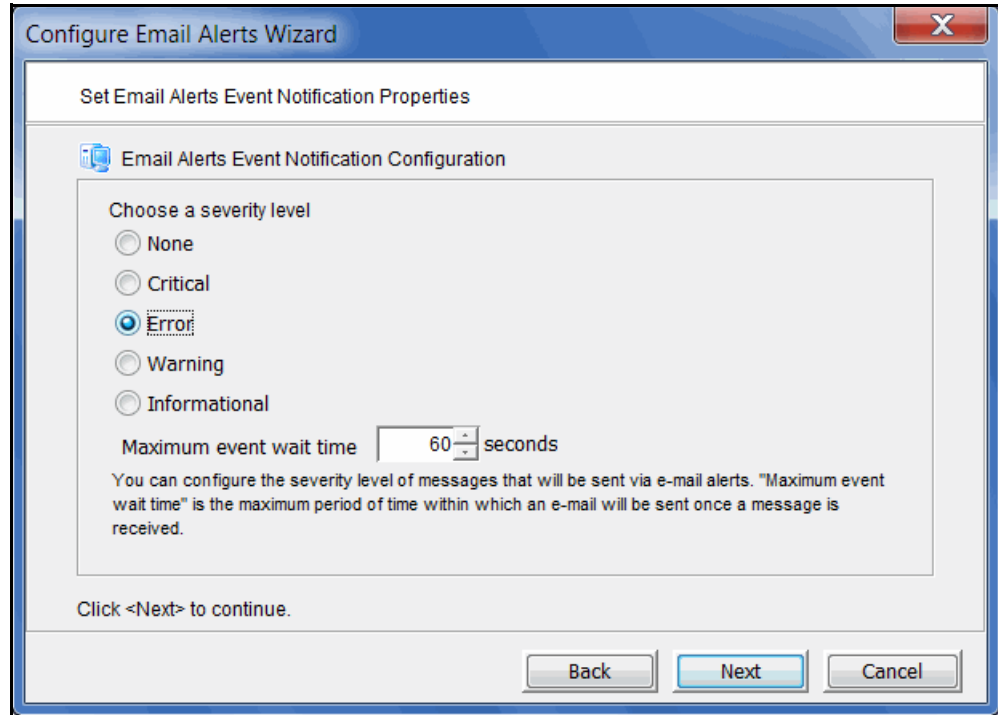


The system log records important events or errors that occur in the system, including those generated by VTL.

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.

- Indicate the severity level of event log messages that should be sent as email alerts.



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.

- 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.

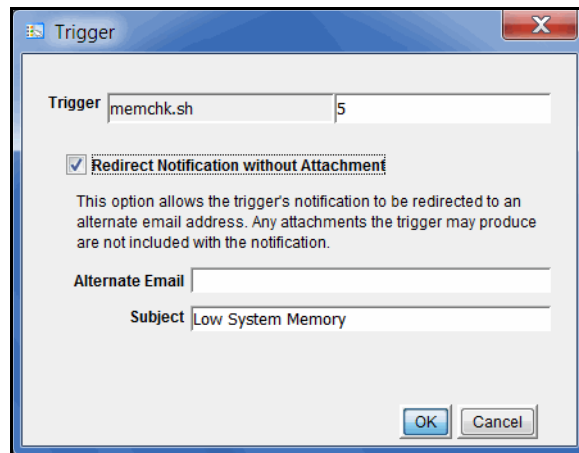
The default value is the same as the Email Alerts interval that was set on the first dialog (or the *General* tab if Email Alerts is already configured).

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*.



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, FalconStor 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, 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 >> $IPSTORCLHMLLOG
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 FalconStor 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.

Usage

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.

Depending upon the settings that were selected when your system was installed, you may only be allowed to run commands on your local server; you may not be allowed to run remote commands. You may also not be permitted to run certain commands, such as `adduser`.

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	--vdevicid	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.

Commands

The commands that are available for VTL servers are described below and on the following pages.

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.

Server 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 removelicence -s <server-name> [-u <username> -p <password>] -k <license-keycode>
[-X <rpc-timeout>]
```

```
iscon removelicence --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.

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 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:scsi 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

If *None* is selected, the device will not be allocated.

`-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

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.

Create storage pool

```
iscon createstoragepool -s <server-name> [-u <username> -p <password>]
-SP <storage-pool-name> [-D <physical-devices>] [-R <resource-list>]
[-X <rpc-timeout>]
```

```
iscon createstoragepool --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--storage-pool-name=<storage-pool-name> [--physical-devices=<physical-devices>]
[--resource-list=<resource-list>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a storage pool.

<storage-pool-name> is a unique storage pool name.

<physical-devices> can be a list of SCSI addresses separated by commas or a file enclosed in <> containing the SCSI address of each physical device to add on a separate line in the following format: adapter:channel:scsi-id:lun

- 99:0:1:0
- 99:0:1:1

Examples for this command argument:

- -D 99:0:1:0,99:0:1:1,99:0:1:2
- -D "<full_path_to_phydev_list_file>"

The device category for all the physical devices in the storage pool have to be the same.

The <resource-list> is used to restrict resource type(s) that can be created on this storage pool in the combination of the following values:

- Tape (VTL pool type)

The default (empty value) is to enable all the resource types, which means there is no restriction.

The resource types can be a list separated by comma, or a file enclosed in <> containing resource type in each line. For example:

- -R Tape
- -R "<full_path_to_res_list_file>"

-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 storage pool

```
iscon deletestoragepool -s <server-name> [-u <username> -p <password>]
-sp <storage-pool-id> | -SP <storage-pool-name> [-X <rpc-timeout>]
```

```
iscon deletestoragepool --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--storage-pool-id=<storage-pool-id> | --storage-pool-name=<storage-pool-name>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes the specified storage pool.

Either <storage-pool-ID> or <storage-pool-name> can 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.

Add physical devices to storage pool

```
iscon addpdevstostoragepool -s <server-name> [-u <username> -p <password>]
-sp <storage-pool-id> | -SP <storage-pool-name> -D <physical-devices>
[-l <lun-reservation>] [-X <rpc-timeout>]
```

```
iscon addpdevstostoragepool --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--storage-pool-id=<storage-pool-id> | --storage-pool-name=<storage-pool-name>
--physical-devices=<physical-devices> [--lun-reservation=<lun-reservation>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds physical devices to an existing storage pool.

Either <storage-pool-ID> or <storage-pool-name> can be specified.

<physical-devices> can be a list of SCSI addresses separated by commas or a file enclosed in <> containing the SCSI address of each physical device to add on a separate line in the following format: adapter:channel:scsi-id:lun

- 99:0:1:0
- 99:0:1:1

Examples for this command argument:

- -D 99:0:1:0,99:0:1:1,99:0:1:2
- -D "<full_path_to_phydev_list_file>"

<lun-reservation> is optional for automatically preparing an unprepared device. This works if the target pool has no defined type or the given reservation type is one of the pool types of the target pool.

The reservation type can be:

- Tape (VTL pool type)

-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 physical devices from storage pool

```
iscon removepdevsfromstoragepool -s <server-name> [-u <username> -p <password>]
-sp <storage-pool-id> | -SP <storage-pool-name> -D <physical-devices> [-X <rpc-timeout>]
```

```
iscon removepdevsfromstoragepool --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--storage-pool-id=<storage-pool-id> | --storage-pool-name=<storage-pool-name>
--physical-devices=<physical-devices> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command removes physical devices from an existing storage pool.

Either <storage-pool-ID> or <storage-pool-name> can be specified.

<physical-devices> can be a list of SCSI addresses separated by commas or a file enclosed in <> containing the SCSI address of each physical device to add on a separate line in the following format: adapter:channel:scsi-id:lun

- 99:0:1:0
- 99:0:1:1

Examples for this command argument:

- -D 99:0:1:0,99:0:1:1,99:0:1:2
- -D "<full_path_to_phydev_list_file>"

-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 storage pools

```
iscon getstoragepools -s <server-name> [-u <username> -p <password>]
[-sp <storage-pool-id> | -SP <storage-pool-name>] [-o <output-format>]
[-X <rpc-timeout>]
```

```
iscon getstoragepools --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--storage-pool-id=<storage-pool-id> | --storage-pool-name=<storage-pool-name>]
[--output-format=<output-format>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves information about all storage pools.

Either <storage-pool-ID> or <storage-pool-name> can be specified.

The default is to display the name of all storage pools if neither option is specified.

<output-format> is an option to choose one of the following formats for the output: *list* (default) or *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.

Rename storage pool

```
iscon renamestoragepool -s <server-name> [-u <username> -p <password>]
-sp <storage-pool-id> | -SP <storage-pool-name> -N <new_pool_name> [-X <rpc-timeout>]
```

```
iscon renamestoragepool --server-name=<server-name>
--storage-pool-id=<storage-pool-id> | --storage-pool-name=<storage-pool-name>
--pool-name=<new_pool_name> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command renames an existing storage pool.

Either <storage-pool-ID> or <storage-pool-name> can be specified to indicate which storage pool is going to be renamed.

<new_pool_name> indicates the new name for the storage pool.

-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 ACL to storage pool

```
iscon adduseracl -s <server-name> [-u <username> -p <password>]
-sp <storage-pool-id> | -SP <storage-pool-name> -ua <user-acl> [-X <rpc-timeout>]
```

```
iscon adduseracl --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--storage-pool-id=<storage-pool-id> | --storage-pool-name=<storage-pool-name>
--user-acl=<user-acl> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds user access control to an existing storage pool.

-ua (--user-acl) is an option to specify the user(s) to be added in the following format: user1,user2,user3

-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 ACL from storage pool

```
iscon removeuseracl -s <server-name> [-u <username> -p <password>]
-sp <storage-pool-id> | -SP <storage-pool-name> [-ua <user-acl>] [-X <rpc-timeout>]
```

```
iscon removeuseracl --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--storage-pool-id=<storage-pool-id> | --storage-pool-name=<storage-pool-name>
[--user-acl=<user-acl>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command removes user access control from an existing storage pool.

-ua (--user-acl) is an option to specify the user(s) in the following format: user1,user2,user3

If the -ua (--user-acl) option is not specified, all user(s) will be removed.

-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.

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 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.

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*.

- *VLIBS* = display virtual tape libraries only.
- *VDRIVES* = display standalone virtual tape drives only
- *VAULT* = display virtual tape vault 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.

-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 library or drive to an FC client

```
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 an option to assign a specific LUN. If this option is not specified, the next available LUN will be assigned. For virtual tape libraries, the option will take effect only when combined with the -y (--vlib-only) 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.

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

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 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 WWPNS is a 16-byte Hex value. Separate initiator WWPNS with commas if more than one initiator WWPNS 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.

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. For existing users, specify the name of the user. To create new users, provide the user name and password separated by a colon. Multiple users must be separated with commas. The password must conform with the password security policy of your organization. By default, the client will allow unauthenticated access. For example: -a iuser1,iuser2:user2password,iuser3

-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.

Get iSCSI client initiators

```
iscon getiscsiclientinitiator -s <server-name> [-u <username> -p <password>]
[-X <rpc-timeout>]
```

```
iscon getiscsiclientinitiator --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command lists the iSCSI client initiator names discovered by the VTL 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.

Delete iSCSI client initiators

```
iscon deleteiscsiclientinitiator -s <server-name> [-u <username> -p <password>]
-I <initiator-name-list> [-X <rpc-timeout>]
```

```
iscon deleteiscsiclientinitiator --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--initiator-name-list=<initiator-name-list> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes the specified iSCSI initiator name. All clients must be unassigned from the initiator prior to running the command.

-I (--initiator-name-list) is required to provide at least one valid initiator name. Multiple names must be 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.

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 with 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.

Get users

```
iscon getusers -s <server-name> [-u <username> -p <password>]
[-l] [-X <rpc-timeout>]
```

```
iscon getusers --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--longlist] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command allows you to get the list of users that can be assigned to a storage pool. This excludes the root user.

-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.

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 with 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 with 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.

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>] -N <auto-repl-mode>
-S <target-name> [-U <target-username> -P <target-password>] [-M <#[D|H|M]>] |
-O <auto-object-storage-migration-mode>
{-Ot <AWS_S3> -Oi <AWS-IAM-access-key-id> | -Ot <HCP> -OU <HCP-username>} [-Y <days>] ]
[-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]
[-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-replication=<auto-repl-mode> --target-name=<target-name>
[--target-username=<target-username> --target-password=<target-password>]
[--delay-delete-time=<#[D|H|M]>] |
--auto-object-storage-migration-mode=<copy|move>
{--object-account-type=<AWS_S3> --key-id=<AWS-IAM-access-key-id> |
--object-account-type=<HCP> --object-hcp-username=<HCP-username>}
[--delay-delete-days=<days>] ] [--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]
[ [--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 or 500 drives for library IBM 03584L22, IBM 03584L32, VTL 03584L.

-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

-O (--auto-object-storage-migration-mode) is an option with one of the following values:

- "copy" (migrate tape data to object storage.)
- "move" (migrate tape data to object storage and convert the source tape to a stub tape once migration is complete.)
- "none" (turn off auto-object-storage-migration-mode.)

-Ot (--object-account-type) is the provider name of the object storage with one of the following values:

- AWS_S3 (for Amazon AWS S3)
- HCP (for Hitachi HCP)

-Oi (--key-id) is required for provider AWS_S3 and is used to specify the key ID for the object storage service provider.

-OU (--object-hcp-username) is required for provider HCP and is used to specify the username for the object storage service provider.

-Y (--delay-delete-days) is an option for move mode to specify the number of days to wait before conversion. The maximum is 30 days. The default is 0 days.

-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 or 500 drives and 1800 GB of tape capacity. The following libraries support max. 500 drives per library: IBM 03584L22, IBM 03584L32, VTL 03584L22, VTL 03584L32, STK L700, STK L700e and STK SL500.

-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.

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 <vdev> [-d]
[-X <rpc-timeout>]
```

```
iscon deletevirtuallibrary --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdev=<vdev> [--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 (--vdev) 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.

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 createstandalone drive -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 createstandalone drive --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 library.

-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.

Enable object storage migration

```
iscon setvirtuallibraryobjectstoragemigration -s <server-name>
[-u <username> -p <password>] -v <vdevid> {-Z <off> | -Z <on> -O <copy|move>
{-Ot <AWS_S3> -Oi <AWS-IAM-access-key-id> |
-Ot <HCP> -OU <HCP-username>}
[-Y <days>]} [-X <rpc-timeout>]
```

```
iscon setvirtuallibraryobjectstoragemigration --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --vdevid=<vdevid>
{--auto-object-storage-migration=<off> | --auto-object-storage-migration=<on>
--auto-object-storage-migration-mode=<copy|move>
{--object-account-type=<AWS_S3> --key-id=<AWS-IAM-access-key-id> |
--object-account-type=<HCP> --object-hcp-username=<HCP-username>}
[--delay-delete-days=<days>]}
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command enables object storage migration for a specific library.

-v (--vdevid) is required to specify the virtual library ID.

-Z (--auto-object-storage-migration) is required to specify whether migration is "on" (enabled) or "off" (disabled) for the virtual library.

-O (--auto-object-storage-migration-mode) is required to specify when migration is enable with following values:

"copy" (migrate tape data to object storage.)

"move" (migrate tape data to object storage and convert the source tape to a stub tape once migration is complete.)

-Ot (--object-account-type) is required to specify the provider name of the object storage with one of the following values when migration is enabled:

- AWS_S3 (for Amazon AWS S3)
- HCP (for Hitachi HCP)

-Oi (--key-id) is required for provider AWS_S3 and is used to specify the key ID for the object storage service provider.

-OU (--object-hcp-username) is required for provider HCP and is used to specify the username for the object storage service provider.

-Y (--delay-delete-days) is an option for move mode to specify the number of days to wait before conversion. The maximum is 30 days. The default is 0 days.

-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> | -SP <storage-pool-name>]
[-n <vtapename>] [-WT] [-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>] [-TSP storage-pool-name]]] [-X <rpc-timeout>]

iscon createvirtualtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--parent-vid=<parent-vid> [--size-gb=<#(GB)>]
[--scsiaddress=<ACSL> | --storage-pool-name=<storage-pool-name>]
[--vdevname=<vtapename>] [--worm-tape]
[--barcode=<barcode | barcode-range>] [--count=<count>]
[--enable-auto-replication --target-name=<target-name>]
[--target-username=<target-username> --target-password=<target-password>]
[--target-storage-pool-name=<storage-pool-name>] [--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:scsi id:lun)

-SP (--storage-pool-name) is an option if the above option is not specified. This will automatically choose the proper physical device in the storage pool to create the virtual resource.

-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: <>"&\$/\'

-WT (--worm-tape) is an option to specify the virtual tape is a write once, read many (WORM) tape.

-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 from 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.

-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 the auto-replication option is specified. The default remote server from the parent library will be used if it is not specified.

-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.

-TSP (--storage-pool-name) is an option to automatically choose the proper physical device in the target storage pool to create the replica resource.

-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>]
[-N <auto-repl-mode> -S <target-name>
[-U <target-username> -P <target-password>] [-M <#[D|H|M]>] ]
-O <auto-object-storage-migration-mode>
{-Ot <AWS_S3> -Oi <AWS-IAM-access-key-id> | -Ot <HCP> -OU <HCP-username>} [-Y <days>] ]
[-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-eject-to-ie] | --auto-replication=<auto-replication-mode>
--target-name=<target-name>
[--server-username=<username> --server-password=<password>]
[--delay-delete-time=<#[D|H|M]>] ]
--auto-object-storage-migration-mode=<copy|move>
{--object-account-type=<AWS_S3>
--key-id=<AWS-IAM-access-key-id> | --object-account-type=<HCP>
--object-hcp-username=<HCP-username>} [--delay-delete-days=<days>] ]

[--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).

-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

-O (--auto-object-storage-migration-mode) is an option with one of the following values:

- "copy" (migrate tape data to object storage.)

- "move" (migrate tape data to object storage and convert the source tape to a stub tape once migration is complete.)

-Ot (--object-account-type) is the provider name of the object storage with one of the following values:

- AWS_S3 (for Amazon AWS S3)
- HCP (for Hitachi HCP)

-Oi (--key-id) is required for provider AWS_S3 and is used to specify the key ID for the object storage service provider.

-OU (--object-hcp-username) is required for provider HCP and is used to specify the username for the object storage service provider.

-Y (--delay-delete-days) is an option for move mode to specify the number of days to wait before conversion. The maximum is 30 days. The default is 0 days.

-N (--auto-replication) cannot be specified if replication is enabled for the 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.

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 <vdevvid> | -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>]
--vdevvid=<vdevvid> | --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 (--vdevvid) 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-vdevvid> -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-vdevvid=<source-vdevvid> --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-vdevvid**) 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.

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 server.

The command does not start a new space reclamation job unless the force argument is used. The *force* argument 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.

-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 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.

-l (--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.

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]>] [-SR]] | [-r <on>]
[[-t <timeout>] [-I <retry-in>] [-C <retry-for>]] [-c <on|off>] [-e <on|off>]
[[-L <#:#:#:#>] | [-TSP storage-pool-name]] [-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]>]]
[--data-change]] | [--repl-first <on>] [ [--replication-timeout=<timeout>]
[--replication-retry-interval=<retry-in>] [--replication-retry-count=<retry-for>]]
[--compression=<on|off>] [--encryption=<on|off>]
[ [--preferred-lun=<#:#:#:#>] | [--target-storage-pool-name=<storage-pool-name>]]
[--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.

Any combination of the following two options can be used in order to set up a replication trigger policy for a virtual tape. 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).

`-SR` (`--data-change`) is an option that will trigger replication when a virtual tape is unloaded from a tape drive and the tape data has changed. `-SR` (`--data-change`) is mutually exclusive with `-w` (`--watermark`) and `-d` (`--date`).

`-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).
- `-l` (`--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`)

`-TSP` (`--storage-pool-name`) is an option that will automatically choose the proper physical device in the target storage pool to create a replica resource.

`-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 <vdev> [-X <rpc-timeout>]
```

```
iscon resumereplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --vdev=<vdev>
[--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-vdev) 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-vdev> [-w <watermark(MB)> | [-d <YYYYMMDDHHMM> -i <#[H|M]>] [-SR]] |
[-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-vdev=<source-vdev> [--watermark=<watermark (MB)>
[--watermark=<watermark (MB)> | [--date=<YYYYMMDDHHMM> --interval=<#[H|M]>]]
[--data-change]] |[--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-vdev) 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 tape.

-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).

-SR (--data-change) is an option that will trigger replication when a virtual tape is unloaded from a tape drive and the tape data has changed. -SR (--data-change) is mutually exclusive with -w (--watermark) and -d (--date).

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.

-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).
- -l (--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.

Set 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.

Data encryption

Enable virtual tape or deduplication repository 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.

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 create a password that will be used for encryption activation.

-C (--confirm-password) is required to confirm the encryption activation password. The two password arguments must match.

The password must conform with the password security policy of your organization.

-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.

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.

-O (--old-password) is required to provide the current password.

-W (--password) is required to provide the new password.

-C (--confirm-password) is required to confirm the new password. The two new password arguments must match.

The password must conform with the password security policy of your organization.

-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.

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.

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.

Object Storage

Add object storage account

```
iscon object-storage-add -s <server-name> [-u <username> -p <password>]
-On <object-storage-name>
{-Ot <AWS_S3> -Oi <AWS-IAM-access-key-id> -Os <AWS-IAM-secret-access-key>
-Ob <bucket-name> [-Or <region>] [-Od <true|false>] [-OS <storage-class>] |
-Ot <HCP> -OU <HCP-username> -OP <HCP-password> -ON <HCP-namespace>
-OT <HCP-tenant> -OD <HCP-domain> [-OL <HCP-protocol>] |
-Ot <IBM_COS> -Oi <HMAC-access-key-id> -Os <HMAC-secret-access-key>
-Ob <bucket-name> -Or <region>}
-Ot <GENERIC_S3> -Oi <HMAC-access-key-id> -Os <HMAC-secret-access-key>
-Ob <bucket-name>}
[-Op <true|false>] [-Ou <uri>] [-Oe <true|false>] [-Ol <true|false>] [-Oc <comment>]
[[-Ps <proxy-server> -Pp <port>] [-Ph <true|false>]
[-PU <proxy-auth-user> -PP <proxy-auth-password>]][-X <rpc-timeout>]
```

```
iscon object-storage-add --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--object-storage-name=<object-storage-name>
{--object-account-type=<AWS_S3> --key-id=<AWS-IAM-access-key-id>
--secret-access-key=<AWS-IAM-secret-access-key> --bucket-name=<bucket-name>
[--region=<region>] [--use-dual-stack=<true|false>]
[--storage-class=<storage-class>] |
--object-account-type=<HCP> --object-hcp-username=<HCP-username>
--object-hcp-password=<HCP-password> --object-hcp-namespace=<HCP-namespace>
--object-hcp-tenant=<HCP-tenant> --object-hcp-domain=<HCP-domain>
[--object-hcp-protocol=<HCP-protocol>] |
--object-account-type=<IBM_COS> --key-id=<HMAC-access-key-id>
--secret-access-key=<HMAC-secret-access-key>
--bucket-name=<bucket-name> --region=<region> |
--object-account-type=<GENERIC_S3> --key-id=<HMAC-access-key-id>
--secret-access-key=<HMAC-secret-access-key> --bucket-name=<bucket-name>}
[--object-use-https=<true|false>] [--object-endpoint=<uri>]
[--object-end-to-end-encryption=<true|false>] [--object-sironly=<true|false>]
[--object-comment=<comment>] [[--proxy-server=<proxy-server> --proxy-port=<port>]
[--proxy-use-https=<true|false>] [--proxy-auth-user=<proxy-auth-user>
--proxy-auth-password=<proxy-auth-password>]] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command allows you to add object storage account access information to the server.

-On (--object-storage-name) is required to define a unique name. This identity will be used in other commands.

-Ot (--object-account-type) is required to specify the provider name of the object storage:

- AWS_S3 (for Amazon AWS S3)
- HCP (for Hitachi HCP)
- IBM_COS (for IBM Cloud Object Storage)
- GENERIC_S3 (for S3 compatible object storage)

-Oi (--key-id) is required for providers AWS_S3, IBM_COS, and GENERIC_S3 and is used to specify the key ID for the object storage service provider.

-Os (--secret-access-key) is required for providers AWS_S3, IBM_COS, and GENERIC_S3 and is used to specify the secret access key of the key ID that is used to access object storage.

-Ob (--bucket-name) is required for providers AWS_S3, IBM_COS, and GENERIC_S3 and is used to specify the bucket name to access object storage.

-Or (--region) is an option for provider AWS_S3 to specify the region of Amazon AWS S3 service:

- us-east-1 (US East (N. Virginia)) (default)
- us-east-2 (US East (Ohio))
- us-west-1 (US West (N. California))
- us-west-2 (US West (Oregon))
- ca-central-1 (Canada (Central))
- ap-south-1 (Asia Pacific (Mumbai))
- ap-northeast-2 (Asia Pacific (Seoul))
- ap-northeast-3 (Asia Pacific (Osaka-Local))
- ap-southeast-1 (Asia Pacific (Singapore))
- ap-southeast-2 (Asia Pacific (Sydney))
- ap-northeast-1 (Asia Pacific (Tokyo))
- cn-north-1 (China (Beijing))
- cn-northwest-1 (China (Ningxia))
- eu-central-1 (EU (Frankfurt))
- eu-west-1 (EU (Ireland))
- eu-west-2 (EU (London))
- eu-west-3 (EU (Paris))
- sa-east-1 (South America (Sao Paulo))

-Or (--region) is required for provider IBM_COS to specify the region of IBM COS service:

Regional Endpoints:

- us-south (US South)
- us-east (US East)
- eu-gb (EU United Kingdom)
- eu-de (EU Germany)
- au-syd (AP Australia)
- jp-tok (AP Japan)

Cross Region Endpoints:

- us (US Cross Region)
- eu (EU Cross Region)
- ap (AP Cross Region)

Single Data Center Endpoints:

- ams03 (Amsterdam, Netherlands)
- che01 (Chennai, India)
- hkg02 (Hong Kong)
- mel01 (Melbourne, Australia)

- mex01 (Mexico City, Mexico)
- mil01 (Milan, Italy)
- mon01 (Montréal, Canada)
- osl01 (Oslo, Norway)
- sjc04 (San Jose, USA)
- sao01 (São Paulo, Brazil)
- seo01 (Seoul, South Korea)
- tor01 (Toronto, Canada)

-Od (--use-dual-stack) is an option to specify "true" or "false" to indicate whether object storage is accessed via IPv6/IPv4 or IPv4 only. The default is "false" (IPv4 only).

-OS (--storage-class) is an option for provider AWS_S3 to specify the storage class:

- standard
- standard_ia
- onezone_ia
- intelligent_tiering
- glacier (default)
- deep_archive

-OU (--object-hcp-username) is required for provider HCP and is used to specify the username for the object storage service provider.

-OP (--object-hcp-password) is required for provider HCP and is used to specify the password for the object storage service provider.

-ON (--object-hcp-namespace) is required for provider HCP and is used to specify the namespace for the object storage service provider.

-OT (--object-hcp-tenant) is required for provider HCP and is used to specify the tenant for the object storage service provider.

-OD (--object-hcp-domain) is required for provider HCP and is used to specify the domain for the object storage service provider.

-OL (--object-hcp-protocol) is an option for provider HCP and is used to specify the protocol/API to access: *rest* (default) or *s3*

-Op (--object-use-https) is an option to specify "true" or "false" to indicate whether object storage is accessed via "https" or "http". The default is "true" (use "https").

-Ou (--object-endpoint) is an option to specify the endpoint/URI used to access object storage.

- For AWS_S3, this is s3.[region].amazonaws.com and the default region is us-east-1
- For HCP, this is [tenant].[domain]
- For IBM_COS, this is s3.[region].cloud-object-storage.appdomain.cloud
- For GENERIC_S3, this option is required. Specify the full URI/URL path

-Oe (--object-end-to-end-encryption) is an option to specify "true" or "false" to enable or disable end-to-end encryption. The default is "true" (use end-to-end encryption).

-OI (--object-sironly) is an option to specify "true" or "false" to use in deduplication (true) or tape migration (false). The default is "false" (use in tape migration)

-Oc (--object-comment) is an option to specify a comment string.

-Ps (--proxy-server) is an option to specify the proxy server to access to object storage.

-Pp (--proxy-port) is an option to specify the port number of the proxy service.

-Ph (--proxy-use-https) is an option to specify "true" or "false" to indicate whether proxy server is accessed via "https" or "http". The default is "true"(use "https").

-PU (--proxy-auth-user) is an option to specify the username for proxy authentication.

-PP (--proxy-auth-password) is an option to specify the password for proxy authentication.

-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.

Modify object storage information

```
iscon object-storage-modify -s <server-name> [-u <username> -p <password>]
-On <object-storage-name> [-Nn <new-object-storage-name>]
{-Ot <AWS_S3> [-Oi <new-IAM-access-key-id> -Os <new-IAM-secret-access-key>]
[-Od <true|false>] [-OS <storage-class>] |
-Ot <HCP> [-OP <HCP-password>] [-OL <HCP-protocol>] |
-Ot <IBM_COS> [-Oi <new-HMAC-access-key-id> -Os <new-HMAC-secret-access-key>] |
-Ot <GENERIC_S3> [-Oi <new-HMAC-access-key-id> -Os <new-HMAC-secret-access-key>]]
[-Op <true|false>] [-Oe <true|false>]
[-Oc <comment>] [[-Ps <proxy-server>] [-Pp <port>] [-Ph <true|false>]
[-PU <proxy-auth-user> -PP <proxy-auth-password>]] [-X <rpc-timeout>]
```

```
iscon object-storage-modify --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--object-storage-name=<object-storage-name>
[--new-object-storage-name=<new-object-storage-name>]
{--object-account-type=<AWS_S3> [--key-id=<new-IAM-access-key-id>
--secret-access-key=<new-IAM-secret-access-key>]
[--use-dual-stack=<true|false>] [--storage-class=<storage-class>] |
--object-account-type=<HCP> [--object-hcp-password=<HCP-password>]
[--object-hcp-protocol=<HCP-protocol>]]
--object-account-type=<IBM_COS> [--key-id=<new-HMAC-access-key-id>
--secret-access-key=<new-HMAC-secret-access-key>] |
--object-account-type=<GENERIC_S3> [--key-id=<new-HMAC-access-key-id>
--secret-access-key=<new-HMAC-secret-access-key>]]
[--object-use-https=<true|false>] [--object-end-to-end-encryption=<true|false>]
[--object-comment=<comment>]
[[--proxy-server=<proxy-server>] [--proxy-port=<port>] [--proxy-use-https=<true|false>]
[--proxy-auth-user=<proxy-auth-user> --proxy-auth-password=<proxy-auth-password>]]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command allows you to modify object storage account access information for the server.

-On (--object-storage-name) is required to specify the unique name defined in object-storage-add command.

-Nn (--new-object-storage-name) is an option to change the unique name.

-Ot (--object-account-type) is required to specify the provider name of the object storage:

- AWS_S3 (for Amazon AWS S3)
- HCP (for Hitachi HCP)
- IBM_COS (for IBM Cloud Object Storage)
- GENERIC_S3 (for S3 compatible object storage)

-Oi (--key-id) is an option for providers AWS_S3, IBM_COS, and GENERIC_S3 and is used to specify the changed key ID for the object storage service provider.

-Os (--secret-access-key) is an option for providers AWS_S3, IBM_COS, and GENERIC_S3 and is used to specify the changed secret access key for the object storage service provider.

-Od (--use-dual-stack) is an option to specify "true" or "false" to indicate whether object storage is accessed via IPv6/IPv4 or IPv4 only. The default is "false"(IPv4 only)

-OS (--storage-class) is an option for provider AWS_S3 to specify the storage class:

- standard
- standard_ia
- onezone_ia
- intelligent_tiering
- glacier (default)
- deep_archive

-OP (--object-hcp-password) is an option for provider HCP and is used to specify the changed password for the object storage service provider.

-OL (--object-hcp-protocol) is an option for provider HCP and is used to specify the protocol/API to access: *rest* (default) or *s3*

-Op (--object-use-https) is an option to specify "true" or "false" to indicate whether object storage is accessed via "https" or "http". The default is "true" (use "https").

-Oe (--object-end-to-end-encryption) is an option to specify "true" or "false" to enable or disable end-to-end encryption. The default is "true" (use end-to-end encryption).

-Oc (--object-comment) is an option to specify a comment string.

-Ps (--proxy-server) is an option to specify the proxy server to access to object storage.

-Pp (--proxy-port) is an option to specify the port number of the proxy service.

-Ph (--proxy-use-https) is an option to specify "true" or "false" to indicate whether proxy server is accessed via "https" or "http". The default is "true"(use "https").

-PU (--proxy-auth-user) is an option to specify the username for proxy authentication.

-PP (--proxy-auth-password) is an option to specify the password for proxy authentication.

-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.

Display object storage information

```
iscon object-storage-show -s <server-name> [-u <username> -p <password>]
[-Ot <AWS_S3|HCP|IBM_COS|GENERIC_S3>] [-l] [-X <rpc-timeout>]
```

```
iscon object-storage-show --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--object-account-type=<AWS_S3|HCP|IBM_COS|GENERIC_S3>] [--longlist]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command allows you to display object storage account access information for the server.

-Ot (--object-account-type) is an option to specify the provider name of the object storage:

- AWS_S3 (for Amazon AWS S3)
- HCP (for Hitachi HCP)
- IBM_COS (for IBM Cloud Object Storage)
- GENERIC_S3 (for S3 compatible object storage)

-l (--longlist) is an option to display information in 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.

Delete object storage

```
iscon object-storage-delete -s <server-name> [-u <username> -p <password>]
-On <object-storage-name> [-f] [-X <rpc-timeout>]
```

```
iscon object-storage-delete --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--object-storage-name=<object-storage-name> [--force] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command allows you to delete object storage account access information from the server.

-On (--object-storage-name) is required to specify the unique name defined in object-storage-add command.

-f (--force) is an option to force the deletion of object storage account access information. The corresponding data/object in the cloud will also 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.

Start object migration

```
iscon startobjectstoragemigration -s <server-name> [-u <username>
-p <password>] -v <tape-vid> [-On <object-storage-name>] [-X <rpc-timeout>]
```

```
iscon startobjectstoragemigration --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<tape-vid> [--object-storage-name=<object-storage-name>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command submits an object storage migration job for a specific tape.

-v (--vdevid) is required to specify the virtual tape ID.

-On (--object-storage-name) is an option to specify the unique name defined in object-storage-add command.

-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.

Recover tape from object storage

```
iscon recovertapefromobjectstorage -s <server-name> [-u <username>
-p <password>] -v <stub-tape-vid> -L <tape-library-vid> [-l <slot-no>] [-O <copy|move>]
[-Og <AWS-Glacier-retrieval>] [-X <rpc-timeout>]
```

```
iscon recovertapefromobjectstorage --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<stub-tape-vid> --tape-library-vid=<tape-library-vid>
[--slot-no=<slot-no>] [--recovery-mode=<copy|move>]
[--glacier-retrieval=<AWS-Glacier-retrieval>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command submits an object storage recovery job for a specific tape.

-v (--vdevid) is required to specify the stub virtual tape ID.

-L (--tape-library-vid) is required to specify the virtual library ID.

-l (--slot-no) is an option to specify the position/slot of the virtual tape library after tape recovery is done.

-O (--recovery-mode) is an option to specify the recovery mode with the following values:

- "copy" (copy tape data from object storage.) (default)
- "move" (copy tape data from object storage and delete the object tape data once the copy is done.)

-Og (--glacier-retrieval) is an option for provider AWS_S3 to specify the retrieval method when the storage class is glacier:

- expedited
- standard
- bulk (default)

-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.

Convert migrated tape to a stub tape

```
iscon convertmigratedtapetostub -s <server-name> [-u <username>
-p <password>] -v <tape-vid> [-X <rpc-timeout>]
```

```
iscon convertmigratedtapetostub --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevid=<tape-vid> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command converts a migrated tape to a stub tape.

-v (--vdevid) is required to specify the 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.

Get object storage job status

```
iscon getobjectstoragejobstatus -s <server-name> [-u <username> -p <password>]
[[-j <job-id-list>] | [[-T <job-type>] [-S <job_status>] [-D <date-range|date>]]]
[-v <tape-id>] [-l][-X <rpc-timeout>]
```

```
iscon getobjectstoragejobstatus --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>]]] [--vdevid=<tape-vid>] [--longlist]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command displays the status of the object storage migration and recovery jobs present in the queue. If no filters are specified, the command displays all of 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 the 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: MIGRATION or RECOVERY

-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, OTHER

"READY" jobs are waiting for the tape/drive to be loaded.

-D (--date-range) is an optional date based filter in the following format: YYYYMMDD-YYYYMMDD or YYYYMMDD

-v (--vdevid) is an option to specify the virtual tape ID.

-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.

Manage object storage jobs

```
iscon manageobjectstoragejobs -s <server-name> [-u <username> -p <password>]  
-a <action> -j <job-id-list> [-X <rpc-timeout>]
```

```
iscon manageobjectstoragejobs --server-name=<server-name>  
[--server-username=<username> --server-password=<password>]  
--action=<action> --job-id-list=<job-id-list> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command manages object storage migration and recovery jobs in the queue.

-a <--action> is the action to perform. The accepted job action values are: RESTART, CANCEL, DELETE.

-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.

Mirroring

Create a mirror

```
iscon createmirror -s <server-name> [-u <username> -p <password>]
-v <vdev> -I <acsl>
[-X <rpc-timeout>]
```

```
iscon createmirror --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdev=<vdev> --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 (--vdev) 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 <vdev>
[-X <rpc-timeout>]
```

```
iscon getmirrorstatus --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdev=<vdev>
[--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 (--vdev) 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 <vdev>  
[-X <rpc-timeout>]
```

```
iscon removemirror --server-name=<server-name>  
[--server-username=<username> --server-password=<password>]  
--vdev=<vdev>  
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command cancels any active mirror synchronization job for the specified virtual device and removes the mirror.

-v (--vdev) 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 <vdev>  
[-X <rpc-timeout>]
```

```
iscon swapmirror --server-name=<server-name>  
[--server-username=<username> --server-password=<password>]  
--vdev=<vdev>  
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command swaps the specified primary device with its mirrored copy.

-v (--vdev) 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 <vdev>  
[-X <rpc-timeout>]
```

```
iscon syncmirror --server-name=<server-name>  
[--server-username=<username> --server-password=<password>]  
--vdev=<vdev>  
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command synchronizes the specified virtual device with its mirroring device. The command does not wait for the operation to finish.

-v (--vdev) 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.

Alarm policies

Set virtual library alarm policy

```
iscon setvlibalarmpolicy -s <server-name> [-u <username> -p <password>]
-L <tape-library-vid> [-t <and> | <or> ] {-up <usedspace-percent> |
-n <empty-tape-number> [-m <empty-tape-sizeMB>]} [-X <rpc-timeout>]

iscon setvlibalarmpolicy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-library-vid=<tape-library-vid> [--trigger=<or|and>]
{--usedspace-percent=<usedspace-percent> | --empty-tape-num=<empty-tape-number>
[--empty-tape-sizeMB=<empty-tape-sizeMB>]} [--rpc-timeout=<rpc-timeout>]
```

Description:

This command sets an alarm policy on the specified virtual library. There are two types of alarm triggers, one based on used space and the other on the number of empty tapes.

-L (--tape-library-vid) is required to specify the virtual library ID.

-t (--trigger) is an option to specify if either alarm occurs (<or>) or both alarms occur (<and>). The default is or.

-up (--usedspace-percent) is an option to specify the percentage of used space. The server will trigger an alarm if used space is above this value.

-n (--empty-tape-num) is an option to specify the number of empty tapes. The server will trigger an alarm if the number of empty tapes is below this value.

-m (--empty-tape-sizeMB) is an option to specify the maximum size of an empty tape in MB when the number of empty tapes is specified. The default is 4.

-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 virtual library alarm policy

```
iscon getvlibalarmpolicy -s <server-name> [-u <username> -p <password>]
[-L <tape-library-vid>] [-X <rpc-timeout>]

iscon getvlibalarmpolicy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--tape-library-vid=<tape-library-vid>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command lists the alarm policies on the VTL system or for a specific virtual library.

-L (--tape-library-vid) is an option to get the alarm policy of a specific virtual 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.

Remove virtual library alarm policy

```
iscon removevlibalarmpolicy -s <server-name> [-u <username> -p <password>]  
-L <tape-library-vid>[-X <rpc-timeout>]
```

```
iscon removevlibalarmpolicy --server-name=<server-name>  
[--server-username=<username> --server-password=<password>]  
--tape-library-vid=<tape-library-vid>[--rpc-timeout=<rpc-timeout>]
```

Description:

This command removes an alarm policy for a specific virtual library.

-L (--tape-library-vid) is required to specify the virtual library 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.

Support utilities

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.

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).

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.

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 creatededupetapeactivityreport -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 creatededupetapeactivityreport --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=<additional 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 policies 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 policies. 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: `DeduplicationTapeUsage-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)` is 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> is 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 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
- ISD[C | R] - Import from standalone drive, Copy or Recycle

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 *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.

Object storage jobs report

```
iscon createobjectstoragejobreport -s <server-name> [-u <username> -p <password>]
[-z <report period>] | [-D <date-range>] [-i <filter>] [-o <filename>] [-X <rpc-timeout>]
```

```
iscon createobjectstoragejobreport --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 the object storage migration 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:

- `MIGM` - Object storage migration, Move
- `MIGC` - Object storage migration, Copy
- `RCVM` - Object storage recovery, Move
- `RCVC` - Object storage recovery, Copy

Job Status:

- `FAIL` - Failed
- `COMP` - Completed
- `CANC` - Cancelled
- `HOLD` - On hold

- RUN - Running

For example: `-i RCV,C,COMP,CANC`

This command will include only *object storage recovery copy, 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: `ObjectStorageMigrationJobReport-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.

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> is 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 createvirtuallibinfoport -s <server-name> [-u <username> -p <password>]
[-o <filename>] [-X <rpc-timeout>]
```

```
iscon createvirtuallibinfoport --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, Migration to Object Storage, and Recovery from Object Storage. 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 createvirtualtapeinfo report -s <server-name> [-u <username> -p <password>]
[-i <filter>] [-o <filename>] [-X <rpc-timeout>]
```

```
iscon createvirtualtapeinfo report --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, the 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)
- RR (replica resources view, includes all replica tapes)
- VV (vault view, includes all tapes from the vault)
- DT (detailed tape view)
- OM (object storage migration 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.

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.

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	Raw replication data transferred, in MB
ReplActualDataTx	Actual amount of replication data transferred, in MB
ReplTotalTapesTx	Number of tapes replicated

falcVtlMonPerformanceInfo

Displays VTL performance.

Object	Description
AvgDedupRatio	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, 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)
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)

falcVtlServerOptionsInfo

Displays VTL server options.

Object	Description
falcVTLServerOptionsInfo	Displays VTL server options (Fibre Channel, iSCSI, VTL database, Email Alerts) 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 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)
falcVtlVirtLibMediaType	Media type
falcVtlVirtLibNumTapes	Number of tapes in library
falcVtlVirtLibSerialNum	Serial number of library
falcVtlVirtLibAutoReplication	Indicates if auto replication 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
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

VtlJob / falcVtlJobQueueInfo

Displays information about the tape import/export job queue.

Object	Description
falcVTLJobID	Job ID
falcVTLJobType	Job type: <ul style="list-style-type: none"> createFromCacheMetaDataModeAndCopy - Create cache in copy meta data mode moveTapeESlot - Moving tape to an IE slot
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 / falcVtIDeduplicationPoliciesInfo

Displays information about deduplication policies.

Object	Description
PolicyName	Policy name
PolicyID	Policy ID
NumberOfTapes	Number of tapes in the policy
TriggerType	Trigger type: <ul style="list-style-type: none"> • endOfBackupTapeFull • endOfBackup • noSchedule • scheduleHourly • scheduleDaily • scheduleSunday • scheduleMonday • scheduleTuesday • scheduleWednesday • scheduleThursday • scheduleFriday • scheduleSaturday
SirCluster	<i>Localcluster</i>
ReplicationStatus	Indicates if replication is enabled or disabled
TurboDeduplication	Indicates if turbo deduplication is enabled or disabled
ReplicationMode	Replication mode
falcVtIDeduplicationPoliciesInfo	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 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.

StoragePools

Information about the hardware specifications and characteristics of each SCSI device storage device.

Object	Description
StoragePoolName	Storage pool name
StoragePoolID	Storage pool ID
StoragePoolType	Storage pool type (Tapes, All)
StoragePoolDevCount	Storage pool device count
StoragePoolCount	Number of storage pools
StoragePoolTotalSize64	Size of the storage pool
StoragePoolUsedSize	Used size of the storage pool
StoragePoolAvailableSize	Available space in the storage pool

falcVtlPhysicalResourcesInfo

Total number of devices.

Object	Description
Adapters	Number of adapters
Devices	Number of physical devices

falcVtlSanClients

Displays information about clients (backup application servers).

SANClient

Displays information about each backup application server.

Object	Description
falcVtlSanClientID	Client ID
falcVtlSanClientName	Client name
falcVtlSanClientType	Client type: Fibre Channel, iSCSI

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
InitiatorWWPN	WWPN of the client's initiator 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 application servers.

Object	Description
falcVtlNumOfSANClients	Total number of backup application 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 server.

Object	Description
Name	<i>Localcluster</i>
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 server. 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 server in the past 24 hours, in MB
NumberOfTapesDeduplicatedIn24Hour	Number of tapes deduplicated by the server in the past 24 hours, in MBs
DataReplicatedIn24Hour	Amount of data replicated by the target replication server in the past 24 hours, in MBs
AverageDedupeRatioIn24Hour	Average deduplication ratio of data processed by the server in the past 24 hours, in MBs
AverageCompressionRatioIn24Hour	Average compression ratio of unique blocks processed by the server in the past 24 hours, in MBs

falcSirMonNodePerformanceInfo

Displays performance information for the server. 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 in the past 24 hours, in MB
OverallDedupeRatio	Overall deduplication ratio
AverageDedupeRatio	Average deduplication ratio in the past 24 hours
AverageCompressionRatio	Average compression ratio 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 deduplication.

VTL

Displays information about VTL server.

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, 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 node.

Object	Description
Name	Server name
IPAddress	Server IP address
Type	Server type
PowerFactor	Number of CPUs x 2. This number determines the maximum number of deduplication and resolver processes that can run on each deduplication node.

SIRHashIndexStorage

Displays information about the deduplication node hash index list.

Object	Description
SirName	Server name
Size	Size of hash index

SIRHashFolderStorage

Displays information about the deduplication node hash folder list.

Object	Description
SirName	Server name
Size	Size of hash folder

SIRHashDataStorage

Displays information about the deduplication node hash data list.

Object	Description
SirName	Server name
Size	Size of hash data

falcSirCluster - falcSirClusterStats

Displays statistics about folders, repository usage, and deduplication results.

Folder

Displays folder information.

Object	Description
Barcode	Barcode of folder
BackupApp	Backup application: ARCServe, Atempo, Bacula, Commvault, DataProtector, IBMiSeries, Legato, MicrosoftTapeFormat, 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, 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 VTL 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 <i>Attention Required</i> 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 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 (<i>/var/log/messages</i>) 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 <i>Physical Resources</i> object and select <i>Rescan</i>). Make sure that the <i>Discover New Devices</i> option is selected. Specify a <i>LUN Range</i> that you reasonably expect will include the LUN. |
| Check system log messages | Check the Event Log or syslog (<i>/var/log/messages</i>) 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: <ul style="list-style-type: none"> • 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 and device connectivity.

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 application 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	<p>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.</p> <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • Windows - For IBM Ultrium devices you can use <code>ntutil</code>, a command line tool that can check the tape device. • Unix systems - You can use the <code>mt</code> or <code>tar</code> commands to access the tape device, for example: <code>mt -f /dev/rmt/0 status</code>
OS cannot access device	<p>If the operating system <i>cannot access</i> the device, you need to troubleshoot virtual device access.</p> <ul style="list-style-type: none"> • 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 (<code>/var/log/messages</code>) for message indicating IO errors. Such messages usually begin with <code>log_scsi_error</code>. • 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	<p>If the operating system <i>can access</i> the device, you need to troubleshoot the backup software. Verify that you have the correct drivers.</p>

Client can no longer access a virtual device (tape library or drive)

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 the VTL 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.

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.

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 FalconStor 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.

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

A tape 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.

Replication process

Replication fails to start If you enable replication/migration on a virtual tape library with existing tapes that have tape-level replication configured, when the tape is ejected to the virtual vault, both the replication/migration job and the tape replication job will be triggered at the same time and the tape replication job will fail.

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).

Deduplication

Post-processing and/or Turbo deduplication jobs fail

Migration to object storage jobs have a higher priority than Post-processing and Turbo deduplication jobs.

If a migration to object storage job is running when a deduplication job is triggered, the deduplication job will fail but will attempt to run at the next retry interval.

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.

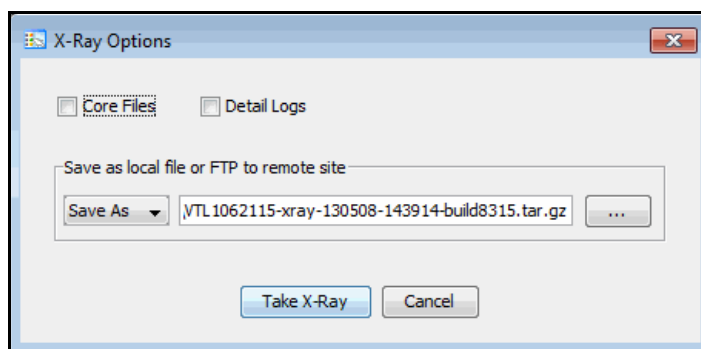
For troubleshooting purposes, `/var/log/messages` keeps track of the last 20 MB of system events. When the file reaches 20 MB, it is renamed to `messages.n`, where *n* is a sequential number between 1 and 30. When it is renamed, it is also compressed to save space.

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*.



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.

Depending upon the settings that were selected when your system was installed, the FTP may not be available.

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 `vtl_xray`) is a directory on the FTP server (i.e., `ftp\vtl_xray`). Do not enter an absolute path like `c:\vtl_xray`.

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 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 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 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.
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 the storage log.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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 be 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.
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 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 the storage log.

Number	Type	Text	Probable Cause	Suggested Action
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 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 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 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 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 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 the storage log.
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 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 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 the storage log.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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 the storage log.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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 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 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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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 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 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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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, 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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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'.
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.

Number	Type	Text	Probable Cause	Suggested Action
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 FalconStor to purchase additional capacity license.
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.

Number	Type	Text	Probable Cause	Suggested Action
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 replicating 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.
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 or software components, restart the server machine.

Number	Type	Text	Probable Cause	Suggested Action
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.
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'.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.
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 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 the storage log. Fix any detected issues and rescan devices to refresh the configuration.
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.

Number	Type	Text	Probable Cause	Suggested Action
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 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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.
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.
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 the storage log. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.
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 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.
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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 the storage log. Fix any detected issues and rescan devices to refresh the configuration.
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.
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.
13710	W	The live trial license expired for %1. Contact FalconStor 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 FalconStor or its representative to purchase a license.	The specified option is not licensed properly.	Obtain proper licenses.
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 the storage log. Fix any detected issues and rescan devices to refresh the configuration.
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 the storage log. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
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 the storage log. Fix any detected issues and rescan devices to refresh the configuration.
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 the storage log. Fix any detected issues and rescan devices to refresh the configuration.
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.
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.
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.	Check the physical device status, device connectivity, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.
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.	Check the physical device status, device connectivity, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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 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.

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 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 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 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, switch zoning.
18001	E	The repository module failed to start because the Deduplication Repository disk is missing.	The server could not connect to a physical device used by the deduplication repository.	Check the underlying physical storage, FC connectivity, switch zoning.
18050	E	The reclamation triggering module stopped because index pruning has failed.	The index disk may be in an invalid state.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
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 Support 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 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 repository was full. At least one VIT tape is required for reclamation to proceed.	Contact Technical Support.
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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 the storage log.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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 the storage log. Fix any detected issues and rescan devices to refresh the configuration.
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 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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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 the storage log. Fix any detected issues and rescan devices to refresh the configuration.
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 the storage log. Fix any detected issues, rescan devices to refresh the configuration, and try again.
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.
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 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 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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.
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 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 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 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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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 the storage log.
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.
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.
40231	W	The object storage migration job %1 failed on tape [%2] ID %3.	This may be due to an object storage network connection failure.	Restart the migration job after the network issues are fixed.

Number	Type	Text	Probable Cause	Suggested Action
40232	W	The object storage migration job %1 failed on tape [%2] ID %3. The job will be retried in %4 minutes.	This may be due to an object storage network connection failure.	Fix the object storage connection issue. The system is configured to retry jobs in the Tape Import/Export queue and the job will be automatically retried.
40238	W	The object storage recover job %1 failed on tape [%2] ID %3.	This may be due to an object storage network connection failure.	Restart the tape reconstruction job after the network issues are fixed.
40239	W	The object storage recover job %1 failed on tape [%2] ID %3. The job will be retried in %4 minutes.	This may be due to an object storage network connection failure.	Fix the object storage connection issue. The system is configured to retry jobs in the Tape Import/Export queue and the job will be automatically retried.
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.
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.
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 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
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 the storage log.
50549	W	Deduplication Policy %1: The policy was not started because the repository %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.
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.
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.
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.
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.
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 the storage log. Fix any detected issues.

Number	Type	Text	Probable Cause	Suggested Action
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 or a storage issue.	Take necessary actions based on the specified reason in the message.
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.
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.	There may not be enough storage on the VTL server.	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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.	Wait for the index cache to complete loading; this may take a while.
50956	E	Deduplication data disk %1 is full.	There is not enough space on the data disk for backup.	Run space reclamation 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 reclamation to free disk space or add more storage.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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 be 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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
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 FalconStor Technical Support.

Best Practices

This section provides sizing guidelines.

Deduplication repository sizing

A simple formula to estimate how large your deduplication repository needs to be is:

Repository size = Amount of data / deduplication ratio

The deduplication ratio is an estimated value based on the nature of the data. The deduplication repository does not grow over time; you must determine its maximum size when you configure the system.

The example below illustrates a Virtual Desktop Infrastructure (VDI) environment using flash array storage. The formula is the following, assuming that each VDI is 40 GB:

Repository size = (Number of VDIs x 40 GB) / deduplication ratio

Number of VDIs (40 GB per VDI)	Expected Deduplication Ratio	Minimum Repository Size (rounded up, in TB)
1000	10:1	5
2000	10:1	10
3000	10:1	15

Index/folder disk sizing

The size of the disk used for the Index/folder for each deduplication node is 4.3% of deduplication repository capacity.

CPU cores

You need at least 2 CPU cores; for large systems requiring high amount of memory based on the deduplication repository size, consider 4 or 8 cores:

- Deduplication repository up to 50 TB: 2 CPU cores
- Deduplication repository up to 100 TB: 4 CPU cores
- Deduplication repository over 100 TB: 8 CPU cores

Deduplication system sizing

The number of nodes and the amount of memory per node determine the maximum size of your deduplication repository.

The following table illustrates the relationship between these parameters:

Dedupe Memory per Node (GB)	Dedupe Repository Size (TB) – 1 Dedupe Node	Dedupe Repository Size (TB) – 2 Dedupe Nodes	Dedupe Repository Size (TB) – 4 Dedupe Nodes
32	1.7	3.4	6.8
48	3	6	12
64	4.4	8.8	17.6
96	7	14	28
128	9.9	19.8	39.6
192	15.3	30.7	61.4
256	20.8	41.6	83.2
384	31.8	63.6	127.2
512	42.7	85.4	170.8

Backup cache sizing

You need external storage as backup cache to hold virtual tape data. The required capacity can be estimated by adding up the following values:

- Size of incoming backup data awaiting inline deduplication for the unique data to be moved to the deduplication repository. If the input rate is high, the cache must be large enough to buffer data for several days. To be safe, you can estimate the amount as one week's worth of data, considered as *weekly ingest data*.
- Size of VITs, which is estimated as 3% of *weekly ingest data* multiplied by the number of weeks to retain data.
- Size of FVITs in case of incoming replication, which is estimated as 6% of *weekly ingest data* multiplied by the number of weeks to retain data.
- Size of largest data to restore.

Memory sizing

VTL uses system memory for deduplication jobs and by default leaves 12 GB of memory for the operating system and for running processes.

This default is good for most cases, but on large systems with many CPU cores and large amounts of memory, it may not be sufficient when many concurrent deduplication jobs run. In these cases, you may need to increase the system memory.

The memory reserved for deduplication depends on the size of the deduplication repository, where 2 GB memory is used for each 1 TB of deduplication repository capacity.

For example, for a system with 48 GB memory and a 12 TB deduplication repository, 24 GB (2 GB*12) memory is recommended. The rest of the system memory (48 GB-24 GB) can be made available to the operating system.

The VTL environment variable `RDE_RESERVED_SYSTEM_MEMORY` defined in `$ISHOME/etc/.isuperm.env` sets the amount of memory reserved for the operating system:

Using our example, the parameter should be set to 24576 (24*1024):

```
[enabled] [RDE_RESERVED_SYSTEM_MEMORY] [24576] [12288]
```

You may still need to adjust the number based on the load on the system. For example, having hundreds of thousands of virtual tapes, may result in the need for additional operating system memory. Refer to the following for the system memory requirements.

Each VTL server needs:

- 16 GB (predefined minimum value)
- 2 GB memory per 1 TB of capacity for deduplication

From the server memory, reserve the sum of the following values for the operating system:

- 12 GB (predefined minimum value)
- The rounded up value of the server `total memory*0.012`
- 150 MB for each inline deduplication stream – equivalent to the number of tape drives assigned to clients with I/O activity; maximum 1024 streams or 153.6 GB
- 150 MB for each post-deduplication stream; maximum 16 streams or 2.4 GB
- 3 KB for each virtual tape
- 200 MB for each tape migration job to object storage (max 20 jobs at a time)

Appendix

This appendix contains information about:

- [Ports used by VTL/SIR](#)
- [IP address and netmask update](#)
- [Storage LUN migration](#)
- [FIPS security](#)
- [Shared library environment variable](#)
- [Linux auditing](#)
- [Block device support](#)

Port usage

VTL 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. For FalconStor appliances, the ports marked ● are enabled by default.

Port	Protocol	Used By	Description
20	TCP/UDP	FTP client	Standard FTP port used for file data transfer.
21	TCP/UDP	FTP client	Standard FTP port used for sending commands.
22 ●	TCP	Host client	Standard Secure Shell (SSH) port used for remote connection to the server.
25	TCP/UDP	Mail server	Standard SMTP port used for Email Alerts.
80 ●	TCP	FalconStor web license server	Standard Internet port for online registration of license key codes. 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.
123	UDP	NTP server	Standard Network Time Protocol (NTP) transport layer used to access external time servers.
161	UDP	SNMP server	Standard Simple Network Management Protocol (SNMP) port used to query VTL MIBs.
199	TCP	SNMP client	Standard SNMP multiplexing (SMUX) protocol port used to query Dell OpenManage system MIBs.

Port	Protocol	Used By	Description
705	UDP	SNMP client	Standard SNMP AgentIX port used to query agents such as Fujitsu ServerView.
3260	TCP	iSCSI client	Communication port between iSCSI clients and the server.
10161	TCP/UDP	SNMP	SNMP communication port.
11576 ●	TCP	CLI and the FalconStor Management Console	Secure RPC communication port between FalconStor Management Console and the server.
11577 ●	TCP/UDP	Replica servers	Communication port for incoming data replication. This port is only open while replication is being performed.
11579	TCP/UDP	Replication servers	Replication authentication.
11582 ●	TCP	CLI	Communication port for used to send CLI commands to the server.
11583	TCP	FalconStor 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	Replication servers	Communication port between replication servers for data replication of deduplicated data.
11676	TCP	VTL server	Communication port for deduplication management and repository health check.
18651	TCP	Replication servers	Communication port between servers for non-encrypted replication.
18652	TCP	Replication servers	Communication port between servers for encrypted replication.

IP address and netmask update

When replication or deduplication are configured, you cannot change the server IP parameters of the VTL server from the FalconStor Management Console. However, you can use the procedure below to make these changes when replication or deduplication are configured without needing to restart servers or services.

This procedure allows you to change the IP address/netmask and replica IP address/netmask.

Environment This procedure is for VTL servers with IPMI power control.

Detailed instructions The following procedure can be performed on one of the servers in your environment. Configuration files are automatically created with the IP parameter changes you plan to make. Using these configuration files, apply the changes to each server.

1. If applicable, remove all servers from multi-node groups.
2. Stop all client IO.
3. Stop all jobs, such as replication, deduplication, migration, recovery, and space reclamation.

4. Open a command prompt and go to the utility directory.

```
# cd $ISHOME/ci
```

5. Run the `ci-cli` command with the following syntax to generate the configuration files:

```
# ./ci-cli --username=<username> --password=<password>
--server=<ServerIP1>[:<port>] [--server=<ServerIP2>[:<port>] ... ]
--fromip=<FromIP1>[/<FromNetmask1>] --toip=<ToIP1> [/<ToNetmask1>]
[--fromip=<FromIP2>[/<FromNetmask2>] --toip=<ToIP2> [/<ToNetmask2>] ...]
```

The `server` parameters indicate the IP address of each server in your environment. SSH is used to access remote servers via standard port 22. If you use a different port, add it after the server IP address, such as,

`ServerIP:PortNum`. You can specify up to four servers; all servers must use the same username/password.

The `fromip` and `toip` parameters specify the IP address/netmask to change; up to 24 `fromip/toip` pairs can be specified.

The `ci-cli` command generates a `ci-Localcluster-<hostname>.txt` file for each server.

6. Run the `ci-updateinfo` command to apply the changes for each configuration file generated in the previous step.

```
# ./ci-updateinfo <configuration file>
```

This command restarts network services. Wait until services are up before running the command for the next configuration file.

Note that the terminal session will be lost after `ci-updateinfo` is run if the IP address of the terminal session used to run the command is one of the IP addresses being changed.

7. Re-add the servers with the new IP addresses into the FalconStor management console.
8. Suspend each deduplication policy with replication enabled and then resume it.

Storage LUN migration

VTL offers a command line tool to migrate data on a storage device to one or more other devices. This is useful if you are upgrading storage devices to newer or larger devices. With this tool, you can specify the target device(s) or let the system auto-select the first available device(s) that have an equal or larger capacity. 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 device, the source device will be unassigned.

- Requirements** The following requirements must be met before running a migration job:
- The source device must have data; you cannot migrate an empty device.
 - The source device should not belong to a storage pool.
 - The target storage must have enough space. Data is copied over by segments. If the source segment is larger than the target segment, migration will fail. For example, the source device is 10 GB and there are two target devices that are 5 GB each. Even though the total of the two target devices is 10 GB, the job will fail if the first segment on the source device is 6 GB since each target device is only 5 GB and neither is large enough to accept the 10 GB segment.
 - VTL database devices cannot be selected as source or target devices.
 - I/O must be stopped on the source and target device(s) while migration is in progress.
 - Space reclamation must be stopped before migration. Be sure to turn it back on after the migration is complete.

Usage `lunmigration.sh <Source ACSL> <-p Target Storage Pool>|<Targets ACSL> [-f]`

Where:

`Source ACSL` specifies the ACSL of the source physical device to migrate in the format `a:c:s:l`

To identify target devices to use for migration, you can specify a storage pool, a specific physical device, multiple physical devices, or let the server automatically select devices. If one of the specified devices is in a pool, that pool will be used for allocation first:

`-p Target Storage Pool` identifies the storage pool to use for the migration target device.

OR

`Targets ACSL` can have one of the following values:

- `AUTO` automatically selects devices to use
- `a:c:s:l` specifies the ACSL of one device to use
- `a:c:s:l, a:c:s:l, ...` specify multiple devices to use

`-f` forces migration without asking for confirmation.

Stop a job To stop a migration job, run the following command:

```
stop_lunmigration.sh a:c:s:l
```

Where:

`a:c:s:l` specifies the ACSL of the source physical device being migrated.

Restart a job To restart an incomplete/failed migration job from where it left off, run the following clean-up command before re-running the migration job:

```
lmclean.sh
```

If errors occur, rescan physical devices from the console and re-run the migration job.

FIPS security

VTL 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.

The option is enabled when VTL software is first installed on an appliance. If it was not enabled and you want to use it, contact Technical Support for assistance.

Shared library environment variable

LD_LIBRARY_PATH is an environment variable containing a list of directories used when searching for libraries. For security reasons, 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.

Linux auditing

The Linux auditing system comprehensively logs and tracks access to files, directories, and resources on your system.

Auditing may have been enabled during installation but may affect system performance. Run the following commands if you want disable it:

```
# chkconfig auditd off
# service auditd stop
```

Block device support

A block-to-SCSI driver enables block devices to be used for downstream physical devices on VTL servers.

Configuration

Block devices must be listed in the `bksc_option.conf` configuration file.

In order to add a block device to this file, you must know the device's major/minor numbers. You can use the `lsblk` command with the block device's special file to determine them. In the following example, `/dev/hioa` is the block device's special file.

```
[jensd@cen2 ~]$ lsblk /dev/hioa
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
hioa 252:0 0 1.1T 0 disk
└─hioa1 252:1 0 1.1T 0 part
```

The major and minor can then be entered to `bksc_option.conf`. The following line in `bksc_option.conf` configures `/dev/hioa1`:

```
-t block -m 252 -n 1
```

Up to 64 block devices are supported.

All block devices appear on adapter number 98 in the console.

Note: You cannot 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.

Index

A

- Accounts
 - Administrator 47
- Activity Log
 - Properties 54
- Administrator
 - Management 47
 - Types 47
- Administrator accounts 47
- Administrator management 47
- Advanced tape creation 38
- Alarm policies
 - Command line 333
 - Get 333
 - Remove 334
 - Set 333
- Attention required
 - Command line
 - Get 336
- Attention Required tab 31, 52
- Auditing 454
- Auto expansion 89, 91
- Auto migration 88
- Auto Replication 87, 168
- Auto Save Config 56
- Autopathing 77

B

- Backup application server
 - Device scan 25
- Backup software
 - Detect devices 26
 - Discover library 25
- Backups
 - Overview 27
 - Run jobs 27
- Best practices 443
 - Deduplication repository sizing 443
 - Deduplication system sizing 444
 - Index/folder disk sizing 443
 - Memory sizing 445
- Block device 455
 - Configuration 455

C

- Clients 11, 36
 - Add 23

- Command line
 - Add iSCSI client 280
 - Assign iSCSI target 281
 - Create iSCSI target 281
 - Delete 278
 - Delete iSCSI client initiators 282
 - Delete iSCSI target 282
 - Get iSCSI client initiators 282
 - Properties 280
 - Rename 278
 - Virtual device list 278
- iSCSI 185
 - Authenticated access 187
 - Unauthenticated access 187
- Multipathing 99
- Rename 23
- Troubleshooting 387, 388
- Unassign virtual tape libraries 102

COD

- Virtual tapes 89, 91
- Command line 258
 - Alarm policies 333
 - Clients 278
 - Commands 260
 - Common arguments 259
 - Data encryption 315
 - Deduplication 301
 - Import/Export 318
 - Licensing 262
 - Login/logout 260
 - Mirroring 330
 - Object Storage 321
 - Physical devices 264
 - Replication 307
 - Reports 337
 - Server info 261
 - Server licensing 262
 - Support utilities 335
 - Usage 258
 - Users 283
 - Virtual devices 274
 - Virtual library and drives 286
 - Virtual tapes 293
- Compression 19
 - Hardware 103
 - Icon 84
 - Software 103

- Virtual tape drive 103
- Virtual tape replication 117
- Configuration
 - Protect 64
 - VTL 15
- Configuration repository 36
- Connect
 - VTL appliance 14
- Console
 - Activities object 32
 - Backup Server object 31
 - Clients object 36
 - Connect to VTL server 14
 - Connection problems 382
 - Dashboard Summary 33
 - Deduplication Repository 33
 - VTL Performance 33
 - VTL Space Usage 33
 - Deduplication Job Queue 32
 - Deduplication Policies object 35
 - Display problems 385
 - Group object 32
 - Group Reports object 36
 - Install 13
 - Multi-Node Group object 32
 - Overview 29
 - Physical Resources object 36, 70
 - Replica Resources object 35
 - Replication Queue 32
 - Reports object 35
 - Repositories object 36
 - Run 13
 - Server
 - Properties 54
 - Set options 38
 - Status object 33
 - Storage Devices object 74
 - Storage Pools object 80
 - System maintenance 40
 - Tape Import/Export Queue 32
 - Troubleshooting 382
 - Unique Replication Queue 32
 - User interface 29
 - Virtual Tape Drives object 34
 - Virtual Tape Libraries object 34
 - Virtual Tape Library System object 34
 - Virtual Vault object 34

D

- Data encryption
 - Command line
 - Activate data encryption 316
 - Change encryption password 316
 - Enable 315
 - Enable virtual tape encryption 315
 - Get data encryption information 316
- Database 19
- Database Repository 36
- Date
 - System 42
- Deduplication 106
 - Command line
 - Add policy 302
 - Add tape to policy 305
 - Delete policy 304
 - Get tape activity 306
 - List policies 301
 - Remove tape from policy 305
 - Start policy 302
 - Stop policy 302
 - Create policy 22
 - Deduplication methods 107
 - Enable 22
 - Expand deduplication repository 139
 - Icons 85
 - Index pruning 136
 - Inline 108, 110
 - Monitor 124
 - Post-processing 108
 - Reclaim disk space 136
 - Reclamation 136
 - Automatic 138
 - Manual 138
 - Requirements 137
 - Schedule 138
 - Thresholds 137
 - Repository statistics 124
 - Repository usage 124
 - Space reclamation 136
 - Statistics 124
- Tape
 - Create policies 109
 - Modify policy 119
 - Overview 106
 - Policies 109
 - Priority 111
 - Replication 113

-
- Tape policy
 - Active policies 130
 - Add/remove tapes 121
 - Delete 121
 - Event Log tab 133
 - General Info tab 128
 - Manage active 123
 - Manually run 121
 - Resume replication on a target 122
 - Run History tab 132
 - Suspend 121
 - Suspend replication on a target 122
 - Tape history 130
 - Tapes tab 128
 - Turbo 108
 - Virtual tapes
 - Deduplication Policies object 127
 - Statistics 127
 - Deduplication - Policy Status Report 213
 - Deduplication Job Queue 134
 - Deduplication Replication Status Report 217
 - Deduplication repository 36
 - Object storage 171
 - Configuration 172
 - Expansion 172
 - Requirements 171
 - Deduplication Repository - Memory and Space Usage Report 234
 - Deduplication Repository - Reclamation Report 220
 - Deduplication repository encryption 140
 - Deduplication Tape Activity Report 214
 - Deduplication Tape Usage Report 233
 - Device scan 25, 26
 - Devices
 - Rescan 73
 - Disk
 - Replace a physical disk 62
 - Disk compression 19
 - Disk Space Allocation for Virtual Tapes in Libraries Report 239
 - Disk Space Usage History Report 235
 - Disks
 - Troubleshooting 386
 - Downloads 12
 - E**
 - Email Alerts 248
 - Configuration 248
 - Include system log entries 252
 - Message severity 253
 - Modifying properties 254
 - System log check 252
 - Triggers 250, 256
 - Customize email 254
 - New script 256
 - Output 256
 - Return codes 256
 - Sample script 257
 - Encryption 20, 140
 - Activate for a server 143
 - Change encryption activation password 143
 - Deduplication repository 140
 - Enable virtual tape 142
 - Icon 84
 - Keys
 - Change 145
 - Create 144
 - Delete 146
 - Export 146
 - Import 147
 - Manage encryption keys 144
 - Virtual tape encryption 141
 - Virtual tape replication 117
 - Encryption at rest 140
 - Encryption keys
 - Change 145
 - Create 144
 - Delete 146
 - Export 146
 - Import 147
 - Manage 144
 - Error codes 394
 - Event Log 50
 - Command line 336
 - Export 51
 - Filter information 50
 - Print 51
 - Sort information 50
 - F**
 - Find
 - Virtual tapes 96
 - FIPS compliance 452
 - Firewall
 - Ports 446
 - Firmware
 - Change 104

G

- Glossary 12
- Groups 66
 - Add servers 67
 - Benefits 66
 - Create 67
 - Management 66
 - Remove a server 69

H

- Halt server 43
- Hostname 18
 - Change 42

I

- IBM System i configuration 196
- Icons
 - Physical resource 70
 - Virtual tape 84
- Import/export
 - Command line
 - Cancel jobs 320
 - Delete jobs 319
 - Job status 318
 - Restart jobs 319
 - Resume jobs 318
 - Suspend jobs 319
- Import/Export Jobs Report 221
- Index pruning 136
- Index/folder repository 36
- Inline deduplication 108, 110
- IP address
 - Change 448
- iSCSI Target Mode 185
 - IBM i 192
 - Initiators 185
 - Linux
 - Add iSCSI client 191, 192
 - Configuration 190
 - Create targets for iSCSI client 191
 - Enable 190, 192
 - Log client onto target 191
 - Prepare iSCSI initiator 190
 - Unassign targets 102
 - Targets 185
 - Users 185
 - Windows 187
 - Configuration 186
 - Enable 186

- Requirements 186

- iSCSI targets
 - Unassign 102

J

- Jumbo frames 17, 41

K

- Keycodes
 - Command line
 - Add 262
 - Get information 262
 - Register 263
 - Remove 262
 - Offline registration 382
 - Troubleshooting 382

L

- LD_LIBRARY_PATH
 - Set 453
- License keycodes
 - Offline registration 382
 - Troubleshooting 382
- Licenses 44
 - Register 44
- Licensing 15
 - Command line
 - Add keycode 262
 - Get keycode information 262
 - Register keycode 263
 - Remove keycode 262

Linux

- Auditing 454
- Location 57
- Logical resources
 - Troubleshooting 386
- LUN reservation 19
 - Command line
 - Change reservation 267
- LUNs Report 237

M

- Migration
 - Libraries to storage pools 81
 - Storage devices 450
- Mirror
 - Fix minor disk failure 62
 - Remove configuration 63
 - Replace a physical disk 62

-
- Replace disk in active configuration 62
 - Replace failed disk 61
 - Repository disks 60
 - Status 61
 - Swap 62
 - Mirroring
 - Command line
 - Create 330
 - Remove 331
 - Status 330
 - Swap 331
 - Sync 332
 - MTU 17, 41
 - Multi-node groups 66
 - Add servers 67
 - Benefits 66
 - Create 67
 - Management 66
 - Remove a server 69
 - Multipathing 99
 - N**
 - NAS
 - Performance 53
 - Space usage 46
 - Netmask
 - Change 448
 - Network configuration 15
 - Network Time Protocol 42
 - NTP 42
 - O**
 - Object recovery jobs 173
 - Object storage 171
 - Account 179
 - Add account 179
 - Command line 321
 - Add account 321
 - Convert migrated tape to stub tape 328
 - Delete 326
 - Display information 326
 - Get job status 328
 - Manage jobs 329
 - Modify information 324
 - Recover tape from object storage 327
 - Start object migration 327
 - Deduplication data repository 171
 - Manage account 184
 - Migration 88
 - Migration icon 84
 - Object storage account
 - Amazon Web Services (AWS) S3 180
 - Generic S3 183
 - Hitachi Content Platform (HCP) 181
 - IBM Cloud Object Storage (COS) 182
 - Microsoft Azure 179
 - Object Storage Migration Jobs Report 222
 - Offline tapes
 - Troubleshooting 386
 - P**
 - Passwords
 - Add/delete administrator password 47
 - Change 48
 - Change administrator password 47
 - Default 14, 48
 - Strong 47
 - Enable 48
 - Unlock 49
 - Patches 12
 - Apply 58
 - Rollback 58
 - Performance 56
 - Performance statistics 53
 - Physical device
 - Change LUN reservation 75
 - Command line
 - Delete 268
 - Get adapter information 265
 - Get device information 264
 - Prepare disk 266
 - Rename 267
 - Rescan 265
 - Storage allocation 273
 - Storage pool
 - Add physical devices 270
 - Assign ACL 272
 - Create 268
 - Delete 270
 - Get list 271
 - Get users 284
 - Remove physical devices 271
 - Rename 272
 - Unassign ACL 272
 - Filter 75
 - Throughput 76
 - Physical devices
 - Rescan 73

- Unassign 71
- Virtualize 71
- Physical Resource Allocation Report 242
- Physical resource icons 70
- Physical resources 70
 - Troubleshooting 385
- Physical Resources Configuration Report 243
- Physical Resources object 70
- Physical storage device
 - Prepare 71
- Ports 446
- Preferred management IP address 56
- Prepare physical devices 71
- Prepare physical storage devices 71
- Prepare virtual devices 19
- Protect
 - Configuration 64
 - Server configuration 56

R

- Reboot server 43
- Reclamation
 - Automatic 138
 - Command Line 138
 - Command line 301
 - Manual 138
 - Requirements 137
 - Run 138
 - Thresholds 137
- Recovery 60
- Release notes 12
- Remote Copy 169
- Replica
 - Preferred management IP address 56
- Replica resources 150, 157
- Replication 149
 - Auto Replication 168
 - Cascaded 113
 - Command line
 - Create replica 307
 - Demote in test mode 313
 - Get properties 311
 - Get status 312
 - Network throttling 314
 - Promote in test mode 313
 - Promote replica 308
 - Remove 309
 - Resume 310
 - Set properties 310

- Start 312
- Stop 312
- Suspend 309
- Concurrent 114
- Deduplicated tapes 160
 - Access data on tape 167
 - Add target server 163
 - Edit target 165
 - Remove configuration 167
 - Requirements 162
 - Setup 163
 - Status 167
 - Stop in progress 167
 - Throttle 166
- FVIT 160
- LVIT 160
- Methods 149
- Parallel 114
- Remote Copy 169
- Serial 114
- SIR Replication tab 31
- Troubleshooting 389
- Virtual tapes 113, 150
 - Advanced 113
 - Cascaded mode 113
 - Change configuration 159
 - Compression 117
 - Concurrent 114
 - Encryption 117
 - Force 159
 - Local 150
 - Parallel mode 114
 - Policies 154
 - Primary tape 150
 - Promote 158
 - Remote 150
 - Remove configuration 159
 - Replica resource 150
 - Requirements 151
 - Resume schedule 159
 - Serial 114
 - Setup 151
 - Single mode 113
 - Start manually 159
 - Status 157
 - Stop in progress 159
 - Suspend schedule 159
 - Throttle 157
 - Troubleshooting 389

Replication Status Report 224
Replication throttle 56
Reports 35, 36
 Command line 337
 Create 204
 Deduplication - Policy Status 213
 Deduplication Replication Status 217
 Deduplication Repository - Memory and Space Usage 234
 Deduplication Repository - Reclamation 220
 Deduplication Tape Activity 214
 Deduplication Tape Usage 233
 Delete 212
 Disk Space Allocation for Virtual Tapes in Libraries 239
 Disk Space Usage History 235
 Email 212
 Export data 212
 Import/Export Jobs 221
 LUNs 237
 Object Storage Migration Jobs 222
 Physical Resource Allocation 242
 Physical Resources Configuration 243
 Properties 210
 Refresh 212
 Replication Status 224
 Retention 210
 Schedule 206
 Storage Pools Configuration 244
 View 209
 Virtual Library and Drive Assignments 225
 Virtual Library Information 226
 Virtual Tape Activity 227
 Virtual Tape Information 229
 VTL Performance 245
Reports object 203
Repository
 Configuration 36
 Database 36
 Deduplication 36
 Index/folder 36
Rescan 73
Rescan physical devices 73
Reserve
 LUNs 19
Restore configuration 64

S

Save configuration 64

Search
 Virtual tapes 96
Security
 Ports 446
 System 446
Server
 Attention Required tab 31
 Event Log 31
 General info 31
 Location 31, 57
 Performance statistics 53
 Properties 54
 Server commands 200
 SIR Replication tab 31
 Stop processes 200
 Version info 31
 VTL server modules 201
Setup
 Confirm successful backup 28
Shared library environment variable
 Set 453
Shred
 Virtual tape 104
Sizing guidelines 443
SNMP
 Deduplication MIBs 371
 Event Log messages 354
 Integration 354
 MIBs 354
 Traps 54, 354
 VTL MIBs 355
Software updates
 Add patch 58
 Rollback patch 58
Sort
 Virtual tapes 97
Space reclamation 136
Space usage 46
Standalone tape drive
 Command line
 Create 290
Storage Devices object 74
Storage LUN migration 450
Storage monitoring 57
Storage pools 80
 Access rights 80
 Configuration report 244
 Create 80
 Migrate existing libraries 81

- Update 81
- Storage Pools Configuration Report 244
- Strong passwords 47
 - Enable 48
- Support portal 12
- System i configuration 196
- System maintenance 40
 - Halt 43
 - Reboot 43
 - Restart network 43
 - Restart server 43
 - Set hostname 42
- System preferred path 79

T

- Tape backup servers
 - Add 23
- Tape capacity-on-demand 89, 91
- Tape expansion
 - Troubleshooting 386
- Tape Import/Export Queue 32, 173
- Tape migration 173
 - Account 179
 - Configuration 175
 - Convert to stub tape 176
 - Jobs 173
 - Manage migrated/stub tapes 177
 - Migrate to object storage 176
 - Recover data 177
- Terminology 12
- Time
 - System 42
- Troubleshooting 382
 - Clients 387, 388
 - Console 382
 - Offline disks 386
 - Offline tapes 386
 - Physical resources 385
 - Replication 389
 - Tape expansion 386
 - VIT 388
- Turbo deduplication 108

U

- Unique Replication Queue 167
- User name
 - Default 14
- Users 47
 - Add or modify 48

- Command line
 - Add 283
 - Delete 283
 - List 284
 - Reset password 285
 - Set password 284
- Unlock account 49

V

- Virtual devices
 - Command line
 - Assign to iSCSI client 275
 - List 274
 - Unassign 276
 - Prepare 19
- Virtual index tape
 - Status 135
- Virtual Library and Drive Assignment Report 225
- Virtual Library Information Report 226
- Virtual Tape Activity Report 227
- Virtual tape drives
 - Assign to clients 99
 - Command line
 - Add 289
 - Get supported 286
 - Compression 103
 - Create standalone 91
 - Unassign from clients 102
- Virtual tape encryption 20, 141
- Virtual tape icons 84
- Virtual Tape Information Report 229
- Virtual tape libraries
 - Assign to clients 21, 99
 - Client multipathing 99
 - Command line
 - Create 287
 - Delete 289
 - Get supported 286
 - Create 21, 86
 - Properties 102
 - Unassign from clients 102
- Virtual tape library/drive
 - Firmware 104
- Virtual tapes
 - Advanced replication 113
 - Advanced tape creation 92
 - Auto load 97
 - Auto Replication 168
 - Auto unload 97

- Barcode 97
- Cascaded mode replication 113
- Command line
 - Copy 299
 - Create 293
 - Delete 296
 - Get tape info 293
 - Move 297
 - Set properties 295
 - Shred 300
- Create 92
- Display 96
- Dynamic LUN Allocation 95
- Filter 96
- Find 96
- How they are allocated 95
- Locate 96
- Move to virtual vault/slot/drive 97
- Parallel mode replication 114
 - Concurrent 114
 - Serial 114
- Properties 97
- Remote Copy 169
- Replication 150
- Round Robin Logic 95
- Shred 104
- Single mode replication 113
- Sort 97
- WORM 12, 92
- Write protect 34, 97
- Virtual vault 34
- VIT
 - Status 135
 - Troubleshooting 388
- VLAN tagging 41
- VTL
 - Components 11
 - Configure 15
 - Performance 53
 - Space usage 46
- VTL appliance
 - Connect 14
- VTL console 11
- VTL info
 - Command line 275
- VTL Performance Report 245
- VTL server
 - Stop processes 200

W

- WORM 12, 92
- Write protection 34

X

- X-ray 392
 - Command line 335