

FALCONSTOR®

FalconStor Software®

StorSafe® with StorSight®

USER GUIDE

FalconStor® StorSafe® with StorSight® User Guide

StorSafe 11.14, StorSight 10.17

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 noted as “updated”.

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

FalconStor Software, Inc.
111 Congress Ave, Suite 500
Austin, TX 78701 USA
Phone: 631-777-5188
Technical Support: 631-777-3332
Web site: www.falconstor.com

Copyright © 2020-2025 FalconStor Software, Inc. All Rights Reserved.

FalconStor®, FalconStor Software®, StorSafe®, StorGuard™, and StorSight® are trademarks of FalconStor, 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, Inc. 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, Inc. 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

Architecture	18
Terminology	20
User access icons in this guide	22
Download software and documentation	23
Additional resources and helpful links	23

StorSight Web Portal

Portal tabs	25
Dashboard	25
Customize the dashboard	26
Widgets	27
Monitor	31
Manage	32
Server info	33
Server icons	33
Port and server module status	33
Analytics	34
Administration	35
Global portal functions	37

Administration

StorSight	39
License	40
Offline registration	40
License renewal	41
Multi-factor authentication	42
Enable strong passwords for StorSight users	43
Patches	44
Apply patch	44
Roll back patch	44
X-ray	44
HTTP proxy	45
Restart/reboot StorSight (updated October 2024)	46
Set automatic logout timeout for the server	46
Portal notification messages	46
Set global notification message	46
Set a server-specific notification message	47
Servers	48
Servers page	48
View capacity information (updated October 2024)	49
Add a server	51
Reset server password	52
Disconnect a server	52
Force rediscovery of a server	53

Lock a server	53
Delete a server	53
Customers	54
Multi-tenant environments	56
Multi-tenancy example	57
Multi-tenancy rules and restrictions	59
Customers page	60
Add a customer	60
Assign resources to a customer	60
Unassign resources from a customer	61
Synchronize users	61
Reset authentication	61
Delete a customer	61
Users	62
User roles	62
Users page	64
Add a user	64
Update a user	65
Synchronize users	65
Delete a user	66
AD/LDAP settings	67
Active Directory	68
Windows 2019 Active Directory configuration example	69
LDAP	71
Local	73
Self-signed certificates	74
Alert notifications	75
Email settings	75
Email format	76
SMS settings	77
Alert scripts	78
Add a script	78
Manage scripts	78
Global settings	79
License usage prediction	80

Configuration

Overview	81

Physical Resources

Storage pools	84
Storage pool access rights	85
Storage Pools page	86
Storage Pools page icons	87
Create a storage pool	87
Update storage pool settings	88

Delete a storage pool	88
Migrate disks to a storage pool	89
Physical devices	90
Physical Devices page	90
Physical Devices page icons	91
Rescan physical devices	92
Set autopathing	93
Rename a physical device	94
Prepare physical storage devices	95
Change the reservation of a physical device	96
Test a physical device	96
Adapters	97
Adapters page	97
Adapters page icons	98
Rescan adapters	98
Update adapter settings	99
Display SNS information	100
Storage web links	101
Storage Web Links page	101
Storage Web Links page icons	101

Virtual Tape Libraries

Tape libraries	102
Libraries page icons	103
Create a virtual tape library	104
Update virtual tape libraries	107
Delete virtual tape libraries	107
Create virtual tapes	107
How virtual tapes are allocated	109
Load virtual tapes	110
Unload virtual tapes	110
Reset a library	110
Add virtual tape drives	110
Configure Cloud Tiering for a virtual tape library	111
Tape drives	113
Drives page icons	114
Create a standalone virtual tape drive	114
Update virtual tape drives	116
Load virtual tapes	116
Delete virtual tape drives	116
Create a virtual tape	116
Vault	117
Vault page icons	118
Manage tapes in the vault	119
Shred a tape	119
Perform tape caching functions	120
Migrate virtual tape data to object storage	120
Convert a virtual tape to a stub tape	120

Recover data from object storage	121
Export data to physical tape	122
Display virtual tape information	122
Tapes	123
Tapes page icons	124
Filter tapes	125
Create virtual tapes	128
Update virtual tapes	128
Manage virtual tapes	128
Delete virtual tapes	129
Perform tape caching functions	129
Configure replication for virtual tapes	130
Start replication for a virtual tape	130
Suspend/resume replication for a virtual tape	130
Display virtual tape information	130

Physical Tape Libraries

Physical tape libraries	131
Physical Tape Libraries page icons	132
Inventory a physical library	132
Enable/disable a physical library	132
Reset a physical library	133
Scan a physical library for stacked tapes	133
Move tapes in a physical library	133
Assign a physical library	133
Unassign a physical library	134
Import data from physical tapes to a virtual library	134
Import data from a tape in another virtual tape library	135
Assign ACSLS/Library Station	136
Hardware configuration	136
Configure StorSafe to work with ACSLS	137
Configure ACSLS to work with StorSafe Failover	138
Set eject policy	138
Filter tapes displayed in the portal	139
Add/remove tapes	139
Replace ACSLS physical drive	140
Physical tape drives	141
Physical Tape Drives page icons	142
Enable/disable a physical tape drive	142
Prepare for stacking	142
Retrieve barcode	143
Scan a standalone physical drive for stacked tapes	143
Import data from physical tapes to a virtual library	143
Physical tapes	145
Physical tape export history	146
Export History page icon	146

Host Clients

Host Clients page	147
Host Clients page icons	148
Add iSCSI clients	149
Create an iSCSI target for an iSCSI client	150
Add Fibre Channel clients	151
Associate World Wide Port Names (WWPN) with clients	152
Manage clients	154
Assign virtual tape libraries and drives to clients	154
Client multipathing	156
Update client properties	157
Unassign virtual tape libraries and drives from clients	157

Fibre Channel Configuration

Configure Fibre Channel on server	160
Ports	160
QLogic drivers	160
Zoning	161
Persistent binding	162
QLogic adapter settings	162
Configure Fibre Channel on storage	166
Load balance paths for downstream storage	166
Configure Fibre Channel on host clients	166

iSCSI Configuration

iSCSI users	167
Windows configuration	168
Requirements	168
Enable iSCSI	168
Prepare client initiators	168
Add an iSCSI client	169
Create targets for the iSCSI client to log onto	169
Assign a virtual tape library to the iSCSI target	169
Log the client onto the target	169
Linux configuration	170
Requirements	170
Enable iSCSI	170
Prepare client initiators	170
Add an iSCSI client	171
Create targets for the iSCSI client to log onto	171
Assign a virtual tape library to the iSCSI target	171
Log the client onto the target	171
IBM i configuration	172
Requirements	172
Enable iSCSI	172

Prepare client initiators	172
Add an iSCSI client	172
Create iSCSI target	172
Assign a virtual tape library to the iSCSI target	175
Log the client onto the target	175
AIX configuration	176
Requirements	176
Enable iSCSI	176
Add an iSCSI client	176
Create targets for the iSCSI client to log onto	177
Assign a virtual tape drive to the iSCSI target	177
Run a backup on a AIX client	177

Failover

StorSafe/StorSafe-S failover	178
Failover triggers	179
Fibre Channel Target failure and Fibre Channel clients	180
Failover terminology	181
Failover behavior	182
IP connectivity	183
Fibre Channel connectivity	184
Resume backups after failover/failback	185
Failover requirements	187
General requirements (updated October 2024)	187
Network requirements	188
Fibre Channel/iSCSI requirements	188
Storage requirements	189
Power control management options	189
Backup application server failover configuration	189
Configure failover	190
Manage FC switch port swapping	194
Check failover status	196
Change your failover configuration	197
Check failover settings	197
Change failover options	197
Change the power control password	197
Start/stop failover or recovery	198
Start takeover	198
Stop takeover	198
Suspend/resume failover	198
Recover from failover	199
Remove a failover configuration	199
SIR cluster failover	200
Failover terminology	200
Failover triggers	201
Failover behavior	201
Server behavior during and after failover	201
Failover requirements	202

Configure failover	204
Start nodes in a deduplication cluster	206
Check failover status	206
Change the power control password	206
Suspend/resume failover	207
Remove a failover configuration	207

Deduplication

How virtual tape deduplication works	208
How NAS deduplication works	209
Deduplication methods - at a glance	210
Deduplication Policies page	211
Deduplication Policies page icons	212
Create a deduplication policy	212
Update deduplication policies	217
Manage deduplication policies	217
Delete deduplication policies	218
Deduplication Tapes page	219
Deduplication Tapes page icon	221
Add tapes to a deduplication policy	221
Remove tapes from a deduplication policy	221
Run tapes in a deduplication policy	222
Schedule Deduplication Policy Tapes report	222
Deduplication Activity page	223
Schedule Deduplication Activity reports	224
Deduplication Run History page	225
Schedule Deduplication Run History report	226
Deduplication Tape History page	227
Schedule Deduplication Tape History report	228
Replication Targets page	229
Add a replication target	229

Replication

Deduplicated tapes	232
Replication requirements for deduplicated tapes	234
Configure replication for deduplicated tapes	235
Access data on a replicated VIT	235
Tapes without deduplication	236
Tape replication	236
Replication requirements for virtual tapes	236
Configure replication for virtual tapes	237
Promote a virtual tape replica	239
Manage replication for virtual tapes	240
Auto Replication	241
Tape Copy	243

Cloud Tiering

Cloud tiering policies	244
Migration and recovery jobs (updated October 2024)	245
Immutable object storage	246
Object storage providers	246
Configuration	247
Object Storage Provider page	248
Object Storage Provider page icons	248
Add an object storage account	249
Manage object storage accounts	258
Cloud Tier page	259
Cloud Tier page icons	260
Convert a virtual tape to a stub tape	260
Recover data from object storage	260
Recover data to a new StorSafe server instance	261

NAS

NAS resources	262
NAS Resources page icons	263
Create a NAS resource	264
Delete a NAS resource	265
Expand a NAS resource	265
Unmount/mount a NAS resource	265
Format a NAS resource	266
Check the file system on a NAS resource	266
Exclude/include a NAS resource from deduplication	266
Exclude/include a NAS resource from global replication	266
Copy NAS files and directories using Fast Copy	267
Display information about a NAS resource	267
Shares	268
Shares page icons	269
Create a NAS share	270
Update a NAS share	272
Delete a NAS share	272
Exclude/include a NAS share from deduplication	272
Exclude/include a NAS share from global replication	273
Set properties for a NAS share	274
Show files for a NAS share	275
Display statistics for a NAS share	275
Map/mount shares	275
Windows clients	275
NFS clients	276
Folders	279
Folders page icons	280
Manage folders and shares	280
Exclude/include a NAS folder or share from deduplication	280
Set properties for a NAS folder or share	281

Display statistics for a NAS folder or share	281
SMB clients - Users	282
SMB User page icons	283
Create an SMB user	283
Update an SMB user	284
Delete an SMB user	284
Set the authentication mode	284
Update the domain controller	287
Sync SMB Clients	287
Set Windows ACLs	288
Set SMB options (updated October 2024)	288
Display SMB connection status	290
Update access rights for SMB client shares	290
SMB clients - Groups	291
SMB Groups page icons	292
Create an SMB group	292
Update an SMB group	293
Delete an SMB group	293
Update access rights for SMB group shares	293
NFS clients	294
NFS Clients page icons	295
Create an NFS user	295
Update an NFS client	296
Delete an NFS client	296
Update access rights for NFS client shares	296
Replication - Configuration	297
Replication Configuration page icons	298
Enable outgoing replication	298
Edit the outgoing replication configuration	300
Set outgoing replication throttle	301
Disable outgoing replication configuration	301
Configure incoming replica	302
Purge replica files	302
View purge results	303
Refresh replication session status	303
Replication - Policies	304
Replication Policies page icons	305
Create a NAS replication policy	305
Update a NAS replication policy	307
Delete a NAS replication policy	307
Start replication of a NAS policy	307
Synchronize files	307
Suspend/resume replication schedule	308
Integrity Check	309
Integrity Check page icons	309
Configure integrity checking	310
Update the integrity check configuration	312
Disable the integrity check configuration	312
Start/suspend integrity check	312

Refresh status of integrity checking	312

Physical Tape Export

Tape Caching	313
Tape caching policies	314
Create/update a tape caching policy	315
Tape caching triggers	315
Space reclamation policy	317
Space reclamation triggers	318
Create a cache for your physical tapes	318
Create uncached virtual tapes	319
Auto Archive	320
Tape Duplication	322
Remote Export	324
Manual Export	325

Encryption

Encryption page	329
Encryption page icons	330
Enable virtual tape encryption for a StorSafe server	330
Activate encryption for a server	331
Change the encryption activation password for a server	331
Encryption keys	332
Encryption Keys page icons	332
Manage encryption keys	333
Create a key	333
Update a key name or password	334
Delete a key	334

Server Settings

Failover	335
Manage deduplication	335
Associate a StorSafe server with an SIR server	335
Discover deduplication data disks	336
Dissociate a StorSafe server from a deduplication cluster	336
Restart network	336
Restart services	337
Reboot	337
Shutdown	337
Take X-ray	337
Download event log	338
Manage Patches	339
Manage network settings	340
Configure network bonding	342
Configure VLAN tagging	344

Change hostname	345
Configuration repository	346
Create the configuration repository	346
Manage the configuration repository	346
Tape database	348
Create the tape database	348
Manage the tape database	349
Deduplication repository	350
Manage the deduplication repository	353
Expand the deduplication repository	353
Reclaim deduplication repository disk space	354
Switch to a flexible storage configuration	357
Mirror a device	357
View the physical layout for a device	357
Delete the deduplication repository	357
Local users	358
Manage iSCSI users	358
Change root password	358
Modify system time	358
Enable Fibre Channel and/or iSCSI target mode	359
Configure SNMP	359
Configure Auto-Save	360
Save server configuration	361
Information and requirements	361
Recover server configuration	362
Information and requirements	362
Recover a standalone server	362
Recover a server in a failover environment	363
After restoring your configuration	364
Set replication throttle	365
Set global virtual tape library system properties	366
Configure Hosted Backup	367
Disable Hosted Backup	367
NDMP	368
Configure NDMP	368
OST Settings	370
Configure the Activity Log	370
Set the storage usage threshold	370
Encryption Support	371
Enable NAS	371
Enable strong passwords	372

Job Status

Deduplication Jobs page	373
Deduplication Jobs page icons	374
Manage deduplication jobs	374
Replication Jobs page	375
Replication Jobs page icons	376

Manage replication jobs	376
Import/Export Jobs page	377
Import/Export Jobs page icons	378
Manage import/export jobs	378
Set retry parameters	378
Schedule Import/Export Jobs report	379

Analytics

Reports	381
Reports page	381
Inventory reports	382
Capacity Management reports	385
Performance Monitoring reports	388
Service Status reports	391
Server Load reports	393
Trend reports	395
Modify a trend report	397
Top N Components reports	398
Custom trend reports	399
Create a report	399
Trend report selection table	400
Capacity	400
Load	402
Performance	403
Scheduled reports	406
Schedule a report	407
Smart rules	408
Smart rule behavior	408
Create a smart rule	408
Modify smart rules	410
Delete a smart rule	410
General rules	411
Create a general rule	412
Update a general rule	415
Delete a general rule	415
Update multiple general rules	416
Set global time exclusions	416
Apply rules to multiple servers	417
Alerts	418
Alert descriptions	419
Manage alerts	421
Capacity prediction	422
Audit log	424
Chargeback	426
Chargeback settings	426
Global settings	426
Customer settings	427
Chargeback report	427

Monitor

Overview page	428
Server view page	429
Overview page	429
Server Overview	429
Storage Usage	430
Usage History	430
Deduplication Usage	431
Virtual Tape Libraries page	432
Libraries page	432
Drives	433
Vault	433
Cloud Tier	434
Tapes	435
Deduplication page	436
Policies	436
Tapes	436
Activity	437
Run History	438
Tape History	439
Replication Targets	439
Jobs page	440
Deduplication	440
Replication	440
Unique replication	440
Import/Export	441
NAS Dedupe/Replication	441
NAS In-Progress Replication	442
Object Storage page	442
NAS page	442
NAS resources	442
Shares	443
NFS Clients	443
SMB Clients	443
Host Clients page	445
Performance	445
Drill Down	446
Physical Resources page	447
Storage Pools	447
Physical Devices	447
Physical Tape Libraries	447
Physical Tape Drives	448
Adapters	448
FC/iSCSI Targets page	449
Client View page	450

IBM i System Configuration

Overview	451
Before you begin	452
Set up the tape library	452
Import cartridges	454
Export cartridges	454

SNMP Integration

Enable SNMP	456
VirtualTapeLibraryMIB	458
falCVtlMonitorMIB - falCVtlMonCapacity	458
falCVtlMonitorMIB - falCVtlMonPerformance	459
falCVtlHistoryMIB	461
falCVtlServer	464
falCVtlLibrarySystem	466
falCVtlPhysicalResources	474
falCVtlSanClients	476
Deduplication Repository MIBs	478
falCSirMonitorMIB	478
falCSirCluster - falCSirClusterConf	481
falCSirCluster - falCSirClusterStats	485
Common MIBs	489
CommonMIBs-Traps	490

Command Line

Usage	491
Common arguments	492
Commands	493
Server login/logout	493
Server info	494
Server licensing	495
Physical devices	497
Virtual devices	507
Clients	511
Users	516
Virtual libraries and drives	519
Virtual tapes	527
Deduplication	535
Replication	542
Data encryption	550
Physical libraries and drives	553
Physical tapes	554
Automated tape caching	557
Import/Export	562
Object Storage	570

Mirroring	579
Hosted Backup	582
Failover	583
Alarm policies	584
NAS devices	586
NAS shares	589
SMB/NFS clients	594
NAS information	598
NAS utilities	600
Support utilities	602
Reports	604

Best Practices

StorSight deployment	625
Storage considerations	625
Security considerations (updated October 2024)	625
Protect and recover the StorSight server (updated January 2025)	626
Deduplication repository sizing	630
Index/folder disk sizing	630
CPU cores	630
Deduplication system sizing	631
Backup cache sizing	631
Memory sizing	632
StorSafe server memory	632
StorSafe-S server memory	633
StorSafe-S with StorSight server memory	633
SIR server memory	634

Troubleshooting

Product registration	635
Physical resources	636
Logical resources	637
Fibre Channel connectivity issues	641
Fibre Channel connectivity and Solaris clients	642
NDMP	643
Replication	644
Replication process	644
Deduplication	646
Import/Export	647
NAS resources/shares	648
Client cannot access NAS resources or shares	648
SMB or NFS client.	648
SMB client (Domain/Active Directory mode).	648
SMB client	648
NFS client	649
NFS client requires remapping user IDs to nobody	650

Domain controller authentication issue	650
StorSafe server cannot get some users/groups from the domain controller	650
Problem migrating data with ACLs and extended attributes	650
Truncated folder names	652
Permission denied to folders	652
File access issues	652
System event messages	653
Error codes	654

Appendix

Port usage (updated October 2024)	731
IP address and netmask update	735
FSOST configuration	738
Storage LUN migration	739
FIPS security	741
Shared library environment variable	742
Linux auditing	743
Block device support	744
Configuration	744
HTTPS connection for StorSight with an SSL certificate	745

Index.	746

Introduction

FalconStor StorSafe® provides high-speed backup/restore, global data deduplication, enterprise-wide replication, and long-term cloud-based archive, without requiring changes to the existing environment.

StorSafe delivers data portability, retention, and cost reduction. It delivers seamless integration with S3 compatible clouds, as well as integration with Hitachi HCP, IBM COS, and other object storage systems.


StorSafe integrates with legacy backup and archive software and provides physical tape emulation and advanced data deduplication technology. Using a FalconStor StorSafe Single Instance Repository (SIR) server for block-level deduplication, the solution removes redundant copies of data, thereby reducing capacity requirements and improving data portability in the cloud.

At the end of a backup job, the tape is ejected to the vault, where the tape data can be compressed and migrated to the object storage of a cloud provider.

If you already have physical tapes that you would like to protect, you can easily import data from physical tapes to StorSafe. This provides an inexpensive solution for preserving data and renewing backups in the cloud where no physical tape libraries, drives, or tapes are no longer required.

Architecture

FalconStor Enterprise Solutions consist of the following components:



FalconStor StorSight

StorSight is a single integrated platform that simplifies the management of data across legacy, modern, and virtual storage environments. StorSight gathers and consolidates information coming from different StorSafe servers into a scalable repository of services, tenants, users, predictive analytics, alert rules, reporting, and historical data. StorSight provides a web-based portal for centralized management and monitoring of multiple storage, backup, and deduplication servers.

StorSight uses Representational State Transfer (REST) Application Program Interfaces (APIs) to:

- Manage StorSight users
- Manage servers for a user
- Provide a proxy REST API to allow for servers residing inside a firewall that are being managed by StorSight
- Provide statistical data retrieved from managed servers

FalconStor StorGuard Server

The FalconStor StorGuard Server is a managed storage server that provides storage virtualization and business continuity services for continuous availability to business data in virtual and physical environments.

FalconStor StorSafe Server

The FalconStor StorSafe Server provides tape library emulation, high-speed backup/restore, data archival to supported S3 clouds for long-term storage, network-attached storage (NAS), global data deduplication, and enterprise-wide replication. A deduplication repository can be on the StorSafe server itself as a standalone server (StorSafe-S) or it can be on a separate SIR server, associated with StorSafe.

FalconStor StorSafe Single Instance Repository Server

The FalconStor StorSafe Single Instance Repository (SIR) server is a standalone server that offers block-level data deduplication to FalconStor StorSafe servers.

Host Clients

Host clients are the backup application servers that are assigned to virtual tape libraries and drives for backup.

Terminology

Auto Archive	Auto Archive writes data to physical tape whenever a virtual tape is moved to an IE slot (the vault) from a virtual library after a backup.
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.
NAS	Network-attached storage provides file-level data storage to host clients using the Server Message Block (SMB) and Network File System (NFS) protocols. Inline deduplication is automatically used for NAS resources.
NAS resource	NAS resource file systems are virtualized disks that are used to create NAS shares.
Remote Export	Remote Export automatically exports the contents of a replicated virtual index tape (VIT) to a physical tape at the remote site
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.
Tape Caching	Tape Caching enhances the functionality of StorSafe by acting as a cache to your physical tape library, providing transparent access to data regardless of its location. With tape caching, tapes always appear to be inside virtual libraries and are visible to the backup application regardless of whether the data is actually on disk or on physical tape.
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/exported to a WORM tape. Importing a WORM tape requires a virtual WORM tape. Exporting a WORM tape requires a physical WORM tape; you cannot export a virtual WORM tape to a non-WORM physical tape. You can export virtual WORM tapes and non-WORM tapes to a physical WORM tape. The non-WORM tapes will inherit the WORM properties of the physical tape. If they are imported, they will become virtual WORM tapes. A WORM tape can be replicated to a target server. The replica tape will maintain the WORM property.

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

WORM tapes cannot be used in a Tape Caching configuration.

User access icons in this guide

For quick at-a-glance recognition, this guide uses icons to identify the user access level required to complete tasks in the portal. The table below describes the access level icons.

Icon	Description
	Super Administrator (green with three stripes)
	Domain Administrator (exclusive mode) (dark blue with two stripes)
	Domain Administrator (shared mode) (light blue with one stripe)
	User (dark gray)
	Viewer (light gray with glasses)

Refer to [“Users”](#) for more information regarding user access.

Download software and documentation

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

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

Additional resources and helpful links

The following links may be helpful:

Site	Link
FalconStor Certification Matrix	https://www.falconstor.com/support/certification-matrix/
FalconStor Technical Support	https://www.falconstor.com/support/technical-support
FalconStor license registration email address	activate.keycode@falconstor.com

StorSight Web Portal

The StorSight web portal allows you to connect to each server for management and monitoring purposes. It is designed to run on any HTML5-capable browser.

To access the portal:

1. Open a browser and type the URL of your FalconStor StorSight Server.

For example, 10.1.1.1

2. Enter your username, domain, and password to log in.

The default login is:

- Username: superadmin
- Domain: <blank>
- Password: freestor

If this is the first time you are connecting to this StorSight server, the end user license agreement will appear. Click *Agree and continue* if you agree.

Upon login, each connected server is queried to gather real-time server status.

Administrators, Users, and Viewers can only see servers that were assigned to them by the Super Administrator.


Portal tabs

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

Dashboard


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

The dashboard is composed of *widgets* (standalone Web application components) that can be added, moved, configured, or removed. Each user can add his or her own tab of widgets in the dashboard.



The dashboard is customizable and can provide monitoring capabilities for the system, displaying current status and historical trends, allowing you to easily identify unhealthy servers, busiest servers, view individual server capacity usage, as well as individual server performance, load, etc. This information is updated every 10 seconds.


Customize the dashboard



You can do the following on the dashboard:

- Add a tab - Click the “+” icon next to the last tab. Widgets can then be added to the tab. You can have up to 10 tabs per dashboard.
- Rename a tab - Click the drop-down icon next to the tab title.
- Delete a tab - Click the drop-down icon next to the tab title. When a tab is deleted, the widgets added to that tab are also deleted.
- Add widgets - Click the “+” icon on the far right to select from a list of available widgets. You can have up to eight widgets on a tab.
- Arrange widgets - Click the *Arrange* icon to reorganize the page. Click the *Arrange* icon again when you are done to exit the mode.
- Change the title of the widget.
- Select the server for which the details of the widget will be displayed.
- Select the properties (IOPS, throughput, or latency) of the widget to be displayed.
- Remove a widget.


To configure a widget, highlight a widget and click the *Settings* icon to access the options for that widget.



Widgets


A variety of widgets are available to be displayed on the dashboard.

Alerts



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

Server Overview



Server Overview displays the name of the selected server, IP address, version, SIR server (if applicable), uptime of the server, operating system, kernel, processor, appliance information, and GUID. Click *Server Ports* to view the status of the server's ports and modules.


The *Options* section displays whether or not encryption is enabled for virtual tapes and the deduplication repository, as well as whether or not encryption is enabled.

The *Inventory* section displays information that is dependent on the type of server.

- StorSafe servers - Number of host clients, virtual tape libraries and drives, physical tape libraries, virtual tapes, FC adapters, physical devices, and storage pools.
- SIR servers - Number of adapters, physical devices, and storage pools.

The *System Resources* section displays CPU usage, number of cores, as well as used and free memory. For SIR and StorSafe-S, used and free memory are calculated based on the system memory left behind at the operating system level after memory pre-allocation for deduplication.

Storage Usage




Storage Usage displays the total managed capacity, which includes reserved/free storage as well as storage utilized for the resources on the selected server or storage pool; unprepared storage is not included.

By default, the server is displayed but you can select to show information for a specific storage pool if you click the drop-down box.

Click the *More Info* link to display the total server capacity (reserved and non-reserved) and the amount of storage reserved for the resources, as well as the amount of storage used by the resources as a percentage of reserved space.


In addition to total and unallocated capacity, you can see how much space is used by your resources, including: deduplication repository, NAS resources, virtual tapes, configuration (the configuration repository and tape database), etc.

Usage History



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


Top StorSafe Servers



Top StorSafe Servers displays the servers with the highest read/write/total throughput for the past five minutes, one hour, or one day.

The information in the widget can be displayed as a bar graph, list, or a pie chart sorted. Information on graphs can be from highest to lowest and vice versa by clicking the appropriate sort button.


Top Physical Devices



Top Physical Devices displays the physical devices with the highest read/write/total IOPS, latency, or throughput for the past five minutes, one hour, or one day.

The information in the widget can be displayed as a bar graph, list, or a pie chart sorted. Information on graphs can be from highest to lowest and vice versa by clicking the appropriate sort button.

Top Virtual Tape Libraries



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

The information in the widget can be displayed as a bar graph, list, or a pie chart sorted. Information on graphs can be from highest to lowest and vice versa by clicking the appropriate sort button.

Monitor

Monitor provides an overview of storage for all servers as well as the ability to view summary, performance, host clients, virtual resources, and physical resources for each server.

Overview page

The *Overview* page allows you to view storage capacity and usage for all servers.

Information can be sorted by name, capacity, free space, free space percentage, and throughput.

Server view page

Select a server from the list of the left side of your screen to view information about that server. Note that clustered servers are displayed individually here as opposed to the *Manage* pages, where clustered servers are displayed within their cluster.

Client view page

The *Client View* page displays information about clients on storage servers managed by the StorSight server.

Manage

Manage allows you to view server overview information as well as configure and manage settings for each server.

The *Server Overview* page displays general server information, names of partner servers (failover, deduplication), inventory, and system usage information. There are also icons to perform server functions, such as restart network/services, reboot, shut down, download x-ray, and manage patches. Refer to the *Server Settings* chapter for your server type for more information.

Note that clustered servers are displayed within their clusters, as opposed to the *Monitor* pages, where clustered servers are displayed individually.

Server search

Sort servers

Server status and info

Feature icons

Alerts

Port and server module status

Health indicator (color indicates the highest severity alert)

Storage space usage (free and total).

Number of each type of server. Click to filter list.

Average downstream read and write IOPS for the last five minutes with physical devices attached to the server.

Server Overview


Server	IP Address	OS
H13-72	172.22.13.72	Red Hat Enterprise Linux Server release 7.6 (Maipo)
H13-82	172.22.13.82	Linux 3.10.0-957.27.2.el7.x86_64 #1 SMP Mon Jul 29 21:21:47 PDT 2019
OBDR190	10.8.25.190	PowerEdge R720

Host Clients	Virtual Libraries	Virtual Tape Drives	Virtual Tapes	Tapes in Vault
4	11	80	57	1

You can sort servers by name, location, or customer (super administrators only). If location or customer are selected, the associated location and customer name will be visible.

Server info

To see more information about a server, click the *info* button.



Server icons

The table below describes the server icons used on the *Manage* pages.

Icon	Feature	Description
	Failover	Provides high availability for two servers, protecting them from a wide variety of problems. The icon color will change depending upon the status (orange=suspended, red=taken over, blue=took over).
	Not connected	The server is not connected.

Port and server module status

Click *Server Ports* to view the status of the server's port and modules.

Server Ports

Port Status | Module Status

Port	Protocol	Status	Usage
20	TCP/UDP	Not Open	FTP clients
21	TCP/UDP	Not Open	FTP clients
22	TCP	Open	HOST clients
25	TCP/UDP	Not Open	Mail servers
80	TCP/HTTP	Open	Web license server, REST API client
111	TCP/UDP	Open	NFS v3 clients

Analytics

Analytics allows you to view reports, system alerts, capacity predictions, audit log, and chargeback reports, as well as configure rules.

Reports/scheduled reports

StorSight provides a wide variety of predefined reports that help you manage your servers. Reports are available in several categories and display information for all servers.

Smart rules

Smart rules allow you to monitor StorSafe servers for a variety of conditions, including virtual tape library/drive performance, physical device performance, and client performance.

General rules

General rules allow you to set a variety of configuration conditions at the server level or the device level for which you want to be notified.

Alerts

The *Alerts* page displays the alerts that get generated by the Smart Rules and General Rules. For each alert, you will see the severity level, time stamp of the first and last occurrence, and details about the alert. If the last occurrence shows the current time, it means that the alert condition has not been resolved yet.

Capacity prediction

The capacity prediction feature uses an algorithm based on the past 30 days of storage data information for a server or storage pool to anticipate future capacity requirements.

Audit log

The audit log tracks and logs system activities for all servers by all administrators and users. This page is only visible for Super Administrators and Administrators.

Chargeback

Track customer storage capacity usage or the number of IOPS in order to allow service providers to easily generate bills for customers.

Administration

Administration allows you to configure and manage StorSight, servers, customers, users, and alert notifications, as well as set Active Directory/LDAP authentication and update client agent software.

StorSight

This page allows you to manage StorSight, including licenses, applying/rolling back maintenance patches to StorSight, downloading StorSight X-rays, setting an HTTP proxy, restarting services, rebooting the server, and setting global or per server portal messages. This page is only visible for Super Administrators.

Servers

This page displays all of the servers that have been added on this StorSight server. Only a Super Administrator can add, update, and delete servers.

Customers

In a multi-tenancy environment, customers are tenants. Typically, each customer represents a company (e.g., companyA, companyB, companyC). The *Super Administrator* acts as the landlord for all tenants.

This page is only visible for Super Administrators. A Super Administrator can add or delete customers and can assign or unassign servers in exclusive or shared mode.

Users

This page displays users. There are two views for this page, depending upon who is logged in as security, which plays a big role in what each logged in user can see and do.

A *Super Administrator* can see and manage all of the users in the system.

An *Administrator* can see and manage the users in their domain. An Administrator can create other administrators, users, and viewers for that domain.

AD/LDAP settings

There are three different methods that can be used to authenticate and authorize StorSight users and administrators:

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

Each customer can use a different method, but only one method can be used per customer.

Alert notifications

Alerts and rules can be configured to send emails and/or SMS text messages when an event is triggered. You can also use custom scripts to send notifications. This page is only visible for *Super Administrators*.

Global portal functions

Some functions are available on multiple pages in the portal.

Column sort



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

To sort additional columns, press the *Shift* button and click on each additional column heading.

▲ Virtual Device	VID	▼ Server	▲ Type	Storage Pool/ Physical Device
R805-10123140-SANDISK-00084	11	R805-10123142	Replica Thick	FALCON:IPSTOR ...
R720-10123110-jamie-00040	41	R720-10123113	Replica Thick	FALCON:IPSTOR ...
R720-10123110-jamie-00041	44	R720-10123113	Replica Thick	FALCON:IPSTOR ...
R720-10123110-jamie-00042	42	R720-10123113	Replica Thick	FALCON:IPSTOR ...
R720-10123110-jamie-00043	43	R720-10123113	Replica Thick	FALCON:IPSTOR ...
R720-10123110-SANDisk-00037	38	R720-10123113	Replica Thick	FALCON:IPSTOR ...

Server sort



Click the drop-down arrow to change the order of the server display. Servers can be sorted by HA pair, name, and location. Super administrators can also sort by customer.

Refresh



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

Export/Send



Click the *Export/Send* icon to export the information on the current page to a PDF or CSV file. If email and/or SMS have been configured by your administrator, you can also send a PDF or CSV file via email or SMS. When you export, you can convert size units of data in tables in order to have a unique unit for all values in a column. For some values, the resulting number may be too small in the desired unit. For example, if you select to convert all values to TBs, throughput numbers (i.e., 70 MB/s) may become too small if converted to TB/s, so they will remain as MB/s.

- To export: Select *Export to PDF* or *Export to CSV* from the menu and then specify the file name.
- To email: Select *Send via Email* from the menu, provide destination email addresses, specify the file type, file name, subject, and message.
- To text: Select *Send via SMS* from the menu, provide the phone number without dashes or spaces, specify the file type, file name, and message.

Save/Reset Layout




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

Filter



Click the *Filter* icon to filter the information displayed in a column. Then, click the icon that appears to the right of the column name.




Type in the *Search* box to search all text in that column, visible or linked (text in a pop-up window); clear the search field to return to the full view. This is useful for searching for data in fields with links. You can also check and uncheck items to filter the view. Checking/unchecking an item is limited to the selected column; when you uncheck an item in one column, it will still appear in the filter list for a different column.

Arrange




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



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

Search



Type a name or value in the *Search for* box to search all visible text in a list; clear the search field to return to the full view of the list. You can specify multiple strings "item1 item2 item3".

As soon as you start typing, only those items that have matching text are displayed.

You can also type in the name of a column heading for columns with checkboxes. For example, if you type "IBM" on a virtual libraries page, only those virtual libraries with "IBM" in the name will be displayed.


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

Administration

When you select *Administration* from the menu bar, you can access *StorSight*, *Servers*, *Customers*, *Users*, *AD/LDAP Settings*, *Alert Notifications*, and *License Usage Prediction*.

StorSight

 This page allows Super Administrators to manage StorSight, including licensing, Multi-factor authentication, strong passwords, applying/rolling back StorSight maintenance patches, downloading StorSight X-rays, setting an HTTP proxy, setting logout timeout for the server, restarting services, rebooting the server, setting global or per server portal messages .



Server Summary displays the name of the StorSight server, IP address, version, appliance type (virtual or physical), and uptime of the server. The *System Resources* section displays the current CPU usage, number of cores, as well as total, used, and free memory. For SIR and StorSafe-S, used and free memory are calculated based on the system memory left behind at the operating system level after memory pre-allocation for deduplication.

System Disk Usage displays the following for each system disk: mount path, size, used size, available size, and usage (%).

Used Memory displays the amount of memory used over the past two hours.

Load Average displays the average amount of CPU load over the past two hours.

Tape Data Load displays tape data, including the amount of data scanned for deduplication and the amount of unique deduplicated data stored in the repository over the past 7, 30, 60, or 90 days.

License

There is one initial license per StorSight; this license covers a preset amount of capacity that is used by all of your servers. If you need additional capacity, you will be given a separate license that must be added. Follow the steps below to add and register licenses. Note that only a Super Administrator can access *Licenses*.

1. Select *Administration* from the menu bar and then select *StorSight*.
2. Click *License*.
3. Click the “+” icon to add a license.
4. Enter the license key and click *Add*.
5. Register your license.

Select *Register Online* if you have an Internet connection. Otherwise, refer to [“Offline registration”](#).

Offline registration

Offline registration is useful when you do not have an Internet connection. To do this:

1. Click *Register Offline info* to create the `license.dat` file that records your server and license information.
2. Copy the `license.dat` file to a computer with an Internet connection and email it to the FalconStor 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.
3. When you receive a reply, copy the contents of the attached signature file (with the `.sig` extension).
4. Back in the portal, click *Register Offline* and paste the contents of the signature file to complete your offline registration.

Notes:

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

License renewal

When it is time to renew your product license, you can send an email to FalconStor that includes a report of your storage capacity usage. The sales department will use this information as part of your license renewal process.

1. Select *Administration* from the menu bar and then select *StorSight*.
2. Click *License*.
3. Click *Email License* and then *Send*.

Multi-factor authentication

Multi-factor authentication (MFA) is a multi-step account process that requires users to provide two verification factors in order to gain access to the portal and/or to delete virtual tapes. In addition to their password, each user will have to enter a verification code that is sent to their email to validate their identity.

MFA requires that an email server be defined in *Alert Notification* and each user must have a valid email address.

Once MFA is enabled, each user, including multi-tenant domain accounts, needs to enter a six-digit verification code in addition to a password when logging into the portal and/or when deleting virtual tapes. The verification code is a one-time code sent via email and is valid for five minutes.

To enable MFA, select *Administration* from the menu bar, select *StorSight*, click *Multi-Factor Authentication*, and click the option(s) to enable, *Enable Multi-Factor Authentication for Login* and/or *Enable Multi-Factor Authentication for Tape Deletion*.

Note: MFA does not apply to Active Directory/LDAP users.

Enable strong passwords for StorSight users

Strong
passwords for
StorSight users

Strong passwords can be enabled for all StorSight users who log into the portal. However, the root user's password can be reset without limitations. For security purposes, the *Strong Passwords* option cannot be disabled once it has been enabled.

Strong passwords must contain a minimum of 14 characters, including at least one lowercase letter, one uppercase letter, and one number, plus at least one special character: @#\$%

Note: Strong passwords and Active Directory/LDAP authentication are mutually exclusive; if AD/LDAP is in use in your organization or those of your tenants, you cannot use strong passwords on any StorSight server.

To enable strong passwords for StorSight, select *Administration* from the menu bar, select *StorSight*, and click the *Strong Passwords* icon.

Strong
passwords for
StorGuard
servers

The *Strong Passwords* option can be enabled for a StorSafe server from the *Manage* --> *Settings* tab. Refer to [“Enable strong passwords”](#) for more information.

Patches

Only a Super Administrator can apply maintenance patches to the StorSight server.

The download patch utility is a Perl program that uses FTP connections. This module supports FTP proxy (not HTTP or SOCKS) via a system variable named "FTP_FIREWALL".

Note: Server upgrade patches must be applied directly on the server and cannot be applied or rolled back via the portal.

Apply patch

1. Upload the patch onto the machine running the browser or a location accessible from that machine.

Patches can be downloaded from the FalconStor support communities.

2. Select *Administration* from the menu bar and then select *StorSight*.
3. Click *Patch*.
4. Locate the patch file and click *Upload Patch*.

The patch will be copied to the server and installed. Afterward, the StorSight server will restart.

Roll back patch

To remove (uninstall) a patch and restore the original files:

1. Select *Administration* from the menu bar and then select *StorSight*.
2. Click *Patch*.
3. Select the patch and click *Roll Back Patch*.

X-ray

An X-ray is a diagnostic tool used by your Technical Support team to help solve system problems. Each X-ray contains technical information about your StorSight server, such as messages from the system event log, a snapshot of your server's current configuration and environment, and multi-tenancy information about servers and storage pools. You should not create an X-ray unless you are requested to do so by your Technical Support representative. Only a Super Administrator can generate a StorSight X-ray.

To create an X-ray from the portal:

1. Select *Administration* from the menu bar and then select *StorSight*.

2. Click *X-ray*.
3. Click *Take X-ray*.
4. Click *Download X-ray*.

The operation is complete when the *.tgz* file has been downloaded.


If the portal is not available, you can create an X-ray by running the following command on the StorSight terminal:

```
# xray <X-ray file name without extension>
```

HTTP proxy

You can set a proxy for HTTP traffic.

1. Select *Administration* from the menu bar and then select *StorSight*.
2. Click *HTTP Proxy*.



The image shows a 'Set HTTP Proxy' dialog box with the following fields and controls:

- Use Proxy:** A toggle switch that is currently turned on (blue).
- Proxy IP Address*:** A text input field containing the placeholder text 'IP Address'.
- Port*:** A spin box with a small arrow icon on the right.
- Use https protocol:** A toggle switch that is currently turned off (grey).
- Username:** A text input field containing the placeholder text 'Username'.
- Password:** A text input field containing the placeholder text 'Password'.

At the bottom of the dialog, there is a message: "You must reboot StorSight after updating proxy settings for changes to take effect." and two buttons: "Save" and "Cancel".

Click the *Use Proxy* button to enable.

Specify the server, port, and protocol that should be used and the user account that will be used to log into the server.

3. Click *Save*.

Restart/reboot StorSight (updated October 2024)

You can restart StorSight services or reboot the server from the portal. Select *Administration* from the menu bar, select *StorSight*, click *Startup*, and select the appropriate operation.

You can also start, restart, and stop StorSight services from a command line. The following commands are available:

- `storsight start` - Starts the StorSight modules.
- `storsight restart` - Stops and then starts the modules.
- `storsight stop` - Stops server modules.

You can run these commands on the server locally using the `root` account. To run a command via a remote shell SSH session for servers installed with a `root` SSH key, connect as `root`. Otherwise, connect with the `fmsadmin` account and the default password `IPStor101`, and then run `su - root` to become the `root` user.

Set automatic logout timeout for the server

You can set an user inactivity timeout that results in an automatic server logout. Select *Administration* from the menu bar, select *StorSight*, click *Startup*, and specify the duration. The range is 1 to 60 minutes. Changes take effect for the next login session.


Portal notification messages

You can set messages to appear in the portal for the StorSight or for specific servers. This is useful if downtime is planned for a server or for StorSight. The message will appear when a user accesses the server and must be acknowledged in order to continue.

Set global notification message

1. Select *Administration* from the menu bar and then select *StorSight*.
2. Click *Message of the Day* and select *Global Message*.

3. Click *Display Message of the Day*.




Specify the message to be displayed.

Lock User Access prevents users from accessing StorSight. Once selected, a list of current StorSight user sessions is displayed so that you can notify them to log out. Once *Lock User Access* is enabled, only the Super Administrator is allowed to log into StorSight.

Set a server-specific notification message

1. Select *Administration* from the menu bar and then select *StorSight*.
2. Click *Message of the Day* and select *Server Specific*.
3. Click the “+” (Create) icon.



Specify the message to be displayed and select one or more servers.

Servers

Each StorSight server is a managed server that provides storage-virtualization and business-continuity services for continuous availability to business data in virtual and physical environments.

The ability to add, update, or delete a server is limited to Super Administrators.

Servers page



Select *Administration* from the menu bar and then click *Servers*. The *Servers* page lists all servers that have been added to this StorSight server.

Either a Super Administrator or Administrator can access the *Servers* page; however, only a Super Administrator can make any configuration changes.

Server	ID	Type	Configuration	IP Address	Version	Build	Status	Total Capacity	Managed Capacity	Total
tsv10	1151695	Backup	Standalone	10.6.1.122	StorSafe-S 11.01	11171	Unreachable			
H13-72	1347149	Backup	Failover	172.22.13.72	StorSafe-S 11.14	12418-03	Connected	10.86 TB	3.8 TB	
H13-82	1841413	Backup	Failover	172.22.13.82	StorSafe 11.14	12418-03	Connected	10.15 TB	3.33 TB	
IPS9-1 (IPS9-1-2)	26765	Storage	IO Cluster	10.6.12.107	StorGuard 9.01	9575-09	Connected	2.07 TB	2.07 TB	
IPS9-2 (IPS9-1-2)	29175	Storage	IO Cluster	10.6.12.139	StorGuard 9.01	9575-09	Connected	2.07 TB	2.07 TB	
StorGuard169 (1014IOC...	1840994	Storage	IO Cluster	172.22.6.169	StorGuard 9.02	9625-02	Connected	172.97 GB	172.97 GB	
StorSafe-Replica-H6210	1834725	Backup	Standalone	172.22.6.210	StorSafe-S 11.14	12418-03	Connected	3.43 TB	3.43 TB	
StorSafe153	2993448	Backup	Standalone	10.6.1.40	StorSafe-S 11.11	12300	Connected	3.24 TB	2.27 TB	
StorSafeH6212	390673	Backup	Standalone	172.22.6.212	StorSafe-S 11.14	12418-03	Connected	4.6 TB	4.6 TB	
VM22-123	1876338	Backup	Standalone	172.22.22.123	StorSafe-S 11.15	12500	Connected	209.96 GB	189.96 GB	
VM22-133	1180164	Backup	Standalone	172.22.22.133	StorSafe-S 11.11	12179-05	Connected	249.93 GB	249.93 GB	1

At the top of the page, you can see the total number of managed servers and total managed capacity as well as the total used primary and replica storage.

The properties for each server are available at a glance:

- Server name
- Server ID
- Server type (Backup, Deduplication)
- Configuration (Standalone, Failover)
- IP address
- Installed version
- Installed build number
- Connection status
- Total storage capacity
- Storage capacity that has been prepared
- Free space
- Free space percentage

- Fibre Channel (FC) protocol (enabled or disabled)
- iSCSI protocol (enabled or disabled)
- User name of the person who added the server to StorSight
- (Super Administrator only) Customer

View capacity information (updated October 2024)



A Super Administrator can view capacity information for all servers managed by this StorSight server. An administrator can view capacity information for all servers used by the customer account.

1. Select *Administration* from the menu bar and then select *Servers*.
2. Click the *Information* icon.

Capacity Information	
Managed Servers	22
Unreachable Servers	12
Customers	2
Managed Capacity	26.22 TB
Managed Capacity for StorGuard Servers	2.24 TB
Managed Capacity for StorSafe Servers	20.4 TB
Used for StorGuard Servers	0.51 TB
Reserved for Tapes	13.15 TB
Used for Tapes	58.05 GB
Reserved for Deduplication Resources	5.34 TB
Used for Deduplication Resources	14.26 GB
Reserved for NAS Resources	1.73 TB
Used for NAS Resources	0 B
Billable Capacity	26.22 TB
Used Licensed Capacity	285.48 GB from 100 TB
Full License Usage Prediction	1000 days
Scanned Tape Data up to Yesterday	85.27 GB
Unique Data	48.15 GB

Here you can see information about managed servers, along with capacity by server type and underlying resources. Other than *Managed Capacity*, a field will not be displayed if its value is zero.

Managed Servers - Number of servers managed by StorSight.

Unreachable Servers - Number of managed servers that cannot be reached.

Disconnected Servers - Number of managed servers that have been temporarily disconnected for maintenance reasons.

Customers - Number of customer domains for multi-tenancy.

Managed Capacity - Total capacity of virtualized storage from physical devices (not object storage).

Managed Capacity for StorGuard Servers - Total capacity of virtualized storage of reachable StorGuard servers; the capacity of each IO cluster node is counted.

Managed Capacity for StorSafe Servers - Total capacity of virtualized storage of reachable StorSafe servers.

Used for StorGuard Servers - Total used capacity of virtualized storage for StorGuard servers.

Reserved for Tapes - Total reserved capacity of virtualized storage for StorSafe tape backup cache.

Used for Tapes - Total used capacity of virtualized storage for StorSafe tape backup cache.

Reserved for Deduplication Resources - Total reserved capacity of virtualized storage for StorSafe deduplication repository data, folder, and index devices.

Used for Deduplication Resources - Total used capacity of virtualized storage for StorSafe deduplication repository data, folder, and index devices.

Reserved Object Storage for Deduplication Repository - Total reserved capacity of object storage for StorSafe deduplication repository data.

Used Object Storage for Deduplication Repository - Total used capacity of object storage for StorSafe deduplication repository data.

Used Object Storage for Tape Export - Total used capacity of object storage for exported StorSafe tapes to the cloud.

Reserved for NAS Resources - Total reserved capacity of virtualized storage for StorSafe NAS resources.

Used for NAS Resources - Total used capacity of virtualized storage for StorSafe NAS resources.

Billable Capacity - Sum of *Managed Capacity*, *Reserved Object Storage for StorSafe Repository*, and 10% of *Used Object Storage for Tape Export*.

Used Licensed Capacity - Sum of used capacity of virtualized storage, *Used Object Storage for StorSafe Repository*, and 10% of *Used Object Storage for Tape Export*.

Full License Usage Prediction - Estimated number of days for the managed capacity to reach the licensed capacity.

Last Month Tape Data - Size of ingest data incoming to StorSafe tapes.

Scanned Tape Data up to Yesterday - Size of scanned data for deduplication of StorSafe tapes.

Unique Data - Size of unique data after StorSafe deduplication.

Deduplication Ratio - Average deduplication ratio for StorSafe servers.

3. Click *Close* when done.

Add a server



Only a Super Administrator can add a server.

Note: Each server must only be added to one StorSight server.

1. Select *Administration* from the menu bar and then select *Servers*.
2. Click the “+” icon.

3. Enter the IP address or hostname of the server.

When Network Address Translation (NAT) is enabled, both the private IP address and public IP address in the NAT must be present on the server being added. An alias network device, for example, `eth0:0`, should be created for the NAT IP address in the system network files. There are several ways to add an alias IP address, for example, you can use the following command:

```
ifconfig <alias name> <NAT IP Address> up
```

For example, the private IP address is 10.178.0.146 and the public IP address from NAT is 195.142.106.184. If 10.178.0.146 is on `eth0`, add the alias IP address 195.142.106.184 on `eth0:0` by running the following command:

```
ifconfig eth0:0 195.142.106.184 up
```

4. Enter a username and password for the server (not a StorSight username/password).
5. Click *Add*.

If you attempt to add a server running an older version of FalconStor software, you receive an error message because the server type could not be identified. To continue, select the type that matches your server.

Reset server password




A Super Administrator can re-establish StorSight connectivity. If you have changed the password for this server elsewhere, click the *Update* icon, enter the new password, and click *Update*.

Disconnect a server



A Super Administrator can temporarily disconnect a server for maintenance reasons. Statistic collection and service alerts stop when a server is manually disconnected.

1. Place a checkmark next to the server to disconnect in the *Administration > Servers* view.
2. Click the  icon.



3. Click the *Disconnect* button from the *Disconnect Server* window.

The server is temporarily disconnected from management services, and the status is changed to *Disconnected* in the server list.

Note: A status of *Disconnected* is not the same as *Unreachable*. *Disconnected* servers have been manually disconnected by a user while *Unreachable* servers cannot be reached for other reasons, such as connectivity issues. StorSight attempts to reconnect to *Unreachable* servers every ten seconds.

To reconnect a server that has been manually disconnected, select the server and click the  icon. Confirm reconnection by clicking the *Reconnect* button.

Force rediscovery of a server



Only a Super Administrator can force rediscovery of a server.

StorSight attempts to communicate with unreachable servers every ten seconds. If there is a connectivity loss between a server and StorSight, select the server and click the *Force Rediscovery* icon to retry connecting to the server and reloading the server configuration, which includes reloading event log messages.

Force Rediscovery does not reconnect a server that has been manually disconnected via the *Disconnect* icon for maintenance reasons. You must use the *Reconnect* icon to bring back a server that has been manually disconnected.

Lock a server



A Super Administrator can temporarily block user access to a server for management, monitoring, or scheduled maintenance purposes. Make sure users are not connected to the server before locking. You can restart services to close all connections.

To lock a server, select the server, click the *Lock* icon, and select which options to apply.

Read-Only - Users are allowed to see servers but no configuration or management changes are allowed. Once set, a yellow lock will appear next to the server name.

Fully Locked - Users are not allowed to see the server resources or perform configuration or management changes. Once set, a red lock will appear next to the server name.

To unlock a server that has been locked, select the server, click the *Lock* icon, and disable the appropriate options.

Delete a server



Only a Super Administrator can delete a server. To delete a server, select the server and click the *Delete* icon.

Customers

StorSight provides a secure multi-tenant architecture that allows Managed Service Providers (MSPs) and large enterprises to offer high-speed backup/restore, global data deduplication, enterprise-wide replication, and long-term cloud-based archive as a service to their customers, called *tenants*.

Each tenant (a business or organization) has its own secure computing environment, called a *domain* (i.e., companyA.com). While all storage and network resources are shared among domains, StorSight logically isolates data by restricting visibility of storage resources to specific customers. Therefore, customers only see their own data and are not aware of other tenants.

The *landlord* in a multi-tenant environment is a Super Administrator, a master role that does not belong to any domain. By default, there is one Super Administrator account per StorSight server; however, the Super Administrator can create additional Super Administrator accounts.


The Super Administrator creates an administrator account for each company or organization, associating each tenant's domain with the account. One Administrator account is automatically created for each domain.

The Administrator for each domain then creates user accounts within their organization. Typically, the Administrator will create Administrator accounts for departments within their organization, possibly creating one Administrator for each department. Each Administrator will then create User and Viewer accounts as needed for their department.

For privacy, access to downstream storage can be isolated via storage pools, access to Fibre Channel targets via separate Fibre Channel zones, and access to networks via separate private subnets.

Servers can be assigned to a customer domain in *exclusive* or *shared* mode:

- In *exclusive* mode, the server is only available to the users in that customer domain.
- In *shared* mode, the server is available to users of multiple customers. Data isolation among those customers is accomplished by assigning different storage pools to each single customer so they can only see resources created from storage pools assigned to them.



An Administrator can also create and manage virtual resources for their domain. An Administrator of a server in exclusive mode has the same rights as the *superadmin* (root) user on the server.

Multi-tenant environments

In order to configure a multi-tenant environment, the following needs to be done:

1. The Super Administrator creates a *customer* for each domain.

A customer is typically created for each company or organization. An Administrator account is created for each customer with the *admin* role.

Refer to [“Add a customer”](#) for more information.

2. The Super Administrator assigns one or more servers/resources to each customer domain.

This provides each customer domain with specific storage resources.

Servers can be assigned in *exclusive* or *shared* mode.



In exclusive mode, only the users in that customer domain will have access to storage resources on that server.



In shared mode, the server is used by multiple customers. The Super Administrator must assign at least one storage pool to a customer; each storage pool can only be assigned to a single customer at any time.

Refer to [“Assign resources to a customer”](#) for more information.

3. Depending upon the authentication and authorization method for StorSight users and administrators, the Administrator for each customer logs in and adds administrators, users, and viewers for the domain.

Refer to [“Add a user”](#) for more information.

4. The Administrator assigns access rights to storage pools, virtual resources, host clients, and groups.

In order to create virtual tape libraries, drives, and tapes, users must have storage pools assigned to them by their administrator. Without access rights, users will not be able to create or manage virtual resources, host clients, and groups.

Refer to [“Storage pools”](#) for more information.

Multi-tenancy example

In our multi-tenancy example, an MSP has two customers, C1 and C2, and plans to configure the following:

Primary site:

- 2 StorSafe servers in a failover configuration
- SIR 2+1 cluster
- Physical tape library

Secondary site:

- 2 StorSafe servers in a failover configuration
- SIR 2+1 cluster

Both customers plan to use virtual tape services for data migration to object storage and replication for disaster recovery. Customer C2 plans to use a physical tape library for import/export.

In this scenario, the following actions are performed where the MSP is the Super Administrator:

1. The *Superadmin* adds all servers to StorSight.
2. The *Superadmin* prepares physical devices on both sites for virtualization, reserving some for configuration, some for tapes, and some for deduplication.
3. The *Superadmin* creates the Configuration Repository and Tape Database.
4. The *Superadmin* configures failover between StorSafe server on both sites.
5. The *Superadmin* configures the SIR 2+1 cluster on both sites.
6. The *Superadmin* adds StorSafe servers from the replica site as replication targets for StorSafe servers on the primary site.
7. The *Superadmin* creates two storage pools P1, P2.
8. The *Superadmin* assigns a physical library to a StorSafe server that is meant to be used by customer C2.
9. The *Superadmin* defines customer C1 and C2 in StorSight with their domains and sets default passwords for their *Admin* accounts.
10. The *Superadmin* assigns StorSafe servers (not SIR servers) and storage pool P1 to customer C1.
11. The *Superadmin* assigns StorSafe servers (not SIR servers), storage pool P2, and the physical tape library to customer C2.

12. The *Superadmin* performs all administration tasks for the StorSight, StorSafe servers, and SIR servers.
13. The *Admin* of each customer sees all StorSafe servers (not SIR servers) and only those storage pools assigned to their domain. Customer C2 sees the physical tape library.
14. The *Admin* creates user accounts within the company domain and host clients running backup software.
15. The *Admin* assigns host clients to users.
16. *Users* create virtual tape libraries from the associated storage pool and assign the library to the client to which the user has access.
17. *Users* create object storage accounts for cloud migration.
18. *Users* create policies for deduplication, replication, migration etc.
19. The *Admin* of customer C2 gets access to the physical library and can create import/export jobs or assign the physical library to a *User* to create those jobs.

Multi-tenancy rules and restrictions

- SIR servers cannot be shared. Management of an SIR server can only be performed by a Super Administrator when it is associated with StorSafe servers assigned in shared mode to customers. In this case, SIR is used for deduplication of data of all customers. SIR servers cannot be assigned to customers in shared mode.

The customer's Administrator can manage an SIR server if SIR is associated with a StorSafe server that is assigned to the customer in exclusive mode. In this case, the SIR server must be assigned to a customer in exclusive mode.

- The deduplication repository on a StorSafe-S server can only be managed by a Super Administrator when the server is assigned in shared mode to customers.

The customer's Administrator can manage the deduplication repository if the StorSafe-S server is assigned to the customer in exclusive mode.

- To support two customer tenants, at least two physical resources reserved for tapes must exist in order to create two storage pools for data isolation.
- A storage pool can only be used by tenants of one customer and cannot be shared among customers. Management of storage pools is performed by the Super Administrator.

The customer's Administrator can manage a storage pool if the storage pool is on a StorSafe server that is assigned to the customer in exclusive mode.

- A physical library/drive can only be used by one customer and cannot be shared among customers. The Super Administrator associates the physical library to a server, which can be assigned in shared or exclusive mode. Management of physical libraries/drives/tapes is performed by the Super Administrator or the customer administrator of an exclusive server.

The customer's administrator can manage a physical library/drive if it is associated with a StorSafe server that is assigned to the customer.

- All server resources (virtual library/drive/tape, FC/iSCSI client, object storage account, deduplication/replication/migration policy, etc.) can be owned by one tenant, except storage pools, which can support multiple tenants within the same customer domain.
- Resources without any user Access Control List (ACL) are only visible to the Super Administrator.

If a domain Administrator assigns a virtual tape library with existing tapes to a domain user, all the tapes will also be owned by that user. The server may become very busy while resetting the user access rights on those tapes when there are many tapes, so it is recommended to stop I/O traffic and suspend failover, if applicable.


Only resources with the *useracl* property are available to a customer administrator so their user ACL can be changed. Other resources are only available to the actual user owning the resource.

Customers page



Select *Administration* from the menu bar and then click *Customers*. The *Customers* page lists customers that have been created on this StorSight server.

Only a Super Administrator can access the *Customers* page.



Click to see which storage pools are assigned to each customer. This is useful for servers in shared mode.

Add a customer

1. Select *Administration* from the menu bar and then select *Customers*.
2. Click the “+” icon.
3. Enter the name and domain of the new customer.

The domain can be an actual domain (i.e., customera.com) or it can be an organizational unit (i.e., finance, engineering).

4. Enter a password for the Admin account that is automatically created for this customer.
5. Click *Save*.

Assign resources to a customer

1. Select *Administration* from the menu bar and then select *Customers*.
2. Click the *Assign Resource* icon.
3. Select the server and the access mode.

Servers can be assigned in *exclusive* or *shared* mode.



In *exclusive* mode, only the users in that customer domain will have access to storage resources on that StorSafe server.



In shared mode, the server is used by multiple customers.

4. If you selected shared mode, select one or more storage pools to assign to the customer.
5. Click *Assign*.

Unassign resources from a customer

1. Select *Administration* from the menu bar and then select *Customers*.
2. Select a customer and click the *Unassign Resource* icon.
3. Select one or more resources to unassign.

You can only unassign resources if there are no resources (virtual resources, clients, etc.) is assigned to a user. The *Allowed* column displays the status. If “No” is displayed, click to see the assigned resources.

If a server is in exclusive mode, you can unassign the server.

If a server is in shared mode, you can unassign specific storage pools or you can unassign the server itself.

4. Click *Unassign*.

Synchronize users

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

To synchronize users, select *Administration* from the menu bar, select *Customers*, highlight the customer, and click the *Sync Users with Managed Servers* icon.

Reset authentication

Use this option if you need to change a customer's authentication method to *Local*. You will only see this option if a customer has Active Directory or LDAP selected as their authentication method.

To reset authentication, select *Administration* from the menu bar, select *Customers*, highlight the customer, and click the *Reset Authentication* icon.

Delete a customer

To delete a customer, select *Administration* from the menu bar, select *Customers*, highlight the customer, and click the *Delete* icon.

Users

StorSight uses a hierarchy of user account types – a Super Administrator, Administrators, Users, and Viewers – to control access to the system.



- **Super Administrator (Superadmin)**- A master role that does not belong to any domain, the Super Administrator has full access to StorSight. In a multi-tenant environment, the Super Administrator represents the landlord and can create customer domains, add and assign servers, virtualize physical resources, reserve resources for specific usage (configuration, tapes, deduplication, NAS), create storage pools of virtualized storage, perform global configuration and maintenance of servers, configure failover, and create the deduplication repository.



- **Administrator (Admin)** - The administrator for a customer domain who can create *User* or *Viewer* accounts and assign these accounts to specific storage pools, virtual resources, and clients. An Administrator can also create and manage virtual resources for their domain. An Administrator of a server in *exclusive* mode has the same rights as the *superadmin* 'root' user on the server.



- **User** - A tenant user with read-write access to resources that have been created by that user or assigned to that user by the domain administrator.



- **Viewer (Read-Only)** - A tenant user with read-only access who can view configuration and reports as well as receive alerts; this account cannot make any configuration changes.

User roles

Using these four account types, system administrators can control access to servers and their resources, including physical devices, storage pools, virtual resources, and clients. Access rights are outlined in the following table:

	Super Admin	Admin (exclusive)	Admin (shared)	User	Viewer
Manage StorSight (configuration, license, administration)	Yes	No	No	No	No
Create customer domains	Yes	No	No	No	No
Assign servers and storage pools to customers	Yes	No	No	No	No
Assign physical libraries/drives to customers	Yes	No	No	No	No
Add Super Administrators	Yes	No	No	No	No
Add Administrators, Users, Viewers	Yes	Yes	Yes	No	No
Add servers	Yes	No	No	No	No
Change server root password	Yes	No	No	No	No

	Super Admin	Admin (exclusive)	Admin (shared)	User	Viewer
Configure server settings (including failover, hostname, network, encryption, hosted backup)	Yes	Yes	No	No	No
Perform server administration tasks (restart services/network, reboot, shutdown, X-ray, patch, delete export history for physical tapes, etc.)	Yes	Yes	No	No	No
Create and manage Configuration Repository	Yes	Yes	No	No	No
Create and manage Tape Database	Yes	Yes	No	No	No
Create and manage SIR repository	Yes	Yes	No	No	No
Create and manage storage pools	Yes	Yes	No	No	No
Assign access rights to storage pools	No	Yes	Yes	No	No
Manage physical libraries/drives	Yes	Yes	No	No	No
Assign access rights to physical libraries/drives	Yes	Yes	Yes	No	No
Assign access rights to virtual resources, clients, and groups	No	Yes	Yes	No	No
Create and manage clients	Yes	Yes	Yes	Yes	No
Create and manage virtual resources	Yes	Yes	Yes	Yes	No
Create and manage object storage accounts	Yes	Yes	Yes	Yes	No
Create and manage policies and run related jobs	Yes	Yes	Yes	Yes	No
Create and modify alerts/rules	Yes	Yes	Yes	No	No
View reports and alerts	Yes	Yes	Yes	Yes	Yes
View audit log	Yes	Yes	Yes	No	No

Users page



Select *Administration* from the menu bar and then click *Users*. The *Users* page lists all system users who have access to this StorSight server.

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

Only a Super Administrator or Administrator can access the *Users* page. A Super Administrator can see all users created on this StorSight server; an Administrator can only see users in their domain.

User	Role	Customer	Domain	Source	Active	Connected	Last Login from	Last Login Time
admin	Admin	C55	c55.com	Local	☒	☐	10.197.255.12	2022-01-11 09:30:43
superadmin	Super Admin			Local	☒	☒	10.197.255.9	2024-10-30 14:09:58
user	User	C55	c55.com	Local	☒	☐	10.197.255.12	2022-08-27 15:58:24
admin	Admin	c56	c56.com	Local	☒	☐		

The *Source* column shows where each user originated; Active Directory, LDAP, or added locally.

Add a user

Add users who can log into the StorSight portal.

Even in a single-tenancy environment, you must create a customer domain for host client authentication. For authentication, the actual user account is composed of the user name, followed by an underscore and the domain name, in which the dot is replaced by an underscore. For example, for *user1* in domain *customer.net*, the user account for authentication becomes: *user1_customer_net*.

Note: User names cannot be used more than once for *each customer*; each customer can have an *Admin1* user, but each customer cannot have multiple *Admin1* users. If Active Directory or LDAP authentication methods are used, a user account will be created automatically the first time the user logs into the StorSight portal. Therefore, you should not create local users for these users before they log in for the first time.

1. Select *Administration* from the menu bar and then select *Users*.
2. Click the “+” icon.

3. Enter information for the new user.

Active - Indicate if this is an active account that can be used.

Username - 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.

Password/Retype Password - Enter a password for this user.

Role - Select the role.

First Name/Last Name - Specify a user name.

Email/Phone Number - Specify an email address and telephone number.

4. Click *Save*.

Each user/administrator becomes a *local user* for this server. These accounts are needed when adding a StorSafe server from a host-side application.

Update a user

Accounts with a higher or equal role can change the properties of other accounts. For example, a Super Administrator can change the password of Super Administrator, Administrator, User, or Viewer accounts; an Administrator can change the password of Administrator, User, and Viewer accounts; Users and Viewers can only change their own password.

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

Super Administrators and Administrators can change a user from *Administration* --> *Users*. Highlight the user and click the *Update* icon.

Users and Viewers can change their password by selecting *Change Password* under the user profile icon on the top right side.

Synchronize users



(This feature is only available for Administrators.) Use this option to synchronize users between the StorSight and StorSafe servers. This is useful if a StorSafe server was down when a new local, Active Directory, or LDAP user was added to the StorSight server.

To synchronize users, select *Administration* from the menu bar, select *Users*, and click the *Sync Users with Managed Servers* icon.

Delete a user

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

The default Super Administrator account and the default Administrator account for a customer domain cannot be deleted.

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

AD/LDAP settings



There are three different methods that can be used for authenticating and authorizing StorSight domain users and administrators in a multi-tenancy environment:

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

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

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

Once the Super Administrator creates customer domains and assigns resources to them, customer domain administrators can set up any of the three authentication methods. Leave the *Local* authentication method for the Super Administrator so both AD/LDAP customer domain users and the local Super Administrator can connect.

Notes:

- Active Directory and LDAP can only be used if strong passwords are not enabled.
- If your Active Directory/LDAP server uses a self-signed certificate, refer to [“Self-signed certificates”](#) before configuring AD/LDAP settings.
- When AD or LDAP are used, a user account is automatically created the first time each AD/LDAP user logs into the StorSight portal. If a user's password is changed on the AD/LDAP server, the user's password must be manually updated in the portal.

Active Directory

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

If you will be using a self-signed certificate, refer to [“Self-signed certificates”](#) BEFORE continuing.

1. Select *Administration* from the menu bar and then click *AD/LDAP Settings*.
2. Click the *Active Directory* button.

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

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

Active Directory Settings

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

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

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

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

Domain - Domain name for the Active Directory server.

Active Directory Authorization

The fields that you will see depend upon the role that you have logged in with. For each, select Active Directory group(s) that contain users for a given role (Super Administrator, Administrator, User, or Viewer). These users will be granted the associated privileges. Note that StorSight only supports administrator/user groups at the parent level in the Active Directory tree.

Super Administrators will see:

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

Administrators for a customer will see:

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

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

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

3. Click *Save*.

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

Windows 2019 Active Directory configuration example

For our example:

- Domain: eve2019ad.test
- Group SuperAdmin: EveSuperAdmin
- Group Admin: EveAdmin
- Group User: EveUser
- Group Viewer: EveViewer

Active Directory Users and Computers [EveW2019AD-7101.eve2019ad.test]		
Active Directory Users and Computers [EveW2019AD-7101.eve2019ad.test]		
Active Directory Users and Computers [EveW2019AD-7101.eve2019ad.test]		
Saved Queries		
eve2019ad.test		
Built-in		
Computers		
Domain Controllers		
ForeignSecurityPrincipals		
Managed Service Accounts		
Users		
	Name	Type
	Administrator	User
	Allowed RODC Password Replication Group	Security Group - Domain Local
	Cert Publishers	Security Group - Domain Local
	Cloneable Domain Controllers	Security Group - Global
	Denied RODC Password Replication Group	Security Group - Domain Local
	DnsAdmins	Security Group - Domain Local
	DnsUpdateProxy	Security Group - Global
	Domain Admins	Security Group - Global
	Domain Computers	Security Group - Global
	Domain Controllers	Security Group - Global
	Domain Guests	Security Group - Global
	Domain Users	Security Group - Global
	Enterprise Admins	Security Group - Universal
	Enterprise Key Admins	Security Group - Universal
	Enterprise Read-only Domain Controllers	Security Group - Universal
	Group Policy Creator Owners	Security Group - Global
	Guest	User
	Key Admins	Security Group - Global
	Protected Users	Security Group - Global
	RAS and IAS Servers	Security Group - Domain Local
	Read-only Domain Controllers	Security Group - Global
	Schema Admins	Security Group - Universal
	EveAdmin	Security Group - Global
	EveUser	Security Group - Global
	EveViewer	Security Group - Global
	EveAdminU	User
	EveUserU	User
	EveViewerU	User
	<	

LDAP

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

If you will be using a self-signed certificate, refer to [“Self-signed certificates”](#) BEFORE continuing.

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

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

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

LDAP Settings

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

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

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

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

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

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

LDAP Authentication

LDAP tree search is performed starting with User Base DN to match username with User Search Attribute.

User Search Base DN *

User Search Attribute *

LDAP Authorization

LDAP tree search is performed starting with Group Search Base DN to match User DN with Group Search Attribute.

Group Search Base DN *

Group Search Attribute *

Group Name Attribute *

Specify user groups on the authentication server to map to FMS user roles.

Group Name for Super Admin Role *

LDAP Authentication

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

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

LDAP Authorization

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

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

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

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

Super Administrators will see:

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

Administrators for a customer will see:

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

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

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

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

3. Click *Save*.

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

Local

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

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

Self-signed certificates

If you are using a self-signed certificate, you must do the following before configuring AD/LDAP settings.

1. Verify that self-certification is enabled in Active Directory/LDAP.
2. Copy the certificate from the Active Directory/LDAP server to any directory on the StorSight server.
3. Run the following command on the StorSight server to register the certificate via Java® Keytool.

This command can be run anywhere. You need to log in as *root* (not *fmsadmin*).

```
$JAVA_HOME/bin/keytool -import -trustcacerts -keystore $JAVA_HOME/
lib/security/cacerts -storepass changeit -noprompt -alias <alias>
-file <path_to_certificate_file.cert>
```

The file path should be the path from step 2 where the certificate was copied.

4. Restart the StorSight tomcat module.

```
storsight restart tomcat
```
5. After StorSight has restarted, follow the steps to configure Active Directory/LDAP.

Be sure to change the port number to 636, select the *Use TLS* or *Use SSL* option, and enter the domain name for the Active Directory server.

Alert notifications



Alerts and rules can be configured to send email and/or SMS text messages when an event is triggered. You can also use custom scripts to send notifications. Alert notification settings are configured for all StorSafe servers.

Only a Super Administrator can configure alert notification settings.

Email settings

1. Select *Administration* from the menu bar, click *Alert Notifications*, and then click *Email/SMS Settings*.
2. Click the *Email* button to enable email.

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

SMTP Server - Specify the mail server that will be used to send emails. You can enter an IP address or hostname consisting of alphabet letters, numbers, "_", "-", or ".". The maximum length is 255 characters. If you are sending emails external to your domain, make sure your SMTP server is an open mail relay SMTP server that allows anyone to send email through it.

SMTP Port - Specify the mail server port that should be used. The default is port 25.

Sender Email - Specify a fully qualified email address that will be used in the "From" field of emails. The *Sender* should be the same as the *SMTP User* or the *Sender* needs to be added as an alias of the *SMTP User*. If this is not done, some SMTP servers will not allow emails to be sent from the specified *Sender* while others will replace the *Sender* with the SMTP user name.

SMTP User/SMTP Password - Specify the user account that will be used to log into the mail server.

3. If desired, click *Send Test Email* to verify your configuration.
You can send a test email message to specified recipients once all required fields are filled in.
4. Click *Save Settings*.

Email format

When an email alert notification is sent, the subject will include the StorSight server name and the severity of the issue. If there are multiple issues in a single email, the total will be displayed.


The email body will include:

- StorSight server host name
- StorSight server IP address
- StorSafe server host name
- StorSafe server IP address
- Device name
- Severity
- Policy name
- Start time

SMS settings

StorSight supports Twilio® for SMS text messages. You will need to create an account with Twilio in order to configure SMS in the StorSight portal.

1. Select *Administration* from the menu bar, click *Alert Notifications*, and then click *Email/SMS Settings*.
2. Click the *SMS* button to enable SMS.



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

Twilio Account SID - Specify an SID.

Twilio Account Token - Specify a token.

Sender Number - Specify the Twilio number for the account.

3. If desired, click *Send Test SMS* to verify your configuration.

You can send a test SMS text message to specified recipients once all required fields are filled in.

4. Click *Save Settings*.

Alternatively, you can send SMS text messages via your email server. Instead of specifying an email address for alerts, you will need to provide the recipient's cell phone number followed by the SMS gateway. The SMS gateways for some common carriers are:

- AT&T: @txt.att.net
- Sprint: @messaging.sprintpcs.com
- T-Mobile: @tmomail.net
- Verizon: @vtext.com
- Virgin Mobile: @vmobl.com

For example, 6311234567@txt.att.net or 6315551234@vtext.com. Check with your carrier if it is not listed above.

Be sure to enable and configure email for this method.

Alert scripts

A custom script can be called when an alert is generated. Custom scripts can be used to integrate with third-party alert management tools or to forward alerts to a centralized server.

Each script should reside under `$MSHOME/etc/` on the StorSight server. We recommend putting all scripts in a subdirectory (e.g., `alertscripts`) so that they will be retained if StorSight is upgraded via USB.

The following parameters are generated by each alert:

- `--server`
- `--severity`
- `--rulename`
- `--resource`
- `--first`
- `--last`
- `--details`

As listed above, each parameter is preceded with a double dash.

The parameter `--first` indicates the first occurrence of the alert and the parameter `--last` indicates the last occurrence of the alert.

Single quotes are added to variable values that can have spaces.

The following is an example of the output:

```
--server=10.1.1.1 --severity=CRITICAL --rulename='Unresponsive Server' --resource='StorSafeVA-90'
-- first='2016-05-12 16:43:49' --last='2016-05-23 19:10:36' --details='Server StorSafeVA-90 is
unresponsive. Check connectivity and make sure the server is up and services are running.'
```

Note that only a Super Administrator can access alert scripts; when configuring rules, customers will not see the script and will not be able to disable the script.

Add a script

1. Select *Administration* from the menu bar, click *Alert Notifications*, and then click *Alert Scripts*.
2. Click the “+” icon.
3. Provide a label/description for the script and the full path to the script.

The label can have up to 32 characters.


Manage scripts

You can modify the name/description or path for a script by clicking the *Update* icon.

If you delete a script, the script name is removed but the actual script is not deleted.

Global settings

1. Select *Administration* from the menu bar, click *Alert Notifications*, and then click *Global Settings*.



The screenshot shows the 'Global Settings' tab in a configuration interface. It contains four settings: 'Alert Frequency' set to 60 minutes, 'Global Emails' with a text input 'Email addresses separated by', 'Global SMS Phone Numbers' with a text input 'Phone numbers separated by c', and 'Global Alert Script' set to '- None -'. Each setting has a help icon (question mark) to its right. A 'Save Settings' button is located at the bottom right of the form.

Alert Frequency - Indicate how much time should elapse between alert notifications if an alert condition is still valid.

Global Emails - If email is enabled, specify the email addresses that should be used for alert notifications. Separate multiple email addresses with commas. These email addresses will be notified in addition to any email addresses defined for an individual rule.

Global SMS Phone Numbers - If SMS is enabled, specify the SMS phone numbers that should be used for alert notifications. Separate multiple phone numbers with commas. These phone numbers will be notified in addition to any phone numbers defined for an individual rule.


Global Alert Script - Select the default script that should be used for all alert notifications.

2. Click *Save Settings*.

License usage prediction



Select *Administration* from the menu bar and then click *License Usage Prediction*. This page displays the current and future data usage for this StorSight server and calculates data usage over time to predict when a new license key will be needed.



The screen displays the current and future data usage.

Gather total used values per day for prediction; 7 day-values give a better prediction.

The total used number includes space used for configuration plus data usage for deduplication resources, object storage for the deduplication repository, virtual tapes, NAS, and a percentage of object storage for tape export.

Configuration

Overview

An administrator must do the following to configure StorSafe servers:

1. Open a browser and type the URL of your FalconStor StorSight Server.

Enter your username, domain, and password to log in. The default login is:

- Username: superadmin
- Domain: <blank>
- Password: freestor

If this is the first time you are connecting to this StorSight server, the end user license agreement will appear. Click *Agree and continue* if you agree.

2. Select *Administration* --> *StorSight* to add and register your StorSight license.

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

3. Select *Administration* --> *Servers* to add your StorSafe server.

Refer to [“Add a server”](#) for more information.

4. Select the *Manage* tab to select your StorSafe server.

5. Select the *Physical Resources* --> *Physical Devices* tab to virtualize physical storage devices so that you can create logical resources.

Devices can be reserved for:

- Configuration - A 20 GB device used for the configuration repository and tape database.
- Deduplication repository - Used for index and folder devices of the deduplication repository.
- Tapes - Used for backup cache devices.

Refer to [“Prepare physical storage devices”](#) for more information.

6. If desired, select the *Physical Resources* --> *Storage Pools* tab to create storage pools.

A storage pool is a group of one or more virtualized physical devices. Refer to [“Storage pools”](#) for more information.

7. Select the *Object Storage* tab to add your object storage account to your StorSafe server.

You must have created at least one object storage account with your object storage provider. To isolate data between different users, you can create separate accounts.

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

8. If you will be using data repository or virtual tape encryption, select the *Settings* tab to enable [“Encryption Support”](#).
9. If you want to enable virtual tape encryption, select the *Encryption* tab to activate it for the server and create encryption keys.

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

10. Select the *Settings* tab to create the configuration repository, tape database, and deduplication repository, as well as to enable Fibre Channel and/or iSCSI target mode. From the *Settings* tab, you can also

The configuration repository contains configuration information for each server. Refer to [“Configuration repository”](#) for more information.

The tape database contains information about virtual tapes. Refer to [“Tape database”](#) for more information.

The deduplication repository stores deduplicated data and includes deduplication index, folder, and data disks. If encryption support was enabled, you can enable encryption for the deduplication repository. Refer to [“Deduplication repository”](#) for more information.

Host clients are the backup application servers that access virtual resources using the Fibre Channel (FC) and/or iSCSI protocol. In order to add a host client to communicate with the server, you must enable the target mode. Refer to [“Enable Fibre Channel and/or iSCSI target mode”](#) for more information.

11. If applicable, select the *Physical Resources --> Physical Tape Libraries* tab to assign physical tape libraries to StorSafe servers.

Refer to [“Assign a physical library”](#) for more information.

12. Select the *Virtual Libraries* tab to create your virtual tape libraries. You can also enable object storage migration for your virtual tape libraries.

Virtual tape libraries are used for backup by third-party tape or disk backup software, database backup utilities, and archiving applications. Each virtual tape library consists of one or more virtual tape drives and can have one or more virtual tapes. Refer to [“Create a virtual tape library”](#) for more information.

Refer to [“Configure Cloud Tiering for a virtual tape library”](#) for more information. about enabling object storage migration for a virtual tape library.

13. Select the *Host Clients* tab to add host clients and assign virtual tape libraries to them.

Host clients are the backup application servers that are assigned to virtual tape libraries and drives for backup. Refer to [“Host Clients”](#) for more information.

14. Select the *Deduplication* tab to use the default deduplication policy or create a new deduplication policy.

Deduplication policies specify which virtual tapes need to have deduplication and when deduplication should occur. Refer to [“Create a deduplication policy”](#) for more information. For general deduplication information, refer to [“Deduplication”](#).

If you will be replicating deduplicated data, designate the replica server as the target server and create deduplication/replication policies. Refer to [“Configure replication for deduplicated tapes”](#) for more information.

15. Optionally, enable a multi-tenant environment.

Refer to [“Multi-tenant environments”](#) for more information.

Physical Resources

Physical resources include:

- [Storage pools](#) - Groups of one or more virtualized physical devices (physical devices that have been prepared for use in creating virtual resources).
- [Physical devices](#) - Physical disks directly accessing an existing local disk or partition.
- [Adapters](#) - Fibre Channel and SCSI physical adapters.

There is one tab for each type of physical resource. Each page displays the resources that have been created on the selected server. There is also a tab for links to the web management interface for storage adapters.

Storage pools

StorSight enables the aggregation of multiple physical storage devices (such as JBODs and RAIDs) with various interface protocols (such as SCSI or Fibre Channel) into logical *storage pools*.

A storage pool is a group of one or more virtualized physical devices (physical devices that have been prepared for use in creating logical resources). Creating a storage pool enables you to provide all of the space needed by your clients in a very efficient manner. You can create and manage storage pools in a variety of ways, including:

- Tiers - Performance levels, cost, or redundancy
- Device categories - Virtual, including primary storage, and configuration
- Specific application use

For example, you can organize your storage by tier (low-cost, high-performance, high-redundancy, etc.) and assign it based on these classifications. Using this example, you may want to have your business critical applications use storage from the high-redundancy or high-performance pools while having your less critical applications use storage from other pools.

Storage pools work with all automatic allocation mechanisms in StorSight. This capacity-on-demand functionality automatically allocates storage space from a specific pool when storage is needed for a specific use.

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 a Super Administrator or Administrator in exclusive mode can manage storage pools.


Only an Administrator can assign access rights to pools.

All administrators and users can create virtual resources and allocate space from the storage pools assigned to them. The storage pool management rights of each type of administrator are summarized in the table below:

Role	Create/manage/delete	Add/remove storage from pools	Assign access rights
Super Administrator	Yes	Yes	No
Administrator (exclusive mode)	Yes	Yes	Yes
Administrator (shared mode)	No	No	Yes

Storage Pools page

When you select *Manage* from the menu bar and then click the *Physical Resources* tab, the *Storage Pools* page is displayed for the selected server by default. This page lists storage pools that have been created on the selected server, as well as properties.



Depending upon the columns that have been selected for display and the server type, the properties for each storage pool are visible at a glance:

- **Storage Pool** - The name that was specified when the storage pool was created.
- **ID** - The storage pool number automatically assigned by the system.
- **Size** - The total size of storage in the pool.
- **Allocated Size** - The total amount of storage that has already been allocated.
- **Free Size** - The amount of free space available in the pool.
- **Resource Types** - The type of resources for which physical space will be allocated from the pool.
- **Physical Devices** - The number of physical devices assigned to the pool.
- **Customer** - (Super Administrators only) The customer to which the storage pool belongs.


Storage Pools page icons

The table below describes the unique icons for this page and briefly describes the associated features:

Icon	Feature	Description
	Create	Create a storage pool.
	Update	Rename storage pool settings, add physical devices.
	Delete	Delete the selected storage pool.
	Migrate disks to pool	Migrate virtual tape library and NAS resource data disks to a storage pool.

Create a storage pool

1. Select *Manage* from the menu bar, select a server, and click the *Physical Resources* tab.
2. Click the “+” icon.



Create Storage Pool

Name*

Resource Types

Search for...

Physical Device	ACSL	ID	Reservation	Size
☐ DELL:MD36xxf	100:0:0...	aed2ea35-2e4e-434c-9600-5b1...	Nas	499.9
☐ DELL:MD36xxf	100:0:2...	aa663176-6f54-4598-b653-84fb...	Nas	523.9
☐ DELL:MD36xxf	100:0:2...	0efb8d4d-3b37-48d8-a712-9c7ff...	Nas	499.9

Showing 3 entries

3. Specify a storage pool name and resource types.
The name can have up to 64 characters excluding < > & \$ " / \'

Specify if you want to limit the type of resources for which physical space will be allocated from the pool:

- *Tapes* - Virtual tapes and virtual tape replicas
- *NAS* - NAS resources
- *All* - Virtual tapes, virtual tape replicas, and NAS resources

4. Select the physical devices to include.

Physical devices must be prepared before they can be used for a storage pool.

Physical devices can only be added to a storage pool with a matching pool type (e.g., only devices reserved for tapes can be added to a storage pool with the type *Tapes*). Devices reserved for tapes and devices reserved for NAS resources can coexist in a storage pool with the type *All*. The server will assign the proper physical device when it is required.

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. Click *Create* to save the storage pool.

Update storage pool settings



To update storage pool settings:

1. Select *Manage* from the menu bar, select a server, and click the *Physical Resources* tab.
2. Select a pool and click *Update* icon.
3. Make your changes and click *Update*.

Delete a storage pool



To delete a storage pool:


1. Select *Manage* from the menu bar, select a server, and click the *Physical Resources* tab.
2. Select a pool and click the *Delete* icon.
3. Click *Delete*.

Migrate disks to a storage pool



You can migrate virtual tape library data disks and NAS resource data disks to an existing storage pool. For virtual tape libraries, after migration, the library will belong to the storage pool and any virtual tapes that are subsequently created will be allocated to the storage pool.

1. Select *Manage* from the menu bar, select a server, and click the *Physical Resources* tab.
2. Select a pool and click the *Migrate Disks to Pool* icon.
3. Select *Virtual Tape Libraries* or *NAS Resources*.



4. Select the virtual tape libraries or NAS resources and click *Migrate*.

Physical devices

Physical devices directly access an existing local disk or partition.

Physical Devices page

Select *Manage* from the menu bar, click the *Physical Resources* tab, and then the *Physical Devices* tab. This page lists physical devices used by the selected server, as well as properties.

The screenshot shows the FalconStor StorSafe with StorSight User Guide interface. The top navigation bar includes Dashboard, Monitor, Manage, Analytics, and Administration. The Physical Resources tab is selected, and the Physical Devices sub-tab is active. The left sidebar shows server details for H13-72, H13-82, and OBD190. The main area displays a list of physical devices with columns for Physical Device, ACSL, Status, Category, Reservation, Size, Allocated Size, Free Size, and Storage. The list shows 23 entries, all with a status of 'Online'.

Physical Device	ACSL	Status	Category	Reservation	Size	Allocated Size	Free Size	Storage
DELL.MD36xof	102:0:8:12	Online	Virtualized, For...	Tapes	N/A	N/A	N/A	
DELL.MD36xof	102:0:9:4	Online	Virtualized, For...	Tapes	N/A	N/A	N/A	
DELL.MD36xof	102:0:8:2	Online	Virtualized	Tapes	1.22 TB	373.38 GB	880.61 GB	
DELL.MD36xof	102:0:8:14	Online	Virtualized	NAS Resources	499.99 GB	1 GB	498.99 GB	
DELL.MD36xof	102:0:8:13	Online	Virtualized, For...	Tapes	N/A	N/A	N/A	
DELL.MD36xof	102:0:9:15	Online	Virtualized	NAS Resources	523.99 GB	0 B	523.99 GB	
DELL.MD36xof	102:0:9:7	Online	Virtualized, For...	Tapes	N/A	N/A	N/A	
DELL.MD36xof	102:0:8:8	Online	Virtualized, For...	Deduplication ...	N/A	N/A	N/A	
DELL.MD36xof	102:0:9:16	Online	Virtualized, For...	NAS Resources	N/A	N/A	N/A	
DELL.MD36xof	102:0:9:1	Online	Virtualized, For...	Deduplication ...	N/A	N/A	N/A	
DELL.MD36xof	102:0:9:11	Online	Virtualized, For...	Tapes	N/A	N/A	N/A	
DELL.MD36xof	102:0:9:0	Online	Virtualized, For...	Tapes	N/A	N/A	N/A	
DELL.MD36xof	102:0:8:18	Online	Virtualized	NAS Resources	499.99 GB	0 B	499.99 GB	
DELL.MD36xof	102:0:8:10	Online	Virtualized	Configuration	29.99 GB	8.79 GB	21.2 GB	
DELL.MD36xof	102:0:8:17	Online	Virtualized, For...	NAS Resources	N/A	N/A	N/A	
DELL.MD36xof	102:0:9:5	Online	Virtualized	Deduplication ...	1.27 TB	1.27 TB	0 B	

Depending upon the columns that have been selected for display, the properties for each physical device are visible at a glance:

- Physical Device Name - The device name.
- ACSL - The adapter number, channel number, SCSI ID, LUN.
- Status - The device status.
- Category - Type of device - virtualized, unassigned, foreign
- Reservation - Storage usage for StorSafe server - configuration, deduplication repository, NAS resources, tapes
- Size - The size in GB.
- Allocated Size - The total amount that has already been allocated.
- Free Size - The amount of free space available.
- Storage Pool - The storage pool in which this physical device is used.
- Protocol - The communication protocol - FC or SCSI
- Queue Depth - The number of concurrent I/O operations for the device.
- GUID - Global unique identifier.

Physical Devices page icons

The table below describes the unique icons for this page and briefly describes the associated features based on the server type.

Icon	Feature	Description
	Rescan	Rescan physical devices.
	Manage paths	Set and manage autopathing.
	Manage	Rename a physical device, prepare/reserve physical devices, and test a physical device.
	Delete	Delete a physical device.
	Information	Display physical device name, status, size, owner, queue depth, ACSL, category, firmware, total number of segments on the device, and physical segment/layout information.

Rescan physical devices



StorSight automatically scans for devices whenever you reboot, but you can always perform a manual rescan to identify new/existing devices. If you have not rebooted your system since attaching new storage, you can rescan to allow the server to identify new devices.

To globally rescan physical devices:

1. Select *Manage* from the menu bar, select a server, click the *Physical Resources* tab, and then the *Physical Devices* tab.
2. Click the *Rescan* icon.

Scan Mode - Select if you want to scan existing devices only. If you want to discover new devices, you can specify additional scan details. If you discover new devices only, existing devices will be omitted, minimizing the processing time of SCSI commands sent to physical adapters.

Adapter range - When discovering new devices, set the range of adapters that you want to scan in the range of 0 to 999.

SCSI ID range - When discovering new devices, set the range of SCSI IDs that you want to scan in the range of 0 to 999.

Use a specific LUN range - When discovering new devices, set the range of LUNs that you want to scan, in the range of 0 to 255. If disabled, the system sends a SCSI request to LUN 0 and asks for a list of LUNs. The system will then scan the list of LUNs provided. If VSA is enabled and/or the LUN is beyond 256, you must disable this option to allow the system to scan beyond 255. Note that this SCSI command is not supported by all devices.

Stop rescan when a LUN without any devices is detected - Scan LUNs sequentially and stop when a LUN without a device is found. When this option is not selected, LUNs will scan to the end. Use this option only if all of your LUNs are sequential.

Automatically detect SCSI ID range - Enable auto detection of SCSI IDs with persistent binding.

Scan inactive paths if inactive paths are unavailable - Force a status check from a path that is not in use. If there is a new device discovered during rescan, it will still be listed although this device is not part of the range of IDs to discover.

3. Click *Scan*.

Set autopathing

Autopathing gives you the ability to balance I/O to multiple Fibre Channel devices used by a StorSafe server by setting preferred paths to use when accessing a device.

To set autopathing:

1. Select *Manage* from the menu bar, select a server, click the *Physical Resources* tab, click the *Physical Devices* tab, click the *Manage Paths* icon, and select *Set Auto Path*.

2. Select the mode to load balance I/O over available paths to storage devices.

Storage Preferred - (Default) Use the preferred paths set by the storage controller to access physical devices.

Active path if no Preferred - Use active paths for devices that do not have paths set as preferred by storage controllers. If you select this option, the system will not trigger paths to trespass (switch).

Active Path - Use the currently active paths.

Any - Use any available path. This can cause paths to trespass.

3. Select FC adapters, SCSI IDs, and LUNs that should be excluded in order to not use their paths for the selected load balancing mode

You may need to do this if you have a specific path that you want to maintain.

4. Click *Set* when done.

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, click the *Manage Paths* icon and select *Save Auto Path*. You will then need to restart StorSafe services.

If you ever need to restore your settings back to the last version that was saved, click the *Manage Paths* icon and select *Restore Auto Path*. This is useful if someone manually changes the settings and you want to revert to the saved version.

Rename a physical device



To rename a physical device:

1. Select *Manage* from the menu bar, select a server, click the *Physical Resources* tab, and then the *Physical Devices* tab.
2. Select a physical device and click the *Manage* icon and select *Rename*.
3. Make your change and then click *Update*.

Prepare physical storage devices



Before you can create virtual resources, you must virtualize physical disks so that you can create logical resources for storage devices. You can also unassign physical disks.

Notes:

- Changing the device category will remove any existing data from the device.
- If a device is in a storage pool, you must remove it from the pool in order to change its category.

To prepare physical devices:

1. Select *Manage* from the menu bar, select a server, click the *Physical Resources* tab, and then the *Physical Devices* tab.
2. Select one or more physical devices, click the *Manage* icon, and select *Prepare*.
3. Select the appropriate category.
 - Virtual - Virtual storage device to be used for StorSafe devices.
 - Unassigned - 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.
4. Select the reservation type.
 - Configuration - A 20 GB device used for the configuration repository and tape database. The configuration repository contains configuration information for each server. The tape database contains information about virtual tapes. A maximum of four devices can be reserved, including mirror devices. After the configuration repository/tape database is created, you cannot change the reservation of devices used for configuration.
 - Deduplication repository - (combined StorSafe with SIR servers and SIR servers) includes deduplication index, folder, and data disks, as well as the associated mirror devices.
 - Tapes - Used only for tape storage (StorSafe servers and combined StorSafe with SIR servers).
 - NAS resources - Used only for NAS resources (StorSafe servers and combined StorSafe with SIR servers).
 - 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.
5. Click *Prepare*.

Change the reservation of a physical device

You can change the 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.

To change the reservation:

1. Select *Manage* from the menu bar, select a server, click the *Physical Resources* tab, and then the *Physical Devices* tab.
2. Select one or more physical devices, click the *Manage* icon, and select *Reserve*.
3. Select a new reservation type.

The following rules apply:

- A device can only be reserved for a single purpose.
- If the device is empty, you can choose any reservation type. If resources have been created on a device with a specific reservation, the reservation can only be changed to *None*.

4. Click *Reserve*.

Test a physical device



You can check the performance of your physical devices to see:

- Sequential and random I/O in MB/s
- Sequential and random I/O in IOPs
- Sequential I/O response time in ns

The system will test the device and then display the throughput results for different spindle locations (start, middle, and end).

To test a physical device:

1. Select *Manage* from the menu bar, select a server, click the *Physical Resources* tab, and then the *Physical Devices* tab.
2. Select a physical device, click the *Manage* icon and select *Test*.
3. Click *Test*.

Adapters

Adapters include the Fibre Channel and SCSI physical adapters on a server.

Adapters page

Select *Manage* from the menu bar, click the *Physical Resources* tab, and then the *Adapters* tab. This page lists adapters used by the selected server, as well as properties.

Adapter	Number	Status	Information	Physical Devices	Mode	WWPN	Target WWPN
FC Adapter 100	100	Link Up	QLogic	28	Dual	21000024f3f62c0	2100000d773f62c0
FC Adapter 101	101	Link Up	QLogic		Dual	21000024f3f62c1	2100000d773f62c1
FC Adapter 102	102	Link Up	QLogic	19	Dual	21000024f349c24	2100000d77349c24
FC Adapter 103	103	Link Down	QLogic		Initiator	21000024f349c25	
SCSI Adapter 0	0		MegaRAID	4			
SCSI Adapter 32	32		Hosted Backup	16			
SCSI Adapter 36	36		vSCSI	3			

Showing 7 entries

The properties for each adapter are available at a glance:

- Adapter - The adapter name.
- Status - The FC adapter status (*Link Up* or *Link Down*).
- Number - The adapter number.
- Information - The adapter vendor.
- Physical Devices - The number of physical devices using the adapter. Click the link for a list of devices.
- Mode - The FC adapter mode - Dual, Initiator
- WWPN - The FC adapter base World Wide Port Name; this field indicates the initiator WWPN for dual-mode ports.
- Target WWPN - List of alias FC WWPNs, including target WWPNs for dual-mode ports.

Adapters page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Rescan	Rescan an adapter.
	Rename	Update adapter name.
	FC Target Mode	Enable/disable FC target mode.
	SNS information	Show or refresh Simple Name Server (SNS) information for FC adapters.


Rescan adapters



StorSight automatically scans for devices whenever you reboot, but you can always perform a manual rescan to identify new/existing devices. If you have not rebooted your system since attaching new storage, you can rescan to allow the server to identify new devices.

To rescan adapters:

1. Select *Manage* from the menu bar, click the *Physical Resources* tab, and then the *Adapters* tab.
2. Select an adapter and click the *Rescan* icon.



The dialog box titled "Rescan Adapter" contains the following fields and controls:

- Adapter:** FC Adapter 100
- Scan Mode:** Two radio buttons: "Rescan existing devices" (selected) and "Discover new devices".
- SCSI ID Range:** Two input boxes with dropdown arrows, containing "0" and "15", separated by the word "to".
- Use a specific LUN range:** A toggle switch (currently off) and a help icon (?)
- Buttons:** "Scan" and "Cancel" at the bottom right.

3. Select rescan options.

Scan Mode - Select if you want to scan existing devices only. If you want to discover new devices, you can specify additional scan details. If you discover new devices only, existing devices will be omitted, minimizing the processing time of SCSI commands sent to physical adapters.

SCSI ID range - When discovering new devices, set the range of SCSI IDs that you want to scan in the range of 0 to 999.

Use a specific LUN range - When discovering new devices, set the range of LUNs that you want to scan, in the range of 0 to 255. If disabled, the system sends a SCSI request to LUN 0 and asks for a list of LUNs. The system will then scan the list of LUNs provided. If VSA is enabled and/or the LUN is beyond 256, you must disable this option to allow the system to scan beyond 255. Note that this SCSI command is not supported by all devices.

Stop rescan when a LUN without any devices is detected - Scan LUNs sequentially and stop when a LUN without a device is found. When this option is not selected, LUNs will scan to the end. Use this option only if all of your LUNs are sequential.

Automatically detect SCSI ID range - Enable auto detection of SCSI IDs with persistent binding.

Scan inactive paths if inactive paths are unavailable - Force a status check from a path that is not in use. If there is a new device discovered during rescan, it will still be listed although this device is not part of the range of IDs to discover.

4. Click *Scan*.

Update adapter settings



You can rename an adapter and enable/disable FC target mode.

To rename an adapter, select *Manage* from the menu bar, click the *Physical Resources* tab, click the *Adapters* tab, select an adapter and click the *Rename* icon.

To enable/disable FC target mode, select *Manage* from the menu bar, click the *Physical Resources* tab, click the *Adapters* tab, select an adapter and click the *FC Target Mode* icon.

Display SNS information



You can display Simple Name Server (SNS) information for FC adapters. This is useful for identifying the WWPN of each of your ports.

To display SNS information:

1. Select *Manage* from the menu bar, click the *Physical Resources* tab, and then the *Adapters* tab.
2. Select an adapter and click the *SNS Information* icon.

Physical Adapter		FC Adapter 100
WWPN	Port ID	Adapter/Client
2100000d772d48e9	011a01	Adapter 100: QLogic (Mode: initiator)
21000024ff37864e	010d00	
2100001b32869802	010c00	
2101001b32a69802	012400	
2100001b32972164	011500	
2101001b32b72164	011100	
2100000d772832ee	011901	
		Refresh Close


3. To rescan, click *Refresh*.
4. Click *Close* when done.

Storage web links

External web links provide links to manage storage adapters on a server.

Storage Web Links page

Select *Manage* from the menu bar, click the *Physical Resources* tab, and then the *Storage Web Links* tab. This page lists links to the web management interface for storage adapters used by the selected server.



The properties for each link are available at a glance:

- **Storage** - The specified label for the storage adapter.
- **Link** - The fully qualified link to the management GUI. Click the link to manage your storage in a separate browser tab.

Storage Web Links page icons

The table below describes the unique icons for this page and briefly describes the associated features:

Icon	Feature	Description
	Create	Create a storage web link.
	Update	Update the storage label and storage link.
	Delete	Delete the selected storage link.

Virtual Tape Libraries

FalconStor StorSafe servers are independent virtual tape library servers that can be managed from the StorSight portal.

Tape libraries



When you select *Manage* from the menu bar and click the *Virtual Libraries* tab, the *Libraries* page lists the virtual tape libraries on the selected StorSafe server. Virtual tape libraries emulate physical tape libraries and are used for backup by third-party tape or disk backup software, database backup utilities, and archiving applications.

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

From this page you can also access the *Drives*, *Vault*, *Cloud Tier*, and *Tapes* pages for your StorSafe servers.

Only a Super Administrator can access the *Libraries* page.

The screenshot displays the FalconStor StorSight interface. The top navigation bar includes 'Dashboard', 'Monitor', 'Manage' (selected), 'Analytics', and 'Administration'. The 'Virtual Tape Libraries' page is active, showing a list of libraries. The left sidebar provides a summary of storage and backup status for several servers, including FBD VM22-190, FBD VM22-139, FBD VM22-140, and FBD VTL01. The main content area shows a table of virtual tape libraries with the following data:

Library	ID	Vendor	Product	Revision	Media Type	Slots	Drives	Tapes
ADIC-Scalar 100-00030	30	ADIC	Scalar 100	2.62	ULTRIUM1	80	6	
ADIC-Scalar 100-00118ee	118	ADIC	Scalar 100	2.62	ULTRIUM1	80	6	
ADIC-Scalar 100-00127-sang	139	ADIC	Scalar 100	2.62	ULTRIUM3	80	12	
ADIC-Scalar 100-02946pra	126	ADIC	Scalar 100	2.62	ULTRIUM3	80	12	
Auto-ADIC-Scalar 100-00045	45	ADIC	Scalar 100	2.62	ULTRIUM1	80	6	
Client-190-ADIC-Scalar 100-00108	108	ADIC	Scalar 100	2.62	ULTRIUM1	80	2	
DSI-TS3500L32-00066	66	DSI	03584L32	2460	ULTRIUM4	253	6	
IBM-TS3500L32-00059	59	IBM	03584L32	2460	ULTRIUM7	253	6	
LSU_mpc-test	5	FALCON	OST	3.15	ULTRIUM1	2000	4	
mp-ADIC-Scalar 100-00080	80	ADIC	Scalar 100	2.62	ULTRIUM1	80	6	
MP-IBM-TS3500L32-00105	105	IBM	03584L32	2460	ULTRIUM6	253	2	
STK-L700-00152 (Frank)	152	STK	L700	3.05	T9840A	678	10	
WORM-IBM-TS3500L32-00089								

Showing 13 entries

Depending upon the columns that have been selected for display, the properties for each virtual tape library are visible at a glance:

- Library name

- Library ID
- Vendor
- Product
- Firmware
- Media type
- Number of slots
- Number of drives
- Number of tapes
- Number of loaded tapes
- Clients (click to see a list)
- Starting/ending barcodes
- Storage pool
- Capacity on Demand usage
- Auto Archive usage
- Auto Replication usage
- Auto migration (tape migration to cloud object storage) usage
- Auto migration mode - Copy or Move
- Tape Caching usage
- Tape Duplication usage
- Tape Encryption usage

Libraries page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Create	Create a virtual tape library.
	Update	Update virtual tape library properties, firmware, tape services (Tape Caching, Auto Archive, Auto Replication, Cloud Tiering), and tape duplication.
	Delete	Delete the selected virtual tape library.
	Manage	Create tapes, create a cache for physical tapes (tape caching), load and unload tapes, reset library, add tape drives.
	Auto Migrate to Cloud Object Storage	Configure or update auto migration to cloud for a virtual tape library.

Create a virtual tape library

To create a virtual tape library to emulate a physical tape library:

1. Select *Manage* from the menu bar, click the *Virtual Libraries* tab, and select a StorSafe server.
2. Click the “+” (*Create*) icon on the *Libraries* page.
3. On the *General* tab, specify general library properties.

The screenshot shows the 'Create Virtual Tape Library' dialog box with the 'General' tab selected. The dialog has three tabs: 'General', 'Tape Services', and 'Tape Duplication'. The 'General' tab contains the following fields and controls:

- Library Type:** A dropdown menu set to 'ADIC-Scalar 100'.
- Name*:** A text field containing 'ADIC-Scalar 100-00193'.
- Number of Slots:** A numeric spinner set to 80.
- Number of I/E Slots:** A numeric spinner set to 4.
- Starting Barcode*:** A text field containing '00AC0000'.
- Ending Barcode*:** A text field containing '00ACZZZZ'.
- Drive Type:** A dropdown menu set to 'IBM-ULTRIUM-TD3'.
- Drive Name Prefix*:** A text field containing 'IBM-ULTRIUM-TD3'.
- Drive Count:** A numeric spinner set to 12.
- Media Type:** A dropdown menu set to 'ULTRIUM3'.
- Maximum Tape Capacity:** A numeric spinner set to 340, followed by a 'GB' unit label.
- Capacity On Demand:** A toggle switch that is turned on, with a help icon (?) to its right.
- Initial Size:** A numeric spinner set to 5, followed by a 'GB' unit label.
- Incremental Size:** A numeric spinner set to 5, followed by a 'GB' unit label.
- Encrypt Data:** A toggle switch that is turned off.

At the bottom right of the dialog are 'Create' and 'Cancel' buttons.

Library Type - Select the physical tape library you are emulating. If you have a physical tape library, you need to create a virtual tape library that resembles it in order for the virtual tapes to use the same format as the physical tapes. This is

important for importing and exporting functions and guarantees that your backup application will accept the tapes.

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.

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

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

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

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

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

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

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

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

Drive Type - Select the drives used in your library.

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

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

Media Type - Informational.

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

- If you will *not* be exporting data to physical tape, you can enter any maximum capacity.

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

Capacity On Demand - 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, StorSafe will allocate each virtual tape at the full size of the tape you are emulating. 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 Size/Incremental Size - Enter the initial size of each resource and the amount by which it will be incremented.

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

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

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

4. On the *Tape Services* tab, specify which services to use for this virtual library.

Note that Tape Caching, Auto Archive, Auto Replication, or Cloud Tiering services are mutually exclusive so only one can be configured per virtual library.

Tape Caching enhances the functionality of StorSafe by acting as a cache to your physical tape library, providing transparent access to data regardless of its location. Refer to ["Tape Caching"](#) for more information.

Auto Archive writes data to physical tape whenever a virtual tape is moved to an Import/Export (IE) slot (the vault) from a virtual library by a backup application or other utility after a backup. Refer to ["Auto Archive"](#) for more information.

Auto Replication replicates data to another StorSafe server whenever a virtual tape is moved to an IE slot from a virtual library by a backup application or other utility after a backup. ["Auto Replication"](#) for more information.

Cloud Tiering exports virtual tapes to object storage when they are ejected by backup software. Refer to ["Configure Cloud Tiering for a virtual tape library"](#) for more information.

5. On the *Tape Duplication* tab, specify if you want to make up to five duplicate copies of a physical tape whenever virtual tape data is exported to a local physical tape.

Refer to ["Tape Duplication"](#) for more information.

6. Click *Create* when done.

Update virtual tape libraries

You can update virtual tape library properties, firmware revision, tape services (Tape Caching, Auto Archive, Auto Replication, Cloud Tiering), and tape duplication.

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, select a virtual tape library, and click the *Update* icon.

Delete virtual tape libraries

You can only delete a virtual tape library if it is not assigned to clients.

You have the option to delete virtual tapes in the library or move them to the vault.

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, select a virtual tape library, and click the *Delete* icon.

Create virtual tapes

New tapes inherit the Auto Archive, Cloud Tiering, Tape Duplication, and Capacity on Demand configuration from the virtual tape library.

1. Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, select a virtual tape library, click the *Manage* icon, and select *Create Tapes*.

You can also create virtual tapes for a library from the *Tapes* tab.

2. On the *General* tab, specify general tape properties.

The screenshot shows the 'Create Tapes' dialog box with the 'General' tab selected. The fields are as follows:

- Virtual Tape Library:** IBM-TS3500L32-00265
- Name Prefix:** VirtualTape-
- Size:** 5 GB
- Storage Type:** Storage Pool (selected), Physical Device
- Storage Pool*:** (Select)
- Number of Tapes:** 1 (Maximum 253)
- Barcode Range:** From 010900 to 0109ZZ
- Starting Barcode:** 01EA0000
- WORM:** Toggle switch is off.

Buttons: Create, Cancel

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

Size - If *Capacity on Demand* is enabled for the library, the default is the initial size but you can change it. Otherwise, the maximum tape capacity is displayed and you cannot change it.

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

Number of Tapes - The maximum number of tapes is the number of slots in the library. You cannot specify if Remote Export is configured for the library or if you are matching physical tapes in an Auto Archive configuration.

Starting Barcode - If desired, specify the starting barcode. The default is to generate a barcode based on the available barcode range for the library. You cannot specify if Auto Archive, Auto Replication, or Remote Export are configured for the library.

WORM - Indicate if you want the virtual tape to be a write-once-read-many (WORM) tape. A WORM tape cannot be overwritten by backup software. Only virtual tapes that support ULTRIUM5 media type and above can be configured as WORM tapes. WORM tapes are not supported with Veritas NetBackup WENCR media (WORM media on which NetBackup encrypts data).

3. On the *Service* tab, specify what services you want to enable for the virtual tape and optionally, add the tape into a deduplication policy.

Remote Export automatically exports the contents of a replicated virtual index tape (VIT) to a physical tape at the remote site. Refer to ["Remote Export"](#) for more information.

If *Auto Replication* is enabled for the virtual library, specify if you want to enable it for the new virtual tapes in order to replicate data to another StorSafe server whenever a virtual tape is moved to an IE slot from a virtual library. Refer to ["Auto Replication"](#) for more information.

If *Auto Archive* is enabled for the virtual library, specify if you want to create virtual tapes that match physical tapes in the physical library. Refer to ["Auto Archive"](#) for more information.

4. Click *Create* when done.

How virtual tapes are allocated

StorSafe uses a sophisticated methodology to determine which LUN to use when allocating space to expand virtual tapes and to create virtual tapes in a storage pool or from multiple physical devices.

Using two algorithms, *Dynamic LUN Allocation* and *Round Robin*, virtual tapes are allocated from the LUN that is currently experiencing the least amount of I/O.

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, StorSafe 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, StorSafe 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, StorSafe continues allocating from another LUN.

When *Capacity on Demand* is enabled for the library and tape expansion is needed, StorSafe 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.

Load virtual tapes

To load tapes into all empty virtual tape drives, select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, select a virtual tape library, click the *Manage* icon, and select *Load Tapes*.

Unload virtual tapes

To unload tapes from all virtual tape drives, select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, select a virtual tape library, click the *Manage* icon, and select *Unload Tapes*.

Reset a library

Resetting a library terminates all active backup/restore jobs and unloads all tapes from their tape drives.

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, select a virtual tape library, click the *Manage* icon, and select *Reset Library*.

Add virtual tape drives

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

1. Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, select a virtual tape library, click the *Manage* icon, and select *Add Drive*.

You can also add virtual tape drives to a virtual tape library from the *Drives* tab.

The screenshot shows a dialog box titled "Add Tape Drives to Library". It contains the following fields and controls:

- Virtual Tape Library:** A text field containing "ADIC-Scalar 100-00127-sang".
- Drive Name Prefix*:** A text field containing "ADIC-Scalar 100-".
- Number of Drives to Create:** A spinner control set to "1". To the right of the spinner is the text "Maximum 243".
- Buttons:** "Create" and "Cancel" buttons at the bottom right.

If desired, specify a drive name prefix up to 59 characters excluding < > & \$ " / \'. Indicate how many tape drives to add.

2. Click *Create* when done.

Configure Cloud Tiering for a virtual tape library

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

In order to use this feature, deduplication must be configured and you must have at least one object storage account with your object storage provider. This is added on the *Object Storage* tab. Refer to ["Add an object storage account"](#) for more information.

You can only use Cloud Tiering if you are not currently using Tape Caching, Auto Replication, or Auto Archive on the virtual tape library.

To configure or update Cloud Tiering for a virtual tape library:

1. Select *Manage* from the menu bar, click the *Virtual Libraries* tab, and select a StorSafe server.
2. Select a tape library on the *Libraries* page and click the *Auto Migrate to Cloud Object Storage* icon.

The screenshot shows a configuration window titled "Auto Migrate to Object Storage". It contains the following fields and controls:

- Virtual Tape Library:** ATL-1500-00137
- Auto Migrate:** A toggle switch that is currently turned on (blue).
- Object Storage Account:** A dropdown menu showing "AKIAIBPQ3S3XBCXVIZKA".
- Mode:** Two buttons, "Copy" (selected) and "Move".
- Delay:** A text input field with "0" and a unit selector set to "Days".
- Buttons:** "Save" and "Cancel" buttons at the bottom right.

You can also configure when you create or update a virtual tape library.

3. Click *Cloud Tiering* to enable.
4. Select an existing object storage account.

Each library can use a different object storage account.

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

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

5. Select the mode.

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

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

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

6. Click *Save* when you are done.


Tape drives



Select *Manage* from the menu bar, click the *Virtual Libraries* tab, and click the *Drives* tab. This page lists virtual tape drives on the selected StorSafe server. Select to view standalone tape drives or tape drives in a specific tape library from the *Tape Library* drop-down list.

If you are using deduplication with your StorSafe servers, “SIR-TapeDrives” will be listed. These tape drives are for deduplication use only and cannot be modified or deleted.

Only a Super Administrator can access the *Drives* page.



Depending upon the columns that have been selected for display, the properties for each virtual tape drive are visible at a glance:

- Drive name
- Drive ID
- Status - Empty, Loaded, Ejected
- Vendor
- Product ID
- Firmware
- Media type
- Element number
- Clients (click to see more info)
- Capacity on Demand usage
- Maximum capacity
- Read and write throughput for the past 60 minutes
- Compressed read and write for the past 60 minutes

Drives page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Create	Create a standalone virtual tape drive or add virtual tape drives to an existing virtual tape library.
	Update	Update virtual tape drive name and firmware.
	Delete	Delete the selected virtual tape drive.
	Load tapes	Load ejected tapes into a virtual tape drive.
	Create tape	Create a virtual tape for an empty standalone virtual tape drive.

Create a standalone virtual tape drive

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

Note: This procedure is for *standalone* virtual tape drives. If you want to add virtual tape drives to an existing virtual tape library, refer to ["Add virtual tape drives"](#).

1. Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Drives* tab.
2. Select *Standalone Drives* as the *Tape Library*.
3. Click the "+" (*Create*) icon.

4. Enter information about the standalone virtual tape drive.

Standalone Tape Drive Type - Select the type of drive.

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

Media Type - Informational.

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

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

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

5. Click *Create* when done.

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

Update virtual tape drives

You can update the virtual tape drive name and firmware revision.

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Drives* tab, select a virtual tape library or standalone drives, select a virtual tape drive, and click the *Update* icon.

Load virtual tapes

To load ejected tapes into a virtual tape drive, select *Manage* from the menu bar, click the *Virtual Libraries* tab, click the *Drives* tab, select a StorSafe server, select a virtual tape drive, and click the *Load Tapes* icon.

Delete virtual tape drives

You can only delete a virtual tape drive if:

- The drive is not assigned to clients.
- No drives with a higher element number remains in the library after deletion.
- At least one tape drive remains in the library.

You can not delete “SIR-TapeDrives”.

You have the option to delete virtual tapes in the drive or move them to the vault.

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Drives* tab, select a virtual tape library or standalone drives, select a virtual tape drive, and click the *Delete* icon.

Create a virtual tape

You can create a virtual tape for an empty standalone virtual tape drive.

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Drives* tab, select a standalone drive, and click the *Create Tape* icon.

Specify a tape name prefix, up to 25 characters excluding spaces or special characters < > & \$ " / \ ' * ? | %

Select a single storage pool or physical device(s) from which to create the tapes.

Vault



Select *Manage* from the menu bar, click the *Virtual Libraries* tab, and click the *Vault* tab. This page lists virtual tapes in the vault on the selected StorSafe server.

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

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

Only a Super Administrator can access the *Vault* page.

Barcode	Tape	ID	Size	Used	Status	VIT	Parent Lib...	Migration Mode	Object Storage...	Migrated	Stub
008B000H	VirtualTape-ora00298	10000298	0 B	0 B	Online	Plain	139				
008B000R	VirtualTape-pra00342	10000342	0 B	0 B	Online	Plain	139				
008B000S	VirtualTape-pra100343	10000343	0 B	0 B	Online	Plain	139				
01040AL2	tsv007-MP-VirtualTape-00253	10000333	1 GB	0 B	Online	Plain					
01650000	VirtualTape-00239	10000239	5 GB	0 B	Online	Plain	357				
01650001	VirtualTape-mp-00265	10000265	5 GB	0 B	Online	Plain	357				
01650001	VTL01-VirtualTape-mp-00265	10000266	5 GB	10 MB	Online	Plain					
01650005	VirtualTape-mp3-00270	10000270	5 GB	20 MB	Online	Plain	357				
01650005	VTL01-VirtualTape-mp3-00270	10000271	5 GB	20 MB	Online	Plain					
01650006	VirtualTape-mp4-00272	10000272	5 GB	30 MB	Online	Plain	357				
01650006	VTL01-VirtualTape-mp4-00272	10000273	5 GB	30 MB	Online	Plain					
01650008	VirtualTape-mp6-00275	10000275	5 GB	40 MB	Online	Plain	357				
01650008	VTL01-VirtualTape-mp6-00275	10000276	5 GB	40 MB	Online	Plain					
01650009	VirtualTape-00277	10000277	5 GB	50 MB	Online	Plain	357				

Showing 35 entries

Depending upon the columns that have been selected for display, the properties for each tape are visible at a glance:

- Barcode
- Tape name
- Tape ID
- Size
- Used size
- Tape status - online, offline, incomplete (missing physical segments)

- Virtual tape index status for deduplicated tapes - pure (all tape data has been deduplicated), mixed (tape contains pointers to deduplicated data as well as data that has not yet been deduplicated), plain (tape not yet deduplicated)
- Parent virtual tape library ID (click for details)
- Cloud Tiering mode - Copy or Move
- Object storage provider ID
- Indication if Cloud Tiering occurred
- Stub - Checked if Cloud Tiering occurred and a stub tape exists
- Write-protected status
- Write-once-read-many (WORM) tape usage
- Encrypted tape indicator
- Deduplication policy
- Deduplication status - Successful, pending, failed
- Replication is/not configured
- Replication server
- ID of replica tape
- Replication source server
- ID of source tape
- Replication of deduplicated tapes - Status on target servers for replication of deduplicated tapes
- Source of replica promoted in test mode - Checked if this is a primary tape promoted from a replica in test mode
- Previously replica promoted in test mode - Checked if this is a replica tape promoted from a replica in test mode
- Tape caching is/not configured
- Unmigrated cache - Checked if tape caching is used and the tape has unmigrated data
- Direct-Link - Checked if tape caching is used and the tape is a direct link tape (a link to the physical tape)

Vault page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Manage	Rename, write protect, move, copy, shred, or stop shredding a tape in the vault.
	Tape caching	Export cached data, export entire tape, reclaim cache space.
	Migrate to object storage	Manually migrate virtual tape data to object storage.

Icon	Feature	Description
	Convert to stub tape	Manually convert a virtual tape to a stub tape.
	Recover from object storage	Recover a stub tape to a virtual tape in order to recover data.
	Export	Export virtual tape data to physical tapes.
	Information	Display tape summary and physical layout information.

Manage tapes in the vault

You can perform the following functions for tapes in the vault:

- **Rename** - The new name can be up to 32 characters excluding spaces or special characters < > & \$ " / \ ' * ? | %
- **Write protect**
- **Move** multiple tapes to a virtual tape library or move a single tape to a standalone virtual tape drive or a virtual tape library.
- **Copy** - Replicates the full contents of a single tape to a local or remote server. Refer to ["Tape Copy"](#) for more information
- **Shred/stop shredding** - Refer to ["Shred a tape"](#) for more information.

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Vault* tab, select a tape, click the *Manage* icon, and select the action you want to perform.

Shred a tape

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

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

Notes:

- Tape shredding may adversely affect backup performance. We recommend that you perform tape shredding when there are no backups running.
- When you shred a VIT, the index information will be erased and data in the repository will be cleaned up during reclamation.

To shred tapes:

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Vault* tab, select one or more tapes, click the *Manage* icon, and select *Shred*. If desired, select the option to delete the tape after shredding it. Note that when a WORM tape is shredded, it will automatically be deleted.

If you want to cancel the shredding process, select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Vault* tab, select one or more tapes, click the *Manage* icon, and select *Stop Shredding*.

Perform tape caching functions

You can perform the following functions for a tape in the vault with tape caching enabled:

- Export cached data - Migrate data in a cache to physical tape. The tape must have unmigrated data and cannot be a stub tape.
- Export entire tape - Migrate an entire tape to physical tape even if there is no unmigrated data. The tape cannot be a stub tape.
- Reclaim cache space - Delete the data that has been migrated to physical tape in order to free up cache disk space.

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Vault* tab, select a tape with tape caching enabled, click the *Tape Caching* icon, and select the action you want to perform.

Migrate virtual tape data to object storage

Migration is automatically triggered when a virtual tape is ejected to the vault. If you enabled Cloud Tiering on a StorSafe system with existing tapes that can be migrated, you can manually trigger migration to object storage. Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Vault* tab, select a tape, and click the *Migrate to Object Storage* icon.

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

Convert a virtual tape to a stub tape

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

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. You need to remove the write-protected flag from the tape for deduplication to trigger.

To recover data from object storage:

1. Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, and click the *Vault* tab.
2. Select a stub tape and click the *Recover from Object Storage* icon.
3. Select the virtual library to which you want the recovered tape to reside.
4. 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.

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

You can select from *Bulk* (default), *Standard*, and *Expedited* (available for *Glacier* only).

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 portal 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.

6. Confirm the information and click *OK*.

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

Export data to physical tape

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Vault* tab, select one or more tapes, and click the *Export* icon to manually export data from a virtual tape to a physical tape in a physical tape library or standalone drive. Refer to "[Manual Export](#)" for more information.

Display virtual tape information

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

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

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

Tapes




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

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

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

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

Only a Super Administrator can access the *Tapes* page.



Depending upon the columns that have been selected for display, the properties for each tape are visible at a glance:

- Barcode
- Tape name
- Tape ID
- Size
- Used size
- Data written
- Tape status - online, offline, incomplete (missing physical segments)
- Location (click to see the tape library name)
- Parent virtual tape library ID (click for details)
- Cloud Tiering usage
- Cloud Tiering mode - Copy or Move
- Object storage account

- Indication if Cloud Tiering occurred
- Stub - Checked if Cloud Tiering occurred and a stub tape exists
- Write-protection usage
- Write-once-read-many (WORM) tape usage
- Encryption usage
- Deduplication policy
- Deduplication usage and status - Successful, pending, failed
- VIT
- Replication usage
- Replica server
- Replica tape ID
- Replication source server
- Replication source tape ID
- Replication of deduplicated tapes - Status on target servers for replication of deduplicated tapes
- Replication start time for replicas
- Source of replica promoted in Test Mode - Checked if this is a primary tape promoted from a replica in test mode
- Previously replica promoted Test Mode - Checked if this is a replica tape promoted from a replica in test mode
- Tape caching usage
- Unmigrated cache - Checked if tape caching is used and the tape has unmigrated data
- Direct-Link - Checked if tape caching is used and the tape is a direct link tape (a link to the physical tape)
- Customer

Tapes page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Filter	Filter tapes.
	Create	Create virtual tapes.
	Update	Update virtual tape properties, tape services (Auto Archive, Auto Replication, Cloud Tiering), and Tape Duplication.
	Delete	Delete tapes. For a replica resource, you can promote the replica before deletion.

Icon	Feature	Description
	Manage	Rename, move, or copy a tape.
	Tape caching	Enable tape caching, export cached data, export entire tape, reclaim cache space, renew cache.
	Replication	Create, update, and delete replication policies for a tape, start and suspend/resume replication, promote a replica.
	Information	Display tape summary, physical layout, and outgoing replication information.


Filter tapes

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

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

To filter tapes:

1. Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Tapes* tab, select a tape location, and click the *Filter Tapes* icon.
2. On the *General* tab, you can indicate the type of tape(s) you are looking for



This includes empty, direct access, and direct link tapes as well as a variety of tape properties, including Auto Archive, Remote Copy, Tape caching, Tape Duplication, Remote Export, write protection, media type (vault), and parent library (vault).

3. On the *Range* tab, you can enter a range of barcodes and/or sizes.

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

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

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

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

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

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

- On the *Deduplication* tab, you can look for tapes associated with/without a deduplication policy or from a specific replica source server.

The screenshot shows a dialog box titled "Filter Virtual Tapes" with a close button (X) in the top right corner. It has four tabs: "General", "Range", "Time", and "Deduplication". The "Deduplication" tab is selected. Inside the dialog, there are six rows of filter settings, each with a label and a dropdown menu set to "No Filter":

- Replication Source: No Filter
- Deduplication Policy: No Filter
- Deduplication: No Filter
- Deduplication-Replication: No Filter
- Deduplication Needed: No Filter
- Deduplication-Replication Needed: No Filter

At the bottom right of the dialog are two buttons: "Apply" (highlighted in blue) and "Cancel".

You can also filter based on whether a tape is deduplicated, deduplicated *and* replicated, in need of deduplication (tape with data that is not a pure VIT), or in need of replication.

You can combine filters and use the *Deduplicated* filter with the *Deduplication-Replication Needed* filter to display tapes that were deduplicated but not replicated.

- Click *Apply*.

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

Create virtual tapes

To create a tape for a virtual tape library, select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Tapes* tab, select a library, and click the *Create* icon. Refer to ["Create virtual tapes"](#) for more information.

You can also create virtual tapes from the *Libraries* tab.

Update virtual tapes

You can change the following for any server in tape library:

- Change a virtual tape barcode
- Write protect a tape
- If enabled, disable Capacity on Demand and expand the tape to its maximum capacity
- Set [Auto Archive](#), [Auto Replication](#), [Auto Migration](#), and/or [Tape Duplication](#)
 - By default, tapes inherit their configuration from the virtual tape library

To update a tape, select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Tapes* tab, select a tape, and click the *Update* icon.

Manage virtual tapes

You can perform the following functions for a virtual tape:

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

- Copy - Replicates the full contents of a single tape to a local or remote server. Refer to ["Tape Copy"](#) for more information.

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Tapes* tab, select a tape, click the *Manage* icon, and select the action you want to perform.

Delete virtual tapes

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

If replication is configured for a virtual tape, deleting the virtual tape removes the replication configuration on the source and replica. You can choose to promote the replica to become a usable virtual tape before deletion. Refer to ["Promote a virtual tape replica"](#) for more information.

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

Select *Promote Replica*, if applicable.

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

Perform tape caching functions

You can perform the following functions for a tape with tape caching enabled:

- Enable Tape Caching - Enable tape caching for existing uncached virtual tapes. This provides an easy way to migrate data on existing uncached virtual tapes to physical tapes without having to back up the data to new tapes. Note that in order to migrate data to physical tapes, physical tape(s) with corresponding barcodes must exist in the physical library.
- Export cached data - Migrate data in a cache to physical tape. The tape must have unmigrated data and cannot be a stub tape.
- Export entire tape - Migrate an entire tape to physical tape even if there is no unmigrated data. The tape cannot be a stub tape.
- Reclaim cache space - Delete the data that has been migrated to physical tape in order to free up cache disk space.
- Renew cache - Renew the cache for a direct link tape (link to a physical tape).

Select *Manage* from the menu bar, click the *Virtual Libraries* tab, select a StorSafe server, click the *Tapes* tab, select a tape with tape caching enabled, click the *Tape Caching* icon, and select the action you want to perform.

Configure replication for virtual tapes

Replication is a process that protects the information on a virtual tape by maintaining a copy of a virtual tape on the same server or on another StorSafe server.

Refer to ["Tape replication"](#) for more information.

Start replication for a virtual tape

Manually start replication for a virtual tape that is not part of a deduplication policy. Only one outgoing replication job can be active at a time; the next job will become active after the current job finishes.

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

Suspend/resume replication for a virtual tape

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

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

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

Display virtual tape information

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

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

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

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

Physical Tape Libraries


FalconStor StorSafe Servers support physical tape library, drives, and tapes that have been assigned to the server.

Physical tape libraries



Select *Manage* from the menu bar, click the *Physical Resources* tab, and then the *Physical Tape Libraries* tab. This page lists the physical tape libraries on the selected StorSafe server.

Only a Super Administrator can access the *Physical Tape Libraries* page.



Depending upon the columns that have been selected for display, the properties for each physical tape library are visible at a glance:

- Library name
- Library ID
- Indicates if the physical tape library has been assigned to the server
- Status: Online or offline
- Vendor
- Product
- Firmware
- Serial number
- ACSL (Adapter:Channel:SCSI ID:LUN)
- Number of slots
- Number of drives - Click for drive information
- Number of tapes
- Number of loaded tapes
- Owner - Server that owns this physical tape library
- Disabled - Indicates if the library has been marked as disabled

Physical Tape Libraries page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Manage	Change the library name, inventory physical tapes, reset library, enable/disable the library, scan for stacked tapes.
	Move tapes	Move tapes between drives, slots, and import/export slots.
	Assign library	Assign a physical tape library to the server.
	Unassign library	Unassign a physical tape library from a server.
	Import tapes	Import data from physical tapes.
	Assign ACSLS/Library Station	Assign an ACSLS/Library station to a physical StorSafe server.

Inventory a physical library

You can perform an inventory of the physical tapes in a physical tape library. This allows you to create virtual tapes that match your physical tapes.

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Libraries* tab, select a physical tape library, click the *Manage* icon, and select *Inventory*.

Enable/disable a physical library

For maintenance purposes, a physical library can be marked as disabled. While it is designated as disabled, the library cannot be used for any import/export functions.

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Libraries* tab, select a physical tape library, click the *Manage* icon, and select *Disable*.

Afterwards, when the library is available, you can enable it.

Reset a physical library

Resetting a library terminates all active import/export jobs and dismounts all tapes from their tape drives. This is useful after a problem (such as a physical tape stuck in a slot) is resolved.

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Libraries* tab, select a physical tape library, click the *Manage* icon, and select *Reset*.

Scan a physical library for stacked tapes

Tape stacking allows multiple virtual tapes to be exported to a single physical tape in a physical tape library, maximizing physical tape usage and allowing the conversion of virtual media with a smaller capacity to physical media with a higher capacity.

You can scan a library to see if a physical tape has stacked virtual tapes. Stacked tapes will appear on the *Tape History* page.

Scanning physical tapes is useful if you have loaded or ejected a tape from a library or if you have different tapes with the same barcode and they are used in the system at different times.

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Libraries* tab, select a physical tape library, click the *Manage* icon, and select *Scan for Stacking*.

Move tapes in a physical library

You can move tapes in a physical tape library between drives, slots, and import/export slots (vault).

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Libraries* tab, select a physical tape library, and click the *Move Tapes* icon. Use the *Move Type* field to select the current location of the tapes and the location type for the destination.

Assign a physical library

You can assign a physical tape library to a server.

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Libraries* tab, select a physical tape library, and click the *Assign* icon.

Unassign a physical library

You can unassign a physical tape library from a server.

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Libraries* tab, select a physical tape library, and click the *Unassign* icon.

Import data from physical tapes to a virtual library

One of the advantages of using a virtual tape library is that you can protect data on your existing physical tapes by importing them into your virtual tape system. If you need to recover files from a physical tape, you can use the import function to directly access the physical tape for immediate recovery. The import function allows you to:

- Copy the contents of a physical tape to a virtual tape
- Directly access a physical tape without copying the entire tape
- Recycle a physical tape
- Import data from virtual tapes used by another FalconStor virtual tape library

You must do the following before you can import tapes:

- Verify that the drive type (i.e. IBM:ULTRIUM-TD3) and the media type (i.e. Ultrium3) of the library you are importing from match the library to which you are importing.
- If you are connecting through a switch, verify that the StorSafe server is in the same zone as the physical tape library or other StorSafe server from which you are importing.

To import tapes:

1. Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Libraries* tab, select a physical tape library, and click the *Import Tapes* icon.
2. Select how you want the data copied.

Copy - Copy the contents of a physical tape onto a virtual tape, leaving the physical tape unchanged. This is the only mode available if you are importing a stacked tape.

Direct Access - Link a physical tape to its virtual counterpart. This gives the backup application immediate access to the tape data without waiting for a complete copy. This is useful when you need to restore a small amount of data from a physical tape. Direct access tapes are write protected. Therefore, you can only read from the tape and not write to it. This mode is not applicable for stacked tapes.

Recycle - Recycle a physical tape after its retention period has been reached. Once the virtual tape is subsequently initialized, the physical tape can be recycled and used for future export operations. This mode is not applicable for stacked tapes.

3. Select the virtual library to which tapes will be imported.

Be sure to pick a library with the same tape size capacity.

Note that stacked tapes can only be imported to virtual libraries that are not enabled with Automated Tape Caching. The maximum number of stacked tapes that can be imported is 1,000.

4. Specify a storage pool to use for creating virtual tapes.
If empty, the first available physical device will be used.
5. Specify whether or not to decrypt data on an encrypted physical tape.
If you select this option, you must select the key to use.

Note: If you select this option and the data was not previously encrypted, an incorrect key is selected, or an invalid password is provided, imported data will be indecipherable.

6. If the tape was exported with stacking, select *Import Stacked Tapes*.
7. Select up to nine physical tapes you want to import.
You can use the same barcode for the virtual tape or you can double-click in the field to enter a new barcode with 4 to 12 characters. The default is the physical barcode
You can also select a slot for the virtual tape.
A WORM tape can be only be imported from a WORM tape.
8. Click *Import* when done.

Import data from a tape in another virtual tape library

In order to import data from a virtual tape in another FalconStor virtual tape library, the two StorSafe servers must be connected or zoned together so that they can communicate with each other. Before running the *Import Tapes* procedure above, you will need to do the following:

1. On the source StorSafe server (where your original tape is located), create a client to represent the target StorSafe server (the server to which the tape will be imported).
2. Rescan adapters on the target StorSafe server to which you are importing tapes and select the *Discover New Devices* option.
3. After re-scanning, verify that the tape library and drives you are importing from are displayed.
4. Assign the physical tape library object.

Select all tape drives that belong to the physical tape library created previously.

Once completed, run the *Import Tapes* procedure described above and select *Copy* as the *Mode*.

Assign ACSLS/Library Station

You can assign an Automated Cartridge System Library Software (ACSL) managed library or a Library Station to the server to manage heterogeneous StorageTek tape libraries. The ACSLS/Library Station option is only available for physical StorSafe systems; it is not available for virtual appliances.

The ACSLS/Library Station option works with ACSLS/Library Station-managed tape libraries, allowing the system to share ACSLS/Library Station-managed libraries among the StorSafe server and your backup application servers. This makes it possible to import data from physical tapes and export data on virtual tapes to physical tapes.

Hardware configuration

1. Physically connect the tape drives that will be assigned to the StorSafe appliance.

Note: Physical tape drives cannot be shared with StorSafe appliances or other applications.

2. (ACSL only) Create at least one storage pool on the ACSLS server for StorSafe and assign tapes to it.


If you have already created a pool, you can use that one.

3. Make a note of the following:
 - ACS IDs, LSM IDs, Panel IDs, and Device IDs of the libraries that hold the tape drives connected to StorSafe. You can run the `cmd_proc` utility on your ACSLS server to determine the IDs. For Library Station users, check with your Library Station administrator to determine the IDs.
 - IP address of the ACSLS/Library Station server.
 - (ACSL only) IDs of the storage pools to be assigned to the StorSafe appliance.
4. Make sure that the StorSafe server and the ACSLS/Library Station server can communicate with each other.

Configure StorSafe to work with ACSLS

The following instructions give you an overview of the steps you must follow to configure your ACSLS/Library Station option.

1. Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Libraries* tab, select a physical tape library, and click the *Assign ACSLS/Library Station* icon.



The dialog box titled "Assign ACSLS/Library Station" contains the following fields and controls:

- Library Type:** Two radio buttons, "ACSLS" (selected) and "Station".
- IP Address:** A text input field with the placeholder text "IP Address".
- ACS ID:** A spin box with a value of "0".
- Pool ID:** A spin box with a value of "0".
- Firewall Support:** A toggle switch currently turned off.
- Buttons:** "Assign" and "Cancel" buttons at the bottom right.

ACSLS

Enter the IP address of the ACSLS, the ACS ID (range 0-126, default 0), and the Pool ID (range 0-65534, default 0).

If applicable, enter your firewall information. You need to specify the ACSLS Server System Interface (SSI) port number (range 0-65536, default 1024) and the Client System Interface (CSI) port number (range 0-65536, default 30031). SSI and CSI handle communication between ACSLS servers and clients, such as StorSafe.

SSI is the network interface component of ACSLS that runs on the client. CSI is the network interface component of ACSLS that runs on the ACSLS platform. When SSI and CSI are used across a firewall, you will need to open firewall ports to allow communication between ACSLS and the StorSafe server. SSI and CSI are independent, so if only one of the sides is behind a firewall, ports need to be opened only that side. The other side will use dynamic ports in the range of 1024-65535.

For example, if firewall security is implemented on the ACSLS server side, you must configure your firewall to allow the client to send requests to the ACSLS server using the port specified by `CSI_INET_PORT` (on the ACSLS server) and `CSI_HOSTPORT` (on the client). Refer to the *Running ACSLS behind a Firewall* section in your ACSLS administration guide for more details.

By design, when a firewall is set for an ACSLS server, only one ACS physical library can be assigned using the same `SSI_INET_PORT` defined in the `acsls_ls_cdk.conf` file. However, if the CSI and SSI INET ports are not set in the configuration file, you can use portmapper to assign several physical libraries. For example, use the following configuration to assign two ACS

physical libraries from an ACSLS behind a firewall. Open all of the ports in your firewall configuration but do not define them in the `acsls_ls_cdk.conf` file.

For the first ACS use:

- SSI inet port: 1024
- CSI host port: 30031

For the second ACS use:

- SSI inet port: 1025
- CSI host port: 30031

If you do not select the *Firewall Support* option, the *portmap* service needs to be running. Otherwise, the system will fail to assign or retrieve the library's status after restarting StorSafe services or rebooting.

To enable the *portmap* service on Linux, run the following command:

```
chkconfig --add portmap
```

Station

Enter the IP address of the Library Station and the ACS ID (range 0-126, default 0).

2. Click *Assign*.

Once completed, the server automatically does an inventory to obtain a list of physical tapes.

Configure ACSLS to work with StorSafe Failover

In order to communicate with a StorSafe server in a failover configuration, the virtual IP address is used. However, ACSLS uses the heartbeat IP address instead.

Therefore, you must identify each StorSafe server by the heartbeat IP address so that ACSLS will continue to function after a takeover (when the secondary server takes over for the primary). Perform the following steps on each StorSafe server:

1. Modify the `/etc/hosts` file and use the heartbeat IP instead of the virtual IP.
2. Run the following command:

```
vtl restart tleupcall
```

Set eject policy

You need to edit the `$ISHOME/etc/acsls_ls_cdk.conf` file to set how many tapes will be ejected after each backup.

Enter your policy information in the following format:

```
RepositoryID:ACSLS/LS IP:acsid:poolid:eject
```

In the following example:


```
4:10.6.2.62:1:2:5
```

- 4 is the Repository/Database ID.
- 10.6.2.62 is the IP address of the ACSLS/Library Station.
- 1 is the ACS ID.
- 2 is the Pool ID.
- 5 is the eject value. In this example, after each backup, the system will eject five tapes at a time. This is usually set to the same value as the *CAP* size.

Filter tapes displayed in the portal

If you pool physical tapes in Library Station and you *only* want to display a range of Library Station barcodes in the portal, you will need to edit the `$ISHOME/etc/acsls_ls_cdk.conf` file.

Locate the following sections:

```
#Library Station barcode filters...
#end Library Station barcode filters
```

Enter your range information **between** these sections (before the `#end` line) in the following format:

```
RepositoryID:ACSLS/LS IP:acsid:startrange-endorange:startrange-endorange
```

The barcode length should be six characters and spaces should not be used. Any number of barcode ranges can be specified. In the following example:

```
4:10.6.2.62:6:000000-004444:AA0000-AA8888
```

- 4 is the Repository ID.
- 10.6.2.62 is the IP address of the ACSLS/Library Station.
- 6 is the ACS ID.
- 000000-004444 is the first range to display.
- AA0000-AA8888 is the next range to display.

Add/remove tapes

Whenever you add or remove tapes from an ACSLS/Library Station pool, you must inventory the tapes through the portal.

Replace ACSLS physical drive

This section explains how to replace ACSLS physical drives that are assigned to StorSafe for tape import/export purposes.

Notes:

- Take an X-ray of the StorSafe server before making any changes.
- Import/export jobs should not access any physical drive that is being replaced. If possible, stop all import/export jobs until the procedure is complete.

1. From the FalconStor portal, find the physical drive and make a note of the following information:
 - SCSI ID
 - ACSLS device ID
 - Serial number of the physical drive

2. Unassign the physical drive.

3. Remove and replace the physical drive in the physical library.

4. Use the ACSLS command interface to configure the new drive.

This can be done using: `config drive 'panel_id'`

For example, if the device ID is 0,2,1,12, you would use the following command:

```
config drive 0,2,1,12
```

Refer to the ACSLS library administrator's guide for more information.

5. From the FalconStor portal, rescan for the new device.

Use the *Discover new devices* option.

6. In the FalconStor portal, assign the new drive to StorSafe.


Add the device ID from step 1.

Afterward, check the physical tape drive to confirm that the serial number of this new drive matches.

Physical tape drives

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

Only a Super Administrator can access the *Physical Tape Drives* page.



ID	Tape Library	Vendor	Product	Firmware	Serial Number	ACSL	Owner
424	ADIC:Scalar 100-0...	IBM	ULTRIUM-TD1	18N2	78ENZ00314	3:0:0:10	VM22-
425	ADIC:Scalar 100-0...	IBM	ULTRIUM-TD1	18N2	78ENZ00315	3:0:0:11	VM22-
426	ADIC:Scalar 100-0...	IBM	ULTRIUM-TD1	18N2	78ENZ00316	3:0:0:12	VM22-
427	ADIC:Scalar 100-0...	IBM	ULTRIUM-TD1	18N2	78ENZ00317	3:0:0:13	VM22-
428	ADIC:Scalar 100-0...	IBM	ULTRIUM-TD1	18N2	78ENZ00318	3:0:0:14	VM22-
429	ADIC:Scalar 100-0...	IBM	ULTRIUM-TD1	18N2	78ENZ00319	3:0:0:15	VM22-

Depending upon the columns that have been selected for display, the properties for each physical tape drive are visible at a glance:

- Drive name
- Drive ID
- Vendor
- Product
- Firmware
- Serial number
- ACSL (Adapter:Channel:SCSI ID:LUN)
- Owner - Server that owns this physical tape drive
- Status: Empty, loaded, ejected, offline, pass through, ready (becoming ready), unloading, disabled
- Disabled - Indicates if the drive has been marked as disabled

Physical Tape Drives page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Manage	Change the drive name and enable/disable the drive. For standalone tape drives, prepare for stacking, retrieve barcode, and scan for stacked tapes.
	Assign drive	Assign a physical tape drive to the server as a standalone physical tape drive or to a physical tape library.
	Unassign drive	Unassign a physical tape drive from a library.
	Import tapes	Import data from physical tapes.

Enable/disable a physical tape drive

For maintenance purposes, a physical drive can be marked as disabled. While it is designated as disabled, the drive cannot be used for any import/export functions.

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Drives* tab, select a physical tape drive, click the *Manage* icon, and select *Disable*.

Afterwards, when the library is available, you can enable it.

Prepare for stacking

You can prepare a standalone physical tape drive for export with tape stacking in order to assign a barcode to the physical tape before export.

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Drives* tab, select a standalone physical tape drive, click the *Manage* icon, and select *Prepare for Stacking*. The physical tape barcode you specify must be 4 to 12 characters.

Retrieve barcode

You can retrieve the barcode of a physical tape in a standalone physical tape drive from the tape header. This is required to record the physical tape barcode in the server memory when you want to import stacked tapes or to export stacked tapes along with appending data.

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Drives* tab, select a standalone physical tape drive, click the *Manage* icon, and select *Retried Barcode*.

Scan a standalone physical drive for stacked tapes

Tape stacking allows multiple virtual tapes to be exported to a single physical tape in a standalone physical tape drive, maximizing physical tape usage and allowing the conversion of virtual media with a smaller capacity to physical media with a higher capacity.

You can scan a standalone tape drive to see if a physical tape has stacked virtual tapes.

Scanning physical tapes is useful if you have loaded or ejected a tape from the drive or if you have different tapes with the same barcode and they are used in the system at different times.

Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Drives* tab, select a standalone physical tape drive, click the *Manage* icon, and select *Scan for Stacking*.

Import data from physical tapes to a virtual library

You can protect data on your existing physical tapes by importing them into your virtual tape system. If you need to recover files from a physical tape, you can use the import function to directly access the physical tape for immediate recovery.

The import function allows you to:

- Copy the contents of a physical tape to a virtual tape
- Directly access a physical tape without copying the entire tape
- Recycle a physical tape

You must do the following before you can import tapes:

- Verify that the drive type (i.e. IBM:ULTRIUM-TD3) and the media type (i.e. Ultrium3) of the library you are importing from match the library to which you are importing.
- If you are connecting through a switch, verify that the StorSafe server is in the same zone as the physical tape library or other StorSafe server from which you are importing.
- Run *Retrieve Barcode* before importing.

To import tapes:

1. Select *Manage* from the menu bar, select a StorSafe server, select the *Physical Resources* tab, select the *Physical Tape Drives* tab, select a standalone physical tape drive, and click the *Import Tapes* icon.

2. Select how you want the data copied.

Copy - Copy the contents of a physical tape onto a virtual tape, leaving the physical tape unchanged. This is the only mode available if you are importing a stacked tape.

Direct Access - Link a physical tape to its virtual counterpart. This gives the backup application immediate access to the tape data without waiting for a complete copy. This is useful when you need to restore a small amount of data from a physical tape. Direct access tapes are write protected. Therefore, you can only read from the tape and not write to it. This mode is not applicable for stacked tapes.

Recycle - Recycle a physical tape after its retention period has been reached. Once the virtual tape is subsequently initialized, the physical tape can be recycled and used for future export operations. This mode is not applicable for stacked tapes.

3. Select the virtual library to which tapes will be imported.

Be sure to pick a library with the same tape size capacity.

Note that stacked tapes can only be imported to virtual libraries that are not enabled with Automated Tape Caching. The maximum number of stacked tapes that can be imported is 1,000.

4. Specify a storage pool to use for creating virtual tapes.

If empty, the first available physical device will be used.

5. Specify whether or not to decrypt data on an encrypted physical tape.

If you select this option, you must select the key to use.

Note: If you select this option and the data was not previously encrypted, an incorrect key is selected, or an invalid password is provided, imported data will be indecipherable.

6. If the tape was exported with stacking, select *Import Stacked Tapes*.

7. Select up to nine physical tapes you want to import.

You can use the same barcode for the virtual tape or you can double-click in the field to enter a new barcode with 4 to 12 characters. The default is the physical barcode

You can also select a slot for the virtual tape.

A WORM tape can be only be imported from a WORM tape.

8. Click *Import* when done.

Physical tapes



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

Only a Super Administrator can access the *Physical Tapes* page.

The screenshot shows the FalconStor StorSafe with StorSight User Guide interface. The top navigation bar includes 'Dashboard', 'Monitor', 'Manage' (selected), 'Analytics', and 'Administration'. The left sidebar shows 'Physical Resources' with a search bar and a list of servers: H13-72, H13-82, and OBD190. The main content area is titled 'Physical Resources' and contains a sub-tabbed interface. The 'Physical Tapes' tab is selected, showing a table of tapes. The table has columns for 'Barcode', 'Location', and 'Stacked'. The 'Tape Library' dropdown is set to 'IBM:03584L32-04912'. The table lists seven tapes with their respective barcodes and locations, all marked as 'Stacked'.

Barcode	Location	Stacked
020703L8	Library Slot 0	☐
020704L8	Library Slot 1	☐
020705L8	Library Slot 2	☐
020706L8	Library Slot 3	☐
020707L8	Library Slot 4	☐
020708L8	Library Slot 5	☐
020709L8	Library Slot 6	☐

Depending upon the columns that have been selected for display, the properties for each tape are visible at a glance:


- Barcode
- Location - Library slot, library drive, import/export slot
- Indicates if the tape is a stacked tape (multiple virtual tapes on a single physical tape)

Physical tape export history



Select *Manage* from the menu bar, click the *Physical Resources* tab, click the *Physical Tapes* tab, and then the *Export History* tab to see a list of physical tapes that were used for export jobs (including stacked tapes). Select the physical tape library from the drop-down list.

Only a Super Administrator can access the *Export History* page.



Depending upon the columns that have been selected for display, the properties for each tape are visible at a glance:

- Physical tape barcode
- Record time
- Virtual tape barcode
- Virtual tape creation time
- Virtual tape modification time
- Virtual tape data size

Export History page icon

The table below describes the unique icon for this page and briefly describes the associated features:

Icon	Feature	Description
	Delete	Delete the selected tape history.

Host Clients

Host clients are the backup application servers that are assigned to virtual tape libraries and drives for backup.

Host Clients page

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

The screenshot displays the FalconStor StorSafe web interface. The top navigation bar includes 'Dashboard', 'Monitor', 'Manage' (selected), 'Analytics', and 'Administration'. The 'Host Clients' tab is active in the sub-menu. The left sidebar shows a list of servers with their names, IP addresses, and build numbers. The main area contains a table of host clients with columns for Client, ID, FC, iSCSI, Devices, iSCSI Initiators, iSCSI Targets, FC WWPNs, VSA, and iSeries.

Client	ID	FC	iSCSI	Devices	iSCSI Initiators	iSCSI Targets	FC WWPNs	VSA	iSeries
10.8.25.13	13	☒	☐	12			2100001b3297d963	☐	☐
10.8.25.150	18	☒	☐				2100001b32920b9e	☐	☐
pra2	28	☐	☒		2			☐	☐
pra3	29	☐	☒		pra1	1		☐	☐

Depending upon the columns that have been selected for display, the properties for each client are visible at a glance. You can click a value displayed under column headings to see a pop-up with additional information.

- Client name
- Client ID
- FC - A check in this box indicates a client with Fibre Channel protocol.
- iSCSI - A check in this box indicates a client with iSCSI protocol.
- Devices - The number of virtual tape libraries and drives assigned to this client. Click the number to see a pop-up that identifies each virtual tape library/drive, Logical Unit Number (LUN), communications protocol, and access rights.
- iSCSI initiators - The number of initiators for iSCSI clients. Click the number to see a pop-up that identifies each initiator IQN.

- iSCSI targets - The number of targets for iSCSI clients. Click the number to see a pop-up that identifies each target name, access rights, and IP address.
- FC WWPNs - The World Wide Port Name (WWPN) or number of WWPNs for FC clients. Click the number to see a pop-up that identifies each WWPN.
- VSA - A check in this box indicates the client has Volume Set Addressing (VSA) enabled.
- iSeries - A check in this box indicates the client emulates IBM tape libraries and drives for an IBM System i server.
- Customer - (Super Administrators only) The customer to which the client belongs.

Host Clients page icons

The table below describes the unique icons for this page and briefly describes the associated features:

Icon	Feature	Description
	Create	Add a new client.
	Update	Update the client name.
	Delete	Delete the selected client.
	Manage	• Update iSCSI initiators • Create/update/delete iSCSI target • Update FC initiators • Update FC device (change LUN)
	Assign	Assign one or more virtual tape libraries or drives to a client.
	Unassign	Unassign one or more devices from a client.

Add iSCSI clients

1. Click *Manage* from the menu bar, select a StorSafe server, and click the *Host Clients* tab.
2. Click the “+” icon.
3. From the *General* tab, specify the client name (maximum of 64 characters) and select iSCSI as the client’s protocol(s).

Create Client

General iSCSI

Name*

Protocols ☐ Fibre Channel ☒ iSCSI

Create Cancel

4. Click the *iSCSI* tab to select the initiators that this client will use.

Create Client

General iSCSI

CHAP Authentication ☐ ?

+Add Initiator Rescan Initiators Search for...

	iSCSI Client Initiator	IP Address	Available	Used by
☐	iqn.1988-12.com.oracle:65176a392bc4	10.7.2.201	No	Frank
☒	YJiscsi		Yes	Javadi
☐	dummy		Yes	Javadi

Create Cancel

If the initiator does not appear, you may need to rescan by clicking the *Rescan iSCSI Initiators* tab. You can also manually add it by clicking the “+”.

5. Enable *CHAP Authentication* to specify if authentication will be used.
If authentication is disabled, the StorSafe server will recognize the client as long as it has an authorized initiator name.
6. If CHAP authentication is enabled, select at least one iSCSI user who can authenticate for this client.
The iSCSI user password is used as the *CHAP Secret* is used.
7. Confirm all information and click *Create* to add this client.

Create an iSCSI target for an iSCSI client

Once an iSCSI client has been added, you can add an iSCSI target. To do this:

1. Click *Manage* from the menu bar, select a StorSafe server, and click the *Host Clients* tab.
2. Select the iSCSI client for which you want to add an iSCSI target.
3. Click the *Manage* icon and select *Create iSCSI Target*.

The screenshot shows a 'Create iSCSI Target' dialog box. It contains the following fields and values:

- Client:** H12-90
- iSCSI Target Name:** iqn.2000-03.com.falconstor.vtl.FBDVA-H16-96.H1
- IP Address:** 172.22.16.96
- Starting LUN:** 0

At the bottom right, there are two buttons: 'Create' (highlighted in blue) and 'Cancel'.

4. Enter a new target name for the client (maximum of 223 characters) or accept the default.

Note: The Microsoft iSCSI initiator can only connect to an iSCSI target if the target name contains less than 221 characters.

5. Select the IP address of the adapter on the StorSafe server.

If you are using failover, be sure to select the server's IP address, not the heartbeat IP address, so that clients can see devices while in failover mode.

6. Specify the starting LUN.

LUN IDs must start with zero; use the default LUN.

Once the iSCSI target is created for a client, LUNs can be assigned under the target using available virtual iSCSI devices.

7. Confirm all information and click *Create*.

Add Fibre Channel clients

1. Click *Manage* from the menu bar, select a StorSafe server, and click the *Host Clients* tab.
2. Click the “+” icon.
3. From the *General* tab, specify the client name (maximum of 64 characters) and select Fibre Channel as the client’s protocol(s).

Create Client

General Fibre Channel

Name*

Protocols ☒ Fibre Channel ☐ iSCSI

Create Cancel

4. Click the *Fibre Channel* tab to select WWPN initiators and set options.

Create Client

General Fibre Channel

Volume Set Addressing ☐ ?

IBM iSeries ☐ ?

+ Add Initiator Rescan Initiators Search for...

	WWPN	Port	Available	Used by
☐	21000024ff31cb48	dc1800	No	10.8.25.190
☐	2100000d7731cb48	dc1801	No	yeg
☒	2100001b32920b9e	dc1400	Yes	
☒	2100000d779186b1	dc1701	Yes	

Create Cancel

Similar to an IP address, the WWPN uniquely identifies a port in a Fibre Channel environment. If necessary, click + *Add Initiator* to add a WWPN for the client.

Select *Volume Set Addressing* (VSA) if your storage device uses VSA.

Select *IBM iSeries* to emulate IBM tape libraries and drives for an IBM System i server.

5. Confirm all information and click *Create* to add this client.

Associate World Wide Port Names (WWPN) with clients

Similar to an IP address, the WWPN uniquely identifies each initiator, target, and client initiator in a Fibre Channel environment. Unlike an IP address, the WWPN is vendor assigned and is hard-coded 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, StorSafe will query the switch for its Simple Name Server (SNS) database and will display a list of all available WWPNs. To see this list, select *Manage* from the menu bar, click the *Physical Resources* tab, click the *Adapters* tab, locate the HBA that matches the recorded target HBA WWPN, and click *SNS Information*. Look for the WWPN that matches the recorded initiator WWPN.
- 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 the boot process or you may have to read it from the physical card. Check with your hardware vendor for their preferred method.

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 StorSafe server, the following message appears on the server screen:

```
FSQLtgt: New Client WWPN Found: 21 00 00 e0 8b 43 23 52
```

Do the following for each client for which you want to assign specific virtual resources:

1. Click *Manage* from the menu bar and select a server.
2. Click the *Host Clients* tab.


3. Select a client, click the *Manage* icon, and select *Update FC Initiators*.



The 'Update FC Properties' dialog box shows the client 'H12-202'. The 'Volume Set Addressing' toggle is off. Below are buttons for '+ Add Initiator' and 'Rescan Initiators', and a search bar. A table lists WWPNs and ports, with the first row selected. The table has columns: WWPN, Port, Available, and Used by. The first row is checked and shows 'H12-202' in the 'Used by' column. The other rows are unchecked and show 'Yes' in the 'Available' column. At the bottom right are 'Update' and 'Cancel' buttons.

WWPN	Port	Available	Used by
☒ 21000024ff349c84	010a00	☒ Yes	H12-202
☐ 2100001b32139190	040900	☒ Yes	
☐ 21000024ff3cddf0	1c0a00	☒ Yes	
☐ 2101001b32bd8636	1c0e00	☒ Yes	
☐ 2100001b329d8636	1c0f00	☒ Yes	
☐ 21000024ff3f62c3	021400	☒ Yes	

4. If necessary, click + *Add* WWPN.



The 'Add FC Initiator WWPN' dialog box prompts the user to 'Enter a valid WWPN in the xx-xx-xx-xx-xx-xx format:'. It features a text input field and 'Add' and 'Cancel' buttons at the bottom right.

5. Click *OK* when done.

Manage clients

Assign virtual tape libraries and drives to clients

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

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

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

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

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

For FC clients, you need to associate the Initiator and Target WWPN(s) belonging to the client.

Assign Devices to Client

Client: 10.8.25.190

Initiator WWPN*: 21-00-00-24-ff-31-cb-48

Target WWPN*: 21-00-00-0d-77-97-d9-63


Assign: Virtual Tape Libraries

Virtual Tape Library	ID	Vendor	Product	Drives
ADIC-Scalar 100-00185noCOD	185	ADIC	Scalar 100	
ADIC-Scalar 100-00192withCOD	192	ADIC	Scalar 100	
ADIC-Scalar 100-00193	200	ADIC	Scalar 100	
ADIC-Scalar 100-00201	207	ADIC	Scalar 100	
ADIC-Scalar 100-00279autorep	335	ADIC	Scalar 100	
ADIC-Scalar 100-00342arjava	342	ADIC	Scalar 100	
ADIC-Scalar 100-02946pra	126	ADIC	Scalar 100	

Showing 18 entries

Assign Cancel

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



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

4. Click *Assign*.
5. Use the backup application server's operating system to discover the StorSafe 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*.

6. Use your backup software to discover the library.

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

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](#).

Update client properties

You can perform the following functions for clients:

- Update the client name
- Delete the client
- Update FC initiators - Enable/disable Volume Set Addressing
- Update FC device - Change the LUN of an assigned device
- Update iSCSI initiators - Enable/disable CHAP authentication
- Update iSCSI targets - Change the name and IP address
- Delete iSCSI targets - Remove the client assignment from an iSCSI target and/or delete the target if there is no client assignment

To update or delete a client, select *Manage* from the menu bar, select a StorSafe server, click the *Host Clients* tab, select a client, and click *Update* or *Delete*.

To update target/initiator information, select *Manage* from the menu bar, select a StorSafe server, click the *Host Clients* tab, select a client, click the *Manage* icon, and select the action you want to perform.

Unassign virtual tape libraries and drives from clients

1. Click *Manage* from the menu bar, select a StorSafe server, and click the *Host Clients* tab.
2. Select a client and click the *Unassign* icon.
3. Select one or more virtual tape libraries or drives and click *Unassign*.


Fibre Channel Configuration

StorSafe supports different protocols to communicate with downstream storage devices such as SCSI, Fibre Channel (FC), and iSCSI. StorSafe also supports different protocols to communicate with upstream clients; Fibre Channel target mode and iSCSI target mode can be enabled on StorSafe so that clients can take advantage of extensive storage capabilities such as virtualization, mirroring, replication, NPIV, and security. Support is offered for all FC topologies, including point-to-point and fabric.

This section provides configuration information for Fibre Channel target mode as well as the associated Fibre Channel SAN equipment (i.e. switch, T3, etc.). An application server can be an iSCSI Client, a Fibre Channel Client, or both. Using separate cards and switches, you can have all types of clients (FC, or iSCSI) on your storage network.

Fibre Channel target mode is supported on various platforms, including Windows, VMware, and Linux.

FC Target Mode



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.

The configuration of Fibre Channel target mode includes the following steps.

1. [Configure Fibre Channel on server.](#)
2. [Configure Fibre Channel on storage.](#)
3. [Configure Fibre Channel on host clients.](#)
4. [Enable Fibre Channel and/or iSCSI target mode.](#)
5. (If applicable) Set HBA ports to target mode.
Refer to ["Update adapter settings"](#).
6. (Optionally) Enable [Failover](#).
7. [Add Fibre Channel clients.](#)
8. (Optionally) [Associate World Wide Port Names \(WWPN\) with clients.](#)
9. [Assign virtual tape libraries and drives to clients.](#)

In order to see the new devices after you have finished configuring your Fibre Channel clients, you must trigger a device rescan or reboot the client machine, depending upon the requirements of the operating system.

Configure Fibre Channel on server

Ports

Your StorSafe will be equipped with several Fibre Channel ports. Some of these ports will interface with storage arrays. Others will interface with physical tape libraries, while the remaining ports will interface with backup (media) servers.

The ports that connect to storage arrays are commonly known as *Initiator Ports*. The ports that will interface with the backup application servers' FC initiator ports will run in a different mode known as *Target Mode*. The ports that are connected to physical tape libraries are known as *Library Connection Ports*.

Fibre Channel connection to downstream storage or upstream client can be via a direct connection or an FC switch.

The following configurations define Fibre Channel topologies:

- Point-to-point: Two devices are connected directly providing dedicated bandwidth.
- Arbitrated loop: All devices are connected in a loop; only one pair of ports can communicate concurrently; failure on one port affects other ports.
- Switched fabric: All devices are connected to Fibre Channel switches; the fabric can scale to numerous ports; failure of a port is isolated and does not affect other ports; a multiple pair of ports can communicate simultaneously.

QLogic drivers

QLogic multi-ID driver Multi-ID is the default driver providing N_Port ID Virtualization (NPIV), which 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. Refer to the FalconStor website certification matrix for a list of supported HBAs.
- The fabric switch must support NPIV.
- If a QLogic FC switch is being used, you must disable "IOStreamGuard" for any switch port that connects to an NPIV target port.

When using the NPIV driver, there are two WWPNs, the *base* port and the *alias*.

Notes:

- With dual mode, clients will need to be zoned to the alias port (called *Target WWPN*). If they are zoned to the base port, clients will not see any devices.
- You will only see the alias port when that port is in target mode.
- You will only see the alias once all of the StorSafe services are started.

QLogic single-ID driver The single-mode driver provides a point-to-point configuration, where each port is either a target or an initiator.

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, Port World Wide Name (pWWN) zoning (also known as soft zoning), and hard zoning (based on port #).

pWWN zoning pWWN zoning is required for the QLogic NPIV driver and uses the WWPN in the configuration. The WWPN remains the same in the zoning configuration regardless of the port location. If a port fails, you can simply move the cable from the failed port to another valid port without having to reconfigure the zoning.

The rules for pWWN zoning between StorSafe and SIR servers using the QLogic NPIV driver are:

- No StorSafe or SIR server target can be zoned with another of its own targets.
- No StorSafe or SIR server target can be zoned with one of its own initiators.
- No StorSafe or SIR server initiators can be zoned with one of its own initiators.

Hard zoning Hard zoning is only supported for QLogic drivers without NPIV. It uses the port number of the switches for zoning. With hard zoning, if a zone has two ports (0 and 1) and port 0 goes down for some reason, you will need to remove the current zoning configuration, move the plug to another valid port, re-zone, and then enable the new zoning configuration.

If hard zoning is used, it is necessary to create zones for each standby target, doubling the number of upstream zones. This extra set of zones is not necessary in the case of soft zoning because zones are defined by WWPN combinations. In a failover event, the standby ports assume the WWPNs of the target ports of the failed StorSafe server. Therefore, the single set of soft zones is still valid.

General zoning requirements StorSafe 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 StorSafe and SIR appliances.

Additionally, make sure that storage devices to be used by StorSafe are not zoned to clients (backup application servers). Ports on storage devices to be used by StorSafe should be zoned to StorSafe's initiator ports while the clients are zoned to StorSafe's target ports. Make sure that from the storage unit's management GUI (such as SANtricity and NaviSphere), the LUNs are re-assigned to StorSafe as the "host". StorSafe will virtualize these LUNS. StorSafe can then define virtual tapes out of these LUNS and further provision them to the clients.

Persistent binding

Persistent binding is automatically enabled for all QLogic HBAs connected to storage device targets upon the discovery of the device (via a Console physical device rescan with the *Discover New Devices* option enabled). However, persistent binding will not be SET until the HBA is reloaded. You can reload HBAs by restarting StorSafe 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).

QLogic adapter settings

The `fshba.conf` file in `$ISHOME/etc` is used to adjust settings for FC adapters installed on the StorSafe appliance.

The table below lists the default settings in the `fshba.conf` file that will override the BIOS settings of the HBA.

Name	Default
frame_size	2 (2048byte)
loop_reset_delay	0
adapter_hard_loop_id	0
connection_option	1 (point to point)
hard_loop_id	0
fibre_channel_tape_support	0 (disable)
data_rate	2 (auto)
execution_throttle	255
LUNs_per_target	256
enable_lip_reset	1 (enable)
enable_lip_full_login	1 (enable)
enable_target_reset	1 (enable)
login_retry_count	8
port_down_retry_count	8
link_down_timeout	45

Name	Default
extended_error_logging_flag	0 (no logging)
interrupt_delay_timer	0
iocb_allocation	512
enable_64bit_addressing	0 (disable)
fibrechannelconfirm	0 (disable)
class2service	0 (disable)
acko	0 (disable)
responsetimer	0 (disable)
fastpost	0 (disable)
driverloadriscocode	1 (enable)
q12xmaxqdepth	255
max_srbs	4096
q12xfailover	0
q12xlogintimeout	20 seconds
q12xretrycount	20
q12xsuspendcount	10
q12xdevflag	0
q12xplogiabsentdevice	0 (no PLOGI)
busbusytimeout	60 seconds
displayconfig	1
retry_gnnft	10
recoverytime	10 seconds
failbacktime	5 seconds
bind	0 (by port name)
qfull_retry_count	16
qfull_retry_delay	2
q12xloopupwait	10

Follow the steps below to change default settings, if required.

10. Back up the `fshba.conf` file:

```
cp fshba.conf fshba.conf.bak
```

11. Modify `fshba.conf` using the `vi` editor and then save the file.

12. Start or restart all StorSafe modules, including the HBA module, in order to recognize new settings:

```
vtl start all
```

Data rate The `data_rate-hbaX` parameter identifies the link speed, which is set to auto-negotiate by default for every FC port. You may need to update this setting to match 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 StorSafe appliance, adjust the link speed for all FC ports (StorSafe and/or tape library). Check with your vendor to obtain any recommended FC HBA settings.

1. Identify the HBA port number.

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 StorSafe appliance may not identify the same devices in the same way. StorSafe 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`.

- a. Run the following command to list the adapter numbers:

```
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.

- b. Run the following command to display the WWPN for that adapter:

```
grep BIOSWWPN /proc/scsi/qla2xxx/###
```

- c. Run the following command to determine which physical port belongs to each adapter number in `fshba.conf`:

```
tail -f /var/log/messages
```

- d. Unplug the FC port to see a loop down message like the one below. 100 is the adapter number in this example:

```
<hostname> kernel: scsi(100): LOOP DOWN detected
```

2. For each adapter to be configured, change the value. For example, change the speed to 8 Gb on `HBA0`:

```
data_rate-hba0=4
```

WWNN If you are using redundant node failover and your storage binds to the World Wide Node Name (WWNN) **and** the World Wide Port Name (WWPN), it may be necessary to set the WWNN so that when failover occurs, the takeover server can maintain its storage connection. To set the WWNN, add a “wwnn” line for each Fibre Channel port on each node in your N+1 cluster. All “wwnn” lines can have identical node names for all ports and all nodes. For example, if you have four ports, you might add the following:

```
wwnn0=2100010001000100  
wwnn1=2100010001000100  
wwnn2=2100010001000100  
wwnn3=2100010001000100
```

Configure Fibre Channel 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`. *Loop* mode is recommended.

Load balance paths for downstream storage

For optimal performance, if you have more than one path available for your storage LUNs, you should evenly distribute I/O between all storage LUNs. Refer to [“Set autopathing”](#) for more information.

Configure Fibre Channel on host 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 only support a direct connection with FC 8 Gb. For 16 Gb or higher HBAs, use a FC switch. Another option is to use iSCSI instead for a direct connection.

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.

iSCSI Configuration

The StorSafe server is protocol-independent and supports multiple outbound target protocols, including iSCSI Target Mode.

iSCSI builds on top of the regular SCSI standard by using the IP network as the connection link between various entities involved in a configuration. iSCSI inherits many of the basic concepts of SCSI. For example, just like SCSI, the entity that makes requests is called an *initiator*, while the entity that responds to requests is called a *target*. Only an initiator can make requests to a target; not the other way around. Each entity involved, initiator or target, is uniquely identified.

By default, when a client machine is added as an iSCSI client of a StorSafe 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 StorSafe server, the two entities need to mutually recognize each other. Use an iSCSI initiator on every backup application server that will access the StorSafe server using iSCSI. This will let you add the StorSafe server as a target portal and log the client onto the iSCSI target you create on the StorSafe server.

iSCSI target mode is supported for the following client platforms:

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

iSCSI users

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. Refer to [“Manage iSCSI users”](#) for information about adding iSCSI users.

Windows configuration

Requirements

- A FalconStor StorSafe server with an Ethernet adapter installed.
- You must install an iSCSI initiator 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.mspx>
- 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.
- If Challenge Handshake Authentication Protocol (CHAP) authentication will be used, create at least one iSCSI user. Refer to ["Manage iSCSI users"](#) for more information.

Enable iSCSI

Note: You cannot enable or disable iSCSI if you have already configured StorSafe failover. If you want to change the state of iSCSI, you need to remove your StorSafe failover configuration first.

In order to add a client using the iSCSI protocol, you must enable iSCSI for your StorSafe server.

Refer to ["Enable Fibre Channel and/or iSCSI target mode"](#) 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 StorSafe server. This enables the StorSafe server to see the initiator.

To do this, launch the iSCSI initiator on the client machine and identify your StorSafe server as the target server. Then, enter the IP address or name (if resolvable) of your StorSafe 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 StorSafe server's IP address or name (if resolvable).
Use the default port (3260) and then click OK to add the client.

Add an iSCSI client

Refer to [“Add iSCSI clients”](#) for more information.

Create targets for the iSCSI client to log onto

Refer to [“Create an iSCSI target for an iSCSI client”](#) for more information.

Assign a virtual tape library to the iSCSI target

Refer to [“Assign virtual tape libraries and drives to clients”](#) for more information.

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 FalconStor StorSafe server with an Ethernet adapter installed.
- You must install an iSCSI initiator 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.
- If Challenge Handshake Authentication Protocol (CHAP) authentication will be used, create at least one iSCSI user. Refer to [“Manage iSCSI users”](#) for more information.

Enable iSCSI

Refer to [“Enable Fibre Channel and/or iSCSI target mode”](#) 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 StorSafe server. This enables the StorSafe 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 StorSafe 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 StorSafe server  
OutgoingUsername=CHAP username  
OutgoingPassword=CHAP password
```

Make a note of the CHAP username and password.

3. Start the initiator by typing:

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

Add an iSCSI client

Refer to [“Add iSCSI clients”](#) for more information.

Create targets for the iSCSI client to log onto

Refer to [“Create an iSCSI target for an iSCSI client”](#) for more information.

Assign a virtual tape library to the iSCSI target

Refer to [“Assign virtual tape libraries and drives to clients”](#) for more information.

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 FalconStor StorSafe server with an Ethernet adapter installed.
- You must install an iSCSI initiator 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.
- If Challenge Handshake Authentication Protocol (CHAP) authentication will be used, create at least one iSCSI user. Refer to [“Manage iSCSI users”](#) for more information.
- Confirm the IBM PTF is installed on the client by checking the *Type-Model*, which should display as 298A-001. This indicates that the iSCSI bus IOP resource is operational on the client. If it does not exist, contact IBM to install the required PTF file.

Opt	Description	Type-Model	Status	Resource Name
	Virtual IOP	298A-001	Operational	CMB01

Enable iSCSI

Refer to [“Enable Fibre Channel and/or iSCSI target mode”](#) 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 portal, add the iSCSI client. You will need to enter the client iSCSI initiator IQN. Refer to [Add iSCSI clients](#).

Create iSCSI target

To create an iSCSI target, you can use SQL commands or the *IBM Navigator for i* GUI that is available on IBM i version 7.3 or higher.

To use the *IBM Navigator for i* GUI, select system *Services* → *iSCSI* tab → *Action* → *Create iSCSI Target*. Enter the target IQN, target IP address or hostname, and CHAP parameters, if necessary.

To use SQL commands, refer to the steps below. For command details, refer to <https://www.ibm.com/support/pages/ibm-i-removable-media-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 StorSafe 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');
```

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

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 StorSafe 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 StorSafe server.

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

AIX configuration

AIX supports tape drives via iSCSI but tape libraries are only supported via Fibre Channel. You can use the `atape` driver on AIX to recognize a tape drive that is attached via iSCSI and use the standard operating system commands, such as `tar` or `dd`, to write to the `/dev/rmt0` device. You will have to manually add the tape to the drive in order to write data.

Requirements

- A FalconStor StorSafe server with an Ethernet adapter installed.
- You must install an iSCSI initiator on each AIX 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 Fibre Channel and/or iSCSI target mode”](#) for more information.

Add an iSCSI client

On AIX, run the following command to collect iSCSI initiator information:

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

The resulting output will look like this:

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

In this example, the iSCSI initiator name is `iqn.localhost.hostid.7f000001`.

From the StorSight portal, add an iSCSI client for AIX. Enter the name of your AIX client and add an iSCSI initiator using the collected IQN as the initiator name, such as `iqn.localhost.hostid.7f000001` from our example above. Do not enable *CHAP Authentication*.

Refer to [“Add iSCSI clients”](#) for more information.

On AIX, run the following command to set the correct discovery policy:

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

Create targets for the iSCSI client to log onto

In the StorSafe portal, create an iSCSI target. Leave the *iSCSI Target Name* set to the default, make a note the target name, and confirm that the IP address is correct.

Refer to [“Create an iSCSI target for an iSCSI client”](#) for more information.

On AIX, run the following command to connect the iSCSI initiator to the iSCSI target:

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

Assign a virtual tape drive to the iSCSI target

From the StorSafe portal, select the iSCSI target and assign a tape drive.

Refer to [“Assign virtual tape libraries and drives to clients”](#) for more information.

On AIX, run the following command to rescan and connect to connect to the tape drive, ignoring the error about the changer device:

```
cfgmgr
```

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

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

Run a backup on a AIX client

From the StorSafe portal, move a tape to the first assigned drive. Select the lowest numbered drive, which should be the `/dev/rmt0` device.

On AIX, run the following command to back up some data to the mounted tape, for example:

```
tar cfv /dev/rmt0 ./var/cache/*
```

Run the following command to list files written to the tape:

```
tar tfv /dev/rmt0
```

Failover

This chapter discusses:

- [StorSafe/StorSafe-S failover](#) - Failover of two StorSafe servers
- [SIR cluster failover](#) - Failover of SIR deduplication clusters

StorSafe/StorSafe-S failover

Failover of standalone StorSafe servers provides high availability, protecting you from a wide variety of problems, including storage device path failure and server failure.

There are two possible configurations:

- *Active-Active Failover (or Mutual Failover)* - Both servers are configured to monitor each other. In the event that either server fails, the other takes over. This configuration requires one StorSafe-S server and one FalconStor StorSafe server.
- *Active-Passive Failover* - The primary server is monitored by the partner server. In the event the primary server fails, the partner takes over. This configuration requires two StorSafe-S servers.

Storage device path failure

A storage device path failure can occur due to a cable or switch/router failure.

You can eliminate this potential point of failure by providing a multi-path configuration, using multiple Fibre Channel switches, and/or multiple adapters, and/or storage devices with multiple controllers. In a multi-path configuration, all paths to the storage devices are automatically detected. If one path fails, there is an automatic switch to another path.

Server failure

The failover feature provides high availability by eliminating the downtime that can occur should an StorSafe/StorSafe-S appliance (software or hardware) fail.

In the failover design, a server is configured to monitor another server. In the event that the server being monitored fails to fulfill its responsibilities to the clients it is serving, the monitoring server will seamlessly take over its identity so the clients will be transparently served by the monitoring server.

A unique monitoring system is used to check the health of the servers. This system includes a self-monitor and a partner heartbeat monitor module.

The *self-monitor* is part of all StorSafe/StorSafe-S servers, not just the servers configured for failover, and provides continuous health status monitoring of the server. It is part of the process that provides operational status to any interested and

authorized parties, including the portal and supported network management applications through SNMP. The self-monitor checks the StorSafe/StorSafe-S processes and connectivity to the server's storage devices.

In a failover configuration, the *heartbeat monitor* continuously monitors the partner server through the same network path that the server uses to serve its clients. The health monitoring results are evaluated to determine if any action is needed.

Failover triggers

Failover will occur under any of the following conditions:

- One or more of the critical server modules on the primary server goes down.
- All network-protected interfaces fails.
- A Fibre Channel target port reports a link down while a virtual library/drive is assigned to a client through that port.
- An attempt to access a storage device by the primary fails and the partner server is able to access it.
- There is a server panic on the primary server.
- The primary server is powered off.
- Deduplication node manager is stopped.

StorSafe/StorSafe-S uses its *self-monitor* and *heartbeat monitor* to determine whether failover should be triggered when one of the above conditions is met.

- The self-monitor detects that a critical server process has failed that is determined to be fatal, yet the error did not affect the network interface. In this case, the server will ask its partner server to take over.
- The self-monitor detects a *storage device connectivity failure* but cannot determine if the failure is local or applies to the partner server also. In this case, the server reports the device error condition to its partner server through the heartbeat. The partner server will check to see if it can successfully access the storage. If it can, it attempts to access all devices. If it can successfully access all devices, the partner server initiates a takeover. If it cannot successfully access all devices, no failover occurs.
- Because the heartbeat uses the same network path that the server uses to serve its clients, the partner server knows when the clients cannot access the partner server when network connectivity is incapacitated. This scenario is considered a *catastrophic failure*. In this case the server will initiate a takeover.
- The partner server must be able to be accessed and reset via a remote power control device (IPMI or HP iLO) in order for automatic failover to occur.

Note: If you want to avoid triggering failover when you need to perform storage maintenance, such as updating the firmware, you should suspend failover. You can resume failover when you are done.

Fibre Channel Target failure and Fibre Channel clients

If a Fibre Channel target port link on a server is down, its partner server immediately takes over to serve the host clients without interruption. This is true regardless of the number of target ports on the server. For example, the server can use multiple targets to provide virtual resources to a client. If a target loses connectivity, even if the client has alternate paths to access those devices, failover occurs. However, the default behavior can be modified by Technical Support.

In a pure Fibre Channel environment, network failure will not trigger failover. Fibre Channel switches can demonstrate different behavior in a multiple path configuration. Before using this configuration, you must verify that the configuration can work on your server *without* the FalconStor software.


To verify:

- Use the hardware vendor's utility or Linux's `cat /proc/scsi/scsi` command to see the devices after the driver is loaded.
- Use the hardware vendor's utility or Linux's `hdparm` command to access the devices.
- Unplug the cable from one device and use the utilities listed above to verify that everything is working.
- Repeat the test by reversing which device is unplugged and verify that everything is still working.

Failover terminology

Active-Active Failover	Also known as <i>Mutual Failover</i> , with Active-Active failover, both servers are configured to monitor each other. Traffic intended for the failed node is passed onto an existing node.
Active-Passive Failover	Provides a fully redundant instance of a node, which is only brought online when its associated primary node fails. Active-Passive failover typically requires the most extra hardware.
Failover/ Takeover	Failover/takeover is the process that occurs when the partner server takes over the identity of the primary.
Failover group	For ease of identification, the primary and partner failover servers in a set are grouped together in the portal.
Heartbeat IP address	A static IP address is used by failover servers to monitor each other's health. This is also referred to as the health monitoring IP address. This address remains with the server in the event of failure so that the server's health can be continually monitored. Clients and the portal should not use this address to connect to the server.
Power control	A hardware-level interface that monitors various hardware functions on a server. IPMI and HP iLO are supported interfaces. Power control is used to prevent servers from competing for access to the same storage. Power control is triggered when a partner server fails to communicate with the primary server over both the network and the configuration repository drive.
Primary/partner servers	The primary and partner servers are separate, independent servers that each have their own assigned clients. The terms <i>Primary</i> and <i>Partner</i> are purely from the client's perspective. Each server is <i>primary</i> to its own clients and <i>secondary</i> to the other's clients. Each server normally services its own clients. In the event one server fails, the other will take over and serve the failed server's clients.
Recovery/ Failback	Recovery/Failback is the process that occurs when the partner server releases the identity of the primary to allow the primary to restore its operation. There is a three minute delay during recovery. During this time, no I/O is permitted.
Virtual IP address	An IP address used for portal and client connections. The Virtual IP address is transferred to the partner server during failover so that portal and client connections using this IP address remain active without any interruptions.

Failover behavior



This diagram illustrates a StorSafe/StorSafe-S failover configuration. When server A fails, server B takes over and serves the clients of server A in addition to its own clients.

When failover occurs, there is a delay of several minutes until the partner server takes over. During this time, no I/O is permitted on the failed server. Therefore, all backup jobs (including import/export and replication) running on the primary server will fail.

When failover occurs, the following operations cannot be performed:

- Change the failover configuration on either server
- Perform system maintenance or change server properties on the failed server
- Manage the storage hardware configuration on the failed server
- Deduplication policies cannot be created in an Active-Active failover configuration.
- Storage pool migration cannot take place in an Active-Active failover configuration.

IP connectivity

When a server fails over to its partner, its virtual IP address is transferred to the partner server so that portal and client connections using this virtual IP remain active without any interruptions. The following example illustrates the steps.

Sample environment	• There are two servers, "A" and "B", which act as failover partners. • Each server has two IP network ports. • Each network port has one heartbeat IP address for failover monitoring: AHIP1 and AHIP2 for server A, BHIP1 and BHIP2 for server B. • Each network port has one alias IP that is used as a virtual IP address for portal and client connections: AVIP1 and AVIP2 for server A, BVIP1 and BVIP2 for server B.
Before failover	Four IP addresses show on each server: • Server A: AHIP1 AVIP1 AHIP2 AVIP2 • Server B: BHIP1 BVIP1 BHIP2 BVIP2
When server A fails over to server B	1. Virtual IP addresses are unloaded from server A. The result is: • Server A: AHIP1 AHIP2 • Server B: BHIP1 BVIP1 BHIP2 BVIP2 2. Virtual IP addresses of server A are added to server B. The result is: • Server A: AHIP1 AHIP2 • Server B: BHIP1 BVIP1 AVIP1 BHIP2 BVIP2 AVIP2
When server A recovers from failover	1. New virtual IP addresses are unloaded from server B. The result is: • Server A: AHIP1 AHIP2 • Server B: BHIP1 BVIP1 BHIP2 BVIP2 2. Virtual IP addresses are added back to server A. The result is: • Server A: AHIP1 AVIP1 AHIP2 AVIP2 • Server B: BHIP1 BVIP1 BHIP2 BVIP2

Fibre Channel connectivity

When a server takes over its partner, it impersonates its initiator standby FC ports as target ports of the failed server by spoofing the initiator WWPNs with the partner's target WWPNs. In this way the Fibre Channel client connections remain active without any service interruption. The following example indicates the spoofing steps.

Sample environment

- There are two servers, "A" and "B", which act as failover partners.
- Each server has two target ports: AT1 and AT2 for server A, and BT1 and BT2 for server B.
- Each server has two standby initiator ports: (AS1 and AS2 for server A, and BS1 and BS2 for server B).

Note that AT1, AT2, BT1, BT2, AS1, AS2, BS1, and BS2 represent WWPNs.

Before failover

Four FC WWPN entries from each server appear on the switch:

- Server A: AT1 AT2 AS1 AS2
- Server B: BT1 BT2 BS1 BS2

When server A fails over to server B

1. Target ports are unloaded from server A and their WWPNs disappear from the switch. The result is:
 - Server A: 00 00 AS1 AS2
 - Server B: BT1 BT2 BS1 BS2
2. Target ports are spoofed on server A with temporary WWPN AX1 and AX2 and re-appear on the switch as initiator ports. The result is:
 - Server A: AX1 AX2 AS1 AS2
 - Server B: BT1 BT2 BS1 BS2
3. Standby initiator ports are unloaded from server B and their WWPNs disappear from the switch. The result is:
 - Server A: AX1 AX2 AS1 AS2
 - Server B: BT1 BT2 00 00
4. Standby initiator ports are spoofed on server B with the original WWPNs from server A, get reloaded as target ports, and re-appear on the switch as target ports. The result is:
 - Server A: AX1 AX2 AS1 AS2
 - Server B: BT1 BT2 AT1 AT2

When server A recovers from failover

1. New target ports are unloaded from server B and their WWPNs disappear from the switch. The result is:
 - Server A: AX1 AX2 AS1 AS2
 - Server B: BT1 BT2 00 00
2. New target ports on server B regain their original WWPN ports and appear on the switch as initiator ports. The result is:

- Server A: AX1 AX2 AS1 AS2
 - Server B: BT1 BT2 **BS1 BS2**
3. Temporary target ports are unloaded from server A and their WWPNs disappear from the switch. The result is:
- Server A: 00 00 AS1 AS2
 - Server B: BT1 BT2 BS1 BS2
4. Target ports on server A regain their original WWPN ports and appear on the switch as target ports. The result is:
- Server A: **AT1 AT2** AS1 AS2
 - Server B: BT1 BT2 BS1 BS2

Resume backups after failover/failback

When failover and recovery occur, there is a three minute delay. During this time, no I/O is permitted and all backup jobs (including import/export and replication) will fail.

While failover and failback are transparent for the server, after failover/failback occurs, you may need to take some action in your backup application in order for it to work properly with StorSafe/StorSafe-S. The action you take varies by backup application and operating system. We have described some of the actions we have used. Your environment may differ. Refer to the documentation that came with your backup application for more details.

Dell NetVault™

- After failover/failback occurs, if the Windows Device Manager hangs, you must reboot your NetVault server. Once devices are visible by both the operating system and NetVault, if the NetVault job hangs while waiting to connect to the tape media, reboot your NetVault server.
- Sometimes after failover or failback, a drive may be online but tapes that were previously in the drive or in the slots are missing. This is caused by an interruption of I/O, which causes the software to lose barcodes. To resolve this, the library, drives, and tapes need a complete rescan. Right-click the library (in the NetVault Device Manager) and select *Open Door*. Wait a few moments and select *Close Door*. You should now see all tapes properly labeled and back in their respective drives and/or slots.

ARCserve

After failover/failback occurs, if the operating system and/or ARCserve losses devices, stop the ARCserve tape engine, rescan the operating system, and then restart the ARCserve tape engine. If you still cannot see devices, reboot your ARCserve server.

Once devices are visible by both the operating system and ARCserve, start ARCserve, eject all tapes from their drives and then re-inventory the library.

HP Data Protector

During failover/failback, backup jobs may fail. When this occurs, the tape will be marked as *poor* quality and it will stay in the tape drive. Manually move the tape back to the slot.

IBM® Tivoli® Storage Manager	After failover/failback occurs, reboot the Tivoli Storage Manager machine to get the devices back before submitting any jobs.
EMC NetWorker®	After failover/failback, NetWorker will mark the tape for the current backup job as <i>full</i> and will use a new blank tape to continue the backup job.
Veritas Backup Exec™	If Backup Exec has stalled jobs, reboot the Backup Exec server. If there are no stalled jobs, but drives are down during failover or failback, everything should recover normally.
Veritas NetBackup™	On Windows, if NetBackup has stalled jobs, reboot the NetBackup server. After rebooting, check the NetBackup tape drive/library status and restart NetBackup services, if needed. On Windows, if drives are failed or missing, bringing up the drives should be sufficient. However, if one or more of the drives cannot be brought up, reboot the NetBackup server. On Solaris, jobs usually fail <i>gracefully</i> and subsequent jobs start without problem.

Failover requirements

General requirements (updated October 2024)

- You must have two servers and both must be added to the StorSight server. For Active-Passive failover, you need two StorSafe-S servers. For Active-Active (mutual) failover, you need one StorSafe-S server and one StorSafe server. The primary server must be a StorSafe-S server.
- Both servers must be running identical Linux operating system versions and must have identical storage configurations.
- Both servers must have exactly the same StorSafe options.
- Both servers must have the configuration repository enabled. Once failover is configured, you cannot enable mirroring of the primary configuration repository.
- Both servers must have the same user name and password in order to manage failover.
- The Deduplication Repository must be created on the primary server but cannot be enabled on the partner server.
- The Tape Database must be created on both servers.
- Either both servers need to be associated with the same deduplication repository server or neither of them can be associated with a deduplication repository server. However, if both servers are associated with a cluster and the device and client ID range on the partner server need adjustment to avoid any conflict between servers, the partner servers need to be unassociated from the cluster before ID range adjustment, and then it should be reassociated with the same cluster as the primary server.
- Both servers must have the same power control type, either IPMI or HP iLO with the same username, password, netmask, and gateway. Refer to '[Power control management options](#)' for more information.
- For Active-Active failover, the first time you set up a failover configuration, the StorSafe server must not have any virtual resources (including virtual tapes, drives, or libraries) or clients.
- For Active-Active failover, the servers must have the same strong password configuration.
- If you will be using FalconStor OpenStorage Option (FSOST), you must enable it on both servers before configuring failover. If you are using the FSOST Direct to Tape Option, you must also enable it on both servers before configuring failover. If you need to enable/disable FSOST (with or without Direct to Tape) in an active failover configuration, you must suspend failover first.
- If you have NDMP enabled on one server, you must enable it on the partner server before configuring failover.
- If you will be using Email Alerts, it is very important to enable them on each server so that an alert will be sent if failure occurs on any server. This should be done before configuring failover.

- Any users added to a server that will become part of a failover group should be added to both servers.

Notes:

- We recommend that you finalize all IP addresses before failover is configured.
- Once failover is configured, you cannot change the host name of the failover nodes and you cannot enable or disable iSCSI.

Network requirements

- Both servers must reside on the same network segment, because in the event of a failover, the partner server must be reachable by the clients of the primary server. This network segment must have at least one other device that generates a network ping (such as a router, switch, or server). This allows the partner server to detect the network in the event of a failure.
- If you are protecting multiple IP addresses, each protected IP address must be in its own subnet.
- You need to reserve an IP address for each network adapter in your failover servers. The IP address must be on the same subnet as the server. These IP addresses are used by the servers to monitor each other's health. The health monitoring IP address remains with the server in the event of failure so that the server's health can be continually monitored. After failover, the health monitoring IP address still exists until the network services are restarted. Note that clients and the portal should not use the health monitoring IP address to connect to a server.
- We recommend using a direct network connection between servers to monitor heartbeats in a failover setup to make sure issues in network infrastructure, routers, and switches do not impact failover.
- Some switches do not automatically swap ports during failover. If your switch does not, you may need to configure port swapping.
- In addition to standard IP addresses, each server must have a power control IP address (used by power control software).
- You must use static IP addresses for your failover configuration. We also recommend that the IP addresses of your servers be defined in a DNS server so they can be resolved.

Fibre Channel/iSCSI requirements

- Both servers must have Fibre Channel Target Mode enabled.
- A standby initiator port needs to be available. This port needs to be connected to the same FC fabric as the target port and should not be zoned to anything else. In a multi-ID configuration, you can configure the target port to act as a standby initiator.
- One StorSafe-S target must be zoned with a StorSafe initiator and vice versa. Each pair must be able to communicate with the other.

- The primary and partner servers should use the exact same Target Port ID scheme on the matching WWPNs. We recommend using the same initiator adapter numbers on both sides to connect to the same storage, so the ACSL of all the devices will look identical on both sides.
- If iSCSI protocol is enabled in one server, the other server must also enable it.

Storage requirements

- Virtual storage devices must be accessible by both servers but devices cannot be owned by both servers. This means that storage devices must be attached in a multi-host SCSI configuration or attached on a Fibre loop or switched fabric. In this configuration, both servers can access the same devices at the same time (both read and write).
- Physical drives cannot be shared among StorSafe/StorSafe-S servers. Physical drives should be visible to the other server but not assigned.
- Both servers must be able to access the StorSafe tape database.

Power control management options

Power control management is a hardware-level interface that monitors various hardware functions on a server. Power control management is required in order for failover to occur. IPMI and HP iLO are supported interfaces.

At times, a server may become unresponsive, but, because of network or internal reasons, it may not release its devices or its IP address, thereby preventing failover from occurring. To allow for graceful failover, IPMI or HP iLO will reset the power of the primary server, forcing the release of the server's devices and IP address.

IPMI is included and configured on all FalconStor appliances.

If you are using a third-party appliance, you must determine if power control management is provided by your hardware vendor. To check for IPMI, use the `dmidecode` command and look for the *IPMI Device Information* section. You must also make sure that the *Interface Type* is defined.

If you determine that power control management is provided, you must follow the vendor's instructions to configure it and you must create an administrative user via your configuration tool. The IP address cannot be the virtual IP address that is set for failover.

Backup application server failover configuration

While failover and failback are transparent for the StorSafe/StorSafe-S server, some configuration may be necessary for your backup application server in order for physical devices to be accessed after failover. The configuration varies by backup application and operating system. To ensure that physical devices can be accessed after failover, we recommend you follow the instructions below.

AIX- All backup software

AIX uses the FC port ID to generate its local ID. StorSafe/StorSafe-S failover will change the FC switch port ID.

If you are using AIX 5.3 or above with a Brocade switch model 3900 or above, you need to enable the dynamic tracking option.

Configure failover

You must know the IP address(es) of the primary server (and the partner server if you are configuring mutual failover). You will also need the health monitoring IP address(es). It is a good idea to gather this information and find available IP addresses before you begin the setup.

1. Select *Manage* from the menu bar, select an StorSafe-S server, click the *Settings* tab, and click the *Failover* icon.

The screenshot shows the 'Failover' configuration page within a management console. At the top, there is a navigation bar with tabs: Virtual Devices, Clients, Groups, TimeMarks, Replication, Physical Resources, and Settings. The 'Settings' tab is active. Below the navigation bar, there is a 'Back' button and the title 'Failover'. The main configuration area contains two radio buttons for 'Failover Mode': 'Mutual' (selected) and 'Non-Mutual'. Below these, there are two text input fields: 'Name' with the value 'R720-60287' and 'Partner Name' with a dropdown menu showing 'R720-60431'. At the bottom of the configuration area is a blue button labeled 'Validate Failover'.

The primary server must be a StorSafe-S server.

2. Select the partner server and failover mode.
Select *Mutual* for Active-Active failover, where both servers monitor each other.
If both servers are StorSafe-S, failover supports the *Active-Passive* mode.
If the primary server is an StorSafe-S and the partner is an StorSafe, failover supports *Mutual* mode.
3. Click *Validate Failover*.
4. If prompted that physical devices do not match, click *Rescan Devices*.

5. Click *Configure Failover*.

6. Enter a group name for the failover pair.

For ease of identification, the primary and partner failover servers in a set are grouped together in the portal.

The group name can contain up to 32 characters, including letters, numbers, a dash, or underscore.

7. Select the IP addresses to use for failover.

Virtual IP(s) - Used for client connections and replication. Each IP address results in creation of a virtual Ethernet interface, (i.e., eth0:0, if the network interface eth0 is used). If this server fails over to its partner, the virtual IP address is transferred to the partner server so that portal and client connections using this IP address remain active without any interruption. Because failover can occur at any time, you should use only those IP addresses that are configured as part of the failover configuration to connect to the server.

Heartbeat IP(s) - Used exclusively by failover servers to monitor each other's health. These IP addresses must not be used for any other purpose. Heartbeat IP addresses remain with the server in the event of failure so that the server's health can be continually monitored. Therefore, it is recommended that you use static IP addresses. After failover, the health monitoring IP address still exists until the network services are restarted.

Partner Virtual IP(s) - Used on the partner server for client connections and replication.

Partner Heartbeat IP(s) - Used on partner server to monitor health of the failover server.

Deduplication Association IP - Local IP address of the partner server used to associate the partner server's deduplication repository.

8. On the *Target WWPNS* tab, specify WWPNS for both servers.

You also need to specify the initiator WWPNS on the partner server that will function as a standby in case the target port on the primary server fails. There should not be any devices attached to the initiator.

9. On the *Power Control* tab, select the type of power control management interface your hardware is using and enter power control information.

The subnet mask, gateway, and IP address information is automatically pre-filled.

By default, the IPMI user name and password are set as *admin* and *IPStor101\$*.

The user name and password you specify here will overwrite the existing information.

Note: After specifying IPMI interface information here, do not change them again using the IPMI-equivalent appliance management console.

10. Confirm all information and click *Configure*.

Manage FC switch port swapping

In a failover setup, the client driver recognizes the drives by port ID instead of WWPN. When failover occurs, as the standby adapters on the partner are connecting to different port IDs, the client will no longer be able to see the drives.

Some switches, such as Cisco, automatically swap ports. However, for some switches, such as Brocade, you need to create the following port swapping scripts:

- A pre-takeover script to switch port IDs when failover occurs.
- A pre-recovery script to change the port IDs back before failback occurs.

Notes:

- Tape drive multi-pathing is not supported with port swapping. Even though you can configure StorSafe-S to assign one drive to two paths, failover will not be transparent, meaning that when the backup job fails, the client will need to be reconfigured to use the device from the second path.
- Port swapping scripts on HP and AIX platforms are not supported with Multi-ID HBAs.
- The portswap command will only work within one switch. You can have multiple pairs of target and standby ports located on different switches but the target-standby ports of each pair need to be located on the same switch.

To configure port swapping after failover is configured:

1. Build SSH host-based authentication between StorSafe-S servers and the switch.
 - SSH to the switch using the "root" user account.
 - On both StorSafe-S servers, run: `ssh root@<switch IP address>`. After logging in to the switch, exit from SSH. This step will populate StorSafe-S servers as "known hosts" to the switch.
2. Copy `preRecovery.portswap` on the primary server as the pre-recovery script.


```
#cd $ISHOME/bin
#cp preRecovery preRecovery.bak
#cp preRecovery.portswap preRecovery
```
3. Copy `preTakeOver.portswap` on the partner server as the pre-takeover script.


```
#cd $ISHOME/bin
#cp preTakeOver.portswap preTakeOver.bak
#cp preTakeOver.portswap preTakeOver
```
4. Use the `switchshow` command to collect the information that needs to be replaced in the `portswap.sh` lines of both scripts.
 - `switch_IP` - IP address of the switch.

- `switch_password` - Password of root user for switch.
- `area id` - Switch port area ID. If no area ID, put any one of the ports to be swapped.
- `slot/port` - All ports to be swapped, one pair after another. Slot is optional.


Note: You will need one portswap line for each switch.

5. Run `$ISHOME/bin/preTakeOver` to make sure port IDs are swapped correctly.
6. Run `$ISHOME/bin/preRecovery` to make sure port IDs are swapped back.

Check failover status

Both servers in a failover configuration will display a *Failover* icon. The icon will change color depending upon the status (green=normal, orange=suspended, red=taken over, blue=took over). Note that the status will only be updated when the page is refreshed.

When you mouse over the icon, the status of the failover pair is displayed.



Failover State can have the following values:

- *Normal* - The server has ownership of its devices only; the server has not failed over to its partner or taken over its partner's devices.
- *Took over* - The server has taken over its partner's devices.
- *Taken over* - The server failed over to its partner and its partner currently has ownership of its devices.
- *Suspended* - Failover has been suspended. This server is not monitoring its partner server.

You can also see the status and configuration on the *Failover Settings* page. There you will also see the failover mode (mutual or non-mutual) and the self-monitor status.

Failover events are written to the primary server's Event Log, so you can check there for status and operational information, as well as any errors. When a failover occurs, the Event Log for the server that failed is available to be downloaded from the portal.

You can also view the failover status by running the `smS -v` command on the server.

Change your failover configuration

If you make hardware changes on a server in a failover setup, the partner server will not be aware of these changes until you rescan. This will synchronize the hardware configuration.

If you make a change to a physical device (such as if you add a network card that will be used for failover), you will need to disable and then reconfigure failover.

Once failover is configured, you cannot change the host name or IP addresses of the failover nodes.

Check failover settings

Select *Manage* from the menu bar, select a failover server, click the *Settings* tab, and click the *Failover* icon to see your failover configuration settings.

Change failover options

Select *Manage* from the menu bar, select a failover server, click the *Settings* tab, click the *Failover* icon, click *Manage Failover*, and select *Update* to change the intervals (heartbeat, self-checking, and auto recovery).

Self-checking Interval - The time interval in seconds for self-health checking on an active server in the range of 1 to 999 seconds; default 2 seconds.

Heartbeat Interval - The time interval in seconds for checking partner's health in the range of 5 to 999 seconds; default 5 seconds.

Auto Recovery - An option to allow the partner server to fail back after it has recovered from failover. A wait time can be specified before giving control to the partner server in the range of 0 to 999 seconds. This option is only applicable to the passive server. It is recommended to check the root cause and health of the server before failback is triggered, otherwise, the error could occur again and cause another server failover.

Note: We recommend keeping the *Self-checking Interval* and *Heartbeat Interval* set to the default values. Changing the values can result in a significantly longer failover and recovery process.

Change the power control password

If you need to change the password that is used for power control, select *Manage* from the menu bar, select a failover server, click the *Settings* tab, click the *Failover* icon, click *Manage Failover*, and select *Update Power Control*.

Start/stop failover or recovery

Start takeover

Select *Manage* from the menu bar, select a failover server, click the *Settings* tab, click the *Failover* icon, click *Manage Failover*, and select *Start Takeover* to take over the other server. You may want to do this if you are taking a server offline, such as when you will be performing maintenance on it.

Stop takeover

Select *Manage* from the menu bar, select a failover server, click the *Settings* tab, click the *Failover* icon, click *Manage Failover*, and select *Stop Takeover* if you want to manually initiate recovery to your primary server.

Suspend/resume failover

Select *Manage* from the menu bar, select a failover server, click the *Settings* tab, click the *Failover* icon, click *Manage Failover*, and select *Suspend* to stop monitoring your servers. The server will NOT take over the primary server's devices even it detects an abnormal condition from the primary server.

Select *Manage* from the menu bar, select a failover server, click the *Settings* tab, click the *Failover* icon, click *Manage Failover*, and select *Resume* to restart the monitoring.

Note: While failover is suspended, changes to the configuration of a library, drive, or client on either failover server will not be synchronized with the partner. Therefore, it is important to not change your library, drive, or client configuration while failover is suspended.

Recover from failover

To recover a server after failover, perform the following steps:

1. Identify and fix the condition that caused the failure.
2. Reboot the server.
3. Run the `sms` command to confirm the server is in a *Ready* state.
4. Check storage and network connectivity.

During failback, the server must receive confirmation from its partner in order to recover its role as the primary server.

If there is a communication problem between servers, the server does not receive confirmation and remains in a *Ready* state but does not recover its role as the primary server.

5. Select *Manage* from the menu bar, select a failover server, click the *Settings* tab, click the *Failover* icon, click *Manage Failover*, and select *Stop Takeover* to trigger recovery.

Remove a failover configuration

Select *Manage* from the menu bar, select a failover server, click the *Settings* tab, click the *Failover* icon, and click *Remove* to eliminate the failover configuration and return the servers to independent servers.

By default, after removing the failover configuration, the health monitoring (heartbeat) IP address will be removed from the server and the original server IP address will be restored. You need to restart network services on the server to clear the heartbeat IP address.

If you want to preserve current network settings, where the heartbeat IP address becomes the server IP address, select the network interfaces to maintain on each server.


If the partner server cannot be reached, you can select to force failover removal.

SIR cluster failover

Deduplication cluster failover provides high availability for your FalconStor StorSafe Single Instance Repository (SIR) servers.

In a deduplication cluster failover configuration, a single standby server monitors the health of all servers in a deduplication cluster. This is referred to as “N+1” failover.

If an active server fails, the standby server sends a command to shut down the failed server. The standby server then assumes the identity of the failed server and the cluster continues to function normally.



When a failed server is ultimately repaired and is turned back on, it assumes the identity of the standby server and begins checking the heartbeats of active servers.

Failover terminology

Active server	An active server is a functioning SIR server. The health of each active server in a cluster is monitored by the standby server.
Standby server	The job of the standby server is to monitor the health of the active servers in a cluster. If an active server fails to respond to the standby server, failover occurs.
Heartbeat	The IP address that identifies each server and is used to verify the health of each server.
Failover	(Also known as <i>Takeover</i>) Failover is the process that occurs when the standby server takes over the identity of an active server.

Failover triggers

Failover will occur under the following conditions when there is a heartbeat failure:

- Server lockup of active node
- Network failure of active node
- Deduplication node manager is stopped

Note: Automatic failover will not occur if IPMI or HP iLO is not accessible.

If the standby server cannot access an active node and it cannot determine the node's power status, it will not automatically take over the failed node. If failover is desired, it can be manually initiated with the `sirtakeover.sh` command. However, in order to prevent potential data integrity problems, you **MUST** confirm that the failed node has been shut down before forcing a takeover.

If you want to avoid triggering failover when you need to perform storage maintenance, such as updating the firmware, you should suspend failover. You can resume failover when you are done.

Failover behavior

During failover, any currently running deduplication jobs will fail. The jobs will restart at their next scheduled time.

If data is being restored from a deduplicated tape at the time failover occurs, the restore job will fail and will need to be restarted afterward.

We recommend running a file system consistency check on any server that has experienced a failure.

Server behavior during and after failover

When a failover occurs, the failed server is powered off by the standby server. The standby server starts the takeover process and assumes the identity (hostname, primary IP addresses, FC targets, and FC initiators) of the failed server.

Once the standby server completes the takeover process, it makes one attempt to turn the failed server back on. If the failed server is successfully powered back on, services will start automatically and the failed server will assume the name and identity of the standby server. If the failed server is not successfully powered on, you must repair the server and turn it back on.

When the failed server is ultimately repaired and is turned back on, it assumes the name and identity of the standby server. All server IP addresses remain the same (excluding heartbeat IP addresses). The new standby assumes the responsibility of checking heartbeats of active servers. The *General* tab for the SIR server shows the current hostname and the original hostname after failover occurred.

If it is necessary to replace the failed node with new hardware, contact Technical Support.

Failover requirements

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

- You must have a standby SIR server. The standby server must have: the same number of FC ports as all active servers, at least one Fibre Channel initiator port.
- All of the SIR servers in the cluster, including the standby server, must be identical (same make, model, and BIOS version) with FC target mode enabled and must all use the same root password. The standby server does not need a configuration repository.
- The QLogic BIOS option that allows the system to boot from a FC LUN must be disabled. This is the default setting for QLogic cards shipped by FalconStor.
- In addition to a standard IP address, each server must have a heartbeat IP address (Identifies each server and is used to verify the health of each server) and a power control IP address (used by power control software).
- All IP addresses (heartbeat and power control) on the active server and the standby server need to be on the same subnet on the same NIC interface.
- If you are protecting multiple IP addresses, each protected IP address must be in its own subnet.
- The role of each port must be the same on each server. For example, if port 1 connects to storage on the active server, port 1 must connect to storage on the standby server.
- Each server must use the exact same Target Port ID scheme on the matching WWPNs. We recommend using the same initiator adapter numbers on both sides to connect to the same storage, so the ACSL of all the devices will look identical on both sides.
- All servers must be configured with a power control management interface on the motherboard. IPMI (Intelligent Platform Management Interface) and HP iLO are supported interfaces.
- If your storage binds to the World Wide Node Name (WWNN) **and** the World Wide Port Name (WWPN), it may be necessary to set the WWNN so that when failover occurs, the takeover server can maintain its storage connection. This is done in the *fshba.conf* file by adding a “wwnn” line for each Fibre Channel port on each node in your N+1 cluster. All “wwnn” lines can have identical node names for all ports and all nodes. For example, if you have four ports, you might add the following:

```
wwnn0=2100010001000100  
wwnn1=2100010001000100  
wwnn2=2100010001000100  
wwnn3=2100010001000100
```


- If you want Email Alerts enabled on the standby server, you must enable it on the standby server before failover is configured.

Notes:

- Once failover is configured, you cannot change the host name or IP addresses of the cluster nodes.
- Once failover is configured, you cannot change the roles of your HBA ports.

Configure failover

Configure failover by adding a standby node to take over when a deduplication cluster node fails. After adding the standby node, services will be restarted on the node. To configure failover:

1. Select *Manage* from the menu bar, select an SIR server, click the *Settings* tab, and click the *Failover* icon.
2. Click *Configure Failover*.

Configure Failover

General Power Control

Standby Node*

Heartbeat NIC*

Standby Node Heartbeat IP*

Standby Node Power Control IP*

Node 1 H13-174 (172.22.13.174)


Power Control IP*

	Interface	IP	Netmask	Heartbeat IP
☐	eth0	172.22.13.174	255.255.254.0	172.22.

Configure Cancel

3. Select the standby node.
4. Select the network interface you want the standby server to use to monitor the heartbeat and trigger failover in the event of a connection loss.
5. Enter the heartbeat IP address and power control IP address for the standby node.
6. Enter the power control IP address for the primary node.
7. Select the network interface and enter the heartbeat IP address for the primary node.

8. Click the *Power Control* tab.



The screenshot shows a 'Configure Failover' dialog box with a close button (X) in the top right corner. It has two tabs: 'General' and 'Power Control'. The 'Power Control' tab is selected. Inside this tab, there is a 'Power Control Type' dropdown menu set to 'IPMI'. Below this are several input fields: 'Power Control Username*' (empty), 'Power Control Password*' (empty), 'Confirm Password' (empty), 'Power Control Netmask*' (containing '255.255.254.0'), and 'Power Control Gateway*' (containing '172.22.12.1'). At the bottom right, there are two buttons: 'Configure' (highlighted in blue) and 'Cancel'.

9. Specify the IPMI user name and password.

The user name and password you specify here will overwrite the existing information.

10. Click *Configure* when you are done.

Start nodes in a deduplication cluster


Once failover has been configured, you must make sure the standby server meets one of the following conditions before starting any active node in a cluster:

- The standby server is powered on with services started.
- The standby server is powered off but power control is accessible. If power control cannot be accessed, contact FalconStor Technical Support.

Check failover status

A cluster with a failover configuration will display a *Failover* icon. The icon will change color depending upon the status (green=normal, orange=suspended, red=taken over, blue=took over). Note that the status will only be updated when the page is refreshed.

When you mouse over the icon, the status of the failover pair is displayed.



Failover State can have the following values:

- *Ready* - All servers are operating normally.
- *Suspended* - Failover has been suspended. There is no server monitoring. the standby server will NOT assume the identity of a failed server

You can also see the status and configuration on the *Failover Settings* page. There you will also see the failover configuration.

Failover events are written to the primary server's Event Log, so you can check there for status and operational information, as well as any errors. When a failover occurs, the Event Log for the server that failed is available to be downloaded from the portal.

You can also view the failover status by running the `sms -v` command on the server.

Change the power control password

Changing the power control password changes it for all servers, including the standby server.

Select *Manage* from the menu bar, select an SIR server, click the *Settings* tab, click the *Failover* icon, click the *Manage* icon, and select *Change Power Control Password*.

Suspend/resume failover

When failover is suspended, the standby server will NOT assume the identity of a failed server. You need to do this when you patch the server or if you have to replace storage. All cluster nodes must be available and with a ready status in order to suspend failover. All cluster nodes must be available and with a suspended status in order to resume failover.

To suspend failover, select *Manage* from the menu bar, select an SIR server, click the *Settings* tab, click the *Failover* icon, click the *Manage* icon, and select *Suspend*.

To resume failover, select *Manage* from the menu bar, select an SIR server, click the *Settings* tab, click the *Failover* icon, click the *Manage* icon, and select *Resume*.

Remove a failover configuration

After a failover configuration is removed, the standby server becomes a regular SIR server without a deduplication repository.


All cluster nodes must be available in order to remove a failover configuration.

Select *Manage* from the menu bar, select an SIR server, click the *Settings* tab, click the *Failover* icon, click the *Remove* icon.

Deduplication

StorSafe deduplication eliminates redundant data used by virtual tapes and NAS. StorSafe deduplication policies specify which virtual tapes need to have deduplication and when deduplication should occur. They may also include replication of virtual tapes.

How virtual tape deduplication works



During deduplication, an intelligent, content-aware tape scanner 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 StorSafe. 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.

In a cluster configuration, deduplication can operate as a single node (one active SIR server) or as a multi-node cluster (two or four active SIR servers).

Deduplication provides replication capability. If replication is configured, deduplication replicates its repository and metadata, effectively performing global data deduplication. Any data duplicated across remote sites is deduplicated at the central site, enabling only globally unique data to be replicated to the disaster recovery site.


How NAS deduplication works

StorSafe uses standard network protocols such as Server Message Block (SMB) or Network File System (NFS) to present a simple, network-based file share as the target for backed-up data. Each NAS file share holds incoming data, acting as a "disk" for disk-to-disk (D2D) backup.

NAS deduplication processes and deduplicates data during backup; there is no need to create a deduplication policy for NAS shares. By default, all NAS files are automatically included for deduplication except for files smaller than 8 KB. Refer to the ["NAS"](#) chapter for more information.

During deduplication, the system analyzes blocks of data and determines whether the data is unique or has already been copied to the repository (virtualized disks that hold deduplicated data). The process then passes only single instances of unique data to the repository and replaces each deduplicated file with a small file (called a *stub* file), that points to the repository and is used to retrieve stored data.

Even though the user interface is file-based, deduplication is performed at the block level, not at a file level. Block-level deduplication examines small sub-blocks, making it far more effective at reducing storage consumption than file-based deduplication.



Because network-based file shares are used for backed-up data, restoring data is fast and easy. The administrator has direct access to all files without the need for a restore job. Even after deduplication occurs, pointers (*stub* files) on the share point to the full file in the repository. Restoring data copied by archiving software is as simple as copying the necessary files from a share back to the appropriate location.

Deduplication methods - at a glance

Typically, backup jobs 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: Inline deduplication, Turbo deduplication, and Post-processing.

NAS file backups use inline deduplication, which processes and deduplicates data during backup.

For tape backups, deduplication can be performed at any time, while backup is running or after each backup job completes, and can be scheduled or can be run on demand. This is set as part of each deduplication policy.

Turbo deduplication pre-processes data during backup but completes the deduplication process at a later time. Before deduplication occurs, the pre-processed, hashed data is stored in a temporary area (which requires approximately 1%-2% of the size of the backed-up data). When deduplication is triggered, the system processes the hashed data stored in the temporary area and the unique data blocks are stored in the repository. After a tape has finished deduplication, the space used by the temporary area is released and the original virtual tape is replaced with a VIT. Turbo deduplication minimizes disk contention by reducing I/O to the disk because data is pre-processed; less has to be read by the deduplication process.

Inline deduplication processes and deduplicates data during backup and then stores the unique data blocks in the repository.

The following table compares the different deduplication methods:

	Inline Deduplication	Turbo Deduplication	Post-Processing (Without Inline or Turbo)
Performance	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.	Pre-processing can significantly improve deduplication performance.	
CPU Processing Power	Requires the most CPU power in the StorSafe server so that incoming backup data can be processed (12 cores of processing power or above recommended).	The more CPUs in the StorSafe server, the less impact to backup performance. With 12 cores of processing power or above, impact will be minimal.	No special requirements.

Deduplication Policies page



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

Only a Super Administrator can access the *Deduplication Policies* page.

Policy	ID	Tapes	Status	Suspended	Last Run	Next Run	Trigger	Replication
Default_Policy	1	11	Idle	☐			Inline	None
H13-62_Cascaded_Policy_test	8	0	Idle	☐			Manual	Cascaded
HCP	4	30	Idle	☐			Tape Eject	None
Spectra-T200	12	400	Idle	☐			Inline	None
STK	3	4	Idle	☐			Inline	None

Depending upon the columns that have been selected for display, the properties for each deduplication policy are visible at a glance:

- Policy name
- Policy ID
- Number of tapes
- Status - Idle, Running, Suspended, Stopping, Stopped, Preparing, Queued
- Indication if the policy is suspended
- Last time the policy was run
- Next time the policy will run
- Deduplication trigger
- Replication mode
- Replica server
- Indication if replication is suspended
- Deduplication cluster to which this policy belongs

Deduplication Policies page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Create	Create a deduplication policy.
	Update	Update deduplication policy name, priority, retry parameters, trigger, Turbo Deduplication usage, replication settings.
	Delete	Delete the selected deduplication policy.
	Manage	Start/stop deduplication policies, suspend/resume deduplication policies, suspend/resume replication in deduplication policies


Create a deduplication policy

Deduplication policies specify which virtual tapes need to have deduplication and when deduplication should occur. In order to configure replication of the tapes in a deduplication policy, the replica server(s) must be added to the StorSight server. You must have also added a target deduplication replication server (from the *Replication Targets* page).

You can have a maximum of 64 deduplication policies.

To create a deduplication policy:

1. Select *Manage* from the menu bar, select a StorSafe server, and click the *Deduplication* tab.
2. Click the “+” (*Create*) icon on the *Policies* page.
3. On the *General* tab, specify the policy name, priority level, and retry parameters.




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

Priority - The order in which deduplication jobs will be launched for all tapes in a policy, *Low*, *Medium*, or *High*.

If you do not specify a priority, the order is determined by the last time the tape was written. Older tapes have higher priority but still less than tapes with a priority.

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.

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 on the Deduplication Jobs page and select *Run Next*, the *Run Next* priority is higher than any existing policy priority that is set.

Retry Count - The number of times to retry a failed deduplication, VIT replication, or resolver job. The range is 0-99999 retries and the default is 48.

Retry Interval - The amount of time between retries. The range is 1-60 minutes and the default is 30 minutes.

The retry policy settings apply to all tapes, regardless of whether or not a priority is set.

4. On the *Policy* tab, specify when deduplication should occur and whether Turbo Deduplication should be used.

The screenshot shows a 'Create Deduplication Policy' dialog box with three tabs: 'General', 'Policy', and 'Replication'. The 'Policy' tab is active. It contains a 'Trigger' dropdown menu currently set to 'Manual', and a 'Turbo Deduplication' toggle switch which is turned off. Both the dropdown and the toggle have a blue question mark icon to their right. At the bottom right of the dialog are 'Create' and 'Cancel' buttons.

Trigger - Deduplication policy trigger:

- **Manual** - Deduplication must be manually initiated.
- **Inline** - Data deduplication occurs while backup is in progress. Indicate if the deduplication method should perform post-processing if inline deduplication is not possible at the start of processing (or at the first write from the backup application). If the post-processing option is not selected and inline is not possible, the current backup job will fail. If the post-processing option is selected, backup data will be written to StorSafe 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 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.
- **Schedule** - Deduplication will occur based on the schedule specified, which can be *Hourly*, *Daily*, or *Weekly*.
- **Tape Eject** - Deduplication starts when a virtual tape is ejected from a drive to a slot, thereby running concurrently with backup. The **Mode** can be set to:
 - **Minimum New Data** - Start deduplication when the amount of new data on the tape reaches the specified threshold. The range is 0 – 2,147,483,647 MB and the default 1 MB.
 - **Tape Full** - Start deduplication when the tape is full.

Turbo Deduplication - Pre-process tapes during backup. *Turbo Deduplication* is available for all triggers except *Inline*. Refer to [Deduplication methods - at a glance](#) for more information about Turbo deduplication.

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

The screenshot shows the 'Create Deduplication Policy' dialog box with the 'Replication' tab selected. The 'Mode' dropdown is set to 'Cascaded'. The 'Source IP Address' dropdown shows '172.22.13.82' and there is an 'Add NAT Address' button. The 'Target Server' dropdown is empty, and the 'Second Target Server' dropdown is also empty. There are 'Create' and 'Cancel' buttons at the bottom right.

Single mode replicates tape data to a single target server.

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

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

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

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

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

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

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

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

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

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

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

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

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

Retry Interval - The amount of time between retries. The range is 0-3600 minutes and the default is 1 minute.

7. Click *Create* when done.

Tapes can be added to the deduplication policy from the *Tapes* tab under *Deduplication Policies*. Refer to [“Add tapes to a deduplication policy”](#) for more information.

Update deduplication policies

Select *Manage* from the menu bar, select a StorSafe server, click the *Deduplication* tab, select a policy, and click the *Update* icon.

You can update a deduplication policy according to the following guidelines:

- You can change deduplication policy priority, retry parameters, trigger, enable/disable Turbo deduplication, enable/disable replication, and change replication properties.
- You cannot change the name of a deduplication policy.
- You can enable or disable replication but you cannot change the replication mode.
- If you disable cascaded replication in a policy on the source server, replication is automatically disabled in the policy on Target 1. You cannot disable cascaded replication in the policy on Target 1.
- You can modify a policy with cascaded replication on the source server as well as on Target 1. However, if you modify the policy on the source server, updates are made in the policy on Target 1 automatically. But, if you modify the policy on Target 1, updates are made only in the policy on Target 1.
- You cannot change the replication target server(s).
- You should not use the StorSight portal to modify deduplication policies that were created for the FalconStor OpenStorage Option (FSOST) via the FSOST `cfglsu` command.

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.

Manage deduplication policies

You can start, stop, suspend, and resume deduplication policies. You can also suspend and resume replication to a target if replication is enabled in a policy.

Select *Manage* from the menu bar, select a StorSafe server, click the *Deduplication* tab, select a policy, click the *Manage* icon, and select the action you want to perform.

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

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

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

Delete deduplication policies

Select *Manage* from the menu bar, select a StorSafe server, click the *Deduplication* tab, select a policy, click the *Delete* icon.

If a 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.

Deduplication Tapes page



Select *Manage* from the menu bar, click the *Deduplication* tab, and click the *Tapes* tab. Select a deduplication policy from the *Deduplication Policy* drop-down list. This page lists virtual tapes in the selected deduplication policy on the selected StorSafe server.

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

Only a Super Administrator can access the *Tapes* page.

Barcode	Tape	ID	Location	Deduplication Status	Deduplication State	Last Deduplication	Dedupe Ratio
00390005	VirtualTape-00018	10000018	Library ID 57 Slot 5	Successful	Writing to VIT Finished	2019-05-08 16:27:38	> 10000:1
00390006	VirtualTape-00019	10000019	Library ID 57 Slot 6	Successful	Writing to VIT Finished	2019-06-05 00:34:52	> 10000:1
00390007	VirtualTape-00020	10000020	Library ID 57 Slot 7	Successful	Writing to VIT Finished	2019-06-02 13:25:08	> 10000:1
00390008	VirtualTape-00021	10000021	Library ID 57 Drive 3	Successful	Writing to VIT Finished	2019-06-03 22:38:21	> 10000:1
00390009	VirtualTape-00022	10000022	Library ID 57 Slot 9	Successful	Writing to VIT Finished	2019-06-03 01:00:39	2236:1
0039000A	VirtualTape-00023	10000023	Library ID 57 Slot 10	Successful	Writing to VIT Finished	2019-05-14 18:24:45	147:1
0039000B	VirtualTape-00024	10000024	Library ID 57 Drive 2	Pending	Writing to VIT Finished	2019-05-04 22:51:47	
0039000C	VirtualTape-00025	10000025	Library ID 57 Slot 12	Successful	Writing to VIT Finished	2019-05-21 17:03:37	> 10000:1
0039000D	VirtualTape-00026	10000026	Library ID 57 Slot 13	Successful	Writing to VIT Finished	2019-05-30 06:48:09	> 10000:1
0039000E	VirtualTape-00027	10000027	Library ID 57 Slot 14	Successful	Writing to VIT Finished	2019-05-31 05:11:47	> 10000:1
0039000F	VirtualTape-00028	10000028	Library ID 57 Slot 15	Successful	Writing to VIT Finished	2019-06-02 01:33:10	30:1
0039000G	VirtualTape-00029	10000029	Library ID 57 Drive 4	Pending	Writing to VIT Finished	2019-05-21 22:45:05	> 10000:1

Depending upon the columns that have been selected for display, the properties for each tape are visible at a glance:

- Barcode
- Tape
- ID
- Location
- Deduplication status - Successful (pure VIT), Pending (or in-progress, not a pure VIT), Failed (not a pure VIT)
- Deduplication state - Completed, Aborted, Empty Tape, Failed, Paused, Preparing, Queued, Running, Stale, Stopped Writing to VIT (writing back the deduplicated data to VIT), Writing to VIT Finished, Writing to VIT Failed
- Last deduplication time

- Deduplication ratio - 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.
- Maximum uncompressed storage capacity (determined when the tape was created)
- Allocated tape size
- Written - Amount of data (before compression) that is written to the tape by backup applications. This amount can be greater than the tape size if the data is compressed.
- Stored - Amount of data (before compression) stored in the deduplication repository.
- Unique data - Actual physical storage used to store tape data after deduplication and compression.
- New data - Amount of data (before compression) that has not yet been deduplicated, including newly appended data to a tape.
- Replication status - Indicated on the source server for a tape in a deduplication policy, on the target server for a deduplicated LVIT. For cascaded or parallel replication, the status is indicated for both first and second replication target servers.
 - Completed
 - Pending - Or in-progress replication on the source server
 - Failed - Can be due to replication failure, tape resolution failure, or in-progress tape resolution
 - Completed Pending - For cascaded or parallel, successful replication on server 1 and pending replication on server 2
 - Completed Failed - For cascaded or parallel, successful replication on server 1 and failed replication on server 2
 - Pending Completed - For cascaded or parallel, pending replication on server 1 and successful replication on server 2
 - Pending Failed - For cascaded or parallel, pending replication on server 1 and failed replication on server 2
 - Failed Completed - For cascaded or parallel, failed replication on server 1 and successful replication on server 2
 - Failed Pending - For cascaded or parallel, failed replication on server 1 and pending replication on server 2
- Replicated - Amount of unique data replicated
- Last replication time
- Last replication status
- Replica server
- Resolution - Status on target server for replication of deduplicated tapes. When resolved, the target server automatically creates an LVIT and puts it in the target server's vault or in a virtual tape library, depending upon how the deduplication policy was configured.

Deduplication Tapes page icon

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Schedule report	Schedule a Deduplication Policy Tapes report.
	Manage	Add, remove, and run tapes.

Add tapes to a deduplication policy

To add tapes to a deduplication policy:

1. Select *Manage* from the menu bar, select a StorSafe server, click the *Deduplication* tab, and click the *Tapes* tab.
2. Select a deduplication policy.
3. Click the Manage icon and select *Add Tapes*.
4. Select a virtual tape library.
5. Select tapes and click *Add*.

Remove tapes from a deduplication policy

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).

To remove tapes from a deduplication policy:

1. Select *Manage* from the menu bar, select a StorSafe server, click the *Deduplication* tab, and click the *Tapes* tab.
2. Select a deduplication policy.
3. Select one or more tapes.
4. Click the Manage icon and select *Remove Tapes*.
5. Confirm and click *Remove*.

Run tapes in a deduplication policy

To deduplicate/replicate one or more individual tapes in a deduplication policy:

1. Select *Manage* from the menu bar, select a StorSafe server, click the *Deduplication* tab, and click the *Tapes* tab.
2. Select a deduplication policy.
3. Select one or more tapes.
4. Click the Manage icon and select *Run Tapes*.
5. Confirm and click *Run*.

Schedule Deduplication Policy Tapes report

The Deduplication Policy Tapes report can be scheduled to be emailed to you on a regular basis. Refer to [“Schedule a report”](#) for more information.

Deduplication Activity page



Select *Manage* from the menu bar, click the *Deduplication* tab, and click the *Activity* tab. Select a deduplication policy from the *Deduplication Policy* drop-down list. This page is divided into two sections and displays information about currently running deduplication policies and replication jobs on the selected StorSafe server.

Only a Super Administrator can access the *Activity* page.

The screenshot shows the FalconStor StorSafe web interface. The top navigation bar includes 'Dashboard', 'Monitor', 'Manage' (selected), 'Analytics', and 'Administration'. The left sidebar shows a list of servers: FBD FBD-61, FBD H13-62, FBD H13-97 (selected), FBD tssir03, and FBD VA22-190. The main content area is titled 'Deduplication Policies' and has tabs for 'Policies', 'Tapes', 'Activity' (selected), 'Run History', 'Tape History', and 'Replication Targets'. Under the 'Activity' tab, there's a 'Deduplication Policy' dropdown set to 'FS 1877' and a 'Tape Scan' section. A table displays the following data:

Barcode	Tape	ID	Serial Number	Parser	Data Size	Scanned	Throughput	Blocks	Status	Progress
00390002		10000015	20GMB0030A	netbackup2	286.93 GB	7.28 GB	275 B/s	119193	Paused	8%

Below this is a 'Tape Replication' section with a similar table structure, including columns for 'Barcode', 'Tape', 'ID', 'Phase', 'Data Size', 'Transmitted', and 'Th'.

Depending upon the columns that have been selected for display, the top section shows the properties for each tape currently being deduplicated:

- Barcode
- Tape
- ID
- Serial number
- Parser
- Size of data
- Amount of data scanned
- Throughput
- Number of blocks
- Status - Queued, Running, Completed, Failed, Paused, Preparing, Scan Paused, Writing VIT (writing the deduplicated data back to the VIT), Writing VIT Paused, Writing VIT Finished, Writing VIT Failed, Writing VIT Preparing, Writing VIT Stopped, Stopped
- Progress - Click *Refresh* to update the data.

While replication is occurring, you will see status displays related to the Index and unique replication phases on the bottom part of the screen.

- Barcode
- Tape
- ID
- Replication phase - Index (transmitting index data) or Unique (transmitting unique deduplication data)
- Data size - Amount of data to be replicated
- Transmitted - Amount of transmitted data
- Throughput
- Deduplication ratio on the replication target server after the data is replicated and resolved
- Progress - Click *Refresh* to update the data.
- Remaining time

Schedule Deduplication Activity reports

The Deduplication Activity Scan report and the Deduplication Activity Replication can be scheduled to be emailed to you on a regular basis. Refer to [“Schedule a report”](#) for more information.

Deduplication Run History page

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

Only a Super Administrator can access the **Run History** page.

The screenshot shows the FalconStor StorSafe interface. The top navigation bar includes 'Dashboard', 'Monitor', 'Manage', 'Analytics', and 'Administration'. The 'Manage' tab is selected, and the 'Deduplication' sub-tab is active. The 'Run History' tab is selected within the 'Deduplication' section. The 'Deduplication Policy' is set to 'FS 1877'. The table displays the following columns: Start Time, Trigger, Tapes, Data to Dedupe, Unique Data, Dedupe Ratio, Dedupe Durat... (Deduplication Duration in Seconds), Data to Replic..., Unique Data to Replicate, and Dedupe Rat on Relica. The table shows 12 rows of deduplication jobs, all triggered by 'Tape Eject'. The 'Showing 3748 entries' message is visible at the bottom of the table.

Start Time	Trigger	Tapes	Data to Dedupe	Unique Data	Dedupe Ratio	Dedupe Durat... Seconds	Data to Replic...	Unique Data to Replicate	Dedupe Rat on Relica
2019-06-24 00:03:39	Tape Eject	1	307.63 GB	2 MB	> 10000:1	1341	0 B	0 B	
2019-06-23 14:57:06	Tape Eject	1	1023.83 GB	1 MB	> 10000:1	5040	0 B	0 B	
2019-06-22 09:43:48	Tape Eject	1	193.74 GB	196 MB	1012:1	962	0 B	0 B	
2019-06-22 06:06:27	Tape Eject	1	1023.83 GB	1 MB	> 10000:1	5017	0 B	0 B	
2019-06-21 12:18:25	Tape Eject	1	1023.83 GB	2 MB	> 10000:1	3969	0 B	0 B	
2019-06-20 17:10:03	Tape Eject	1	1023.83 GB	10 MB	> 10000:1	6373	0 B	0 B	
2019-06-20 07:04:09	Tape Eject	1	1023.83 GB	2 MB	> 10000:1	5088	0 B	0 B	
2019-06-19 16:24:36	Tape Eject	1	716.2 GB	2 MB	> 10000:1	3383	0 B	0 B	
2019-06-18 23:22:33	Tape Eject	1	214.33 GB	1 MB	> 10000:1	836	0 B	0 B	
2019-06-18 19:29:19	Tape Eject	1	1023.83 GB	4 MB	> 10000:1	8447	0 B	0 B	
2019-06-18 03:02:34	Tape Eject	1	1023.83 GB	1 MB	> 10000:1	6518	0 B	0 B	
2019-06-17 22:00:44	Tape Eject	1	626.55 GB	0 B	> 10000:1	2549	0 B	0 B	
2019-06-17 11:44:40	Tape Eject	1	1023.83 GB	110 MB	9531:1	5532	0 B	0 B	

Depending upon the columns that have been selected for display, the properties for each deduplication job are visible at a glance:

- Start time
- Trigger - Manual, Inline, Schedule, Tape eject, Tape caching
- Number of tapes
- Total amount of data to deduplicate
- Total amount of unique data written to the repository
- Deduplication ratio - 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
- Deduplication duration (in seconds)
- Total amount of data to replicate
- Unique data to replicate
- Deduplication ratio on the replication target server after the data is replicated and resolved
- Replication duration (in seconds)

- Status - Completed, Canceled, Failed

Schedule Deduplication Run History report

The Deduplication Run History report can be scheduled to be emailed to you on a regular basis. Refer to [“Schedule a report”](#) for more information.

Deduplication Tape History page



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

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

Only a Super Administrator can access the *Tape History* page.

The screenshot shows the FalconStor StorSafe with StorSight interface. The main content area displays the 'Deduplication Tape History' page for policy 'FS 1877'. The table lists the following columns: Barcode, Tape, ID, Parser, Start Time, Status, Data to Dedupe, Unique Data, Dedupe Ratio, and Dedupe Duration. The table contains 16 rows of data, all with a status of 'Completed'. The dedupe ratio for all entries is '> 10000:1'. The dedupe duration for most entries is '> 10000:1', except for one which is '2237:1'. The bottom of the page shows '1 to 50 of 3,748' and 'Tapes per page: 50'.

Barcode	Tape	ID	Parser	Start Time	Status	Data to Dedupe	Unique Data	Dedupe Ratio	Dedupe Duration
00390005	VirtualTape-00018	10000018	netbackup2	6/10/2019, 2:28:27 PM	Completed	1023.83 GB	1 MB	> 10000:1	> 10000:1
00390006	VirtualTape-00019	10000019	netbackup2	6/5/2019, 12:34:43 AM	Completed	1023.83 GB	18 MB	> 10000:1	> 10000:1
00390007	VirtualTape-00020	10000020	netbackup2	6/2/2019, 1:25:02 PM	Completed	235.93 GB	1 MB	> 10000:1	> 10000:1
00390007	VirtualTape-00020	10000020	netbackup2	6/1/2019, 12:01:04 PM	Completed	787.9 GB	5 MB	> 10000:1	> 10000:1
00390008	VirtualTape-00021	10000021	netbackup2	6/3/2019, 10:38:13 PM	Completed	677.02 GB	1 MB	> 10000:1	> 10000:1
00390008	VirtualTape-00021	10000021	netbackup2	6/2/2019, 8:04:18 AM	Completed	346.81 GB	28 MB	> 10000:1	> 10000:1
00390009	VirtualTape-00022	10000022	netbackup2	6/18/2019, 7:29:19 PM	Completed	1023.83 GB	4 MB	> 10000:1	> 10000:1
00390009	VirtualTape-00022	10000022	netbackup2	6/3/2019, 1:00:33 AM	Completed	2.18 GB	1 MB	2237:1	> 10000:1
00390009	VirtualTape-00022	10000022	netbackup2	6/3/2019, 12:18:11 AM	Completed	1 MB	0 B	> 10000:1	> 10000:1
0039000A	VirtualTape-00023	10000023	netbackup2	6/19/2019, 4:24:36 PM	Completed	715.2 GB	2 MB	> 10000:1	> 10000:1
0039000A	VirtualTape-00023	10000023	netbackup2	6/11/2019, 10:07:06 AM	Completed	307.63 GB	4 MB	> 10000:1	> 10000:1
0039000B	VirtualTape-00024	10000024	netbackup2	6/5/2019, 5:47:13 AM	Completed	1023.83 GB	4 MB	> 10000:1	> 10000:1
0039000C									

Depending upon the columns that have been selected for display, the properties for each tape are visible at a glance:

- Barcode
- Tape
- Tape ID
- Parser
- Deduplication start time
- Status - Completed, Canceled, Failed
- Total amount of data to deduplicate
- Total amount of unique data written to the repository
- Deduplication ratio - The ratio between the scanned data and the 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
- Deduplication duration (in seconds)

- Total amount of data to replicate
- Unique data to replicate
- Deduplication ratio on the replication target server after the data is replicated and resolved
- Replication duration (in seconds)

Schedule Deduplication Tape History report


The Deduplication Tape History report can be scheduled to be emailed to you on a regular basis. Refer to [“Schedule a report”](#) for more information.

Replication Targets page



Select *Manage* from the menu bar, click the *Deduplication* tab, and click the *Replication Targets* tab. This page lists the replication targets on the selected SIR server. You do not have replication targets for FalconStor StorSafe servers without deduplication.

Only a Super Administrator can access the *Replication Targets* page.



Add a replication target

Before you can configure replication for tapes in a deduplication policy, you must enable replication between the source and target SIR servers. If the replica is a StorSafe-S appliance, that appliance is your target.

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.

To add a replication target server:

1. Select **Manage** from the menu bar, select a SIR server, click the **Deduplication** tab, click the **Replication Targets** tab, and click the “+” (Create) icon.

The screenshot shows the 'Add Replication Target' dialog box for a single SIR server. The 'Source Server' field is populated with 'h21-66'. The 'Source IP Address' field is empty, with an 'Add NAT Address' button to its right. The 'Target Server' field is empty. The 'Protocol' section has 'TCP' selected and 'FC' as an option. The 'TCP Connections' field is set to 5, and the 'Replication Timeout' field is set to 5 minutes. The 'Encrypt Data' toggle is turned off. At the bottom right, there are 'Create' and 'Cancel' buttons.

SIR server

The screenshot shows the 'Add Replication Target' dialog box for a deduplication cluster. It features two 'Source Server' sections. The first section, 'Source Server Node 1', has 'H13-174' as the source server and an empty 'Source IP Address' field with an 'Add NAT Address' button. The second section, 'Source Server Node 2', has 'H13-176' as the source server and an empty 'Source IP Address' field with an 'Add NAT Address' button. The 'Target Server' field is empty. The 'Protocol' section has 'TCP' selected and 'FC' as an option. The 'TCP Connections' field is set to 5, and the 'Replication Timeout' field is set to 5 minutes. The 'Encrypt Data' toggle is turned off. At the bottom right, there are 'Create' and 'Cancel' buttons.

Deduplication cluster

2. Select the IP address of the source server.

For a cluster, select for both servers.

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

3. Select the target server and select an IP address for replication.
4. Select the protocol to use for replication.

TCP is the default. Both servers must be able to communicate with each other and the following ports must be open on the source server: port 11781 for unencrypted data and port 11782 for encrypted data.

Fibre Channel (FC) protocol cannot be used when repository encryption is enabled or if you are using object storage for your data repository.

To use FC, all FC targets on each server in the source cluster must be seen by the FC initiators on each server in the replica cluster, which means that each FC adapter target port on the source server must be zoned to each initiator port on the replica server and vice versa; each FC initiator port on the source server must be zoned to each target port on the replica server. If FC initiator ports on the replica are not zoned to target ports on the source server, the unique data replication job will fail.

5. If you are using TCP, set TCP options.

TCP Connections - The number of TCP connections to use to transfer hash data from the source to the target. The default value is 5 and the range is 1 to 16. For slow WAN connections, a low value can save memory without affecting throughput. For fast connections utilizing encryption, a higher value will use more memory and CPUs to increase throughput. A value greater than the number of CPU cores available will not increase throughput.

Replication Timeout - The amount of time that a connection fails to respond before the connection is considered broken. The default value is 5. and the range is 1 to 40 minutes. If a single connection between a source and target is broken, all resolver jobs between them are terminated.

Encrypt Data - An option to encrypt data with 256-bit AES encryption.

6. Click *Create* when you are done.

The target SIR server is configured to be a replication target.

You can now configure replication for tapes when you create a deduplication policy.

Replication

Data on virtual tapes can be replicated from one server to a target server for disaster recovery purposes. This chapter discusses replication of:

- [Deduplicated tapes](#) - Replication of unique data from tapes with a deduplication policy.
- [Tapes without deduplication](#) - Replication of data on tapes without deduplication.

Deduplicated tapes

The replication configuration for deduplicated virtual tapes is defined in the deduplication policy.

If you configure replication 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 or Fibre Channel. In either case, an IP connection between StorSafe servers associated with the deduplication repositories is required. The replication protocol is set when you designate a deduplication server as a replication target.

FC protocol can be used for unique data replication only when deduplication repository encryption is not enabled and when object storage is not used for the deduplication data repository. To use FC, all FC targets on each server in the source cluster must be seen by the FC initiators on each server in the replica cluster, which means that each FC adapter target port on the source server must be zoned to each initiator port on the replica server and vice versa; each FC initiator port on the source server must be zoned to each target port on the replica server. If FC initiator ports on the replica are not zoned to target ports on the source server, the unique data replication job will fail. If there is more than one FC connection between sites, all of the paths will be used to transmit the data. If one path goes down, the remaining paths will be used to continue transmitting data.

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

Each cluster can contain one, two, or four deduplication nodes. The clusters can each have a different number of deduplication nodes. For example, a two-node cluster can replicate to a one-node cluster, and vice versa.

Replication can occur in several configurations:


- One-to-one configuration - one cluster replicates to one cluster.
- One-to-many configuration - the source server replicates data to multiple targets.
 - *Cascaded* replication - the source server replicates data to Replication

Target 1, which in turn replicates data to Replication Target 2.

- *Parallel* replication - the source server replicates data to two target servers, 1 and 2. This can occur concurrently (replication to both target servers at the same time), or serially (replication first to Replication Target 1 and then to Replication Target 2).
- Many-to-one configuration - multiple clusters replicate to a single cluster, which would be the case where remote offices replicate data to a corporate data center.

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:

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).
- 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 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 vault or in the virtual tape library. The name of the LVIT corresponds to the name of the FVIT.
- The FVITs are listed when you select *Replica* tapes on the *Tapes* page.
- Note that this final step may not occur immediately after the initial replication of data and can take some time to complete, depending on the availability of 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 StorSafe servers.
- You must finalize all IP addresses before replication is configured.
- You must have at least two functional deduplication clusters.
- You must enable replication between the deduplication clusters by adding the target deduplication cluster to the primary cluster.
- You must have administrative rights on all servers.
- You must have enough space on the target server for the replica data.
- Each virtual tape you want to replicate must be included in a deduplication policy.
- Just as you have to associate your SIR server with your StorSafe server before you can create deduplication policies, the SIR server on the target side must be associated with your local SIR server before you can configure replication.
- *At the time of configuration*, each virtual tape that will be configured for replication must be in a slot, not a virtual library tape drive.
- While you can configure replication for a virtual tape that has not been deduplicated, replication will not run until at least one deduplication has taken place.
- An IP connection is required. Fibre Channel can be used to replicate unique data but you still need an IP connection between VTL servers associated with the deduplication repositories.
- On the target server, you must prepare physical resources for deduplication use and enable deduplication.
- Network firewalls should allow access through TCP ports 11782 and 11781 for replication encryption and unencryption.

Configure replication for deduplicated tapes

You must do the following to configure replication with deduplication:

1. Associate each StorSafe server with a deduplication cluster.

The source StorSafe server should be associated with the source deduplication cluster and the target StorSafe server should be associated with the target deduplication cluster.

2. Add a target deduplication replication server to enable replication between the source and target SIR 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.

Refer to [“Add a replication target”](#) for more information.

3. Create a deduplication policy and enable replication.

When you enable replication, you will have to add a target StorSafe server. This is either the target server that is associated with the *target* deduplication cluster or it is a combined backup & SIR server. Refer to [“Create a deduplication policy”](#) for more information.

Access data on a replicated VIT

If a replicated virtual tape is needed (due to a failure at the primary site), the administrator can do one of the following so that the data can be accessed by backup software:

- Move the virtual tape from the virtual vault to a virtual library on the target server.
- If the primary cluster has been repaired, you can replicate the entire deduplicated virtual tape back to that server. Depending upon how much data exists in the cluster’s repository, this may be time consuming. However, if all of the data is there, only the VIT will need to be replicated back.
- Export the data on the virtual tape to a physical tape. Refer to the [“Physical Tape Export”](#) for more information.
- If you move a local VIT out of the vault, replication of this VIT will be discontinued until the tape is moved back to the vault. **It is important to note** that any new data added to the tape while it is not in the vault will be overwritten when the tape is returned to the vault and replication proceeds.

Tapes without deduplication

For virtual tapes without deduplication, replication can be accomplished in the following ways:

- [Tape replication](#) - Replicates *changed* data from a primary virtual tape to the same server or another server at prescribed intervals, based on user defined policies.
- [Auto Replication](#) - Replicates data whenever a virtual tape is moved to an IE slot (the vault) from a virtual library.
- [Tape Copy](#) - Replicates the contents of a single tape *on demand*.

Tape replication

Replication is a process that protects the information on a virtual tape by maintaining a copy of a virtual tape on the same server or on another StorSafe 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 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 StorSafe 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 StorSafe server to another over the IP network. With remote replication, the replica disk is located on a separate target StorSafe server.
- Local replication of non-deduplicated tape data allows fast, data synchronization of storage volumes on the same StorSafe 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.

Note: Tape replication is mutually exclusive with Auto Replication, Auto Archive, and Cloud Tiering.

Replication requirements for virtual tapes

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

- You must have enough space on the target server for the replica resource.
- You must finalize all IP addresses before replication is configured.
- For remote replication, you must have two StorSafe servers.
- For remote replication, you must have administrative rights on both servers.

- For remote replication, 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

1. Select *Manage* from the menu bar, select a StorSafe server, click the *Virtual Libraries* tab, and select the *Tapes* tab.
2. Select one or more tapes, click the *Replication* icon, and select *Create Replication Policy*.
3. On the *General* tab, identify the replica (target) server, specify how to create the replica resource, and specify replication rules.

The screenshot shows the 'Create Replication Policy' dialog box with the 'General' tab active. The 'Virtual Tape' field is set to 'VirtualTape-00219'. The 'Enable Local Replication' toggle is turned off. The 'Target Server*' and 'Storage Pool*' are dropdown menus, with the storage pool currently showing '(Select)'. The 'Schedule Replication' toggle is turned on. The 'Start Time' is set to '2019-12-13' at '00:00'. The 'Replication Interval' is set to '10' minutes. The 'Replicate by Watermark' and 'Replicate by Data Change' toggles are both turned off. At the bottom right, there are 'Create' and 'Cancel' buttons.

Indicate whether you want to use local or remote replication.

- Local replication - Specify a storage pool and a name for the replica resource that is up to 32 characters, excluding spaces or special characters `< > & $ " / \ ' * ? | %`
- Remote replication - Select the server that will contain the replica, specify the management and transmission IP addresses, and specify a storage pool for the replica resource.

Specify when replication should occur.

- **Schedule** - Indicate the date and time 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.
- **Watermark** - When a watermark value is reached, replication of changed data will begin as soon as the virtual tape is ejected from the tape drive.
- **Data change** - Trigger replication whenever data changes on a virtual tape. Replication based on data change is mutually exclusive with the schedule and watermark triggers and is not available when Tape Caching is enabled.
- **Migration** - When Tape Caching is enabled, replicate before or after Tape Caching export.

4. On the *Options* tab, specify timeout, retry, and transmissions options.

The screenshot shows the 'Create Replication Policy' dialog box with the 'Options' tab selected. The settings are as follows:

Setting	Value	Unit	Help Icon
Replication Timeout	60	Seconds	Yes
Retry Count	1		Yes
Retry Interval	60	Seconds	Yes
Compress Data	☐		No
Encrypt Data	☐		No

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.

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

Encryption 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 and encryption are not valid for local replication and 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.

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

5. Click *Create* when done.

Promote a virtual tape replica

If a replica resource is needed, you 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.

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

Unless you select the *Test Mode* option, 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. Select *Manage* from the menu bar, select a StorSafe server, click the *Virtual Libraries* tab, select the *Tapes* tab, set the location to *Replica*, select a tape, click the *Replication* icon, and select *Promote*.

Replica Tape	tsvtl08-VirtualTape-00011
Source Server	UNKNOWN
IP Address	10.6.1.200
Promote in Test Mode	☐ ?
Force Promote	☐ ?

Promote Cancel

2. Specify if you want to promote the replica in test mode.

Test mode promotes replica storage *without* breaking the replication configuration. This can be useful for disaster recovery testing.

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.

3. Specify if you need to force promotion, which is needed if the source server or tape is not available.
4. Click *Promote* when done.

Manage replication for virtual tapes

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)

You can also:

- Manually start replication for a virtual tape. Only one outgoing replication job can be active at a time; the next job will become active after the current job finishes.
- Suspend the schedule of a replication policy or resume a suspended policy. Suspending replication prevents future replications from automatically being triggered by your replication policies (watermark, interval, time). This will not stop replication that is currently in progress. You can still manually start the replication process while the schedule is suspended.
- Delete replication. This removes the replication configuration on the primary and either deletes or promotes the replica resource on the target server at the same time.

Note: If you overwrite or relabel a tape with a replication policy, the replica tape will also get the same date after replication. If you want to update, delete, or shred a tape without impacting the replica, remove the replication configuration for that tape before performing these operations.

Select *Manage* from the menu bar, select a StorSafe server, click the *Virtual Libraries* tab, select the *Tapes* tab, select a tape, click the *Replication* icon, and select the action you want to perform.

Auto Replication

Auto Replication replicates data to another StorSafe server whenever a virtual tape is moved to an IE slot (the vault) from a virtual library (such as from a backup application or other utility).

Note: Tape Caching, Auto Archive, Auto Replication, and Cloud Tiering services are mutually exclusive so only one can be configured per virtual library.

You can configure Auto Replication when you create a virtual tape library or you can update an existing virtual tape library for which Auto Replication was not previously enabled. This is done on the *Tape Services* tab for a virtual tape library, which are located on the *Libraries* page. You must also configure Auto Replication for virtual tapes in a library with Auto Replication enabled.

For the library, you must determine if you want virtual tapes moved (removed) or copied (retained) after the data is replicated. If you select *Move*, indicate how long to wait before deleting a tape.

Select the remote server from the list of existing target servers.

The screenshot shows the 'Create Virtual Tape Library' dialog box with the 'Tape Services' tab selected. The 'Service' dropdown menu is set to 'Auto Replication'. Below it, the 'Mode' section has two buttons: 'Move' and 'Copy'. The 'Delay Tape Removal' is set to '1' day. The 'Target Server' dropdown menu is set to 'H13-97'. At the bottom right, there are 'Create' and 'Cancel' buttons.

Notes:

- 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.

For virtual tapes, you must select the transmission IP address of the replica server, and a storage pool for the replica resource.

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 to compress and/or encrypt data. Compression option provides enhanced throughput during replication by compressing the data stream. Encryption secures data transmission over the network during replication.

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

The screenshot shows the 'Create Tapes' dialog box with the 'Service' tab selected. The 'General' tab is also visible. The 'Service' dropdown is set to 'Auto Replicate'. The 'Replica Server' dropdown is set to 'H13-97'. The 'Transmission IP Address' dropdown is empty. The 'Replica Storage Pool' dropdown is set to '(Select)'. The 'Replication Timeout' is set to 60 seconds. The 'Retry Count' is set to 1. The 'Retry Interval' is set to 60 seconds. The 'Compress data' and 'Encrypt data' toggle switches are both turned off. The 'Create' and 'Cancel' buttons are at the bottom right.

Field	Value
Service	Auto Replicate
Replica Server	H13-97
Transmission IP Address	
Replica Storage Pool	(Select)
Replication Timeout	60 Seconds
Retry Count	1
Retry Interval	60 Seconds
Compress data	Off
Encrypt data	Off

Tape Copy

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

Notes:

- Copy is mutually exclusive with Auto Replication, Auto Archive, and Cloud Tiering.
- 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.

You can copy a tape whenever needed. This is done from the *Vault* tab or *Tapes* tab of a StorSafe server. Select a tape, click the *Manage* icon, and select *Copy*.

The screenshot shows a 'Copy Tape' dialog box with the following fields and values:

- Virtual Tape: VirtualTape-00175
- Copy to: Local Server (selected), Remote Server
- Server: VTL01
- IP Address: 10.8.25.130
- Transmission IP Address: 10.8.25.130
- Location: Library
- Remote Library*: ADIC-Scalar 100-00127-sang
- Storage Pool: StoragePool-22

Buttons: Copy, Cancel

If you select to copy to a remote server, you will have to select the server and specify the server's IP address and IP address to transmit data.

Tapes can be copied to a library, standalone drive, or the 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.

Cloud Tiering

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 *Cloud Tiering* can be used to archive virtual tape data to the cloud. At the end of each backup job, the tape is ejected to the vault, where the tape data gets compressed and migrated to the object storage of a cloud provider. You can export virtual WORM tapes and they will maintain the WORM property. If they are imported, they will become virtual WORM tapes.

Cloud tiering policies

Cloud tiering is set at the virtual tape library level to automatically migrate tape data to object storage as soon as a backup is complete and a tape is ejected to the 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.

Cloud Tiering to Object Storage, Automated Tape Caching, Auto Archive, Auto Replication, and replication without deduplication are mutually exclusive.

If replication has been enabled in the deduplication policy of a tape that is configured for migration to object storage, data replication occurs according to the policy as long as the tape has not been moved to the vault. Once the tape is moved to the vault, it is added to the migration job queue. When migration starts, replication will be performed once migration is complete, before the tape becomes a stub tape. If you are using *Move* mode instead of *Copy* mode, set a grace period that is large enough, for example two days, to give you time to replicate the data before VITs becomes stub tapes.

The replica tape will not be automatically migrated. Migration only applies when a tape is ejected to the vault.

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

Migration and recovery jobs (updated October 2024)

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

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

As migration occurs, the allocated space in the cloud object storage bucket will grow based on the required space for the migrated tape. With a compression ratio of 2-to-1, you can estimate the bucket size.

Data upload to S3 is via multiple parts of 16 MB. Incomplete multi-part uploads will be cleaned up after seven days if there are no new request to those parts. The URI is used as prefix for objects. The S3 authentication method supports AWS v4.

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

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

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

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

Similarly, the migration process will keep compressed StorSafe 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 StorSafe, the migration process will compress the data in flight and at rest.

In order to configure tape migration to object storage, you must create at least one object storage account with your object storage provider and add it on the *Object Storage* page. To isolate data between different users, you can create separate accounts. You must also enable Cloud Tiering for a virtual tape library.

Immutable object storage

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

From the IBM COS management GUI, you can configure the retention policy of your bucket and set the minimum and default values to zero, and the maximum retention period to the number of days you want to lock the bucket data.

From the FalconStor portal, you can set a tape migration policy for a virtual tape library or drive. At the end of each backup job, the tape is ejected to the vault, where the tape data gets packaged and migrated to the object storage of a cloud provider. This copy of data remains intact during the retention period and is protected from malicious activity or accidental deletion.

Object storage providers

Object storage is a scalable, cost-effective data storage architecture to store data in its native format and manage the data as objects.

StorSafe supports popular object storage providers, such as Amazon Web Services (AWS) S3, Microsoft Azure, Hitachi Content Platform (HCP), IBM COS, Wasabi, Google, and Generic S3.

Configuration

An administrator must do the following to configure migration to cloud 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 StorSafe server.
Refer to [“Add an object storage account”](#) for more information.

3. Enable Cloud Tiering for a virtual tape library.

This is typically done when you create a virtual tape library. You can also enable it for an existing library on the *Virtual Tape Libraries* page.

Refer to [“Configure Cloud Tiering for a virtual tape library”](#) for more information.

Each library can use a different object storage account.

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

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


If you are using Cloud Tiering on a StorSafe system with existing tapes that need to be migrated, you can manually trigger migration to object storage from the *Vault* tab. Refer to [“Vault”](#) for more information.

Object Storage Provider page



Select *Manage* from the menu bar and click the *Object Storage* tab. This page lists the object storage provider accounts on the selected server.

Only a Super Administrator can access the *Object Storage* page.



Object Storage Provider page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Add	Add an object storage account.
	Update	Update an object storage account.
	Delete	Delete an object storage account.

Add an object storage account

You must create at least one object storage account with your object storage provider before you can add it to your backup server here.

1. Select *Manage* from the menu bar, select a backup server and click the *Object Storage* tab.
2. Click the “+” (*Add*) icon.
3. Select a provider and provide configuration information.

Amazon Web Services (AWS) S3

The screenshot shows a 'Add Object Storage Account' window. The 'Provider' is set to 'Amazon Web Services (AWS)'. The 'Label*' field is empty. The 'Access Key*' field is empty. The 'Secret Key*' field is empty. The 'Region' is set to 'us-east-1: US East (N. Virginia)'. The 'Storage Class' is set to 'Glacier'. The 'S3 Bucket*' field is empty. The 'Use Dual-Stack' toggle is off. The 'Protocol' is set to 'https'. The 'Proxy Server IP Address' is set to 'Server'. The 'Proxy Port' is set to '1'. The 'Proxy Protocol' is set to 'https'. The 'Proxy Username' field is empty. The 'Proxy Password' field is empty. The 'Encrypt Data' toggle is off. The 'Comments' field is empty.

Label - Specify a name for this account.

Access Key - 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.

Region - Specify the region of Amazon AWS S3 service. The default is *us-east-1*.

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

S3 Bucket - Specify the name of an existing bucket. Objects created by tape migration will be kept in this bucket.

Use 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.

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

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

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

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

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

Encrypt Data - Enable or disable end-to-end encryption. When enabled, all data is encrypted before saving it to object store; 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.

Microsoft Azure

The screenshot shows a 'Add Object Storage Account' dialog box. The 'Provider' is set to 'Microsoft Azure'. The 'Label' field is empty with a placeholder 'Label to name the account'. The 'Storage Account' field is empty with a placeholder 'User account'. The 'Access Key' field is empty with a placeholder 'Account password'. The 'Blob Container' field is empty with a placeholder 'Blob object storage'. The 'Endpoint Suffix' field is empty with a placeholder 'Custom suffix for the endpoint'. The 'Use for Deduplication Repository' toggle is turned off. The 'Protocol' is set to 'https'. The 'Proxy Server IP Address' field is empty with a placeholder 'Server'. The 'Proxy Port' field is empty with a placeholder '1'. The 'Proxy Protocol' is set to 'https'. The 'Proxy Username' field is empty with a placeholder 'Username'. The 'Proxy Password' field is empty with a placeholder 'Password'. The 'Encrypt Data' toggle is turned off. The 'Comments' field is empty with a placeholder 'Optional description for the account'. There are 'Add' and 'Cancel' buttons at the bottom right.

Label - Specify a name for this account.

Storage Account - Specify a user account.

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

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

Endpoint Suffix - Optionally, specify an endpoint suffix.

Use for Deduplication Repository - Specify if this account will be used to host data disks of the deduplication repository on object storage instead of being

used to archive virtual tape data to object storage. You cannot use the same account for both.

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

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

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

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

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

Encrypt Data - Enable or disable end-to-end encryption. When enabled, all data is encrypted before saving it to object store; 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.

Hitachi Content Platform (HCP)

The screenshot shows a web-based form titled "Add Object Storage Account" with a close button (X) in the top right corner. The form is for configuring a Hitachi Content Platform (HCP) account. It contains the following fields and controls:

- Provider:** A dropdown menu set to "Hitachi Content Platform (HCP)".
- Label:** A text input field with the placeholder "Label to name the account".
- Domain:** A text input field with the placeholder "Domain name".
- Tenant:** A text input field with the placeholder "Tenant name".
- User:** A text input field with the placeholder "User account".
- Password:** A text input field with the placeholder "Account password".
- Namespace:** A text input field with the placeholder "Tenant bucket space".
- Transport:** A dropdown menu set to "REST".
- Use for Deduplication Repository:** A toggle switch currently turned off.
- Protocol:** A dropdown menu set to "https".
- Proxy Server IP Address:** A text input field with the placeholder "Server".
- Proxy Port:** A text input field with the value "1".
- Proxy Protocol:** A dropdown menu set to "https".
- Proxy Username:** A text input field with the placeholder "Username".
- Proxy Password:** A text input field with the placeholder "Password".
- Encrypt Data:** A toggle switch currently turned off.
- Comments:** A text input field with the placeholder "Optional description for the account".

Label - Specify a name for this account.

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

Tenant - Specify the tenant for the provider.

User - Specify the username for the provider.

Password - Specify the password for the user.

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.

Transport - 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*, the HS3 protocol 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.

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

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

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

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

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

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

Encrypt Data - 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.

IBM Cloud Object Storage (COS)

The screenshot shows a web form titled "Add Object Storage Account". The form contains the following fields and controls:

- Provider:** A dropdown menu with "IBM Cloud Object Storage (IBMCOS)" selected.
- Label:** A text input field with placeholder text "Label to name the account".
- Access Key:** A text input field with placeholder text "Access key ID".
- Password:** A text input field with placeholder text "Account password".
- Region Type:** A dropdown menu with "Regional" selected.
- Endpoint Type:** A dropdown menu with "Public" selected.
- Location Endpoint:** A dropdown menu with "Australia" selected.
- Immutable:** A toggle switch, currently turned off.
- Bucket Name:** A text input field with placeholder text "Existing bucket name".
- Protocol:** A dropdown menu with "https" selected.
- Proxy Server IP Address:** A text input field with placeholder text "Server".
- Proxy Port:** A text input field with the value "1".
- Proxy Protocol:** A dropdown menu with "https" selected.
- Proxy Username:** A text input field with placeholder text "Username".
- Proxy Password:** A text input field with placeholder text "Password".
- Encrypt Data:** A toggle switch, currently turned off.
- Comments:** A text input field with placeholder text "Optional description for the account".

Note: 'Smart Tier' is the default storage class when creating a new bucket on IBM Cloud. You can use 'Smart Tier' or 'Standard' classes for StorSafe to migrate tape data to COS. These classes perform better than 'Vault' or 'Cold Vault' for StorSafe activities of backup with inline deduplication (read/write), replication (read) and tape data migration to COS (write).

Label - Specify a name for this account.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Wasabi

The screenshot shows a 'Add Object Storage Account' dialog box. The 'Provider' is set to 'Wasabi'. The 'Label' field has a placeholder 'Label to name the account'. The 'Access Key' field has a placeholder 'Access key ID'. The 'Secret Key' field has a placeholder 'Account password'. The 'Region' field is empty. The 'Bucket Name' field has a placeholder 'Existing bucket name'. The 'Use for Deduplication Repository' toggle is turned off. The 'Proxy Server IP Address' field has a placeholder 'Server'. The 'Proxy Port' field has a placeholder '1'. The 'Proxy Protocol' is set to 'https'. The 'Proxy Username' field has a placeholder 'Username'. The 'Proxy Password' field has a placeholder 'Password'. The 'Encrypt Data' toggle is turned off. The 'Comments' field has a placeholder 'Optional description for the account'. There are 'Add' and 'Cancel' buttons at the bottom right.

Label - Specify a name for this account.

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

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

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

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

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

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

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

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

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

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

Google

The screenshot shows a 'Add Object Storage Account' dialog box. The 'Provider' is set to 'Google'. The 'Label*' field contains 'Label to name the account'. The 'Access Key*' field contains 'Access key ID'. The 'Secret Key*' field contains 'Account password'. The 'Region*' field is empty. The 'Bucket Name*' field contains 'Existing bucket name'. The 'Use for Deduplication Repository' toggle is turned off. The 'Proxy Server IP Address' field contains 'Server'. The 'Proxy Port' field contains '1'. The 'Proxy Protocol' is set to 'https'. The 'Proxy Username' field contains 'Username'. The 'Proxy Password' field contains 'Password'. The 'Encrypt Data' toggle is turned off. The 'Comments' field contains 'Optional description for the account'. There are 'Add' and 'Cancel' buttons at the bottom right.

Label - Specify a name for this account.

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

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

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

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

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

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

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

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

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

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

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

Generic S3

The screenshot shows a 'Add Object Storage Account' dialog box with the 'Generic S3' provider selected. The form contains the following fields and controls:

- Provider:** A dropdown menu set to 'Generic S3'.
- Label:** A text input field with placeholder text 'Label to name the account'.
- Access Key:** A text input field with placeholder text 'Access key ID'.
- Password:** A text input field with placeholder text 'Account password'.
- Endpoint:** A text input field with placeholder text 'Endpoint URI'.
- Region:** A text input field with placeholder text 'Enter region name'.
- Bucket Name:** A text input field with placeholder text 'Existing bucket name'.
- Use for Deduplication Repository:** A toggle switch currently turned off.
- Proxy Server IP Address:** A text input field with placeholder text 'Server'.
- Proxy Port:** A text input field with placeholder text '1'.
- Proxy Protocol:** A dropdown menu set to 'https'.
- Proxy Username:** A text input field with placeholder text 'Username'.
- Proxy Password:** A text input field with placeholder text 'Password'.
- Encrypt Data:** A toggle switch currently turned off.
- Comments:** A text input field with placeholder text 'Optional description for the account'.
- Buttons:** 'Add' and 'Cancel' buttons at the bottom right.

Label - Specify a name for this account.

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

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

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

Region - Optionally, specify the region for clouds where region cannot be included in the endpoint.

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

Use for Deduplication Repository - Specify if this account will be used to host data disks of the deduplication repository on object storage instead of being used to archive virtual tape data to object storage. You cannot use the same account for both. A Generic S3 account can be used for different cloud providers, where the endpoint identifies the destination, for example, for AWS in

US east region, the endpoint URL is `https://s3.us-east-1.amazonaws.com`; for IBM COS in US east region, the endpoint URL is `s3.us-east.cloud-object-storage.appdomain.cloud`; for Google, the endpoint URL is `https://storage.googleapis.com`.

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

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

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

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

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

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

4. Confirm the information and click the *Add* button.

Manage object storage accounts

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

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

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

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

Notes:

- Changing the HCP user name/password can result in an IO disruption. Be sure to suspend all active IO on StorSafe before changing the HCP credentials. Restart StorSafe services afterward.
- The *localfilesystem* account is an internal account reserved for StorSafe servers. It cannot be modified.

Cloud Tier page



When you select *Manage* from the menu bar and select the *Virtual Libraries* tab, the *Cloud Tier* page lists all virtual tapes in the vault that have been migrated from the selected StorSafe server.

Only a Super Administrator can access the *Cloud Tier* page.

Depending upon the columns that have been selected for display, the properties for each tape are visible at a glance:

- Source tape name
- Barcode
- Tape ID
- Size
- Used size
- Virtual tape index status - pure (all tape data has been deduplicated), mixed (tape contains pointers to deduplicated data as well as data that has not yet been deduplicated), plain (tape not yet deduplicated)
- Cloud Tiering migration mode - Copy or Move
- Object storage provider
- User access key
- Object storage provider ID (click for object storage account details)
- Encrypted tape indicator
- Stub - Checked if Cloud Tiering occurred and a stub tape exists
- Write-once-read-many (WORM) tape usage
- Parent virtual tape library ID (click for library details)

Cloud Tier page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Recover from object storage	Recover stub tape data to a virtual tape from cloud object storage.
	Convert to stub tape	While migration is triggered automatically after backup is complete, you can manually convert a virtual tape to a stub tape.

Convert a virtual tape to a stub tape

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

Recover data from object storage

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

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

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

You must make sure there is enough backup cache capacity available for the recovered data before you begin the recovery process. Reconstructed tapes are not deduplicated even if the original tape was in a deduplication policy. In addition, you must have at least two times the size of the VIT available for the recovery process.

The recovered tape will be write-protected after recovery. You need to remove the write-protected flag from the tape for deduplication to trigger.

To recover data from object storage:

1. Select *Manage* from the menu bar, select a StorSafe server, click the *Virtual Libraries* tab, click the *Cloud Tier* tab, select a stub tape and click the *Recover from Object Storage* icon.

2. Select the virtual library to which you want the recovered tape to reside.
3. Select the virtual tape library slot and recovery mode.

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

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

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

You can select from *Bulk* (default), *Standard*, and *Expedited* (available for *Glacier* only).

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 portal 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.

5. Confirm the information and click *OK*.

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

Recover data to a new StorSafe server instance

This section explains how to recover data from tapes that were exported to object storage to a replacement StorSafe server (version 11.13). While this procedure does not recover the actual StorSafe server, it does require that you had previously saved the server configuration from your StorSafe server.

1. Prepare and install a new StorSafe server with new storage.
Be sure to set the host name to the original host name.
2. Create the tape database and configuration repository for the StorSafe server.
3. Run the following command at the command line:

```
# recoverobjstube.sh -s [saved configuration]
```

You can check the syslog (/var/log/messages) to confirm that the procedure was successful.


NAS

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

NAS resources

When you select *Manage* from the menu bar and click the *NAS* tab, the *NAS Resources* page lists the NAS resources on the selected StorSafe server. NAS resources are virtualized disks that are used to create NAS shares and are used for standard file-based backup and for OpenStorage backups when the FalconStor OpenStorage (FSOST) option is configured. Inline deduplication is used for NAS resources; there is no need to create a deduplication policy for NAS data.

From this page you can also access the *Shares*, *Folders*, *SMB Clients*, *NFS Clients*, *Replication*, and *Integrity Check* pages for your StorSafe servers.



Depending upon the columns that have been selected for display, the properties for each NAS resource are visible at a glance:

- Resource name
- Resource ID
- Mount point - The directory from which the share can be accessed.
- Mount options - Options for mounting the file system.
- Format options - Options for formatting the file system.
- File system status

- Capacity of NAS resource
- Used space of NAS resource
- Indication of whether or not the NAS resource is deduplicated
- Indication of whether or not the NAS resource is replicated
- Indication of whether or not OST is configured for the NAS resource

NAS Resources page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Create	Create a NAS resource.
	Delete	Delete the selected NAS resource.
	Manage	Expand, unmount/mount, format, and check file system.
	Exclusion	Exclude a NAS resource from deduplication or global replication.
	Fast Copy	Copy files and directories between shares on NAS resources using Fast Copy.
	Info	Display physical layout information for a NAS resource.

Create a NAS resource

The deduplication repository must be created before you can create NAS resources. To create a NAS resource for NAS shares:

1. Select *Manage* from the menu bar, click the *NAS Resources* tab, and select a StorSafe server.
2. Click the “+” (*Create*) icon.
3. Enter information about the NAS resource.

Name - Specify the NAS resource name up to 64 characters, excluding < > & \$ " / '. The name cannot start with a dot and cannot contain consecutive dots. By default, a NAS- prefix is used before the device ID if FSOST is not configured; an LSU- prefix is used when FSOST is configured. This is controlled by the *Use for OST* field.

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

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

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

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

- `-b size=4096` - Specify the block size of 4 KB, which determines the minimum amount of space to use for each file.

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

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

Use for OST - Specify if you want to create the NAS resource for OpenStorage so that it will be represented as a Logical Storage Unit (LSU) via the FalconStor FSOST agent.

4. Click *Create* when done.

Delete a NAS resource

To delete a NAS resource, select *Manage* from the menu bar, click the *NAS Resources* tab, select a StorSafe server, select a NAS resource, and click the *Delete* icon.

Expand a NAS resource

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

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

Unmount/mount a NAS resource

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

To unmount/mount a NAS resource, select *Manage* from the menu bar, click the *NAS Resources* tab, select a StorSafe server, select a NAS resource, click the *Manage* icon and select the appropriate action.

Format a NAS resource

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

To format a NAS resource, select *Manage* from the menu bar, click the *NAS Resources* tab, select a StorSafe server, select a NAS resource, click the *Manage* icon and select *Format*. You can specify the format options and mount options to use.

Check the file system on a NAS resource

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

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

Exclude/include a NAS resource from deduplication

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

Note: NAS resources using OST should not be excluded from deduplication; excluding OST resources can cause duplication jobs to hang.

To exclude a specific NAS resource from deduplication, select *Manage* from the menu bar, click the *NAS Resources* tab, select a StorSafe server, select a NAS resource, click the *Exclusion* icon, and select *Exclude Deduplication*.

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

Exclude/include a NAS resource from global replication

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

To exclude a specific NAS resource from replication, select *Manage* from the menu bar, click the *NAS Resources* tab, select a StorSafe server, select a NAS resource, click the *Exclusion* icon, and select *Exclude Global Replication*.

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

Copy NAS files and directories using Fast Copy

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

Select *Manage* from the menu bar, click the *NAS Resources* tab, select a StorSafe server, click the *Fast Copy* icon, and select *Start*. Specify the source and target paths. You can click *Browse* and expand a folder to find a directory. Specify if you want to overwrite target data if it exists. The job will fail if the target directory exists and this option is not selected.


Display information about a NAS resource

You can display physical layout info for a NAS resource. Select *Manage* from the menu bar, click the *NAS Resources* tab, select a StorSafe server, select a NAS resource, and click the *Info* icon.

Shares

Select *Manage* from the menu bar, click the *NAS* tab, and click the *Shares* tab. This page lists NAS shares on the selected NAS resource. Select which NAS resource to view from the *NAS Resource* drop-down list.

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



Depending upon the columns that have been selected for display, the properties for each share are visible at a glance:

- Share name
- NFS Clients (click to see the list with options)
- SMB clients (click to see the list with options)
- SMB share
- SMB comment
- Maximum number of SMB users
- SMP properties
- Indication of whether or not the NAS share is deduplicated (fully or partially)
- Indication of whether or not the NAS share is replicated (fully or partially)

Shares page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Create	Create a NAS share.
	Update	Update client settings for a NAS share.
	Delete	Delete the selected NAS share.
	Exclusion	Exclude a NAS share from deduplication or global replication.
	Set properties	Set NAS directory properties.
	Show files	Show folders and files for a NAS share.
	Statistics	Show deduplication and replication statistics for a NAS share.

Create a NAS share

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab and select a NAS resource from the drop-down box.
2. Click the “+” (*Create*) icon.
3. On the *General* tab, specify share properties.

Create NAS Share/Folder

General | SMB Clients | SMB Properties

NAS Resource: NAS-00254

Parent Directory: Path

Share/Folder Directory*: Name

NFS Clients

☐	Name	Access	Squash	Insecure	Kerberos	Kerberos with Integrity	Kerberos with Privacy	System Protocol
☐	13-180	Read-Write	root	☐	☐	☐	☐	☒
☐	test	Read-Write	root	☐	☐	☐	☐	☒
☐	client1	Read-Write	root	☐	☐	☐	☐	☒
☐	client2	Read-Write	root	☐	☐	☐	☐	☒
☐	client7-1	Read-Write	root	☐	☐	☐	☐	☒

Create Cancel

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

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

4. Select NFS clients who will access the share and specify permissions for each client.

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

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

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

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

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

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

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

System Protocol - Option to enable the system protocol for client security.

5. On the **SMB Clients** tab, select SMB clients who will access the share and specify share name and options.

Create NAS Share/Folder

General | **SMB Clients** | SMB Properties

SMB Clients

Search for...

☐	Name	Type	Access rights
☐	Administrator	User	Read-Write
☐	DefaultAccount	User	Read-Write
☐	falcon101	User	Read-Write
☐	Guest	User	Read-Write
☐	henry	User	Read-Write

Share Name: Smbshare

Description: Comment

Max Number of Connections: 0

Create Cancel

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

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

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

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

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

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

Property	Value
nt acl support	no

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

7. Click *Create* when done.

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

Update a NAS share

You can update client settings for a NAS share. Refer to ["Create a NAS share"](#) for more information.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab, select a share, and click the *Update* icon.

Delete a NAS share

You can delete a NAS share if it does not have any files or subfolders.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab, select a share, and click the *Delete* icon.

Exclude/include a NAS share from deduplication

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

To exclude a specific NAS share from deduplication, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab, select a share, click the *Exclusion* icon, and select *Exclude Deduplication*.

To include a specific NAS share that was previously excluded from deduplication, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab, select a share, click the *Exclusion* icon, and select *Include Deduplication*.

Exclude/include a NAS share from global replication

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

To exclude a specific NAS share from replication, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab, select a share, click the *Exclusion* icon, and select *Exclude Global Replication*.

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

Set properties for a NAS share

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

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab, select a share, and click the *Set Properties* icon.

Set NAS Directory Properties

NAS Resource: NAS-00254

Share/Folder: ellshare

Owner*: root

Group*: nasgrp

Access	Read	Write	Execute
Owner	☒	☒	☒
Group	☒	☒	☒
Others	☒	☐	☒

Sticky Bit: ☐ ?

Apply changes to subdirectories: ☐

Set Cancel

2. Set the appropriate properties.

Owner - Directory owner user account name.

Group - User group with permissions on the directory.

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

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

3. Click *Set* when done.

Show files for a NAS share

You can display the folders and files on a NAS share. You will see the size of each and deduplication and replication status, if enabled.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab, select a share, and click the *Show Files* icon.

Display statistics for a NAS share

You can display deduplication and replication statistics, if enabled. This includes data size, number of files and folders, number of deduplicated files, space used by deduplicated files, number of files awaiting deduplication, space used by files awaiting deduplication, number of replicated files, number of files awaiting replication, space use by files awaiting replication, replicated data size, and size of unique replicated data.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab, select a share, and click the *Show Files* icon.

Map/mount shares

Windows clients

- Map a share You should map a share for each Windows client to allow access to the share. Do the following on each Windows client:
1. Open *Windows Explorer* (or *My Computer*).
 2. Select *Tools --> Map Network Drive*.
 3. Set the path to the shared folder.
The path is: `\\hostname\sharename`
where *hostname* is the StorSafe server's name or IP address and *sharename* is the name of the shared folder. For example: `\\server1\Data1`
 4. Enter login information.
For *User* mode, enter the user name and password that was set when the user was created.
For *Domain* mode, enter the user's full account name (*Domain\username*) and the user's password that is defined at the Active Directory level.

NFS clients

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

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

For example: `/mnt/share`

2. Locally, mount the share.

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

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

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

For example:

```
mount STORSAFE35:/nas/NAS-00002/fds/Data1 /mnt/share
```

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

Then, unmount and re-mount the filesystem.

Refer to the following link for more information:

<https://access.redhat.com/solutions/2791811>

Client mount options

We recommend using the following mount options for NFS clients:

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

AIX 6.x/7.x

NFSv4 (with krb5i authentication):

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

Refer to the following link for additional information about mounting NFSv4 on an AIX client:

<https://www.ibm.com/support/pages/ibm-aix-how-setup-nfsv4-mount-clinet-and-server>

NFSv3:

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

Solaris 10 and 11

NFSv4:

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

NFSv3:

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

Note: We highly recommend using the “llock” mount option with Solaris NFS clients.

Linux

NFSSv4.1:

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

NFSv4.1 (with krb5 authentication):

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

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

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

NFSv3:

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

Note: We highly recommend using the “nolock” mount option with Linux NFS clients.

Oracle Recovery Manager (RMAN) (updated October 2024)


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

Operating System	Options
AIX	<code>cio,bg,rsize=65536,wspace=65536,noac,timeo=600</code>
Solaris	<code>bg,rspace=32768,wspace=32768,noac,forcedirectio,suid</code>
Linux	<code>rspace=524288,wspace=524288,actimeo=0,timeo=600</code>

Folders

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

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



Depending upon the columns that have been selected for display, the properties for files in each folder are visible at a glance:

- Name
- Last modified time for a file
- File size
- Size of stub file for a deduplicated file
- Size of unique data in the deduplication repository for a deduplicated file
- Deduplication ratio for a deduplicated file
- Deduplication status for a file - Completed, pending, excluded, in progress
- Replication status for a file - Completed, pending, excluded, in progress

Folders page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Manage	Create, manage, or delete folders.
	Exclusion	Exclude a folder from deduplication.
	Set properties	Set NAS directory properties.
	Statistics	Show deduplication and replication statistics for a NAS share.

Manage folders and shares

You can perform the following functions for folders and shares:

- Create a NAS share or folder - Refer to ["Create a NAS share"](#) for more information.
- Update client settings - Refer to ["Create a NAS share"](#) for more information.
- Delete a NAS share or folder if it does not have any files or subfolders.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Folders* tab, select a share, click the *Manage* icon, and select the action you want to perform.

Exclude/include a NAS folder or share from deduplication

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

To exclude a specific NAS share from deduplication, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab, select a share, click the *Exclusion* icon, and select *Exclude Deduplication*.

To include a specific NAS share that was previously excluded from deduplication, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Shares* tab, select a share, click the *Exclusion* icon, and select *Include Deduplication*.

Set properties for a NAS folder or share

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

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Folders* tab, select a share, and click the *Set Properties* icon. Refer to ["Set properties for a NAS share"](#) for more information.

Display statistics for a NAS folder or share

You can display deduplication and replication statistics, if enabled. This includes data size, number of files and folders, number of deduplicated files, space used by deduplicated files, number of files awaiting deduplication, space used by files awaiting deduplication, number of replicated files, number of files awaiting replication, space use by files awaiting replication, replicated data size, and size of unique replicated data.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Folders* tab, select a share, and click the *Show Files* icon.

SMB clients - Users

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

The screenshot shows the FalconStor StorSafe web interface. The top navigation bar includes 'Dashboard', 'Monitor', 'Manage' (selected), 'Analytics', and 'Administration'. The left sidebar shows storage metrics for various clients. The main content area has tabs for 'Server Overview', 'Virtual Libraries', 'Deduplication', 'Jobs', 'Object Storage', 'NAS' (selected), 'Host Clients', 'Physical Resources', and 'Settings'. Under the 'NAS' tab, there are sub-tabs for 'NAS Resources', 'Shares', 'Folders', 'SMB Clients' (selected), 'NFS Clients', 'Replication', and 'Integrity Check'. The 'SMB Clients' sub-tab is active, showing a table of users and their assigned NAS shares.

User	ID	Description	Groups	Customer
Administrator	1	Built-in account for administering...		
DefaultAccount	3	A user account managed by the s...		
falcon101	7			
Guest	2	Built-in account for guest access ...		
henry	6			

Below the user table, there is a 'User' drop-down menu set to 'henry' and a 'NAS Shares' section. The 'NAS Shares' section displays a table of shares assigned to the selected user.

Path	Access	Share Name	Description	Max Connections
/nas/NAS-00254/fds/share2	Read-Write	share2	share2	0
/nas/NAS-00254/fds/test2	Read-Write	test2		0

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

- User name
- ID
- Description
- List of groups to which this client belongs

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

- Full shared directory path
- Access rights - read-write or read-only
- Share name
- Share comments
- Maximum number of user connections

SMB User page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Create	Create an SMB user.
	Update	Update the description or the password for an SMB user or update share access rights.
	Delete	Delete an SMB user.
	Manage	Set authentication mode, set SMB options, display SMB connection status. For domain authentication, you can also update the domain controller, sync SMB clients, set Windows ACLs.

Create an SMB user

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

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, and click the *SMB Clients* tab.
2. Click the “+” (*Create*) icon.



The image shows a 'Create SMB Client' dialog box with the following fields and controls:

- Name***: Text input field containing 'Userabc'.
- Password***: Password input field with masked characters (dots).
- Confirm Password***: Text input field containing 'Retype New Password'.
- Description**: Text input field containing 'Comment'.
- Buttons**: 'Create' and 'Cancel' buttons at the bottom right.

3. Enter information about the SMB user.

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

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

Description - Optionally, specify a description for the Windows user up to 256 characters without consecutive colons.

4. Click *Create* when done.

Update an SMB user

You can update the description and change the password for an SMB user if the authentication mode is *User*.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, select a user, and click the *Update* icon.

Delete an SMB user

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

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, select a user, and click the *Delete* icon.

Set the authentication mode

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

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

Notes:

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

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, click the *Manage* icon, and select *Set Authentication Mode*.

2. Enter information on the *General* tab.

The screenshot shows the 'Set Authentication Mode' dialog box with the 'General' tab selected. The 'Authentication Mode' is set to 'Domain'. The 'Server Description' is 'VTL Server 192'. The 'Primary Domain Controller*' and 'Secondary Domain Controller' are both set to 'Domain Controller'. The 'Windows ACL' is set to 'Disabled'. The 'User Name*' is 'superadmin' and the 'Password*' is masked with dots. The 'Kerberos User Name' is 'Login name' and the 'Kerberos Password' is 'Password'. There are blue question mark icons next to the 'Authentication Mode', 'Server Description', 'Primary Domain Controller*', 'Secondary Domain Controller', 'User Name*', and 'Kerberos User Name' fields. The 'Set' and 'Cancel' buttons are at the bottom right.

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

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

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

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


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

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

Kerberos User Name/Password - (Domain mode) Optionally, specify a user name and password with AD administrative privileges when Kerberos

authentication is required for NAS clients. This user name/password is used one time to create a computer account for the server in the domain so the server can become a trusted member in AD. By default, the server tries to join the domain as an AD member server and if that fails, it will try to join the domain as a legacy Windows NT 4 server.

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




The image shows a screenshot of the 'Set Authentication Mode' dialog box, specifically the 'UID/GID Range' tab. The dialog has three tabs: 'General', 'UID/GID Range', and 'Organizational Units'. The 'UID/GID Range' tab is active. It contains two sections: 'Reserved User IDs (UID)' and 'Reserved Group IDs (GID)'. Each section has a '+ Add' button, a '- Remove' button, and a search field labeled 'Search for...'. Below each section is a table with three columns: 'Begin', 'End', and 'Count'. The 'Reserved User IDs (UID)' table shows a range from 45436 to 65435 with a count of 20000. The 'Reserved Group IDs (GID)' table also shows a range from 45436 to 65435 with a count of 20000. At the bottom right of the dialog are 'Set' and 'Cancel' buttons.

Begin	End	Count
45436	65435	20000

Begin	End	Count
45436	65435	20000

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

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



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

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

5. Click *Set* when done.

Update the domain controller

(Domain mode) To update the primary or secondary domain controller and specify a new user/password, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, click the *Manage* icon, and select *Update Domain Controller*. Refer to ["Set the authentication mode"](#) for more information.

Sync SMB Clients

(Domain mode) To synchronize SMB clients with Active Directory, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, click the *Manage* icon, and select *Sync SMB Clients*.

Set Windows ACLs

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

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

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, click the *Manage* icon, and select *Set Windows ACL*. Select users with administrative roles to support ACL permissions of users on NAS resources.

Set SMB options (updated October 2024)

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, click the *Manage* icon, and select *Set SMB Options*.
2. Enter information on the *General* tab.

Set SMB Options

General | UID/GID Range | Admin Users

Server Description: VTL Server 139

Alias Names: as

Domain Sync Interval: 40 Minutes

Windows ACL: Enabled

Samba Options: +Add -Delete Search for...

Samba Option	Value
allow trusted domains	no
browseable	no
client ntlmv2 auth	yes
create mask	0771
deadtime	15
debug pid	yes

Double-click the option value to edit.

Set Cancel

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

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

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


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

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

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

Refer to ["Set the authentication mode"](#) for more information.

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



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

Display SMB connection status

To display active connections, shares, and locked files, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, click the *Manage* icon, and select *SMB Connection Status*.

Update access rights for SMB client shares

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

SMB clients - Groups

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

Group	ID	Description	Members	Customer
Allowed RODC Password Replic...	39	Members in this group can have t...		
Cert Publishers	28	Members of this group are permit...		
Cloneable Domain Controllers	43	Members of this group that are d...		
Denied RODC Password Replic...	40	Members in this group cannot ha...		
DnsAdmins	47	DNS Administrators Group		
DnsUpdateProxy	48	DNS clients who are permitted to...		
Domain Admins	29	Designated administrators of the...		

Path	Access	Share Name	Description	Max Connections
/nas/NAS-00254/fds/share2	Read-Write	share2	share2	0

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

- Group name
- ID
- Description
- List of users in the group

Select a group from the *Group* drop-down box to view properties for each NAS share assigned to the group:

- Full shared directory path
- Access rights - read-write or read-only
- Share name
- Share comments
- Maximum number of connections

SMB Groups page icons


The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Create	(User mode) Create an SMB group.
	Update	(User mode) Update the group description, specify group members, or update share access rights.
	Delete	(User mode) Delete an SMB group.

Create an SMB group

(User mode) To create an SMB group:

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, and click the *Groups* tab.
2. Click the “+” (*Create*) icon.



The dialog box titled "Create SMB User Group" contains three input fields: "Name*" with a placeholder "Name", "Members*" with a placeholder, and "Description" with a placeholder "Comment". At the bottom right are "Create" and "Cancel" buttons.

3. Enter information about the SMB group.

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

Members - Specify users to include in this group.

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

Update an SMB group

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

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, click the *Groups* tab, select a user, and click the *Update* icon.

Delete an SMB group

(User mode) You can delete an SMB group.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *SMB Clients* tab, click the *Groups* tab, select a user, and click the *Delete* icon.

Update access rights for SMB group shares

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

NFS clients

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

The screenshot shows the FalconStor StorSafe web interface. The top navigation bar includes Dashboard, Monitor, Manage, Analytics, and Administration. The left sidebar shows storage metrics for various VMs. The main content area has tabs for NAS Resources, Shares, Folders, SMB Clients, NFS Clients, Replication, and Integrity Check. The NFS Clients tab is selected, showing a table of users and a section for NAS Shares.

User	ID	Description	Groups	Customer
Administrator	1	Built-in account for administering...		
DefaultAccount	3	A user account managed by the s...		
falcon101	7			
Guest	2	Built-in account for guest access ...		
henry	6			

Below the user table, there is a section for NAS Shares. A dropdown menu shows the selected user is 'henry'. The NAS Shares table is as follows:

Path	Access	Share Name	Description	Max Connections
/nas/NAS-00254/fds/share2	Read-Write	share2	share2	0
/nas/NAS-00254/fds/test2	Read-Write	test2		0

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

- Client user name
- ID
- Name or IP address of the client machine
- Description

Select a client from the *Client* drop-down box to view properties for each NAS share assigned to the client:

- Full shared directory path
- Access rights - read-write or readonly
- Permissions - Refer to ["Create a NAS share"](#) for more information


NFS Clients page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Create	Create an NFS client.
	Update	Update the name, machine, or comment for an NFS client or update share access rights.
	Delete	Delete an NFS client.

Create an NFS user

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, and click the *NFS Clients* tab.
2. Click the “+” (*Create*) icon.



The image shows a 'Create NFS Client' dialog box with the following fields and buttons:

- Name***: A text input field with the placeholder text 'Name'.
- Machine***: A text input field with the placeholder text 'Machine'. A blue question mark icon is located to the right of this field.
- Comment**: A text input field with the placeholder text 'Comment'.
- Create**: A blue button.
- Cancel**: A white button.

3. Enter information about the NFS client.
 - Name** - Specify a unique NFS client name up to 36 ASCII characters.
 - Machine** - Specify a unique resolvable name or IP address of the client machine. It can also be a subnet and netmask to grant NAS share access to an entire subnet using the following format: `subnet/netmask`, where `netmask` can be an IP or a prefix, for example, `192.168.0.0/255.255.255.0` or `172.30.123.0/24`.
 - Comment** - Optionally, specify a description for the NFS client up to 128 ASCII characters.
4. Click *Create* when done.

Update an NFS client

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *NFS Clients* tab, select a client, and click the *Update* icon. Refer to ["Create an NFS user"](#) for more information.

Delete an NFS client

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *NFS Clients* tab, select a client, and click the *Delete* icon. You can force the deletion if NAS shares are assigned to the client.

Update access rights for NFS client shares

To update the access rights for NFS clients, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *NFS Clients* tab, select a client in the *NAS Shares* section of the page, and click the *Update* icon. For each path, you can set permissions. Refer to ["Create a NAS share"](#) for more information.

Replication - Configuration

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

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

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

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

The screenshot displays the FalconStor StorSafe web interface. The top navigation bar includes 'Dashboard', 'Monitor', 'Manage' (selected), 'Analytics', and 'Administration'. The left sidebar shows a list of servers: 'storsafe1110', 'TJ-Test', 'VM22-133', 'VM22-139' (selected), and 'VTL01'. The main content area is titled 'NAS' and contains a sub-menu with 'NAS Resources', 'Shares', 'Folders', 'SMB Clients', 'NFS Clients', 'Replication' (selected), and 'Integrity Check'. The 'Replication' tab is active, showing 'Configuration' and 'Policies' sections. The 'Configuration' section displays 'OUTGOING REPLICATION' for target server 'VM22-139' with details: IP Address '172.22.22.133', Username 'root', Transmission IP '172.22.22.133', Encryption 'Enabled', and Global Replication Policy 'Disabled'. Below this, the 'REPLICATION SESSION' section shows 'No Session in Progress'. The 'INCOMING REPLICATION' section shows 'Incoming Replica' and 'Replica Resource' as 'Specific Resource NAS-00254, ID 254'. A table lists the source servers and their IP addresses:

Source Server	IP Address
VM22-133	172.22.22.133
StorSafe-S-VA22-192	172.22.22.192

Replication Configuration page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Outgoing replication	Enable outgoing replication, edit outgoing replication, set outgoing replication throttle, disable outgoing replication.
	Incoming replication	Configure incoming replica, purge replica files, view purge results.
	Refresh replication sessions	Refresh replication status.

Enable outgoing replication

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

- You must have two StorSafe servers with NAS enabled.
- You must have administrative rights on both servers.
- You must have enough space on the target server for the replica data.
- The target server must have been added into the portal.

In this section, replication applies to data on NAS resources that do not use OST. Replication of OST resources is handled by the FalconStor OpenStorage Option plug-in.

Replication of NAS shares and OST resources can co-exist. However, when you configure replication through FSOST, it will overwrite any existing outgoing replication configuration so that both types of replication will share the same parameters.

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, and click the *Replication* tab.

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

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

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

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

Global Replication Policy - Specify if you want to define a global replication policy that replicates all NAS directories instead of individual policies for some directories. For a global policy you will need to manually exclude directory paths that you do not want to replicate. A global policy is applicable only for NAS resources not using OST. In a multi-tenancy environment, the global replication policy does not apply to servers assigned in shared mode to customer domains.

- For a global replication policy, configure policy settings.

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

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

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

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

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

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

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

In the *Exclude Hours* field, you can:

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

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

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

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

5. Click *Enable* when done.

Edit the outgoing replication configuration

You can edit the outgoing replication configuration. Refer to ["Enable outgoing replication"](#) for more information.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, click the *Outgoing Replication* icon, and select *Edit Outgoing Replication*.

Set outgoing replication throttle

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

To set a throttle for outgoing replication, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, click the *Outgoing Replication* icon, select *Set Outgoing Replication Throttle*, and select *Throttle Replication*.

Click *Add* to specify the time range during which throttling should occur and the maximum number of Kbs per second that should be used for replication.

Transmission will not exceed the set value. This is a global server parameter and affects all resources. Valid input for bandwidth is 10-100,000,000 Kb/s.

Disable outgoing replication configuration

To disable outgoing replication, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, click the *Outgoing Replication* icon, and select *Disable Outgoing Replication*.

Configure incoming replica

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

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, and click the *Replication* tab.
2. Click the *Incoming Replication* icon and select *Configure Incoming Replica*.

Use Any Resource - Specify if you want the system to automatically select any available NAS resource for incoming replication the first time replication is started. You must select this option if you are using the FalconStor OpenStorage option.

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

3. Click *Configure* when done.

Purge replica files

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

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, and click the *Replication* tab.
2. Click the *Incoming Replication* icon and select *Purge Replica Files*.

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

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

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

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

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

3. Click *Start* when done.

View purge results

If you selected to save purge results when running *Purge Replica Files*, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, click the *Incoming Replication* icon, and select *Get Purge Results* to download a CSV file containing the purge files.

Refresh replication session status

To refresh the status of active replication session, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, click the *Refresh Replication Sessions* icon.

Replication - Policies

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

The screenshot shows the FalconStor StorSafe with StorSight interface. The left sidebar displays server details for VM22-139, including its IP address (172.22.22.139) and build number (12146). The main panel shows the 'Replication' tab under the 'NAS' section. The 'Policies' sub-tab is active, displaying a table of replication policies. The table has columns for Policy, ID, Method, Paths, Only Deduplicated Files, and Replicate Back. Below the table, the configuration for 'policy1-1' is shown, including its replication status (Idle) and schedule (Normal).

Policy	ID	Method	Paths	Only Deduplicated Files	Replicate Back
policy1-1	2	Schedule	/nas/NAS-00254/fds	☐	☐
policy2	3	Schedule	2	☒	☐
policy3	4	Manual	/nas/NAS-00254/fds/folder4	☒	☒
policy4	5	Schedule	2	☐	☐
policy5	6	Schedule	2	☐	☒
policy5-1	7	Schedule	2	☐	☒
policy6	8	Manual	/nas/NAS-00254/fds/folder4	☐	☐

Below the table, the configuration for 'policy1-1' is shown:

Policy: policy1-1

Replication Status: Idle

Schedule: Normal

Depending upon the columns that have been selected for display, the properties for each NAS replication policy are visible at a glance:

- Policy name
- Policy ID
- Replication method
- Paths to replicate (click for more information)
- Replication options in policy

Replication Policies page icons


The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Create	Create a NAS replication policy. Not valid if a global replication policy is configured.
	Update	Update a NAS replication policy. Not valid if a global replication policy is configured.
	Delete	Delete a NAS replication policy. Not valid if a global replication policy is configured.
	Manage	Start replication, synchronize files, suspend/resume replication schedule.

Create a NAS replication policy

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

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, and click the *Policies* tab.
2. Click the “+” (*Create*) icon.
3. On the *General* tab, specify policy properties.



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

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

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

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

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

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

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

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

4. For scheduled replication, on the *Exclusion* tab, select when replication should not occur.

The screenshot shows a 'Create Replication Policy' dialog box with the 'Exclusion' tab selected. The 'General' tab is also visible. The 'Exclusion' tab contains the following fields:

- Exclude Hours:** A text input field with a help icon (?) to its right.
- Exclude Days:** A text input field with a help icon (?) to its right.
- Exclude Weekdays:** A text input field.
- Exclude Month:** A text input field.
- Replication Window:** Two time pickers showing '00:00' and '23:59'.

At the bottom right of the dialog are 'Create' and 'Cancel' buttons.

In the *Exclude Hours* field, you can:

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

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

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

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

5. Click *Create* when done.

Update a NAS replication policy

You can update a NAS replication policy as long as replication of the policy is not in progress. You cannot update a global replication policy. Refer to [Auto Migration](#) for more information.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, click the *Policies* tab, select a policy, and click the *Update* icon.

Delete a NAS replication policy

You can delete a NAS replication policy as long as replication of the policy is not in progress. You cannot delete the global replication policy. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, click the *Policies* tab, select a policy, and click the *Delete* icon.

Start replication of a NAS policy

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

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, click the *Policies* tab, select a policy in the bottom section of the screen, click the *Manage* icon, and select *Start Replication*.

Synchronize files

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

Only one outgoing replication or file synchronization job can be active at a time. For a global replication policy, you can select a NAS resource, share, or folder to be synchronized instead of synchronizing all resources.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, click the *Policies* tab, select a policy in the bottom section of the screen, click the *Manage* icon, and select *Synchronize Files*.


Suspend/resume replication schedule

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

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Replication* tab, click the *Policies* tab, select a policy in the bottom section of the screen, click the *Manage* icon, and select *Suspend Replication Schedule* or *Resume Replication Schedule*.

Integrity Check

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



Integrity Check page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Configure	Enable, update, or disable integrity checking.
	Manage	Start or stop integrity checking.
	Refresh integrity status	Refresh status.

Configure integrity checking

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

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

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

Notes:

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

1. Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, and click the *Integrity Check* tab.
2. Click the *Configure* icon and select *Enable Integrity Check*.

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

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

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

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

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

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

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

The screenshot shows a dialog box titled "Enable Integrity Check" with a close button (X) in the top right corner. It has two tabs: "General" and "Exclusion". The "Exclusion" tab is selected. Inside the "Exclusion" tab, there are four input fields with labels to their left: "Exclude Hours", "Exclude Days", "Exclude Weekdays", and "Exclude Month". Each input field has a placeholder text that matches its label. To the right of the "Exclude Hours" and "Exclude Days" fields, there are blue circular help icons (question marks). At the bottom right of the dialog, there are two buttons: "Enable" (in blue) and "Cancel" (in white).

In the *Exclude Hours* field, you can:

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

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

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

4. Click *Enable* when done.

Update the integrity check configuration

You can update the integrity check configuration. Refer to ["Configure integrity checking"](#) for more information.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Integrity Check* tab, click the *Configure* icon and select *Update Integrity Check*.

Disable the integrity check configuration

You can disable the integrity check configuration.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Integrity Check* tab, click the *Configure* icon and select *Disable Integrity Check*.

Start/suspend integrity check

You can manually run an integrity check. If running, you can suspend it.

Select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Integrity Check* tab, click the *Manage* icon and select *Start* or *Suspend*.

Refresh status of integrity checking

To refresh the status of integrity checking, select *Manage* from the menu bar, click the *NAS* tab, select a StorSafe server, click the *Integrity Check* tab, and click the *Refresh Integrity Status* icon.

Physical Tape Export

Data on virtual tapes can be exported to physical tapes via:

- [Tape Caching](#) - Acts as a cache to your physical tape library and automatically triggers data export based on a schedule or policy.
- [Auto Archive](#) - Automatically writes data to physical tape whenever a virtual tape is moved to an Import/Export (IE) slot from a virtual library after backup.
- [Tape Duplication](#) - Automatically makes up to five duplicate copies of a physical tape whenever virtual tape data is exported to a local physical tape.
- [Remote Export](#) - Automatically exports the contents of a replicated virtual index tape (VIT) to a physical tape at the remote site.
- [Manual Export](#) - Manually export data from a virtual tape to a physical tape.

Tape Caching

Tape Caching enhances the functionality of StorSafe by acting as a cache to your physical tape library, providing transparent access to data regardless of its location.

With tape caching, tapes will always appear to be inside virtual libraries and will be visible to the backup application regardless of whether the data is actually on disk or on physical tape.

Tape Caching provides advanced flexibility that allows you to set up policies that automatically trigger data export to physical tapes based on criteria, such as the number of days that data has been on disk or the amount of used disk space.

With Tape Caching, you can not only determine which events will activate the action, but also when it will occur. For example, you can set the policy to export the data immediately or at a specific time or day. This enables data to be written to physical tapes as a background process without impacting production servers.

You can also set up a reclamation policy that allows you to specify how and when the data that has been exported to physical tape can be deleted from the disk to make space for new backups.

In order to use Tape Caching, you must enable the feature for your virtual library, set your export and reclamation policies, and create a cache for each of your physical tapes. This can be done when you create a virtual tape library. You can also create a tape caching policy for an existing virtual tape library for which Tape Caching was not previously enabled.

Notes:

- Tape Caching, Auto Archive, Auto Replication, and Cloud Tiering services are mutually exclusive so only one can be configured per virtual library.
- Tape Caching and inline deduplication are mutually exclusive.
- WORM tapes cannot be used in a Tape Caching configuration.

Tape caching policies

A tape caching policy contains the data export triggers and reclamation triggers for a virtual tape library. The tape caching policy affects how data will be read/written from/to tapes.

Scenario 1: Data on virtual tape. Data not written to physical tape.

If the data has not been written to physical tape, reads will be from the virtual tape. Writes will either append or rewrite the virtual tape.

Scenario 2: Data written to physical tape. Virtual tape not reclaimed.

If the data has been written to physical tape but is still retained on the virtual tape, reads will be from the virtual tape. Writes to the tape will either append or rewrite the virtual tape (and restarts the clock on the export policy).

Scenario 3: Data written to physical tape. Virtual tape reclaimed.

If the data has been written to physical tape and the virtual tape has been reclaimed, reads will be directly from the *direct link* tape. A *direct link* tape is not an actual tape but a link to a physical tape. If you overwrite the beginning of the tape, StorSafe will create a new virtual tape, which breaks the *direct link* tape and restarts the clock on the export policy. If you try to append to the tape, StorSafe will append data on the physical tape.

Create/update a tape caching policy

You can configure a tape caching policy when you create a virtual tape library or you can update an existing virtual tape library for which Tape Caching was not previously enabled. This is done on the *Tape Services* tab for a virtual tape library, which are located on the *Libraries* page.

Tape caching triggers

Triggers control when data in the cache will be copied to physical tape.

Note: Regardless of which triggers you set, there must be at least 1 MB of data on the tape in order to trigger data export.

For a *Schedule* trigger, specify when data export should actually occur.

The screenshot shows the 'Create Virtual Tape Library' dialog box with the 'Tape Services' tab selected. The 'Service' dropdown is set to 'Tape Caching'. The 'Trigger' dropdown is set to 'Schedule'. The 'Export Schedule' is set to 'Hourly' with a value of '1'. The 'Space Reclamation Policy' is set to 'Delete Cache'. The 'Space Reclamation Trigger' is set to 'Never'. There are 'Create' and 'Cancel' buttons at the bottom right.

- *Hourly* - Export data every n hours.
- *Daily* - (Default) Export data at a specific time of day. Specify the *Start Time*. Note that if the specified time has passed, the export will start at that time on the next day.
- *Weekly* - Export data on a specific day of the week. Specify the *Start Time*. Note that if the specified time has passed, the export will start at the next scheduled day and time.

For a Policy trigger, determine what criteria will trigger export.

You can set multiple policy triggers:

- **Data Age** - Export data when the data has been on the virtual disk for a specified number of hours or days. Specify the desired number of hours/days in the list box.
- **Disk Usage** - Export data when the used disk space exceeds the specified disk capacity. The actual percentage is a global variable which is set for each virtual tape library in the *Global Settings* on the *Settings* tab for StorSafe servers.
- **Tape Eject** - Export data when a backup has completed and the virtual tape is ejected to a slot. If you select the option *Copy ejected tape when full*, the copy will only occur if the tape is full.

Relation Between Trigger Policies - If multiple triggers are set, select *And* if all the triggers must be met to initiate data copy or select *Or* if meeting any one of them will initiate data copy. For example, if you select both *Data Age* and *Disk Usage*, and you select *And*, data copy will occur only when both the specified number of days/hours has elapsed and the specified disk capacity has been reached. If you select *Or*, the occurrence of either one of those events will trigger the data copy.

Delay Copy Until - Copy will be delayed until the time you specify after one of the above policies has been triggered. You may want to select a time when system usage is very light.

Space reclamation policy

A *Space Reclamation Policy* determines the method to use to free up cache disk space after data has been copied to physical tape.

Delete cache - After the cache is deleted, tapes become *direct link* tapes. A direct link tape is not an actual tape but a link to a physical tape. If your backup application ever overwrites a direct link tape, StorSafe will automatically start caching the physical tape. Direct link tapes are only deleted if you disable tape caching on the library or if you delete the tape. If you unassign a physical library, the direct link tape will not be deleted.

Deduplicate Data - Deduplication is triggered and single instances of unique data are copied to the deduplication repository. After deduplication, the virtual tape is replaced with a virtual index tape (VIT) pointing to deduplication storage. In order to use this feature, you must have a deduplication cluster associated with this StorSafe server. After selecting this option, you will be prompted to create a deduplication policy. By default, this policy will have the same name as the tape library and all tape caching tapes created under this library will automatically be added to the policy. Regular virtual tapes will not be added to the policy and cannot be added manually to this policy.

Convert deduplicated tape to direct link - VIT tapes will become direct link tapes after the specified number of days since the last successful copy. If the tape is in the middle of being deduplicated, the system will wait until deduplication completes.

Note: A physical tape can be ejected to a mail slot by a backup application only if it is associated with a direct link tape that was created using the *Delete cache* method. If the physical tape is associated with a tape reclaimed using the *Deduplicate Data* method, it cannot be ejected by the backup application.

Space reclamation triggers

Space reclamation triggers determine when to trigger reclamation.

- **Immediate** - Cache disk space is freed up as soon as the data copy is complete.
- **Space Usage** - (Only available for *Delete cache*) Cache disk space is freed up when the used space reaches the specified threshold. The actual percentage is a variable which is set for each virtual tape library on the *Settings* tab for virtual tape libraries.
- **Retention Period** - Cache disk space is freed up after a specified number of days has elapsed. Specify the number of days that the data should be retained.
- **Never** - (Only available for *Delete cache*) Cache disk space is never freed up.

Create a cache for your physical tapes

Tape caching stores data on disk before it is migrated to physical tape. In order for this to happen, you must create a cache for each of your physical tapes. This is typically done after you create a virtual tape library. If this has not been done, you must sync the library to create a cache for your physical tapes. You also need to sync the library if a direct link tape was deleted and the associated physical tape is needed to recover data.

1. Select **Manage** from the menu bar, select a StorSafe server, click the *Virtual Libraries* tab, click the **Manage** icon, and select *Create Tapes via Tape Caching Sync*.

Create Tapes via Tape Caching Sync

Virtual Tape Library: ADIC-Scalar 100-00153ee

Physical Tape Library*:

Physical Tapes*: ?

Mode: ☒ Cache ☐ Metadata Copy ☐ Direct Link ?

Tape Name Prefix:

Tape Initial Size:

Physical Devices: ?

Encrypt Data: ☐

Create Cancel

2. Select the appropriate physical library.

A physical tape library can only be synchronized to one virtual tape library at a time.

3. Select the physical tape(s) for which you want to create a cache.

4. Select a mode for creating the cache tapes.

Metadata Copy - Create cache tapes with the same header metadata as the physical tapes. Select this option if your backup application requires a tape header to identify a tape.

Direct Link - Create direct link tapes, which are not actual tapes but are links to directly access physical tapes.

5. (Not valid for direct link tapes) If applicable, specify a tape name prefix and initial size (default 5 GB).

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

6. (Not valid for direct link tapes) Optionally, specify the physical devices to use.

If you do not specify a device, the first available physical device will be used.

7. Indicate if you want to encrypt the data on tapes.

You can select this option only if at least one key has been created. If you select this option, you must select the key to use. All the data on the tape will be indecipherable until is imported back to a virtual tape library and decrypted using the same key. For more information about encryption, refer to [“Manage encryption keys”](#).

Create uncached virtual tapes

Even though you are using Tape Caching for your tape library, you can still create *uncached* virtual tapes that will not be migrated to physical tapes. This can be useful for a single backup that is not part of your normal backup routine.

Select *Manage* from the menu bar, select a StorSafe server, click the *Virtual Libraries* tab, click the *Manage* icon, and select *Create Tapes*. Refer to [“Create virtual tapes”](#) for more information.

Note that if you create virtual tapes, they cannot match the barcodes of your physical tapes.

Auto Archive

Auto Archive writes data to physical tape whenever a virtual tape is moved to an IE slot (the vault) from a virtual library by a backup application or other utility after a backup. In order to use *Auto Archive*, the physical tape library or standalone tape drive must support barcodes because when StorSafe attempts to export to physical tape it must find a matching barcode.

When there are no physical tape libraries, the first available tape drive will be used. A physical tape must be manually loaded in the physical standalone tape drive.

Note: Tape Caching, Auto Archive, Auto Replication, and Cloud Tiering services are mutually exclusive so only one can be configured per virtual library.

You can configure Auto Archive when you create a virtual tape library or you can update an existing virtual tape library for which Auto Archive was not previously enabled. This is done on the *Tape Services* tab for a virtual tape library, which are located on the *Libraries* page. You must also configure Auto Archive for virtual tapes in a library with Auto Archive enabled.

For the library, determine if you want virtual tapes moved (removed) or copied (retained) after the data is transferred to physical tape. If you select *Move*, indicate how long to wait before deleting it.

Indicate if you want to eject your *physical* tapes to the library's import/export slots after archiving.

You can encrypt the data while exporting as long as at least one key has been created. Refer to [“Create a key”](#) for more information.

The screenshot shows the 'Create Virtual Tape Library' dialog box with the 'Tape Services' tab selected. The 'Service' dropdown is set to 'Auto Archive'. The 'Mode' section has 'Move' and 'Copy' buttons, with 'Move' being the active selection. The 'Delay Tape Removal' is set to '1' with a unit of 'Days'. The 'Eject Physical Tape' and 'Encrypt Data' options are both disabled (unchecked).

For virtual tapes, specify the physical library and physical tapes for which you want to create matching virtual tapes.

Create Tapes

General

Service

Service

Match Auto Archive

Remote Physical Library

Library ID

Physical Tapes

Barcode

Create

Cancel

Tape Duplication

Tape Duplication allows you to make up to five duplicate copies of a physical tape whenever virtual tape data is exported to a local physical tape. The number of physical libraries controls the number of duplicate copies; you need six physical libraries to make five duplicates.

You must have at least two physical tape libraries (same model, same number of drives, same tapes with the same barcodes) connected to your local system in order to enable Tape Duplication for this library.

You can configure Tape Duplication when you create a virtual tape library or you can update an existing virtual tape library for which Tape Duplication was not previously enabled. This is done on the *Tape Duplication* tab for a virtual tape library, which are located on the *Libraries* page.

You can select whether the physical tape will have the same barcode as the virtual tape or will use a specific prefix that replaces the first character of the barcode. For example, if your virtual tape barcode is “123456” and you specify that the prefix is “A”, the system will look for a physical tape with the barcode “A23456”.

The duplication job will look for a tape with the correct barcode in one of the physical libraries. If one is found, the data is duplicated to that physical library. If an appropriate tape is not found, but there are additional physical libraries, the system will continue to look for a match.

When data is exported, separate export jobs will be created for each physical library and each job will have a unique job ID. Multiple duplication jobs will run concurrently.

If this library is using Tape Caching or if you selected the *Move* option for Auto Archive on the *Tape Services* tab, the virtual tape data will not be deleted until the duplication job finishes successfully.

Notes:

- You should not have duplicate physical tape barcodes in your system *unless* you are using tape duplication.
- Once you configure tape duplication, you should not unassign any of the physical libraries from StorSafe.
- If this library is using Tape Caching, once the virtual tape has been reclaimed and a direct link tape exists, tape duplication will no longer occur and only the primary physical tape can be updated.

Remote Export

Remote Export automatically exports the contents of a replicated virtual index tape (VIT) to a physical tape at the remote site, providing a copy of a physical tape at the remote site.

If remote export is enabled for a virtual tape, the virtual tape must be added into a deduplication policy.

Once replication is successfully completed, an export job is submitted at the replica site.

If a deduplication policy includes cascaded replication, the tape will be added to the policy only if the target server is the same as one of the two target servers in the policy. This tape will be remotely exported to physical tape(s) on the target servers only if physical tapes with the same barcodes exist on the target servers.

The target server you select must have an available physical tape library and the media type must match that of the virtual tape library.

You must also select a physical library and physical tapes on the target server that should be used to synchronize barcodes. The number of virtual tapes that will be created will be the same as the number of physical tapes selected.

Manual Export

You can manually export data from a virtual tape to a physical tape in a physical tape library or standalone drive.

For deduplicated tapes that have been replicated, exporting a tape is useful for disaster recovery purposes because data can be accessed from the replicated tape. Whenever a deduplicated tape is exported, a complete physical tape is created. This means that the entire, original data set (prior to deduplication) is written to the physical tape. This applies to all VITs, pure or mixed.

Because some third-party backup applications alter what they write to the tape depending on the type of cartridge used, you can only export tapes to *like* media. You cannot export to a dissimilar physical tape. This guarantees that the backup application will accept the tape as valid; from the backup application's point of view, there is no difference between the virtual and physical tape.

Exporting a WORM tape requires a physical WORM tape.

You should inventory the physical tapes in your library before you begin. Refer to ["Inventory a physical library"](#) for more information.

Export with stacking

While exporting, you can stack tapes, which allows multiple virtual tapes to be exported to a single physical tape in a physical tape drive or library. Stacking maximizes physical tape usage and allows the conversion of virtual media with a smaller capacity to physical media with a higher capacity (i.e. DLT to LTO).

Tape stacking is a tool for archival purposes only. Backup applications will not have direct access to stacked tapes. Restoring from a stacked tape will require importing data from a stacked tape back to virtual tape.

You can use tape stacking only if you are not currently using Tape Caching on the selected virtual tape library.

When exporting to a physical tape library, if you want tape stacking to append data to a physical tape (rather than overwrite the data on that physical tape), you must scan the tape before exporting.

When exporting to a standalone physical tape drive, you must prepare the tape to assign it a barcode before exporting. If you want to append data to a physical tape in the drive, you must retrieve the barcode before exporting.

You can export virtual WORM tapes and non-WORM tapes to a physical WORM tape. The non-WORM tapes will inherit the WORM properties of the physical tape. If they are imported, they will become virtual WORM tapes.

You cannot export a virtual WORM tape to a non-WORM physical tape.

To manually export data:

1. Select *Manage* from the menu bar, select a StorSafe server, click the *Vault* tab, select one or more tapes, and click the *Export* icon.

Export Virtual Tapes to Physical Tapes

Target: **Physical Library** Standalone Drive

Physical Tape Library*: ADIC:Scalar 100-00423

Export Mode: **Move** Copy

Delay Days: 365 Days

Eject Physical Tape: ☐

Encrypt Data: ☐

Tape Duplication: ☐

Stack Tapes: ☐

Virtual tapes to export are listed below:

Virtual Tape	Barcode	Same Barcode	Physical Tape
VirtualTape-00051	05Q00001	☒	(Slot:0) 00640000
VirtualTape-00052	05Q00002	☒	(Slot:5) 00640001

By default, a virtual tape is exported to a physical tape that has the same barcode as the virtual tape barcode unless you select a specific physical tape barcode and slot number.

Export Cancel

2. Select the local physical tape library/drive to which you want to export.
By default, a virtual tape is exported to a physical tape with the same barcode. Alternatively, you can select a tape in another slot.
3. Select how you want the data exported and if you want the physical tape exported to the import/export slots.

Move - Copies the contents of the virtual tape to its physical counterpart and then removes the virtual tape from the system. Specify a grace period if you want to keep the virtual tape for a time before deleting it. The default is 365; 0 means immediate tape removal. If you enable *Tape Duplication* below, the virtual tape data will not be deleted until the duplication job finishes successfully.

Copy - Copies the contents of the virtual tape to its physical counterpart and retains the virtual tape after the data is transferred.

Eject Physical Tape - Move physical tapes to I/E slots after exporting.

4. Indicate if you want to encrypt the data on the tape.

You can select this option only if at least one key has been created and you must select the key to use. All data on the tape will be indecipherable until it is imported back to a virtual tape library and decrypted using the same key.

You can use a single key to all virtual tapes when you export them or you can create a unique key for each one. Creating multiple keys provides more security; in the unlikely event that a key is compromised, only the tapes that use that key would be affected. However, if you use multiple keys, you must keep track of which key applies to each tape so that you use the correct key to decrypt the data when you import the physical tape back to virtual tape. For more information about encryption, refer to [“Manage encryption keys”](#).

5. Indicate if you want to use tape duplication.

Tape Duplication makes a duplicate copy of the physical tape when data is exported. You must have at least two identical physical libraries (same model, same number of drives, same tapes with the same barcodes). When data is exported, separate export jobs will be created for each physical library and each job will have a unique job ID. Refer to [“Tape Duplication”](#) for more information.

6. Select the virtual tape(s) you want to export.

If your physical tape drive/library uses barcodes, we highly recommend that you select to use the same barcode as the physical tape.

If your physical tape drive/library does not use barcodes, you can then select which slot to use for the physical tape.

7. Indicate if you want to stack tape data.

Select the physical tape to which you want to export and specify if you want to append or overwrite data on the physical tape.

The default is to append data to the physical tape unless you select *Overwrite Physical Tape*.

8. Click *Export* when done.

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

Encryption

Encryption can be used to ensure that tape data is confidential and secure.

Encryption can protect:

- Deduplicated data stored in the deduplication repository
- Data backed up on virtual tapes
- Data exported to physical tapes

Encryption utilizes the Advanced Encryption Standard (AES) 256-bit key CBC algorithms (Secure Tape) published by the National Institute of Standards and Technology (an agency of the U.S. government) and is FIPS-140-2-compliant.

With deduplication repository encryption, deduplicated data from all cluster servers is encrypted in the deduplication repository. Data in the repository will not be accessible unless the activation password has been entered.

To enable deduplication repository encryption, encryption must be enabled from the [“Encryption Support”](#) icon on the *Settings* page. Once this is done, you will have an option to enable deduplication repository encryption when you create your deduplication repository. Refer to [“Deduplication repository”](#) for more information.

Once *Encryption Support* is enabled, there will be an *Encryption* page where you can configure virtual tape encryption as well as create encryption keys

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 services are started on the server. 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. The Tape Database must exist in order to enable virtual tape encryption.

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.

Encryption page



Select *Manage* from the menu bar and click the *Encryption* tab. This page shows whether virtual tape encryption or deduplication repository encryption are enabled on the selected StorSafe server. It also shows if encryption has been activated. From this page you can access the *Encryption Keys* page.

Note that you will only see this page once *Encryption Support* is enabled on the *Settings* page.

Only a Super Administrator can access the *Encryption* page.

The screenshot displays the FalconStor StorSafe web interface. The top navigation bar includes 'FALCONSTOR', 'Dashboard', 'Monitor', 'Manage' (active), 'Analytics', and 'Administration'. A notification bell shows 59 alerts. The left sidebar lists servers: tssir03, VM22-139, VM22-140, VTL01 (selected), and VTL1002. The main content area is titled 'Encryption' and shows settings for VTL01. It includes tabs for 'Settings' and 'Encryption Keys'. Under 'Settings', the following encryption status is shown:

Setting	Status
Encryption Activation	Disabled
Virtual Tape Encryption	Enabled
Deduplication Repository Encryption	Enabled

Encryption page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Activate encryption	Enter the password to activate encryption for the server.
	Change password	Change encryption password.
	Enable tape encryption	Enable virtual tape encryption.

Enable virtual tape encryption for a StorSafe server

The Tape Database must exist in order to enable virtual tape encryption, which can be enabled as follows:

1. Enable *Encryption Support* on the *Settings* page.
Refer to [“Encryption Support”](#) for more information.
2. Select *Manage* from the menu bar, select a StorSafe server, and click the *Encryption* tab.
3. Click the *Enable Tape Encryption* icon.
4. Create the encryption activation password that will need to be entered each time the services are started on the server.
You must also enter a hint (0–32 characters) to help you remember the password.
If deduplication repository encryption is already enabled, you need to enter the same encryption password for virtual tape encryption.
5. Create one or more encryption keys.
Refer to [“Create a key”](#) for more information.
6. Enable encryption for virtual tape libraries.

You can enable encryption when you create a new virtual tape library.

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. 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.

7. Activate encryption for the server.

This is necessary in order to use and load encrypted tapes. Refer to [“Activate encryption for a server”](#).

Activate encryption for a server

In order to use and load encrypted tapes, encryption must be activated for the server. To activate encryption, select *Manage* from the menu bar, select a StorSafe server, click the *Encryption* tab, and click the *Enable Tape Encryption* icon. You will need to provide the current encryption activation password.

.Change the encryption activation password for a server


To change the encryption activation password, select *Manage* from the menu bar, select a StorSafe server, click the *Encryption* tab, and click the *Change Password* icon. You will need to provide the current encryption activation password.

Encryption keys



Select *Manage* from the menu bar, click the *Encryption* tab and then the *Encryption Keys* tab. Note that you will only see this page once *Encryption Support* is enabled on the *Settings* page.

Only a Super Administrator can access the *Encryption Keys* page.



Encryption Keys page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Create	Create an encryption key.
	Update	Update key name, password, and password hint.
	Delete	Delete the selected encryption key.

Manage encryption keys

Keys can be created and managed for virtual tape encryption and physical tape export encryption. Deduplication repository encryption uses an internal key that is created when encryption is enabled; it is not visible here.

Create a key

Each key can be used for both virtual and physical tapes.

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 the key.

To create a key:

1. Select *Manage* from the menu bar, select a StorSafe server, click the *Encryption* tab and then the *Encryption Keys* tab.
2. Click the *Create* icon.

In the *Key Name* text box, type a unique name for the key (1–32 characters).

In the *Secret Phrase* and *Confirm Secret Phrase* text boxes, 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.

In the *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: ~!@#\$%^&*()-_+=\|[]{};:","<.>/?

In the *Password Hint* text box, type a hint (0–32 characters) to help you remember the password.

3. Click *Create* when done.

Update 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.

Note that an encryption key used for Auto Archive export of a library cannot be modified.

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, select *Manage* from the menu bar, select a StorSafe server, click the *Encryption* tab, click the *Encryption Keys* tab, select a key, and click the *Update* icon. You will need to provide the current encryption activation password.

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, select *Manage* from the menu bar, select a StorSafe server, click the *Encryption* tab, click the *Encryption Keys* tab, select a key, and click the *Delete* icon. You will need to provide the current encryption activation password.

Server Settings

Select *Manage* from the menu bar for most of your StorSafe server management needs. Click the *Settings* tab for quick access to server administration tasks as well as global options.

You must be an StorSight Super Administrator or a domain Administrator (in exclusive mode) to perform most administrative functions.

Failover

Select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab. Click the failover icon to configure/manage your failover environment.

From here, you can check storage consistency, sync standby devices, set the storage exception list, manage power control, etc.

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

Manage deduplication

For pure StorSafe servers (not StorSafe-S), select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, and click the *Deduplication* icon.

From here, you can:

- Associate a StorSafe server with an SIR server
- Discover deduplication data disks
- Dissociate a StorSafe server from a deduplication cluster

Associate a StorSafe server with an SIR server

(Required for all StorSafe servers that will use deduplication) Associating a StorSafe server with an SIR server enables deduplication. This cannot be performed if the server has been taken over by its failover partner.

Select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, and click the *Deduplication* icon.

Enter the IP address, user name, and password of any server in the deduplication cluster. You can select any SIR server in a cluster. Multiple StorSafe servers can be associated with the same SIR server.

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

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

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

Notes:

- You should not associate a StorSafe server while any I/O is running on the StorSafe system. As part of the configuration, the SIR server will rescan all of the Fibre Channel HBA ports on the system and this will cause the I/O jobs running on the clients to fail.
- If you get a connectivity error message, check the storage connectivity between your StorSafe and SIR servers and then rescan.

Discover deduplication data disks

Select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, click the *Deduplication* icon, and click the *Discover deduplication data disks* icon to discover newly added data disks to the SIR server. This operation can take a while to complete and might disturb I/O activity.

Dissociate a StorSafe server from a deduplication cluster

Select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, click the *Deduplication* icon, and click the *Dissociate from Deduplication Cluster* icon to disable deduplication.

This cannot be performed if failover has occurred (the server has been taken over by or has taken over its failover partner). This operation is not applicable to the passive server in an active-passive failover configuration, or to a server that is in an active-active failover configuration with a StorSafe-S server.

Restart network

Restarting network services will bring all network interfaces down and then up. Communication with the server will be temporarily unavailable.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Restart Network* icon.
3. Confirm by selecting the *Restart Network Services* checkbox and clicking *Restart*.

Restart services

Restarting services will bring all services down and then up. Communication with the server will be temporarily unavailable.

1. Select *Manage* from the menu bar and select the server.
2. Click the *Settings* tab and click the *Restart Services* icon.
You can select to restart all modules except FC HBA modules.
3. Confirm by selecting the *Restart Services* checkbox and clicking *Restart*.

Reboot

Rebooting the server will power the server off and then on. Communication with the server will be temporarily unavailable.

1. Select *Manage* from the menu bar and select the server.
2. Click the *Settings* tab and click the *Reboot* icon.
3. Confirm by selecting the *Reboot Server* checkbox and clicking *Reboot*.

Shutdown

Shutting down the server will power the server off. Communication with the server will be lost.

1. Select *Manage* from the menu bar and select the server.
2. Click the *Settings* tab and click the *Shutdown* icon.
3. Confirm by selecting the *Shutdown Server* checkbox and clicking *Shutdown*.

Take X-ray

An X-ray is a diagnostic tool used by your Technical Support team to help solve system problems. Each X-ray contains technical information about your server, such as messages from the system event log 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.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *X-ray* icon.
3. Select *Include Core Files* to retrieve memory dump files when programs fail and generate core files.
4. Click *Take X-ray*.
5. When prompted, click *Take X-ray*.

Download event log


The server retains information about the health and behavior of the physical and virtual resources on the server. The event log records these system events and errors.

There are four types of events:

- Informational (I) - Day-to-day activities, such as when replication was completed, require no action.
- Warning (W) - Indicates that something occurred that may require maintenance or corrective action; however, the system is still operational.
- Error (E) - Indicates that a failure has occurred. For example, a device was not available or an operation failed.
- Critical (C) - Issues that may stop the system from operating properly.

To download the event log:

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Event Log* icon.



3. Specify if you only want events since a specific date and time or turn off *Get Latest Events* and specify a starting and ending date and time.
4. Click *Get Event Log*.
5. When prompted, click *Download Event Log*.

The file format is csv and can be opened in an application that supports csv, such as Microsoft Excel.

Manage Patches

This section describes how to apply and roll back StorSafe server patches from the portal.

1. To apply a patch, download it onto the machine running the browser or a location accessible from that machine.

Patches can be downloaded from the FalconStor support community sites.

2. Select *Manage* from the menu bar and select a StorSafe server.
3. Click the *Settings* tab and click the *Patches* icon.
4. Select the appropriate function.

To apply a patch, locate the patch file and click *Upload Patch*. The patch will be copied to the server and installed.

To roll back a patch, select the patch and click *Roll Back Patch*.

Manage network settings

Use *Network Settings* if you need to change StorSafe server network information. Using third-party utilities will not update the information correctly.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Network Settings* icon.

3. Click *Change Global Settings* and make configuration changes.

Domain - Internal domain name.

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

Default Gateway IP Address - IP address of your default gateway.

Enable SSH - Use Standard Secure Shell (SSH) for remote connection to the server.

NTP Servers - Allows you to keep the date and time of your StorSafe server in sync with Internet NTP servers. Click the “+” icon to add IP addresses.

Restart NTP - Restart the NTP service for time synchronization after changing NTP servers.

DNS IP Address - Click the “+” icon to add IP addresses of your DNS servers.

Click *Change Interface Settings* and specify the following for each NIC:

- Specify if you are using dynamic (*DHCP*) addresses.
- *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.

Note: If the MTU is changed from 9000 to 1500, a performance drop will occur. If you then change the MTU back to 9000, the performance will not increase until the server is restarted.

- *IP Address* - Static IP address
- *Netmask* - Network mask

4. Click *Change Network Settings* when you are done.

Configure network bonding

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

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


You can think of this group as a single virtual adapter that is made up of multiple physical adapters. To the system and the network, it appears as a single interface with one IP address.

However, throughput is increased by a factor equal to the number of adapters in the group. Also, network bonding detects faults anywhere from the NIC out into the network and provides dynamic failover in the event of a failure.

Notes:

- Network bonding cannot be set if failover is configured.
- Network bonding is not applicable for virtual appliances.
- You cannot use network cards with different speeds (e.g., 1Gb and 10Gb) in the same bonding group.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Network Bonding* icon.
3. Click *Create Network Bond*.
4. Click *Add Group*.



The dialog box titled "Add Network Bonding Group" contains the following fields and buttons:

- Bonding Mode***: A dropdown menu with "(Select)" as the current selection.
- IP Address***: A text input field with the placeholder text "(Enter IP Address)".
- Netmask***: A text input field with the placeholder text "(Enter Netmask)".
- NICs***: A dropdown menu with "(Select)" as the current selection.
- Buttons**: "Add" (blue) and "Cancel" (white) buttons located at the bottom right.

5. Select the bonding mode.

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

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

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

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

6. Enter an IP address and netmask to represent the bonding group and select which NICs to include.
7. Click *Add*.

Depending upon how many NIC ports you have, you can create multiple bonding groups.

You can also update and delete bonding groups from here.

8. Click *OK* when done.


Network services will be restarted.

Configure VLAN tagging

Set Virtual LAN (VLAN) Tagging to associate network interfaces with VLAN IDs defined on network switches. This identifies the VLAN to which data belongs and determine which ports to use for communication when a VLAN spans multiple switches.

VLAN tagging cannot be set if failover is configured.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *VLAN Tagging* icon.
3. Click *Set VLAN Tagging*.
4. Click *Add*.

A screenshot of the 'Add VLAN Tagging' dialog box. The dialog has a title bar with a close button. Inside, there are several fields: 'VLAN' with the value 'eth', 'NIC*' with a dropdown menu showing '(Select)', 'VLAN ID' with a text box containing '1' and a spinner icon, 'IP Address*' with a text box containing '(Enter IP Address)', 'Netmask*' with a text box containing '(Enter NetWork)', and 'MTU' with a text box containing '1500' and a spinner icon. At the bottom right, there are two buttons: 'Add' (blue) and 'Cancel' (white).

5. Select a NIC and specify the VLAN ID that was set in the network switch.
6. Specify the IP address, subnet mask, and maximum transfer unit (MTU), of each IP packet.

The MTU range is 1 to 9000. Set the MTU of each IP packet. If your network supports it, set this value to 9000 for jumbo frames.

7. Click *Add*.

You can also update and delete groups from here.

8. Click *OK* when done.

Network services will be restarted.

Change hostname

After changing the hostname the server modules will be restarted. The hostname cannot be changed if the server is part of a failover configuration.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Hostname* icon.

The hostname can be up to 32 characters. Valid characters are letters, numbers, or dash. The services will automatically restart when the hostname is changed.

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 will be marked offline and seen as foreign devices.

Configuration repository

The *Configuration Repository* maintains a continuously updated version of your system configuration. This is useful for disaster recovery purposes, such as if a server is out of commission but you have the storage devices and want to use them to build a new server. Failover configurations require a configuration repository. As a best practice, you should also create a configuration repository on a standalone server.

This feature works seamlessly with failover to provide business continuity in the event that a server fails. For additional redundancy, the configuration repository of a server can be mirrored to another device.

Create the configuration repository

To create a configuration repository, make sure there is at least 10 GB of continuous space since the repository cannot be spanned across multiple devices.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Configuration Repository* icon.
3. Click *Create Configuration Repository*.
4. Select a physical device for the configuration repository and click *Create*.

You can select multiple devices but the repository will be created on the first device with 10 GB of continuous space.

Manage the configuration repository

Select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, and click the *Configuration Repository* icon. Here you will see information about your configuration repository, including its size, status, and mirror status (if mirrored).

You can mirror the configuration repository, view its physical layout, or delete the configuration repository.

To mirror the configuration repository, you must have at least one physical device that has been reserved for *Configuration*. Once mirroring is configured, you can perform the following functions for a mirror:

- Sync - Manually synchronize and restart mirroring once the mirror is synchronized. This is useful if the mirrored devices has a minor failure, such as a power loss.
- Swap - Manually reverse the roles of the primary device and the mirror. You must do this before removing or performing maintenance on the primary device.
- Delete - Delete the mirror.

When you view the physical layout of the configuration repository, you will see physical segment and layout information along with the device size. If the configuration repository is mirrored, you will also see that information for the mirror device.

Tape database

The tape database contains information about virtual tapes.

Create the tape database

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Tape Database* icon.
3. Click *Create Database*.

4. Select a physical device.
5. Specify if you will use software compression.

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 in your backup application.

Compression is a global setting, which means that it will apply to all tapes in your system. However, if compression is enabled for the 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 application that run 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 compression is disabled, it is possible to enable individual drive compression, but it will have no effect.

Note: If you are already using software compression that is performed by your backup application, you should not use StorSafe's compression. Using both types of compression will cause StorSafe to try to compress already-compressed data and this can slow down your backups.

6. Set Tape Caching thresholds.

These thresholds are global variables that are set for tape caching policies for all virtual tape libraries.

Tape Caching Reclamation Threshold % - Space reclamation is triggered when used space reaches this threshold. The range is 50 to 99% and the default is 90%.

Tape Caching Migration Threshold % - Migration to physical tape will occur when the used disk space exceeds the specified disk capacity. The range is 50 to 99% and the default is 90%.

7. Click *Create* when done.

Compression and Tape Caching thresholds can be modified in *Global Settings*.

Manage the tape database

Select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, and click the *Tape Database* icon. Here you will see information about your tape database, including its size, status, and mirror status (if mirrored).

You can mirror the tape database, view its physical layout, or delete the tape database.

To mirror the tape database, you must have at least one physical device that has been reserved for *Configuration*. Once mirroring is configured, you can perform the following functions for a mirror:

- Sync - Manually synchronize and restart mirroring once the mirror is synchronized. This is useful if the mirrored devices has a minor failure, such as a power loss.
- Swap - Manually reverse the roles of the primary device and the mirror. You must do this before removing or performing maintenance on the primary device.
- Delete - Delete the mirror.

When you view the physical layout of the tape database, you will see physical segment and layout information along with the device size. If the tape database is mirrored, you will also see that information for the mirror device.

Deduplication repository

The deduplication repository stores deduplicated data and includes deduplication index, folder, and data disks. Index and folder disks always are located on physical resources. Deduplication data disks can be located on physical devices or on HCP, Azure, or Generic S3 object storage.

The deduplication repository must be created before you can create NAS resources.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Deduplication Repository* icon.
3. Click *Create Deduplication Repository*.

Create Deduplication Repository

Data Device Type: Physical Disk | Object Storage

Data Disk Configuration: Flexible | Classic

Select devices to be used as data disks to store unique data as a result of deduplication, with a maximum size of 16 TB and devices to be used as index/folder disks to store pointers to unique data and information related to deduplication sessions, with a maximum size of 2 TB:

☐	Physical Device	ACSL	GUID	Type
☐	VMware:Virtual disk	2:0:3:0	543d8197-e9a9-4fc9-bf86-2bb22...	Data
☐	VMware:Virtual disk	2:0:4:0	81751473-9e9c-4842-b1e0-2e3...	Data

Once the Deduplication Repository is created, data encryption cannot be enabled or disabled. Encryption will have a performance impact.

Encrypt Data ☐

Create Cancel

4. Indicate if you are using physical FC/SCSI disks or object storage for the deduplication repository.
5. **Object storage** - Indicate the size of the data disk in the range of 400 GB to 3,932,160 GB (240 data devices of 16 TB) on the object storage. On StorSafe-S, select the object storage account that you had created with the *Use for Deduplication Repository* option. On SIR, create a new object storage account.
6. **Physical disk** - Determine which method to use to configure deduplication data disks, *Flexible* or *Classic*.

Flexible - This method allows you to create a cluster with fewer, yet a more flexible number of data drives (255 max per cluster). This method does not require organizing your disks into columns and rows and allows for the use of all deduplication data drives in a cluster, rather than a row-by-row utilization of drives. With *Flexible* storage, the number of data disks is independent from the

number of nodes in the cluster; each node in a cluster has a single hash range and all data drives for the node store data for the node's hash range. Note that if you select this method, you will not be able to add deduplication nodes in the future to increase your cluster size. This method does not require downstream LUNs to be in certain multiples and is mainly meant for easier capacity expansion.

Classic - There are limits to the number and size of the data disks that can be selected. Each data disk is used for a specific hash range. The number of data disks depends on the number of nodes in the cluster. This method requires organizing your disks into columns/rows where all data disks in the same row have the same size; this method allows for future deduplication cluster expansion by adding nodes.

In *Classic* mode, capacity expansion has to go by multiples of column numbers; new capacity will be fulfilled by adding more rows. The deduplication hash range is spread among all cluster nodes where the greater columns per node provides better concurrent IOs.

If you select *Classic* mode, indicate the number of *Stripe Columns per Node*, which is the number of columns for storage striping to spread I/O between data disks. The number of data disks must be a multiple of the number of stripe columns for each node. Valid values for the number of columns are 1, 2, 4, 8, or 16, divided by the number of nodes (2 power 0 and 2 power 4).

How do the number of columns affect performance?

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

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

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

7. **Physical disk** - Select the device(s) that will be used to storage unique data.

The maximum size is 16 TB.

At least two physical devices with a reservation type of *Deduplication Repository* are required since the same disk cannot be used for both Data and Index.

For *Classic* mode, the number of data disks must be a multiple of columns and it must be the same for each cluster node.

8. **Object storage and physical disk** - Select the device(s) that will be used for the deduplication index and folder.

The index is used to store pointers to unique deduplicated data.

The folder is used to store information related to each deduplication session.

The selected device(s) must be of sufficient size to store these resources. The recommended minimum size disk is 0.6% of the total size of the deduplication data disks for index and 4.3% for folders. The maximum size is 2 TB.

9. **Object storage and physical disk** - If encryption support has been enabled, specify if you want to encrypt data in the deduplication repository.

Encryption must be enabled from the [“Encryption Support”](#) icon on the *Settings* page beforehand in order to have the option to encrypt data in the deduplication repository.

Encryption can be used to ensure that the data stored in the deduplication repository is confidential and secure. In order for data on the repository to be accessible, the cluster servers must have encryption activated with the specified password each time services are started on the server. All servers in the cluster use the same encryption activation password; data in the repository will not be accessible unless the activation password has been entered.

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

In the *Password Hint* text box, type a hint (0–32 characters) that will help you remember the password.

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

Notes:

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

10. Click *Create* when done.

Manage the deduplication repository

Select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, and click the *Deduplication Repository* icon. Here you will see general information about your deduplication repository, such as the data disk configuration and an indication if encryption is enabled. You will also see information about your data, index, and folder disks, including size, status, and mirror status (if mirrored).

You can also do the following:

- Expand the deduplication repository - Add deduplication data and index/folder storage
- Reclaim deduplication repository disk space - Set a space reclamation policy, start space reclamation and index pruning
- Switch to a flexible storage configuration
- Mirror a device
- View the physical layout for a device
- Delete the deduplication repository

Expand the deduplication repository

If your SIR server has available space, you can create additional logical resources for storage of deduplication data, index, or folders.

If you have physical disks available, prepare them, connect them to your server, and re-scan before beginning the expansion.

Add deduplication data storage

All active members of a deduplication cluster, not including the standby server, must be available when adding new disks.

The newly added disks will not be usable until all nodes in the deduplication cluster have been restarted. If your SIR server is an active node in a redundant node failover configuration, you must suspend failover before restarting the SIR server to prevent the standby node from taking over.

To add deduplication data storage, select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, click the *Deduplication Repository* icon, click the *Manage* icon, and select *Add Deduplication Data Storage*.

When the deduplication repository data is on physical storage, select from devices that have been reserved for the deduplication repository. If the deduplication repository is configured in classic mode, the same amount of storage must be added to each server in a cluster if the number of columns is greater than one.

When the deduplication repository is on object storage, specify the expansion size of object storage in GB, in the range of 400 GB to 3,932,160 GB (240 data devices of 16 TB). Note that you must have enough available capacity with your storage provider; the system cannot check whether there is sufficient free space.

Add deduplication index/folder disks

Adding index/folder disks applies to the local server and not all nodes in a deduplication cluster.

All members of a deduplication cluster, including the standby server, must be available when adding new disks. The same amount of storage should be added to each active server in a cluster.

Select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, click the *Deduplication Repository* icon, click the *Manage* icon, and select *Add Deduplication Index Disks*.

Select from physical devices that have been reserved for the deduplication repository.

Reclaim deduplication repository 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 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)

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 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.

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 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.

In order for reclamation to occur, the StorSafe server must have access to the deduplication data drives and there must be at least one VIT on any of the StorSafe servers associated with the deduplication cluster.

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

You cannot modify the threshold values.

Set space
reclamation
policy

1. Select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, click the *Manage* icon, and select *Set Space Reclamation Policy*.

2. Select when reclamation should run.
 - *Space Reclamation by Usage* - Specify how often the system should check to see if disk usage has reached the thresholds. The range is 10 to 1,440 minutes and the default is 10 minutes. Reclamation will run automatically whenever a threshold is met.
 - *Space Reclamation by 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 *Space Reclamation by Usage* or *Space Reclamation by Schedule*, reclamation will not run automatically; you will need to run it manually.

3. Click *Set* when you are done.

Start space
reclamation

To manually start space reclamation, select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, click the *Manage* icon, and select *Start Space Reclamation*.

Start index pruning	To manually start index pruning, select <i>Manage</i> from the menu bar, select a StorSafe server, click the <i>Settings</i> tab, click the <i>Manage</i> icon, and select <i>Start Index Pruning</i> .

Switch to a flexible storage configuration

If you configured your deduplication cluster via the *Classic* method (which requires organizing your disks into columns/rows), you can change your storage configuration to use the *Flexible* method. This method allows you to add one or more LUNs, rather than needing to add an equal amount per column.

Notes:

- Once you switch to the *Flexible* method, it is not reversible.
- You will need to restart all nodes in the cluster in order for the change to take effect.

To switch to a flexible storage configuration, select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, click the *Deduplication Repository* icon, click the *Manage* icon, and select *Switch to Flexible Mode*.

Mirror a device

To mirror a deduplication repository device, select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, click the *Deduplication Repository* icon, highlight a device, and click the *Mirror* icon. You must have at least one physical device that has been reserved for *Configuration*.

Once mirroring is configured, you can perform the following functions for a mirror:

- Sync - Manually synchronize and restart mirroring once the mirror is synchronized. This is useful if the mirrored devices has a minor failure, such as a power loss.
- Swap - Manually reverse the roles of the primary device and the mirror. You must do this before removing or performing maintenance on the primary device.
- Delete - Delete the mirror.

View the physical layout for a device

To view the physical segment and layout information for a device, select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, click the *Deduplication Repository* icon, highlight a device, and click the *Physical Layout* icon.

If the device is mirrored, you will also see that information for the mirror device.

Delete the deduplication repository

To delete the deduplication repository, select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, click the *Deduplication Repository* icon, and click the *Delete Repository* icon.

You can only delete the deduplication repository if there are no tapes with a deduplication policy.

Local users

Local users are StorSafe administrators that have full access to StorSafe functions.

Select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab, and click the *Local Users* icon.

Manage iSCSI users

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.

You can add, update, or delete iSCSI users from the *iSCSI Users* icon.

To add an iSCSI user:

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *iSCSI Users* icon.
3. Click the “+” icon.
4. Enter the iSCSI user name (up to 32 characters) and password (12-16 characters) and click *Add*.

Change root password

You must know the existing root user password in order to change it.

1. Select *Manage* from the menu bar and select a server.
2. Click the *Settings* tab and click the *Change Root Password* icon.

Modify system time

To set the server’s time:

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *System Time* icon.

Warning: Changing the system time may interrupt ongoing job processes.

Enable Fibre Channel and/or iSCSI target mode

Fibre Channel clients are the file and application servers that access virtual resources using the Fibre Channel (FC) protocol. In order to add a client that will use the FC protocol to communicate with the server, you must enable FC for your server.

iSCSI target mode is supported for iSCSI initiators on various platforms, including Windows, VMware, and Linux. Refer to the Certification Matrix for all support information. In order to add a client that will use the iSCSI protocol to communicate with the server, you must enable iSCSI for your server.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Target Mode* icon.
3. Select the options you want to use.

Configure SNMP

Simple Network Management Protocol (SNMP) is used to monitor systems for fault conditions, such as device failures, threshold violations, etc. on servers. Hierarchies, and other metadata (such as type and description of the variable), are described by Management Information Bases (MIBs). Refer to [“SNMP Integration”](#) for more information.

Configure Auto-Save

You can 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. To configure Auto-Save:

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Auto-Save Configuration* icon.

3. Complete the following fields:
 - *Secure FTP* - Indicate if you want to use secure FTP (SFTP).
 - *IP Address* - Enter a valid IPv4 address for the FTP server. This server must have FTP server installed and enabled.
 - *Username/Password* - The name and password the system will log in as. You must create this user on the FTP site. The user must have read/write access to the directory named here.
 - *Target Directory* - An existing directory on the FTP server where the files will be stored. The directory name you enter here (such as storsafeconfig) is a directory on the FTP server (for example ftp\storsafeconfig). You should not enter an absolute path like c:\storsafeconfig.
 - *FTP Port* - The FTP port number (default 21 for FTP, 22 for SFTP).
 - *Copies* - The maximum number of copies to keep on FTP. The oldest copy will be deleted as each new copy is added.
 - *Save Interval Unit* - Frequency to copy the configuration to the FTP server in days (1-7), hours (1-24), or minutes (1-59). Depending upon how frequently you make configuration changes to StorSafe, set the interval accordingly.
4. Click *Change Settings* when you are done.

Save server configuration

StorSafe includes a portal option and a command line utility (*vtlrecover*) that enables you to protect your databases and recover your backup or deduplication server configuration in case of the following:

- The Linux boot disk of the appliance is lost or corrupted.
- The file system where the StoreSafe software is installed is lost.

Information and requirements

- All appliance hardware, including FC HBAs and network adapters, storage, physical libraries/drives, and connectivity must be intact and functioning properly.
- Patches are not backed up.
- The tape import/export queue and scheduled reporting will not be saved and cannot be restored.
- In a SIR failover environment, we do not recommend saving/recovering the standby node because it is simpler to just remove the standby node and reconfigure failover after the problem is fixed. Moreover, since the standby node is used to automatically recover from a failed active node, we do not recommend saving/recovering active nodes.

Portal To save the configuration from the portal:

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Save Configuration* icon.
3. Click *Save*.
4. When prompted, click *Download*.
5. If the server is in a failover configuration, repeat on each node.

This must be done when failover is in a normal mode of operation and both nodes are functioning correctly.

Command line To save the configuration from the command line:

1. Run the following command: `$ISHOME/bin/vtlrecover save archive.tar`
The output file includes all configuration information needed for recovery.
2. If the system is a failover configuration, repeat the first step on each node.
This must be done when the failover system is in a normal mode of operation (i.e., both nodes are functioning correctly).
3. Copy the output file to a safe remote location outside the appliance.

Recover server 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. The server will be restarted after the configuration is restored.

Information and requirements

- All appliance hardware, including FC HBAs and network adapters, storage, physical libraries/drives, and connectivity must be intact and functioning properly.
- Patches are not 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/export queue and scheduled reporting will not be restored.

Recover a standalone server

To restore the configuration for a standalone (non-failover) server:

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 StorSafe 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.

5. Restore the configuration:

From the portal, select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab and click the *Recover Configuration* icon. Locate the saved configuration file and click *Recover*.

From the command line, run the following command:

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

Depending on how many tapes are present on a backup server, it may take up to 10 minutes to restore the system.

6. Verify that all configuration information has been restored.

All virtual tapes, including direct link tapes, will be automatically moved to the appropriate libraries.

Recover a server in a failover environment

There are two scenarios for restoring a configuration in a failover environment:

- Takeover was successful - Follow the instructions below to restore your configuration.
- Takeover was not successful because the failed server's configuration and database were corrupted - Contact Technical Support for help restoring the configuration.

To restore the configuration on a failed server when the surviving server has taken over successfully:

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

For a 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 StorSafe 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.

5. Restore the configuration:

From the portal, select *Manage* from the menu bar, select a StorSafe server, click the *Settings* tab and click the *Recover Configuration* icon. Locate the saved configuration file and click *Recover*.

From the command line, run the following command:

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

Depending on how many tapes are present on a backup server, it may take up to 10 minutes to restore the system.

6. Select *Manage* from the menu bar, select the secondary failover server, click the *Settings* tab, click the *Failover* icon, click *Manage Failover*, and select *Stop Takeover*.

7. Verify that all configuration information has been restored.

All virtual tapes, including direct link tapes, will be automatically moved to the appropriate libraries.

After restoring your configuration

- Check the *Tapes* page for tapes that are not *Online*. This will include tapes that were deleted after saving the configuration and direct link tapes that were reclaimed after saving the configuration.
- If you created any tapes after saving your configuration, those tapes will go to the vault.
- 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.

Set replication throttle

You can limit the amount of available network bandwidth that is used for replication of deduplicated tapes (VITs) on the source server 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. If throttling is not used, replication will use the maximum bandwidth that is available.

1. Select *Manage* from the menu bar, select a StorSafe server.
2. Click the *Settings* tab and click the *Replication Throttle* icon to set a throttle speed in the range of 10 to 1,000,000 KB/s, specified in KB/s, MB/s, or GB/s.

Server Overview Virtual Devices Host Clients TimeMarks Replication Physical Resources Settings

◀ Back Manage Global Replication Throttle

Maximum allowable replication throughput in order to limit network bandwidth usage, in the range of 10 to 1,000,000 KB/s

Global Replication Throttle

Activate Global Replication Throttle ☒

Throttle Speed 10 MB/s

Change

3. Click *Change* when you are done.

Set global virtual tape library system properties

To set global options for all virtual tape libraries:

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Global Settings* icon.

Update VTL01 Global Settings

Enable compression mode ☐

Retain tape properties when the tape is overwritten ☐

Tape Caching Reclamation Threshold % ?

Tape Caching Migration Threshold % ?

Update Cancel

3. Select the options you want to use.

Enable compression mode - 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 StorSafe server.

If you are already using software compression that is performed by your backup application, you should not use StorSafe compression. Using both types of compression will cause StorSafe to try to compress already-compressed data and this can slow down your backups.

Compression is a global setting, which means that it will apply to all tapes in your system. However, if compression is enabled on the StorSafe 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.

Retain Tape Properties when tape is overwritten - Allows tape properties to persist even when the tape is overwritten.

Tape Caching Reclamation Threshold - Cache disk space is freed up when the used space reaches this threshold.

Tape Caching Migration Threshold - Migration will occur when the used disk space exceeds the specified disk capacity.

4. Click *Update*.

Configure Hosted Backup

Hosted Backup eliminates the need for a dedicated backup application server by making virtual tape libraries and drives available to the local system, allowing certified backup applications to be installed directly onto the StorSafe appliance.

Note: Your backup application should be installed on the `/apps` partition on your StorSafe appliance.

With Hosted Backup, data transfer between application servers and the backup application is accelerated because the backup application is hosted on the StorSafe appliance itself, enhancing backup performance.

Hosted Backup is used by NDMP and by combined backup & deduplication servers and may already be enabled. To configure Hosted Backup:

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Hosted Backup* icon.
3. If necessary, click *Enable Hosted Backup*.
4. Click the *Assign Devices* icon.
5. Select the virtual libraries and drives that Hosted Backup will use and click *Assign*.

The selected devices will be listed. Resources pre-allocated for use with deduplication will also be listed and cannot be unassigned.

Disable Hosted Backup

In order to disable Hosted Backup, you must unassign all virtual libraries and drives assigned to it.

However, Hosted Backup is used by NDMP and by combined backup & deduplication servers and cannot be disabled if NDMP or deduplication are being used.

NDMP

Network Data Management Protocol (NDMP) allows certified backup applications and industry standard NAS filers (i.e. NetApp filers) to perform backup and restore using the NDMP protocol over an IP network.

With NDMP, the StorSafe appliance acts as an NDMP server, centralizing management and eliminating the need for locally attached tape devices on NAS filers. When a backup occurs, data is moved from the NAS filer directly to a virtual tape. Conversely, during restore, data is moved from a virtual tape directly to a NAS filer.

The NDMP supports the following:

- NDMP v2, 3, 4
- The following data transfer models:
 - Filer to direct-attach tape drive (local filer)
 - Filer to another filer attached tape drive (filer-to-filer)
 - Filer to *NetVault* Client/Server attached tape drive (filer-to-server)
 - *NetVault* Client/Server to filer attached tape drive (server/client-to-filer)
- Library and tape sharing
- Direct Access Restores (DAR)

Notes:

- If you want to use StorSafe as an OpenStorage server and NDMP host, NetBackup requires you to enable the *Direct to Tape* option first. Refer to the *FalconStor OpenStorage Option User Guide* for more information.
- NDMP is not needed when presenting a virtual tape library over FC to a NAS filer as a replacement for a physical library.
- Before you begin configuration, you must define the hostname of the StorSafe server in the `/etc/hosts` file in the format "IPAddress Hostname". For example: 10.7.2.41 Server41
- Because some backup applications use NDMP, if you are running backup software on the StorSafe server, it should be started after StorSafe has started and should be stopped before StorSafe is stopped. Otherwise, the NDMP service that is loaded by the backup software may interfere with StorSafe's NDMP service.

Configure NDMP

To configure NDMP support:

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *NDMP* icon.
3. Enter the user name and password the backup application server will use to communicate with NDMP.

You must enter the same user name/password into the NDMP module in your backup application.

4. Click *Update* when done.

The *Hosted Backup* option automatically gets enabled as well.

5. Click the *Hosted Backup* icon and assign virtual libraries.

Hosted Backup can have any number of virtual libraries assigned to it. Conversely, libraries assigned to the *Hosted Backup* can also be assigned to other clients.

The backup application will now see the devices as local devices. You can use Linux's `cat /proc/scsi/scsi` command on your StorSafe appliance to see the library and all of the drives in the library listed as local devices.

Important notes about NDMP devices:

- Any time you add or remove physical or virtual resources on the StorSafe server and then reboot the server, you will need to check that the paths of the assigned NDMP devices are in the correct sequence in Veritas NetBackup. If they are not in the correct sequence, you will need to reconfigure the paths.
- If you are using a physical library assigned as an NDMP device to your client, make sure that this library is not assigned for any other StorSafe functions.

OST Settings

FalconStor® OpenStorage Option (OST) is a backup option that interfaces with Veritas™ NetBackup™ or Backup Exec™ servers to perform high-speed backup and recovery to intelligent disk devices.

The OST Direct to Tape feature exports OST backup images directly to physical tapes. It requires the OST and NDMP options to be enabled.

Refer to the *FalconStor OpenStorage Option User Guide* for more information.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *OST Settings* icon.
3. Enable *OST* and the *OST Direct to Tape* option, if needed.

Configure the Activity Log

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 StorSafe server reports.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Activity Log Settings* icon.
3. Set the maximum file size and number of days to track activity.

The default values are 50 MB and 365 days.

Set the storage usage threshold

In order to monitor storage usage, you can set a usage threshold. 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 also be monitored from there.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Storage Usage Threshold* icon.
3. Enter the maximum percentage of storage that can be used by the StorSafe server before you should be alerted.

The default is 90%.

Encryption Support

Encryption can be used to ensure that the data stored in the deduplication repository is confidential and secure. It can also be used to ensure that data backed up on virtual tapes is confidential and secure. Before you can enable encryption on either the deduplication repository or for virtual tapes, you must enable encryption support.

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Encryption Support* icon.

This option unlocks encryption. When you create your deduplication repository, you will now have an option to enable deduplication repository encryption. Refer to [“Deduplication repository”](#) for more information.

Once *Encryption Support* is enabled, you must create encryption keys. Refer to [“Encryption”](#) for more information.

Enable NAS

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

Before you can enable NAS, you must have prepared at least one physical device with the *NAS Resources* reservation type and the deduplication repository must have been created. NAS cannot be enabled on servers associated with an encrypted deduplication repository.

To enable NAS:

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *NAS* icon.

Enable strong passwords

For security purposes, the *Strong Passwords* option cannot be disabled once it has been enabled. When enabled, strong passwords are required for all users. However, the root user's password can be reset without limitations.

Strong passwords must contain a minimum of 14 characters, including at least one lowercase letter, one uppercase letter, and one number, plus at least one special character: @#\$%

The password cannot be the Administrator name or the Administrator name in reverse order. It cannot contain part of the Administrator name. Any time a password is changed, the old and new password must differ by at least five characters.

The *Strong Passwords* option includes a locking mechanism for user accounts that repeatedly fail to log in because of an incorrect password; by default, a user is locked out after entering the wrong password three times. 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 before trying again. The default lockout period is five minutes.

Note: Strong passwords and Active Directory/LDAP authentication are mutually exclusive; if AD/LDAP is in use in your organization or those of your tenants, you cannot use strong passwords on any StorSight server.

To enable strong passwords:

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *Strong Passwords* icon.

If you need to change the server password to comply with strong password requirements, use the *Change Root Password* option on the *Settings* tab. Then, to re-establish connection with the server using this new password, select the server on the *Administration -> Servers* page and click the *Update* icon to reset the password.

Job Status

When you select *Manage* from the menu bar, select a StorSafe server, and click the *Jobs* tab, the *Deduplication Jobs* page appears. From here you can also access the *Replication Jobs* and *Import/Export Jobs* for your StorSafe servers.

Deduplication Jobs page



Select *Manage* from the menu bar and click the *Jobs* tab. This page lists the virtual tapes that are currently being deduplicated for all deduplication policies on the selected StorSafe server.

Only a Super Administrator can access the *Deduplication Jobs* page.

The screenshot displays the FalconStor StorSafe web interface. The top navigation bar includes 'Dashboard', 'Monitor', 'Manage' (selected), 'Analytics', and 'Administration'. The left sidebar shows a list of servers, with 'H13-97' highlighted. The main content area is titled 'Jobs' and contains a sub-menu with 'Deduplication', 'Replication', and 'Import/Export'. The 'Deduplication' tab is active, showing a table of jobs. The table has columns for Barcode, Tape, ID, Policy, State, Unduplicated..., Throughput, and Progress. Two entries are shown: one for VirtualTape-01931 in a 'Deduplicating' state with 85% progress, and another for VirtualTape-01932 in a 'Queued' state with 0% progress.

Barcode	Tape	ID	Policy	State	Unduplicated...	Throughput	Progress
0039000Q	VirtualTape-01931	10001931	HCP	Deduplicating	320.83 GB	139 MB/s	85%
0039000R	VirtualTape-01932	10001932	HCP	Queued	352.83 GB	0 B/s	0%

Depending upon the columns that have been selected for display, the properties for the virtual tapes that are being processed are visible at a glance. The following information is displayed for each tape:

- Barcode
- Tape
- Tape ID
- Deduplication policy
- Status - Suspended, Queued, Tape in drive, Preparing for deduplication, Deduplicating

- Amount of undeduplicated data
- Throughput
- Progress percentage - Click *Refresh* to update the data.

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

Deduplication Jobs page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Manage	Change the priority of tapes in the queue or cancel a deduplication job.
	Automatic refresh	Allow the tape job status to automatically refresh every 10 seconds or 1 minute.

Manage deduplication jobs

You can change the priority of tapes in the queue or cancel processing for a tape that is being deduplicated.

Select *Manage* from the menu bar, select a StorSafe server, click the *Jobs* tab, the *Deduplication* tab, and then click the *Manage* icon.

To change the priority of tapes, select *Run Later* or *Run Next* from the *Manage* menu. Priority was set when the deduplication policy was created. Changing the priority in the queue is useful for disaster recovery (DR) purposes when a tape has replication configured and is needed at the DR site.

To cancel processing for a tape, select *Cancel* from the *Manage* menu.

Replication Jobs page



Select *Manage* from the menu bar, click the *Jobs* tab and then the *Replication* tab. This page lists the virtual tapes that are currently being replicated on the selected StorSafe server.

Only a Super Administrator can access the *Replication Jobs* page.

Barcode	Tape	ID	Type	Mode	State
004F06L7	VirtualTape-WORM-00425	10000425	Without Deduplication	Without Deduplication	Running

Showing 1 entries

Depending upon the columns that have been selected for display, the properties for the virtual tapes that are being processed are visible at a glance. The following information is displayed for each tape:

- Barcode
- Tape
- Tape ID
- Type - With or without deduplication
- Replication mode
- Status - Running, Waiting for retry, Waiting for a slot.

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



Replication Jobs page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Suspend globally	Suspend all virtual tape replication jobs.
	Cancel job	Cancel the currently running job.
	Automatic refresh	Allow the tape job status to automatically refresh every 10 seconds or 1 minute.

Manage replication jobs

You can cancel a replication job that is running or suspend all virtual tape replication jobs (without deduplication) on this server.

Select *Manage* from the menu bar, select a StorSafe server, click the *Jobs* tab, the *Replication* tab, click the *Manage* icon and select a tape and select the action you want to perform.

Import/Export Jobs page



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

Only a Super Administrator can access the *Import/Export Jobs* page.

The screenshot shows the FalconStor StorSafe web interface. The top navigation bar includes tabs for Dashboard, Monitor, Manage, Analytics, and Administration. The 'Manage' tab is selected, and the 'Jobs' sub-tab is active. The main content area displays a table of jobs with columns for Job ID, Type, Mode, Virtual Tape ID, Virtual Tape Barcode, Physical Tape Barcode, Status, Start Time, and Elapsed Time. The table lists 22 entries, all of which are 'Completed'.

Job ID	Type	Mode	Virtual Tape ID	Virtual Tape Barcode	Physical Tape Barcode	Status	Start Time	Elapsed Time
612	Recover from Object Storage	Move	10000397	0079000G		Completed	2019-03-19 04...	
613	Recover from Object Storage	Move	10000398	0079000H		Completed	2019-03-19 05...	
614	Recover from Object Storage	Move	10000401	0079000I		Completed	2019-03-19 05...	
604	Export to Library	Move	10000424	004F05L7	004F05L7	Completed	2019-02-26 05...	
605	Export to Library	Move	10000426	004F07L7	004F07L7	Completed	2019-02-26 04...	
609	Recover from Object Storage	Copy	10000397	0079000G		Completed	2019-03-13 02...	
610	Migrate to Object Storage	Move	10000363	00790000		Completed	2019-03-13 16...	
611	Recover from Object Storage	Move	10000361	00760000		Completed	2019-03-19 04...	
615	Recover from Object Storage	Move	10000403	0079000J		Completed	2019-03-19 05...	
616	Recover from Object Storage	Move	10000404	0079000K		Completed	2019-03-19 05...	
617	Recover from Object Storage	Copy	10000405	0079000L		Completed	2019-03-19 16...	
618	Recover from Object Storage	Copy	10000406	0079000M		Completed	2019-03-20 03...	

Depending upon the columns that have been selected for display, the properties for each job are visible at a glance. The following information is displayed for each job:

- Job ID
- Job type - Export to standalone drive/library, import from standalone drive/library, create cache and copy metadata, move to import/export slot, export to standalone drive/library with stacking, import from library with stacking, move to import/export slot with stacking, scan physical tape in standalone drive/library, export tape to object storage, import from object storage
- Cloud Tiering migration mode - Copy or Move
- Virtual tape ID
- Virtual tape barcode
- Physical tape barcode or object storage
- Status - Waiting for tape/drive, Preparing for Object Storage Import/Export, Running, Completed, Cancelled, Failed, Canceling, Waiting for I/E slot, Waiting to be deleted, Hold, Resumed.

- Start time
- Elapsed time in seconds for running jobs
- Amount of data transferred
- Progress percentage

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

Import/Export Jobs page icons

The table below describes the unique icons for this page and briefly describes the associated features.

Icon	Feature	Description
	Schedule report	Schedule an Import/Export Jobs report.
	Set retry parameters	Set rules for retrying failed jobs.
	Manage	Restart, cancel, or delete a job.
	Automatic refresh	Allow the status to automatically refresh every 10 seconds or 1 minute.

Manage import/export jobs

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

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

Set retry parameters

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

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

2. Click *Set Retry Parameters*.

A screenshot of a dialog box titled "Set Job Retry Parameters". It contains three settings: "Retry failed jobs" with a toggle switch turned on, "Retry Count" with a text input field containing the number 10, and "Retry Interval" with a text input field containing the number 0. To the right of the "Retry Interval" field is the label "Minutes". At the bottom right of the dialog are two buttons: "Set" and "Cancel".

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

3. Click the *Set* button.

Schedule Import/Export Jobs report

The Import Export Jobs report can be scheduled to be emailed to you on a regular basis. Refer to [“Schedule a report”](#) for more information.

Analytics

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

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

[Reports](#) - StorSight provides a wide variety of predefined reports that help you manage your servers.

[Scheduled reports](#) - Manage and view scheduled reports.

Rules - You can create rules for a variety of conditions that you may want to be made aware of via alerts, email, or SMS, such as when a server is unresponsive, when memory is low, or when a network error occurs. There are two types of rules:

- [Smart rules](#) - offer an additional level of flexibility by allowing you to set thresholds, use comparison operators, and define a combination of conditions.
- [General rules](#) - configuration rules that can be set at a server or device level.

[Alerts](#) - display as a result of the rules configured. When the condition occurs, an alert is triggered.

[Capacity prediction](#) - uses an algorithm based on historical storage data usage for a server or storage pool to forecast future capacity demands.

[Audit log](#) - tracks and logs system activities.

[Chargeback](#) - Track customer storage capacity usage or the number of IOPS in order to allow service providers to easily generate bills for customers.

Reports



StorSight provides a wide variety of reports that help you manage your system and resources. Some reports focus on server conditions and individual hardware component usage, errors, and performance. Others are related to data services and gather comprehensive status information about the different types of resources (i.e., virtual tapes), as well as replication.

To view a report, select *Reports* from the *Analytics* menu bar.

Note that you will see reports listed that are available only for FalconStor StorGuard servers. These reports will not be discussed.

Reports page

The *Reports* page opens with *Inventory Reports*. You can click the report you want to display from the list of reports in the left panel. You can also search for report name. If applicable, you can use the *Filter* or the *Arrange* icons to select which parameters you want displayed in each report and in which order. You can also sort the information in the columns.

Reports are available in the following categories:

- [Inventory reports](#)
- [Capacity Management reports](#)
- [Performance Monitoring reports](#)
- [Service Status reports](#)
- [Server Load reports](#)
- [Trend reports](#)
- [Top N Components reports](#)
- [Custom trend reports](#) - User-defined trend reports

Super Administrators can view reports for all customers. Each customer can view reports for its own domain.

Inventory reports

Inventory reports provide details about your servers, physical devices, clients, virtual tape libraries and drives, physical tape libraries and drives, and NAS resources. Super Administrators can also view a customer report. You can locate a component across all of your servers by a specific property, such as GUID or serial number in inventory reports. The following inventory reports are available:

Server List of all servers. For each server, this report displays:

- IP address
- Server type
- Configuration - Standalone, Failover
- Version
- Build
- Protocol(s) - FC and/or iSCSI
- Total capacity
- Total used
- Number of clients
- Number of storage pools
- Number of FC adapters
- Number of physical devices
- Number of prepared/unprepared/virtualized/physical devices
- Number of virtual tape libraries
- Number of virtual tape drives
- Number of virtual tapes
- Number of loaded virtual tapes
- Number of standalone virtual tape drives
- Number of tapes in the vault
- Is tape compression is enabled
- Number of NAS resources and NAS shares
- Customer - Super Administrators only

You can click the *Arrange* icon to show read/write IOPS, read/write latency, and read/write throughput for each server.

Physical Devices List of physical devices on all servers. For each physical device, this report displays:

- ACSL
- Server
- Server IP address
- Server type
- Size - Total, allocated, and free
- Storage pool
- Category - Virtualized, foreign, unassigned

- Reservation type - Tapes, NAS resources, deduplication repository, configuration repository
- Owned
- Connection - FC or SCSI
- Queue depth
- GUID

You can click the *Arrange* icon to show read/write IOPS, read/write latency, and read/write throughput for each physical device.

Clients List of FC and iSCSI clients on all servers. For each client, this report displays:

- Client ID
- Server
- Server IP address
- Server type
- Number of virtual resources (click to see a list)
- Protocol(s) - FC and/or iSCSI
- VSA usage
- iSeries usage
- Is the client is clustered
- Customer - Super Administrators only

You can click the *Arrange* icon to show read/write IOPS, read/write latency, and read/write throughput for each client.

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

- ID
- Server
- Server IP address
- Vendor
- Product
- Revision
- Media type
- Number of slots
- Number of drives
- Number of tapes
- Number of loaded tapes
- Number of clients (click to see a list)
- Starting/ending barcodes
- Storage pool
- Tape capacity on demand usage - If tape COD is used, the initial allocation size, increment size, and maximum capacity are displayed.
- Auto archive usage
- Auto replication usage

- Cloud Tiering usage
- Automated tape caching usage
- Tape encryption usage
- Tape duplication usage

Virtual Tape Drives	List of all virtual tape drives in virtual tape libraries and all standalone virtual tape drives on StorSafe servers. For each virtual tape drive, this report displays: • ID • Server • Server IP address • Status - Empty, loaded, ejected • Loaded tape - Name of tape if the status is loaded • Barcode of loaded tape • Element number • Virtual tape library • Number of clients (click to see a list) • Vendor • Product • Firmware
Physical Tape Libraries	List of all physical tape libraries on StorSafe servers. For each physical tape library, this report displays: • ID • Server • Vendor • Product • Firmware • Serial number • ACSL (Adapter:Channel:SCSI ID:LUN) • Number of drives (click to see a list)
Physical Tape Drives	List of all physical tape drives in physical tape libraries. For each physical tape drive, this report displays: • ID • Server • Physical tape library • Vendor • Product • Firmware • Serial number • ACSL (Adapter:Channel:SCSI ID:LUN)

NAS Resources	List of all NAS resources on StorSafe servers. For each NAS resource, this report displays: • ID • Status - Online/offline • Server • Server IP address • Mount point - The directory from which the share can be accessed • Mount status - Mounted, unmounted • Capacity • Allocated and free size • Deduplication - Included, excluded, partial (some sub-folders are excluded) • Replication - Included, excluded, partial (some sub-folders are excluded), empty (not configured or configured with multiple policies) • Number of shares • Storage pool • Physical device (click to see the ACSL)
Customers 	(Super Administrators only) List of all customers. For each customer, this report displays: • Number of active servers (click to see a list of servers and access modes) • Number of storage pools • Number of virtual resources • Number of clients (click to see a list) • Capacity - Total available capacity for this customer • Used Capacity - Total size of all virtual resources and supporting resources • Used Capacity StorSafe Dedupe Repo - Amount of data stored in the deduplication repository • Written Data StorSafe Tapes - Amount of incoming digest data on all tapes • Dedupe Ratio - Ratio between written and stored data

Capacity Management reports

Capacity management reports provide real-time statistics for primary and supporting devices at the device, group, server, and client level. These reports can be used to identify near-capacity components and provide an indication of storage growth rates.

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

For all reports, Super Administrators will see a *Customer* column.

The following capacity management reports are available:

Server Physical Storage	Physical storage usage. For each server, this report displays: • IP address • Server Type • Configuration - Standalone, Failover • Protocol(s) - FC and/or iSCSI • Capacity - Prepared, reserved, and unprepared • Total used • StorSafe servers: • Amount used/reserved for virtual tapes • Capacity for NAS and amount used/reserved • Amount used for reserved/used for the deduplication repository (for StorSafe servers with deduplication) • Amount reserved for configuration (includes StorSafe database) • Amount used for object storage • Total reserved capacity • SIR servers: • Amount used for configuration repository • Amount reserved/used for the deduplication repository • Total reserved capacity
Storage Pool	Storage pool usage for StorSafe servers. For each storage pool, this report displays: • Server • Server IP address • Server type • Resource types (if a number is displayed, click to see a list) - All, tape, NAS resource • Size - Total, allocated, and free • Number of physical devices (click to see a list with ACSL and GUID) • Amount used for NAS • Amount used for tapes
Deduplication Repository	Deduplication repository usage for SIR servers and StorSafe server with deduplication. For each server, this report displays: • Server IP address • Deduplication ratio • Amount of scanned data • Amount of unique data • Automatic reclamation usage • Data disk: • Size, used size, free size - Capacity of the deduplication data disk(s), how much space has been used, how much space is available • Reclamation threshold percentage • Used before reclamation value/percentage - 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 value/percentage - The amount of deduplicated data added after the last successful reclamation. If reclamation has never been run, *Data Disk Used* and *Data Disk Used Since Reclamation* will be the same.
- Folder disk:
 - Size, used size, free size - Capacity of the folder disk, how much space has been used, how much space is available
 - Reclamation threshold percentage
 - Used before reclamation value/percentage - 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 value/percentage - The amount of folder capacity used after the last successful pruning. If reclamation has never been run, *Folder Disk Used* and *Folder Disk Used Since Reclamation* will be the same.
- Index disk:
 - Size, used size, free size - Capacity of the index disk, how much space has been used, how much space is available.
 - Pruning threshold percentage
 - Used before pruning value/percentage- 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 value/percentage - The amount of index space used after the last successful pruning. If pruning has never been run, *Index Disk Used* and *Index Disk Used Since Reclamation* will be the same.
- Index memory:
 - Used and free size - The amount of total deduplication index cache that is used and available.
 - Reclamation threshold percentage
 - Used before reclamation percentage- 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 percentage - The amount of index cache capacity used after the last successful reclamation.

NAS SMB Shares

This report shows disk space usage for a specific SMB share on a specific date or range of dates. For each share, this report displays:

- Time
- Server
- Server IP address
- NAS SMB share name
- Used size of SMB share (in KB)

Performance Monitoring reports

Performance monitoring reports provide real-time information about I/O performance and average load for upstream and downstream components, as well as replication and FC target and pending commands/errors for clients and devices.

These reports can be used to identify overall performance and load differences between different components.

The performance and throughput reports provide real-time information. The FC target, device, and client commands reports allow you to set a time frame for each report. By default, data for the last 24 hours is shown.

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

For all reports, Super Administrators will see a *Customer* column.

The following performance monitoring reports are available:

Physical Device Performance	Physical device performance for all servers. For each physical device, this report displays: • ACSL • Server • Server IP address • Server type • Storage pool • Connection - FC or SCSI • Queue depth • Downstream read, write, and total throughput • Downstream read, write, and total IOPS • Downstream read, write, and mean latency
Virtual Tape Library Performance	Virtual tape library performance for StorSafe servers. Library performance is the sum of performance values of each tape drive in the library. If compression is enabled, the read/write throughput and read/write compressed values will be different. If compression is not enabled, the comparable values (i.e., write throughput and compressed write) will be the same. For each virtual tape library, this report displays: • ID • Server • Server IP address • Read and write throughput • Compressed read • Compressed write

Virtual Tape Drive Performance	Virtual tape drive performance for StorSafe servers. Tape drive performance values are cumulative from the time a tape is loaded on the drive. Unloading or reloading a tape will reset the value to 0; if the current values are less than the previous values, a tape has been reloaded or unloaded in the drive. If compression is enabled, the read/write throughput and read/write compressed values will be different. If compression is not enabled, the comparable values (i.e., write throughput and compressed write) will be the same. For each virtual tape drive, this report displays: • ID • Virtual tape library • Server • Server IP address • Read and write throughput • Compressed read • Compressed write
NAS Resource Performance	NAS resource performance for StorSafe servers. For each NAS resource, this report displays: • ID • Server • Server IP address • Read, write, and total throughput
Client Performance	Client performance for all servers. For each client, this report displays: • Client ID • Client IP address • Type - Operating system • Server • Server IP address • Server type • Read, write, and total throughput
Fibre Channel Target Performance	Fibre Channel target performance for all servers. For each target, this report displays: • Server • Server IP address • Server type • Upstream read and write throughput

iSCSI Target Performance

iSCSI target performance for all servers. For each target, this report displays:

- Server
- Server IP address
- Server type
- Upstream read and write throughput
- Upstream read and write IOPS

Server Throughput

Server throughput for all StorSafe servers. For each server, this report displays:

- Server IP address
- Server type
- Configuration - Standalone, Failover
- Read, write, and total throughput - This is the total performance for all virtual tapes (uncompressed I/O), physical tapes, and NAS resources.
- Read, write, and total IOPS

Service Status reports

Service status reports provide information and status for failover, deduplication, and replication.

Super Administrators will see a *Customer* column.

The following service status reports are available:

Failover	Failover status for all servers. For each server, this report displays: • Server IP address • Server type • Failover configuration - Not configured, 1+1, 2+1, or 4 +1 redundant node configuration (SIR servers) • Status • Partner server - The standby server.
Deduplicated Tape Replication History	Information about completed replication of tapes in deduplication policies for the specified period of time. For each server, this report displays: • Server IP address • Size of ingest data incoming on all tapes • Size of data for replication, based on compressed data if compression is enabled • Size of replicated unique data • Number of tapes (includes tapes waiting to be replicated) • Number of replicated tapes
Tape Deduplication and Replication	Deduplication and replication statistics for tapes in deduplication policies. For each server, this report displays: • Server IP address • Total data size of all tapes that have been deduplication • Data size of all tapes to be deduplicated; for inline deduplication, this may be 0 • Number of deduplicated tapes • Number of tapes to be deduplicated - zero-data tapes (either zero-size tapes or zero-used tapes with an allocation size) are not counted • Amount of data on all tapes to be replicated • Number of tapes to be replicated - zero-data tapes (either zero-size tapes or zero-used tapes with an allocation size) are not counted • Number of zero-data tapes - zero-size tapes or zero-used tapes with an allocation size
NAS Deduplication and Replication	Deduplication and replication statistics for NAS resources. For each NAS resource, this report displays: • NAS resource ID

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

NAS Replication Sessions

Summary information for NAS incoming and outgoing replication sessions for the specified period of time. For each NAS replication session, this report displays:

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

Server Load reports

Server load reports provide real-time data about server usage of system resources such as CPU, memory, system disk, and network.

These reports can be used to identify the saturation level and compare load differences between servers.

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

For all reports, Super Administrators will see a *Customer* column.

The following server load reports are available:

CPU Load CPU load for all servers. For each server, this report displays:

- Server IP address
- Server type
- Configuration - Standalone, Failover
- Average load
- Number of cores

Memory Usage Memory usage for all servers. For SIR servers and StorSafe servers with deduplication, usage is based on the system memory left behind at the operating system level after memory is pre-allocated for deduplication.

For each server, this report displays:

- Server IP address
- Server type
- Configuration - Standalone, Failover
- Amount of total, used, and free memory
- Buffer size
- Amount of cached memory
- Amount of used and free swap memory
- Deduplication index memory - Pre-allocated memory for deduplication on SIR servers.


**Network
Throughput**

Network throughput for all servers. For each network interface, this report displays:

- Network interface IP address
- Server
- Server IP
- Server type
- Configuration
- Speed
- Duplex
- Read, write, and total throughput
- Number of received and transmitted packets
- Number of received and transmitted bytes
- Number of erroneous received and transmitted packets
- Number of dropped received and transmitted packets

Trend reports

Trend reports show performance and capacity trends based on real-time or historical data. These reports can be used to analyze trends.



Trend reports can be presented in line or area format.




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

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

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


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

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



The screenshot shows a dropdown menu with the following options: Real-time, Last Hour, Last 3 Hours, Last 6 Hours, Last 12 Hours, Last Day, and Last 7 Days. Below these is a 'Date Range' section with 'From' and 'To' date pickers, both set to 2016-07-27. At the bottom are 'Apply' and 'Cancel' buttons.

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



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

The following trend reports are available:

- **Physical Device Performance Trend** - Displays real-time read throughput, latency, and IOPS for a physical device on a server.
- **Virtual Tape Library Performance Trend** - Displays real-time read and write throughput including compressed activity, when compression is enabled for a virtual tape library. Unloading or reloading a tape will reset the values to 0.
- **Virtual Tape Drive Performance Trend** - Displays real-time read and write throughput including compressed activity, when compression is enabled for a virtual tape drive. Unloading or reloading a tape will reset the values to 0.
- **NAS Resource Performance Trend** - Displays real-time throughput for a NAS resource.
- **Client Performance Trend** - Displays real-time read throughput, latency, and IOPS for a FC or iSCSI client on any server.
- **Storage Capacity Trend** - Displays used capacity for the selected server, as well as capacity used for virtual resources. The used capacity for virtual resources includes primary virtual resources, their mirrors, replicas, and replica mirrors.

- **Deduplication Repository Trend** - Displays real-time data, folder, or index disk statistics for a deduplication repository on an SIR server or a StorSafe server with deduplication.
- **Tape Deduplication and Replication Trend** - Displays trend graphs for deduplication and replication status of tapes in deduplication policies on a StorSafe server.
- **Deduplicated Tape Replication History Trend** - Displays trend graphs for incoming data and replicated tapes in deduplication policies on a StorSafe server.

Modify a trend report

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

1. Select a trend report.
2. Click the arrow by a series to modify it or click the “+” icon to add a series.
3. Make your selections.

Refer to the [“Trend report selection table”](#).

4. Specify a *Series Name* or accept the default name.
5. Click *OK*.

6. Repeat for the other series in this report.

You can mix and match categories, components, and objects in a single trend report.

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

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

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

Top N Components reports

Top N Reports show the Top 10 servers, physical devices and virtual resources. Reports are generated for Throughput (throughput), IOPS and Latency; devices are arranged according to these values. Each report displays the average value for the selected time period.

The following reports are available:

- Top StorSafe Servers by Throughput
- Top Virtual Devices by Throughput
- Top Physical Devices by Throughput
- Top Physical Devices by IOPS
- Top Physical Devices by Latency
- Top Virtual Tape Libraries by Throughput

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

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

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

Custom trend reports

Custom reports contains saved trend reports that were modified or custom created trend reports. For more information about modifying an existing trend report, refer to [“Modify a trend report”](#).

Create a report

You can create custom reports to display the specific information you need.

1. Click the + *Create Report* icon.
2. Specify a name for the report and click *Create*.
3. Click the “+” icon to add a series.
4. Make your selections.

You will need to select:

- *Category* - Type of report: Capacity, Load, Performance.
- *Component* - Type of system component to report on (e.g., physical device, client, network, or virtual tape library). The available options are related to the selected category.
- *Metric* - Criteria to measure.
- *Server* - The server to report on. This field is only available if an *Object* is also required.
- *Object* - The specific component to report on (e.g., specific physical device or client). If a specific component is not required, you can select the server in this field.

Refer to the [“Trend report selection table”](#) to see the available options.

5. Specify a *Series Name* or accept the default name.
6. Click *OK*.
7. Repeat the steps above to add other series to this report.

You can mix and match categories, components, and objects in a single trend report.

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

Saved reports can be viewed under *Custom* in the left panel.

Trend report selection table

Capacity The following *Capacity* reports are available:

Component	Metric	Object
Customer Storage Usage	• Total Capacity • Total Usage (%) • Total Used	Customer
NAS Resource	• Free • Size • Used	NAS resource
Server Deduplication Repository	• Data Disk Free • Data Disk Reclaim Threshold • Data Disk Size • Data Disk Used • Data Disk Used % Before Reclamation • Data Disk Size Used % Since Reclamation Percent • Data Disk Used Before Reclamation • Data Disk Used Since Reclamation • Folder Disk Free • Folder Disk Reclaim Threshold • Folder Disk Size • Folder Disk Used • Folder Disk Used % Before Reclamation • Folder Disk Used % Since Reclamation • Folder Disk Used Before Reclamation • Folder Disk Used Since Reclamation • Index Disk Free • Index Disk Prune Threshold • Index Disk Size • Index Disk Used • Index Disk Used % Before Pruning • Index Disk Used % Since Pruning • Index Disk Used Since Pruning • Scanned Data • Unique Data	Server

Component	Metric	Object
Server Physical Storage	Metrics are for all physical resources on a server. • Capacity for NAS • Prepared Capacity • Prepared Capacity as Virtual • Reserved for Configuration Repository • Reserved for Deduplication Repository • Reserved for NAS • Reserved for Tape • Total Allocated • Total Capacity • Total Reserved Capacity • Total Used • Unprepared Capacity • Used for Configuration Repository • Used for Deduplication Repository • Used for NAS • Used for Tape • Used for Virtual Headers	Server
Storage Pool	• Prepared Capacity as Virtual • Reserved for NAS • Total Capacity • Total Used • Used as Virtual • Used for Configuration Repository • Used for NAS • Used for Tape • Used for Virtual Headers	Storage pool
Unpooled Storage	• Prepared Capacity • Total Capacity • Total Usage (%) • Total Used • Unprepared Capacity • Used for Configuration Repository • Used for NAS • Used for Tape • Used for Virtual Headers	Server

Load The following *Load* reports are available:

Component	Metric	Object
CPU	- Average Load - Usage	Server
Memory	- Buffers - Cached Memory - Deduplication Memory - Free Memory - Free Swap Memory - Total Memory - Total Swap Memory - Used Memory - Used Swap Memory	Server
Network	- Collisions - Dropped Received Packets - Dropped Transmitted Packets - Erroneous Received Packets - Erroneous Transmitted Packets - Received Bytes - Received Bytes/s - Received Packets - Transmitted Bytes - Transmitted Bytes/s - Transmitted Packets	NIC connected to the network
Network Throughput	- Read Throughput - Total Throughput - Write Throughput	NIC connected to the network

Performance The following *Performance* reports are available:

Component	Metric	Object
Client Pending Commands/Errors	- Maximum Pending Commands - Minimum Pending Commands - Pending Commands - Read Error - Write Error	Client
Device Client Errors	- Maximum Pending Commands - Minimum Pending Commands - Pending Commands - Read Error - Write Error	Virtual resource
NAS and Tape Resources	- Read Throughput - Total Throughput - Write Throughput	Server
NAS Resource	- Read Throughput - Total Throughput - Write Throughput	NAS resource
Physical Device	Metrics are for a single physical device. - Mean Latency - Read IOPS - Read Latency - Read Throughput - Total IOPS - Total Throughput - Write IOPS - Write Latency - Write Throughput	Physical device
Server NAS Resource	- Read Throughput - Total Throughput - Write Throughput	Server
Server Throughput	Metrics are for the server. - Read IOPS - Read Throughput - Total IOPS - Total Throughput - Write IOPS - Write Throughput	Server

Component	Metric	Object
Server Virtual Device	Metrics are for all virtual resources on a server. • Mean Latency • Read IOPS • Read Latency • Read Throughput • Total IOPS • Total Throughput • Write IOPS • Write Latency • Write Throughput	Server
Server Virtual Tape Drive	• Compressed Read • Compressed Write • Read (uncompressed) • Total Compressed • Total Uncompressed • Write (uncompressed)	Server
Server Virtual Tape Library	• Compressed Read • Compressed Write • Read (uncompressed) • Total Compressed • Total Uncompressed • Write (uncompressed)	Server
Virtual Device	Metrics are for a single virtual resource and can be displayed for the primary virtual resource or its mirror. • Mean Latency • Mirror Latency • Mirror Read Latency • Mirror Write Latency • Read IOPS • Read Latency • Read Throughput • Total IOPS • Total Throughput • Write IOPS • Write Latency • Write Throughput	Virtual resource

Component	Metric	Object
Virtual Tape Drive	• Compressed Read • Compressed Write • Read (uncompressed) • Total Compressed • Total Uncompressed • Write (uncompressed)	Virtual tape drive
Virtual Tape Library	• Compressed Read • Compressed Write • Read (uncompressed) • Total Compressed • Total Uncompressed • Write (uncompressed)	Virtual tape library

Scheduled reports

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

- Capacity Management - Server Physical Storage
- Capacity Management - Storage Pool
- Capacity Management - Deduplication Repository
- Server Status - Deduplicated Tape Replication History
- Server Status - Tape Deduplication and Replication
- Server Status - NAS Deduplication and Replication
- Trend
- Custom Trend
- Deduplication Tapes
- Deduplication Activity
- Deduplication Run History
- Deduplication Tape History
- Import/Export Jobs

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

- Report name
- Type of report
- Indication of whether the schedule is suspended
- Time period for the report
- Frequency of the report
- Start time
- Email recipient
- Date last sent
- Customer who created the report

From this page, you can:

- Update the report schedule and email settings.
- View the latest report.
- Suspend the report schedule.
- Delete the report and its schedule.

Schedule a report

Some reports are scheduled from the *Reports* page. Reports can also be scheduled from the *Deduplication Tapes*, *Deduplication Activity*, *Deduplication Run History*, *Deduplication Tape History*, and *Import/Export Jobs* tabs for a StorSafe server. To schedule a report:

1. Select a report.
2. Click the *Schedule Report* icon. If you are starting from the *Deduplication* or *Jobs* tab, click the *Schedule Report* button.

Specify a report name, PDF report file name, destination email address/subject/message, along with the following:

- *Server* - (For summary replication reports only) Server to include.
- *Period* - Number of days of data to include in the report, between 1 and 31.
- *Frequency* - Email frequency - daily, weekly, monthly. If weekly or monthly is selected, select the day/date.
- *Start Time* - Time of day to send the email.

If you do not specify a PDF file name, the report name will be used.

3. Click *Create*.

Smart rules



Smart rules allow you to monitor StorSafe servers for a variety of conditions, including virtual tape library/drive performance, physical device performance, and client performance. While general rules provide a default set of rules, smart rules offer an additional level of flexibility by allowing you to set thresholds, use comparison operators, and define a combination of conditions.

When a smart rule is triggered, an alert appears on the *Alerts* page. In addition, smart rules can be set to notify users via email, SMS, or custom script. Notification parameters can be set for an individual rule or can be set globally for all rules by a super administrator.

Select *Smart Rules* from the *Analytics* menu bar to see a list of existing rules, along with severity. You can also see which rules are enabled. Super Administrators will also see a *Customer* column to identify the owner of each rule. Super Administrators have access to all rules they created as well as the rules created by customer administrators.

Smart rule behavior

When a rule is created, monitoring begins immediately (assuming new data is available for the components in the rule) and the rule definition is compared against the average values for the last five minutes of metrics.

When a rule condition is reached for first time (e.g., read throughput for a virtual resource falls below 40 MB/s), an alert is created. The *Alerts* page will show the *First Occurrence* as the current time. The *Last Occurrence* will not show a time until the condition no longer exists. Monitoring occurs continuously and there is no update to the alert until the condition is resolved. At that time, the *Last Occurrence* will be updated with current time. If the rule condition is reached again, a new alert is created.

Create a smart rule

In addition to Super Administrators, the administrator for each customer can set smart rules for their servers/components.

1. Select *Smart Rules* from the *Analytics* menu bar.
2. Click the “+” icon.

The screenshot shows a dialog box titled "Add Rule" with a close button (X) in the top right corner. It contains two input fields: "Name" with the text "Server Vir Dev Performance" and "Rule Type" with a dropdown menu showing "Server - Virtual Device Performance". At the bottom right, there are "OK" and "Cancel" buttons.

3. Specify a name and select the *Rule Type*:
 - *Physical Device Performance Server Average* - Monitors all physical devices on a server. You can select multiple servers.
 - *Client Performance* - Monitors specific clients. You can select multiple server/client pairs.
 - *Physical Device Performance* - Monitors specific physical devices. You can select multiple server/physical device pairs.
 - *Virtual Tape Drive Performance Server Average* - Monitors all virtual tape drives on a server. You can select multiple servers.
 - *Virtual Tape Library Performance Server Average* - Monitors all virtual tape libraries on a server. You can select multiple servers.
 - *Virtual Tape Library Performance* - Monitors virtual tape libraries. You can select multiple server/virtual tape library pairs.
4. Click *Create*.


The rule template displays.

5. Enable the rule and define rule conditions as appropriate for the rule type.

- *Field* - The field options change depending upon the metric selected, including: Write/Compressed Write, Read/Compressed Read, Total Compressed/Uncompressed, Read/Total/Write IOPS, Mean/Read/Write Latency, Read/Total/Write Throughput
- *Operator* - =, !=, <, >, <=, >=, between, not between
- *Value* - The value field changes depending upon the metric selected, including: KBps, MBps, GBps, and ms.

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

Apply To allows you to specify the server and, depending upon the rule type, the device or client to which to apply this rule.



6. Set the severity level.

The alert severity level can be *Information*, *Warning*, or *Critical*.

7. Set notification information.

Alerts can be sent via email or SMS. You must have global email settings configured to send emails and/or SMS text messages. This is done by a Super Administrator from the *Administration* menu. A script can also be used. Refer to [“Alert scripts”](#) for more information

Modify smart rules

Select *Smart Rules* from the *Analytics* menu bar, select a rule, and select *Update* to modify existing smart rules.

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

Delete a smart rule

Select *Smart Rules* from the *Analytics* menu bar, select a rule, and select *Delete* to delete a smart rule.

General rules



General rules allow you to set a variety of configuration conditions at the server level or the device level for which you want to be notified.

When a rule is triggered, an alert appears on the *Alerts* page. In addition, rules can be set to notify users via email, SMS, or custom script. Notification parameters can be set for an individual rule or can be set globally for all rules by a Super Administrator.

Each customer can configure general rules for its own domain. All administrators for a domain share the same set of configurable general rules. The Super Administrator will not receive alerts for rules set at the customer level unless the Super Administrator has also set the same rule. This is also true in reverse; an administrator will not receive alerts set by the Super Administrator.

Server-related rules (CPU, memory, network, etc.) and service failure rules can only be enabled by Super Administrators and customer administrators if the server is in exclusive mode.

In shared mode, enabling/disabling a rule for a server being shared by multiple customers only impacts the customer who has enabled/disabled the rule; other customers will not be impacted.

Select *General Rules* from the *Analytics* menu bar to see a list of default system rules, along with any rules created for your domain. For each rule, you will see the parameters, rule checking frequency, excluded time windows, and notifications. You can also see which rules are enabled. Refer to [“Alert descriptions”](#) for more information about each rule.

Select a server to view

☐	▲ Rule Name	Active	Parameters	Checking Frequency	Excluded Time Windows	Notification
☐	Deduplication Repository Data Disk Usage	☒	Data Disk Usage: 90%	Every hour		
☐	Deduplication Repository Index Disk Usage	☒	Index Disk Usage: 90%	Every hour		
☐	Deduplication Repository Index Memory Us...	☒	Index Memory Usage: 75%	Every hour		
☐	High CPU Usage	☒	Threshold: 95%, CPU Cores: All	Every hour		
☐	High Storage Pool Usage	☒	Threshold: 90%	Every hour		
☐	Low Memory	☒	Threshold: 5%	Every minute		
☐	NAS Low Free Space	☒	Threshold: 200 GB, Minimum Size: 1024 GB	Every hour		
☐	Network Error	☒	Network Interface: eth0	Every hour		
☐	Offline Physical Device	☒	None	Every minute		
☐	Server Failover	☒	None	Every minute		
☐	Service Failure	☒	None	Every minute		
☐	Unresponsive Server	☒	None	Every minute		

Create a general rule

You can set network error and system log error rules. An alert will be generated when a specific network interface fails.

1. Select *General Rules* from the *Analytics* menu bar and select a server.
2. For Super Administrators, if this server is assigned to a customer, select a role.
To create a rule for all administrators at the customer level, select the customer.
To create a rule for Super Administrators, select *SuperAdmin*.
3. Click the “+” icon.
4. Select the type of rule:

Network Error

An alert will be generated when a specific network interface fails.

The screenshot shows a 'Create General Rule' dialog box with the following fields and controls:

- Alert Policy:** Two tabs, 'Network Error' (selected) and 'System Log Errors'.
- Active:** A blue toggle switch is turned on, with a question mark icon to its right.
- Network Device:** A dropdown menu with a downward arrow, currently empty.
- Exclusion Time Windows:** A grey toggle switch is turned off, with a question mark icon to its right.
- Notification:** A text input field with a '+' icon inside a small box on the left.
- Buttons:** 'Add' and 'Cancel' buttons at the bottom right.

System Log Errors

An alert will be generated if a message in the server's system log matches a specified pattern. The system log records important events or errors that occur in the system; it does not apply to customers. The Service Failure alert rule must also be active in order to get alerts for this rule.

The screenshot shows the 'Create General Rule' dialog box with the 'System Log Errors' tab selected. The 'Alert Policy' is set to 'System Log Errors'. The 'Active' toggle is turned on. A blue informational banner states: 'The Service Failure rule needs to be active to receive alerts for errors in the system log.' The 'Log Checking Frequency' is set to 60 minutes, and 'Memorize Incidents' is enabled. The 'Duration' is also set to 60 minutes. Below these settings is a table of log patterns:

Alert Label	Log Pattern
FC Firmware ISP Error	ISP System Error
FC Command Aborted	ABTS_RECV_IOCB
FC Port Logged Out	Port logged out
FC I/O Queue Full	QUEUE_FULL
FC Busy	TgtId+ PN\lw(2)\lw(2)\lw(2)\lw(2)\lw(2)\lw(2)\lw(2)\lw(2). * F

At the bottom, there are options for 'Exclusion Time Windows' (disabled) and 'Notification' (add button). 'Add' and 'Cancel' buttons are at the bottom right.

This rule looks at the last 20 MB of messages. If the system log is rotated prior to the *Log Checking Frequency*, the previous log is checked as well as the current log.

To limit the number of alerts for the same event, enable *Memorize Incidents* and set the *Duration*, in minutes, to remember each event. If an event is detected several times during the current timeframe, only the first occurrence is reported and the number of repetitions is indicated with the last occurrence of the message.

The *Alert Label* and *Log Pattern* message patterns that are tracked in the system log are displayed. If needed, you can add, update, or delete log patterns.

5. Select the network device and decide whether you want the rule to be active.
When a rule is active, its conditions will be checked and an alert will be generated when a network interface fails.
6. Optionally, define time periods, months, and/or days during which the policy for this rule will not be checked.

If global exclusion windows are also configured, they are combined so that both exclusion times are in effect for the rule. Up to four time windows can be defined per rule.

7. Specify notification information.

When an alert is triggered, notification is sent.


Alerts can be sent via email or SMS. You must have global settings configured to send emails and/or SMS text messages. This is done by a Super Administrator from the *Administration* menu. A script can also be used. Refer to [“Alert scripts”](#) for more information.

8. Click *Add*.

Update a general rule

This operation allows you to update a single general rule. For batch operations, refer to [“Update multiple general rules”](#).

1. Select *General Rules* from the *Analytics* menu bar and select a server.
2. For Super Administrators, if this server is assigned to a customer, select a role.
To modify a rule for all administrators at the customer level, select the customer.
To modify a rule for Super Administrators, select *SuperAdmin*.
3. Select a rule, and select *Update*.



General rules can be modified in the following ways:

- Enable/disable the rule.
- Specify time periods during which checking will be skipped per rule condition. Up to four time windows can be defined per rule.
- For rules with parameters (threshold, etc.), you can customize those parameters to fit your environment.
- Enable/disable email and SMS in order to send an email and/or SMS text message when the alert is triggered.

Delete a general rule

You can only delete a general rule that you created. Select *Analytics*, click *General Rules*, select the rule, and click *Delete*.

Update multiple general rules

You can enable, disable, or delete multiple general rules.

Note that you can only delete a general rule that you created.

1. Select *General Rules* from the *Analytics* menu bar and select a server.
2. Select one or more rules.
3. Click the *Batch Actions* icon and select an operation.

Set global time exclusions


Select *General Rules* from the *Analytics* menu bar and select *Global Exclusion Time Windows* to specify times, months, and days during which policies for general rules will not be checked.

Click the “+” icon to add the start and end time, months, and/or days to exclude checking. If exclusion windows are also configured for an individual general rule, they are combined so that both exclusion times are in effect for the rule.

Apply rules to multiple servers

You can apply one or more rules to other servers. This is convenient if you are managing multiple servers and want to set rules globally.

1. Select *General Rules* from the *Analytics* menu bar and select a server.
2. Select one or more rules, and select *Copy Rules to Servers*.



3. Select the servers to which the rules will be copied and click *Copy*.

Alerts



Alerts display as a result of the rules configured. When the condition occurs, an alert is triggered.

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

Server	Severity	Rule Name	Resource	First Occurrence	Last Occurrence	Details
StorSafe1101-h6141	Error	Offline Physical Device	IBM:ULTRIUM-TD4	2020-11-30 21:46:07	2020-12-15 13:59:09	The Physical Device IBM:ULTRIUM-TD4 (ID: fa...
VTL01	Error	Offline Physical Device	DELL:PERC 6/i	2020-11-30 21:46:07	2020-12-15 13:59:06	The Physical Device DELL:PERC 6/i (ID: 56fat...
VTL01	Error	Offline Physical Device	DELL:PERC 6/i	2020-11-30 21:46:07	2020-12-15 13:59:06	The Physical Device DELL:PERC 6/i (ID: abd9...
StorSafe1101-h6141	Error	Offline Physical Device	XIOTECH:ISE3401	2020-11-30 21:46:07	2020-12-15 13:59:09	The Physical Device XIOTECH:ISE3401 (ID: fa...
StorSafe1101-h6141	Critical	Unresponsive Server	StorSafe1101-h6141	2020-12-03 03:17:05	2020-12-15 05:51:43	Server StorSafe1101-h6141 is unresponsive. C...
StorSafe1101-h6141	Error	Offline Physical Device	XIOTECH:ISE3401	2020-11-30 21:46:07	2020-12-15 13:59:09	The Physical Device XIOTECH:ISE3401 (ID: 0f...
StorSafe1101-h6141	Error	Offline Physical Device	XIOTECH:ISE3401	2020-11-30 21:46:07	2020-12-15 13:59:09	The Physical Device XIOTECH:ISE3401 (ID: 3f...

By default, only unacknowledged alerts are listed. You can change the display criteria to show all alerts or only acknowledged alerts. You can also customize your view by selecting a particular server, level of severity, rules, occurrence date, etc. Super Administrators can filter the list by customer.

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

Each alert includes the server name, severity, rule name, resource, first and last occurrence, as well as details about the alert condition. If the condition continues, the *Last Occurrence* column will contain the current date and time.

When a smart rule condition is reached, the *Last Occurrence* column will show the alert with a “Still Active” state. The state remains that way until the condition no longer exists. When this occurs, the time of last occurrence will change to the date and time the condition was no longer met.

Occurrence dates can be sorted in ascending or descending order by clicking the column header. The sort arrow faces up or down depending upon the sort direction.

For smart rules, alerts will only display the rule condition that triggered the alert. For example, in an “or” rule, whichever condition triggered the alert is shown. In an “and” rule, the last condition will be displayed since that is actually the trigger.

To see additional information about an alert, highlight the alert and click the “*i*” *Information* icon.

If the event description is missing for an alert, use the *Force Rediscovery* option for that server in the *Administration* menu to reload event log messages.

Alert descriptions

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

Rule Name - Alert	Description	Interval	Level
Deduplication Repository Data Disk Usage	Deduplication Repository data disk usage is over a specified percentage (default 90%).	Hourly	Warning
Deduplication Repository Index Disk Usage	Deduplication Repository Index disk usage is over a specified percentage (default 90%).	Hourly	Warning
Deduplication Repository Index Memory Usage	Deduplication Repository Index memory usage is over a specified percentage (default 75%).	Hourly	Warning
High CPU Usage	Server CPU usage is over a specified percentage. Parameters: All (average usage of all CPU cores is above the threshold value) or Any (any of the CPU cores' usage is above threshold value), Threshold (default 95%).	Hourly	Warning
High Storage Pool Usage	Available disk space usage in a storage pool is over a specified threshold. Parameter: Threshold (default 90%)	Hourly	Warning
Low Memory	Available system memory is below a specified percentage or amount. Parameters: Threshold (default 5%), Unit (% or MB)	Every minute	Warning
NAS Low Free Space	Free disk space on a NAS resource is below a specified threshold (default 200 GB) for any resource larger than a specified size (default 1024 GB)	Hourly	Warning
Network Error	Network errors, overruns, dropped frames, or network collisions have occurred since the last check. Parameters: Network interface name	Hourly	Warning
Offline Physical Device	A physical device is offline.	Every minute	Critical
Out-of-sync Mirror	A mirror of a virtual device is out-of-sync.	Every minute	Error
Server Failover	A StorSafe server failed over to its partner.	Every minute	Error

Rule Name - Alert	Description	Interval	Level
Service Failure	A critical event or an error is recorded in the event log indicating a service failure since the last check. The following critical events are reported: • Out-of-sync mirror • Critical failure on server • Failed fakeover • Configuration repository failure • Tape database inconsistency • Inconsistent deduplication repository • Deduplication index disk failure • Invalid deduplication index • Read-only deduplication repository • Deduplication folder disk failure • Insufficient deduplication index • Near-full deduplication index • Near-full deduplication folder • Near-full deduplication data disk • Near-full deduplication index cache • NAS failure The following errors are reported: • Replication failure • Reclamation failure • Remote Copy failure • Export/import job failure • Shredding job failure • Stacking job failure • Deduplication failure • File system error	Every minute	Error
Unresponsive Server	Main server modules (including the alert module) are not running resulting in the server becoming unresponsive.	Every minute	Critical

Manage alerts

You can acknowledge, unacknowledge, or delete an event from the *Manage* icon.

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

Acknowledged alerts are retained for 7 days; non-acknowledged alerts are retained for 14 days. You can manually delete alerts at any time.

To manage alerts:

1. Select *Alerts* from the *Analytics* menu bar.
2. Select one or more alerts.
3. Click the *Manage* icon and select an operation.

To acknowledge, unacknowledge, or delete all alerts on the displayed page, select the top checkbox, click the *Manage* icon, and then select an operation.

Capacity prediction



The *Capacity Prediction* feature uses an algorithm based on the past 30 days of server resources to anticipate future capacity requirements.

To see a trend analysis for the selected server:

1. Select *Capacity Prediction* from the *Analytics* menu bar.
2. From the drop-down menu, select the server and then an object in the table for which you would like an analysis.

You can view information for the server, a specific storage pool, or server-level information for all tapes and the deduplication repository.


For all objects, you can view the capacity and trend prediction graph for past, present or future time periods. In addition, you can mouse over graphs and predictions for more details.

Keep in mind that sufficient data is required for accurate trend predictions. If you have less than 30 days of data, a forward prediction will be made using simulated data points based upon current usage.

Predictions for each server and storage pool are recalculated every six hours, increasing the number of valid data points. The more valid data points, the more accurate the prediction.

When adding a new server, predictions start displaying after five minutes. However, accurate predictions cannot be made until an adequate amount of historical data is available.


The *Capacity Trend and Prediction* graph illustrates usage based upon current usage. Beneath the graph, you can click on the legend to change the view. *Current* displays current data usage. *Future* displays predictive data usage graph. For resources with less than seven days of activity, you can also display *Extrapolated History* for historical data.



The *Current Allocation* area shows the current usage as well as the daily growth rate, which is calculated as usage divided by capacity. The growth rate can be impacted by the addition of storage or the deletion of data.

The *Prediction Based on Usage Thresholds* area shows the predicted number of days until the storage will be full.

The *Prediction Based on Time Thresholds* shows the amount of space available. You can hover over the blue area of the ring to see the usage percentage or hover over the black area of the ring to see percentage of available space.



Audit log



The audit log tracks and logs system activities for all servers by all administrators and users. Information in the audit log is maintained for 365 days. Only a Super Administrator or Administrator can access the audit log.

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

Date/Time	Customer	User	User IP	Server	Action	Component Type	Component ID	Component Name	Result
2019-01-29 21:01:07		superadmin	10.197.255.10		Login	FMS			✓
2019-01-29 14:59:38		superadmin	10.5.3.119	IPS9-3	Update	Physical Device			✓
2019-01-29 14:55:30		superadmin	10.5.3.119	IPS9-3	Add Server	Server			✓
2019-01-29 14:53:54		superadmin	10.5.3.119	IPS9-L3	Add Server	Server			✓
2019-01-29 14:50:13		superadmin	10.5.3.119		Login	FMS			✓
2019-01-29 09:14:31		superadmin	10.5.3.119		Login	FMS			✓
2019-01-28 17:13:53		superadmin	10.5.3.112		Login	FMS			✓
2019-01-28 14:42:27		superadmin	10.5.3.112	2018Q49557-1497	Create	Virtual Device	711	SANDisk-000711	✓
2019-01-28 14:41:09		superadmin	10.5.3.112		Login	FMS			✓
2019-01-28 13:45:19		superadmin	10.5.3.112		Login	FMS			✓
2019-01-28 10:05:09		superadmin	10.5.3.112		Login	FMS			✓
2019-01-28 06:00:33		superadmin	10.197.255.10		Login	FMS			✓
2019-01-28 00:29:49		superadmin	10.197.255.10		Login	FMS			✓

You will see the following for each entry:

- Date/time of occurrence
- Customer - In a multi-tenancy environment, the Super Administrator will see all customers managed by the StorSight server; administrators will only see entries for their domain.
- User who performed the activity
- User IP address
- Server name, if its configuration was changed; the field will be blank for StorSight activities such as log in, log out, and add server.
- Action performed
- Component type - Client, client agent, configuration repository, deduplication job, deduplication policy, object storage, patch, physical adapter, physical device, physical tape drive, physical tape library, replica tape, replication job, server, storage pool, StorSight, tape import/export job, user, virtual resource, virtual tape, virtual tape drive, or virtual tape library
- Component ID for client, consistency group, virtual resource
- Component name
- Result - Click in the *Result* field to see more details. Expand the *Object* field to see detailed information.

You can customize your view by selecting a particular server, component type and component ID, action, start/end date, or user. Super Administrators can also select a specific customer. To clear the search criteria, click the *Clear Search* icon.

Chargeback



Chargeback is a feature that tracks tape ingest data per customer for all managed StorSafe servers. Chargeback reports enable service providers to easily invoice customers. Chargeback settings and reports are only available for Super Administrators.

Chargeback settings

Chargeback settings define the price for incoming backup data on all tapes, as well as the currency used, discounts, and the maximum cash back limit.

Chargeback settings can be global or can be set for a specific customer.

Global settings

Global settings define usage and discounts for all customers. A chargeback policy defined for a specific customer will override these global settings.

To access global settings, select *Analytics* from the menu bar, click *Chargeback*, click *Settings*, and then click the *Global Settings* tab.

The screenshot shows the 'Global Settings' tab for Chargeback. It includes the following settings:

- Currency:** A dropdown menu set to 'US Dollar (\$)'. A question mark icon is next to it.
- StorGuard Used Storage:** A toggle switch that is turned on. A question mark icon is next to it.
- Storage Price:** A text input field containing '1' and a dropdown menu set to 'TB'. A question mark icon is next to it.
- StorGuard IO:** A toggle switch that is turned off. A question mark icon is next to it.
- StorSafe Tape Ingest Data:** A toggle switch that is turned on. A question mark icon is next to it.
- Ingest Price:** A text input field containing '10' and a dropdown menu set to 'TB'. A question mark icon is next to it.
- Discount:** A toggle switch that is turned off. A question mark icon is next to it.

A 'Save' button is located at the bottom right of the settings panel.

Currency - Select the currency used.

StorSafe Tape Ingest Data - Click to enable. Specify the charge for incoming backup data on all tapes in the *Ingest Price* field.

Discount - Click to enable a discount. Set the percentage customers will receive. You can set the *Minimum Tape Ingest Data* field as criteria for receiving the discount. You can also set a maximum cash back limit (zero means unlimited).

Customer settings

Customer settings define currency, charges, and discount for a specific customer. A chargeback policy defined for a specific customer will override the global settings.

To access these settings, Select *Analytics* from the menu bar, click *Chargeback*, click *Settings*, and then click the *Customer Settings* tab.

To add a policy for a customer: click the “+” icon, select a customer, and then configure chargeback settings for this customer. Refer to [“Global settings”](#) for information about the fields.

Chargeback report

The chargeback report displays information per customer for all managed servers. The report enables service providers to easily generate bills for customers.

To access the report, Select *Analytics* from the menu bar, click *Chargeback*, and then click *Report*.

You can select to view the report for all customers (the default) or for a specific customer. You can also specify a time period (the default is yesterday).

For each StorSafe customer, you will see the cost based on global or customer settings, including any discount, and the amount of ingest data.

Monitor

Monitor provides an overview of information about StorSafe servers, including storage for all servers and performance, host clients, virtual resources, and physical resources for each server.

Overview page




The *Overview* page allows you to view storage capacity and usage for all servers.

Servers can be sorted by name, capacity, free space, free space percentage, and throughput.

Servers can be filtered by type.

Hover your mouse over the graph to see more detailed information.

Click to see the alerts for this server.




Server view page



Select a server from the list of the left side of your screen to view information about that server.

Note that clustered servers are displayed individually here as opposed to the *Manage* pages, where clustered servers are displayed within their cluster.



Overview page

The *Overview* page displays current memory usage, CPU usage, storage usage, usage history, and deduplication usage for the selected server. There are tabs for the server that vary by server type.

Server Overview

Server Overview displays the name of the selected cluster/server, IP address, version, appliance type (virtual or physical), uptime of the server, cluster ID (if applicable), operating system, kernel, processor, and FalconStor appliance serial number. Click *Server Ports* to view the status of the server's ports and modules.

The *Inventory* section displays the number of each resource type.

The *System Resources* section displays CPU usage, number of cores, as well as used and free memory. For SIR and StorSafe-S, used and free memory are calculated based on the system memory left behind at the operating system level after memory pre-allocation for deduplication.


Storage Usage

Storage Usage displays the total managed capacity, which includes reserved/free storage as well as storage utilized for the resources on the selected server or storage pool; unprepared storage is not included.

By default, the server is displayed but you can select to show information for a specific storage pool if you click the drop-down box.


Click the *More Info* link to display the total server capacity (reserved and non-reserved) and the amount of storage reserved for the resources, as well as the amount of storage used by the resources as a percentage of reserved space.

Usage History



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

Deduplication Usage



(SIR servers and StorSafe servers with deduplication) *Deduplication Usage* graphically displays the current state of deduplication storage for the cluster. Values are based on all scans performed during the life span of the selected server.

- Deduplication ratio
- Scanned data - Amount of data scanned for deduplication
- Unique data - Amount of deduplicated data stored in the repository
- Folder disk - Capacity and how much space has been used. Highlight the chart to see the reclamation threshold percentage, as well as folder disk usage before and since reclamation.
- Index disk - Capacity and how much space has been used. Highlight the chart to see the pruning threshold percentage, as well as index disk usage before and since pruning.
- Data disk - Capacity and how much space has been used. Highlight the chart to see the reclamation threshold percentage, as well as data disk usage before and since reclamation.
- Index memory - Highlight the chart to see the reclamation threshold percentage, as well as index memory usage before and since reclamation.

Virtual Tape Libraries page

Libraries page

The *Libraries* page displays all virtual tape libraries for this server along with:

- ID
- Vendor
- Product
- Firmware
- Serial number
- Media type
- Number of slots
- Number of drives
- Number of tapes
- Number of loaded tapes
- Number of clients (click to see a list)
- Starting/ending barcodes
- Storage pool
- Capacity on Demand, Auto Archive, Auto Replication, Cloud Tiering usage and mode
- Tape Caching, Tape Duplication , Encryption usage
- Read and write throughput
- Compressed read and write - If compression is enabled, the read/write throughput and read/write compressed values will be different. If compression is not enabled, the comparable values (i.e., write throughput and compressed write) will be the same.
- Configured options - Tape capacity on demand, auto archive, auto replication, auto tape caching, tape encryption, tape duplication

Click on a link to display a dynamic graph that shows information for the last 15, 30, 45, and 60 minutes. Mouse over any point on the graph for more information. Use the slider at the bottom to adjust the granularity of the graph.

Drives

The following information is displayed for each tape drive in a virtual tape library:

- ID
- Status - Empty, loaded, ejected
- Vendor
- Product
- Firmware
- Serial number
- Media type
- Element number
- Number of clients (click to see a list)
- Read and write throughput
- Compressed read and write - If compression is enabled, the read/write throughput and read/write compressed values will be different. If compression is not enabled, the comparable values (i.e., write throughput and compressed write) will be the same.

Click on a link to display a dynamic graph that shows information for the last 15, 30, 45, and 60 minutes. Mouse over any point on the graph for more information. Use the slider at the bottom to adjust the granularity of the graph.

Vault

The following information is displayed for each virtual tape in the vault:

- Barcode
- Tape name
- Tape ID
- Size
- Used size
- Tape status - online, offline, incomplete (missing physical segments)
- Virtual tape index status for deduplicated tapes - pure (all tape data has been deduplicated), mixed (tape contains pointers to deduplicated data as well as data that has not yet been deduplicated), plain (tape not yet deduplicated)
- Parent virtual tape library ID (click for details)
- Cloud Tiering mode - Copy or Move
- Object storage provider ID
- Indication if Cloud Tiering occurred
- Stub - Checked if Cloud Tiering occurred and a stub tape exists
- Write-protected status
- Write-once-read-many (WORM) tape usage
- Encrypted tape indicator
- Deduplication policy

- Deduplication status - Successful, pending, failed
- Replication is/not configured
- Replication server
- ID of replica tape
- Replication source server
- ID of source tape
- Replication of deduplicated tapes - Status on target servers for replication of deduplicated tapes
- Source of replica promoted in test mode - Checked if this is a primary tape promoted from a replica in test mode
- Previously replica promoted in test mode - Checked if this is a replica tape promoted from a replica in test mode
- Tape caching is/not configured
- Unmigrated cache - Checked if tape caching is used and the tape has unmigrated data
- Direct-Link - Checked if tape caching is used and the tape is a direct link tape (a link to the physical tape)

Cloud Tier

The following information is displayed for each virtual tape in the vault that has been migrated:

- Source tape name
- Barcode
- Tape ID
- Size
- Used size
- Virtual tape index status - pure (all tape data has been deduplicated), mixed (tape contains pointers to deduplicated data as well as data that has not yet been deduplicated), plain (tape not yet deduplicated)
- Cloud Tiering migration mode - Copy or Move
- Object storage provider
- User access key
- Object storage provider ID (click for object storage account details)
- Encrypted tape indicator
- Stub - Checked if Cloud Tiering occurred and a stub tape exists
- Write-once-read-many (WORM) tape usage
- Parent virtual tape library ID (click for library details)

Tapes

The following information is displayed for each virtual tape:

- Barcode
- Tape name
- Tape ID
- Size
- Used size
- Tape status - online, offline, incomplete (missing physical segments)
- Location (click to see the tape library name)
- Parent virtual tape library ID (click for details)
- Cloud Tiering mode - Copy or Move
- Object storage provider ID
- Indication if Cloud Tiering occurred
- Stub - Checked if Cloud Tiering occurred and a stub tape exists
- Write-protected status
- Write-once-read-many (WORM) tape usage
- Encrypted tape indicator
- Deduplication policy
- Deduplication status - Successful, pending, failed
- Replication is/not configured
- Replica server
- Replica tape ID
- Replication source server
- Source tape ID
- Replication of deduplicated tapes - Status on target servers for replication of deduplicated tapes
- Source of replica promoted in Test Mode - Checked if this is a primary tape promoted from a replica in test mode
- Previously replica promoted Test Mode - Checked if this is a replica tape promoted from a replica in test mode
- Tape caching usage
- Unmigrated cache - Checked if tape caching is used and the tape has unmigrated data
- Direct-Link - Checked if tape caching is used and the tape is a direct link tape (a link to the physical tape)

Deduplication page

Policies

The following information is displayed for each deduplication policy:

- Policy name
- Policy ID
- Number of tapes
- Status - Idle, Running, Suspended, Stopping, Stopped, Preparing, Queued
- Indication if the policy is suspended
- Last time the policy was run
- Next time the policy will run
- Deduplication trigger
- Replication mode
- Replica server
- Indication if replication is suspended
- Deduplication cluster to which this policy belongs

Tapes

The following information is displayed for each virtual tape in a deduplication policy:

- Barcode
- Tape
- ID
- Location
- Deduplication status - Successful (pure VIT), Pending (or in-progress, not a pure VIT), Failed (not a pure VIT)
- Deduplication state - Completed, Aborted, Empty Tape, Failed, Paused, Preparing, Queued, Running, Stale, Stopped Writing to VIT (writing back the deduplicated data to VIT), Writing to VIT Finished, Writing to VIT Failed
- Last deduplication time
- Deduplication ratio - 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.
- Maximum uncompressed storage capacity (determined when the tape was created)
- Allocated tape size
- Written - Amount of data (before compression) that is written to the tape by backup applications. This amount can be greater than the tape size if the data is compressed.
- Stored - Amount of data (before compression) stored in the deduplication repository.

- Unique data - Actual physical storage used to store tape data after deduplication and compression.
- New data - Amount of data (before compression) that has not yet been deduplicated, including newly appended data to a tape.
- Replication status - Indicated on the source server for a tape in a deduplication policy, on the target server for a deduplicated LVIT. For cascaded or parallel replication, the status is indicated for both first and second replication target servers.
 - Completed
 - Pending - Or in-progress replication on the source server
 - Failed - Can be due to replication failure, tape resolution failure, or in-progress tape resolution
 - Completed Pending - For cascaded or parallel, successful replication on server 1 and pending replication on server 2
 - Completed Failed - For cascaded or parallel, successful replication on server 1 and failed replication on server 2
 - Pending Completed - For cascaded or parallel, pending replication on server 1 and successful replication on server 2
 - Pending Failed - For cascaded or parallel, pending replication on server 1 and failed replication on server 2
 - Failed Completed - For cascaded or parallel, failed replication on server 1 and successful replication on server 2
 - Failed Pending - For cascaded or parallel, failed replication on server 1 and pending replication on server 2
- Replicated - Amount of unique data replicated
- Last replication time
- Last replication status
- Replica server
- Resolution - Status on target server for replication of deduplicated tapes. When resolved, the target server automatically creates an LVIT and puts it in the target server's vault or in a virtual tape library, depending upon how the deduplication policy was configured.

Activity

This page is divided into two section and displays information about currently running deduplication policies and replication jobs: In the top section, the following information is displayed for each tape currently being deduplicated:

- Barcode
- Tape
- ID
- Serial number
- Parser
- Size of data
- Amount of data scanned

- Throughput
- Number of blocks
- Status - Queued, Running, Completed, Failed, Paused, Preparing, Scan Paused, Writing VIT (writing the deduplicated data back to the VIT), Writing VIT Paused, Writing VIT Finished, Writing VIT Failed, Writing VIT Preparing, Writing VIT Stopped, Stopped
- Progress - Click *Refresh* to update the data.

While replication is occurring, you will see status displays related to the Index and unique replication phases on the bottom part of the screen.

- Barcode
- Tape
- ID
- Replication phase - Index (transmitting index data) or Unique (transmitting unique deduplication data)
- Data size - Amount of data to be replicated
- Transmitted - Amount of transmitted data
- Throughput
- Deduplication ratio on the replication target server after the data is replicated and resolved
- Progress - Click *Refresh* to update the data.
- Remaining time

Run History

The following historical information is displayed for each deduplication policy:

- Start time
- Trigger - Manual, Inline, Schedule, Tape eject, Tape caching
- Number of tapes
- Total amount of data to deduplicate
- Total amount of unique data written to the repository
- Deduplication ratio - 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
- Deduplication duration (in seconds)
- Total amount of data to replicate
- Unique data to replicate
- Deduplication ratio on the replication target server after the data is replicated and resolved
- Replication duration (in seconds)
- Status - Completed, Canceled, Failed

Tape History

The following historical information is displayed for each tape in a deduplication policy:

- Barcode
- Tape
- Tape ID
- Parser
- Deduplication start time
- Status - Completed, Canceled, Failed
- Total amount of data to deduplicate
- Total amount of unique data written to the repository
- Deduplication ratio - The ratio between the scanned data and the 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
- Deduplication duration (in seconds)
- Total amount of data to replicate
- Unique data to replicate
- Deduplication ratio on the replication target server after the data is replicated and resolved
- Replication duration (in seconds)

Replication Targets

The following information is displayed for each replication target on the selected SIR server:

- Replica server
- Encryption usage

Jobs page

Deduplication

The *Deduplication* page displays all of the virtual tapes that are being processed. The following information is displayed for each tape:

- ID
- Barcode
- Deduplication policy
- Status - Suspended, queued, tape in drive, preparing for deduplication, deduplicating
- Size of undeduplicated data
- Throughput
- Progress percentage - Click *Refresh* to update the data.

Replication

The *Replication* page displays information for replication jobs for virtual tapes. The following information is displayed for each tape:

- ID
- Barcode
- Type - Deduplicated or without deduplication
- Mode - Non-deduplicated tape to a single server, deduplicated tape to a single server, parallel concurrent for deduplicated tapes, parallel serialized for deduplicated tapes, cascaded for deduplicated tapes
- Status - Running, waiting for retry, waiting for a slot

Unique replication

The *Replication* page displays information for incoming replication jobs for deduplicated tapes. It includes tapes currently replicating (after the index has been replicated) and those awaiting replication. The following information is displayed for each tape:

- Source StorSafe server
- Source tape barcode
- Deduplication policy
- Replica barcode
- Start time
- Status - Running, waiting for retry

Import/Export

The *Import/Export* page lists current import, export, and automated tape caching jobs. The following information is displayed for each job:

- ID
- Type of job - Export to standalone drive/library, import from standalone drive/library, create cache and copy metadata, move to import/export slot, export to standalone drive/library with stacking, import from library with stacking, move to import/export slot with stacking, scan physical tape in standalone drive/library, export tape to object storage, import from object storage
- Mode
- Virtual tape ID
- Virtual tape barcode
- Physical tape barcode
- Start time
- Status - Waiting for tape/drive, running, completed, cancelled, failed, canceling, waiting for I/E slot, waiting to be deleted, hold, resumed

NAS Dedupe/Replication

The *NAS Dedupe/Replication* page lists deduplication and outgoing replication statistics. The following is displayed for each NAS resource:

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

NAS In-Progress Replication

The *NAS In-Progress Replication* page lists current outgoing replication sessions at the server level for all NAS resources. The following is displayed for each file being processed:

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

Object Storage page

The *Object Storage* page displays object storage provider accounts for this server along with:

- Account label
- ID
- Provider
- User access information
- Bucket
- Endpoint region
- Encryption usage
- Amount of data exported
- Comment

NAS page

NAS resources

The *NAS Resources* page displays all NAS resources for this server. The following information is displayed for each:

- ID
- Status - Online/offline
- Mount point - The directory from which the share can be accessed
- File system status
- Total, allocated, and free size
- Data size - Actual size of the data written to the share before deduplication
- Number of shares
- Number of folders
- Storage pool
- Deduplication - Whether the NAS resource is included for deduplication

- Replication - Whether the NAS resource is included for replication
- Read, write, and total throughput - Click on a link to display a dynamic graph that shows information for the last 15, 30, 45, and 60 minutes. Mouse over any point on the graph for more information. Use the slider at the bottom to adjust the granularity of the graph.

Shares

The *Shares* page displays all NAS shares for this server. The following information is displayed for each:

- Share name
- NFS Clients (click to see the list with options)
- SMB clients (click to see the list with options)
- SMB share
- SMB comment
- Maximum number of SMB users
- SMP properties
- Indication of whether or not the NAS share is deduplicated (fully or partially)
- Indication of whether or not the NAS share is replicated (fully or partially)

NFS Clients

The *NFS Clients* page displays NFS clients on the selected StorSafe server and the NAS shares assigned to a client. The following information is displayed for each:

- Client user name
- ID
- Name or IP address of the client machine
- Description

Select a client from the *Client* drop-down box to view properties for each NAS share assigned to the client:

- Full shared directory path
- Access rights - read-write or readonly
- Permissions

SMB Clients

The *SMB Clients* page displays SMB clients on the selected StorSafe server and the NAS shares assigned to a client. The following information is displayed for each:

- User name
- ID
- Description
- List of groups to which this client belongs

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

- Full shared directory path
- Access rights - read-write or read-only
- Share name
- Share comments
- Maximum number of user connections

Host Clients page

The *Overview* page displays all clients for this server along with:

- Client ID
- Type (operating system)
- Protocol (FC and/or iSCSI)
- Number of iSCSI initiators and targets (click to see a list)
- Number of FC WWPNS (click to see a list)
- Number of assigned devices (click to see a list with access mode)
- Number of read and write IOPS - The number of operations in bits per second.
- Read and write throughput - The amount of throughput consumed by operations in bits per second.
- User for authentication

Click on a link to display a dynamic graph that shows information for the last 15, 30, 45, and 60 minutes. Mouse over any point on the graph for more information. Use the slider at the bottom to adjust the granularity of the graph.

Performance

(FC clients only) The *Performance* page displays throughput, and IOPS performance trends in real time, for the last 30 minutes, or for the last hour for a selected client. You can highlight a point on a graph to see more detailed information.

You can select different clients using the drop-down box.

Drill Down

The *Drill Down* page displays performance for a host client to its assigned virtual resources and underlying physical devices. You can select different clients using the drop-down box.

The screenshot shows the 'Drill Down' page for client 6U5ESX6211. The interface includes tabs for Overview, Performance, Virtual Devices, Host Clients, Physical Resources, and FC/iSCSI Targets. The 'Host Clients' tab is active, and the 'Drill Down' sub-tab is selected. The client details table shows one client with ID 29, IP 172.22.6.211, and VMWare type. Below this, the 'Virtual Devices' section lists five virtual devices, all of type Thin, using the 6u5VMDK-G1 storage pool and connected to the XIOTECH.ISE3401... physical device. The 'Physical Devices' section at the bottom shows one physical device, XIOTECH.ISE3401.001, which is virtualized and has a size of 3 TB.

Client	ID	IP Address	Type	FC	iSCSI	Devices	iSCSI Initiators	iSCSI Targets
6U5ESX6211	29	172.22.6.211	VMWare	☐	☒	5	iqn.1998-01.com.vm...	1

Virtual Device	ID	Type	Group	Storage Pool	Physical Device	Total Size	Allocated Size	Clients
ESX6U5-VMDK1-01681	1681	Thin	6u5VMDK-G1	Pri	XIOTECH.ISE3401...	20 GB	11 GB	2
ESX6U5-VMDK1-01682	1682	Thin	6u5VMDK-G1	Pri	XIOTECH.ISE3401...	20 GB	20 GB	2
ESX6U5-VMDK1-01683	1683	Thin	6u5VMDK-G1	Pri	XIOTECH.ISE3401...	20 GB	20 GB	2
ESX6U5-VMDK1-01684	1684	Thin	6u5VMDK-G1	Pri	XIOTECH.ISE3401...	20 GB	11 GB	2
ESX6U5-VMDK1-01685	1685	Thin	6u5VMDK-G1	Pri	XIOTECH.ISE3401...	20 GB	1 GB	2

Physical Device	ACSL	Status	Category	Reservation	Size	Allocated Size	Free Size	Storage Pool
XIOTECH.ISE3401.001	100:0:1:1		Virtualized		3 TB	560.74 GB	2.45 TB	Pri

The page is divided into three sections, displaying information about the selected client, the virtual resources assigned to the client, and the physical devices used by the virtual resources. It includes IOPS, latency, and throughput performance information in real time, for the past 30 minutes, or for the past hour.

Physical Resources page

Storage Pools

The *Storage Pools* page displays all storage pools on this server along with the following default fields:

- Name
- Full, allocated, and free size
- Storage pool resource type - All, tape, NAS resource
- Number of physical devices (click to see a list)
- Customer - Super Administrators only

Refer to [“Storage pools”](#) for more information.

Physical Devices

The following information is displayed by default for each physical device:

- Name
- ACSL
- Status
- Device category - Virtualized, foreign, unassigned
- Full, allocated, and free size
- Reservation type - Tapes, NAS resources, deduplication repository, configuration repository
- Storage pool (StorSafe servers)
- IOPS - The number of operations in bits per second.
- Latency
- Throughput - The amount of throughput consumed by operations in bits per second.
- Owner
- GUID

Click on a link to display a dynamic graph that shows information for the last 15, 30, 45, and 60 minutes. Mouse over any point on the graph for more information. Use the slider at the bottom to adjust the granularity of the graph.

Refer to [“Physical devices”](#) for more information.

Physical Tape Libraries

The following information is displayed by default for each physical tape library:

- Name
- ID
- StorSafe server

- Vendor
- Product model
- Firmware
- Serial Number
- ACSL
- Number of slots, tapes, and drives

Refer to [“Physical Tape Libraries”](#) for more information.

Physical Tape Drives

The following information is displayed by default for each physical tape drive:

- Name
- ID
- StorSafe server
- Vendor
- Product model
- Firmware
- Serial Number
- ACSL
- Status - online or offline
- Enabled/disabled

Refer to [“Physical tape drives”](#) for more information.

Adapters

Adapters include the Fibre Channel and SCSI physical adapters on a server. The following information is displayed by default for each adapter:

- Name
- Status - Link up or link down
- Number
- Adapter information
- FC adapter mode - Dual, initiator, target, standby
- WWPN and target WWPN

Refer to [“Physical Resources”](#) for more information.

FC/iSCSI Targets page

The *FC/iSCSI Targets* page displays all FC and iSCSI targets on this server along with performance information.

- Read IOPS - The number of read operations in bits per second.
- Write IOPS - The number of write operations in bits per second.
- Read Throughput - The amount of throughput consumed by read operations in bits per second.
- Write Throughput - The amount of throughput consumed by write operations in bits per second.

Click on a link to display a dynamic graph that shows information for the last 15, 30, 45, and 60 minutes. Mouse over any point on the graph for more information. Use the slider at the bottom to adjust the granularity of the graph.

Client View page




The *Client View* page displays information about clients on storage servers managed by the StorSight server.

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 StorSafe 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 StorSafe 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 a Fibre Channel or iSCSI connection between the IBM i system host and the StorSafe appliance. For more information about iSCSI configuration, refer to [“iSCSI Configuration”](#).
- IBM i systems and the StorSafe appliance must use certified FC HBAs or iSCSI initiators. Refer to the StorSafe release notes for a list that are currently certified.

Set up the tape library

With StorSafe 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.

In addition, for StorSafe import/export, you must use one of the supported tape drives in the physical tape library. This ensures that you have 1:1 data transfer between the virtual volume and the physical tape media.

Once you have created a virtual tape library and assigned it to the host, you should perform the following tasks to ensure that the 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 portal, 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 portal, you can see the tape in the virtual vault.

SNMP Integration

StorSafe provides Simple Network Management Protocol (SNMP) support to integrate management of StorSafe servers into an existing enterprise management solution, such as HP OpenView, CA Unicenter, IBM Tivoli NetView, or BMC Patrol.

StorSafe 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 StorSafe to send certain types of messages as traps to your SNMP manager. Refer to [“Enable SNMP”](#) to configure this.
- MIBs - Each MIB (Management Information Base) monitors information and processes in StorSafe. You will need to compile the StorSafe MIBs into your SNMP manager. The procedure to do this will vary by SNMP manager. You will need to compile the following StorSafe 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, StorSafe 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.

Enable SNMP

1. Select *Manage* from the menu bar and select a StorSafe server.
2. Click the *Settings* tab and click the *SNMP Settings* icon.

Server Overview Virtual Libraries Deduplication Jobs Object Storage Host Clients Physical Resources Encryption **Settings**

◀ Back SNMP Settings

Simple Network Management Protocol (SNMP) is used to monitor the server and send traps for fault conditions. After modifying the settings on this page, click Update.

Event Severity for Traps: ⓘ

Location: ⓘ

Contact: ⓘ

SNMP Server	SNMP Commu...	Version	User	Authentication
Showing 0 entries				

Event Severity for Traps - 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.

Location - Enter the location of your system.

Contact - Enter contact information. This can be a name or an email address.

3. Click the “+” icon to add an SNMP server (maximum five).

If you are configuring SNMP version 3 (default), you will see the following:

The screenshot shows a dialog box titled "Add SNMP Server" with a close button (X) in the top right corner. The dialog contains the following fields and controls:

- Version:** A dropdown menu with the value "3" selected.
- SNMP Server:** A text input field with a blue question mark icon to its right.
- Username:** A text input field.
- Authentication Type:** A dropdown menu with "None" selected.
- Encryption:** A dropdown menu with "None" selected.
- Engine ID:** A text input field.
- Buttons:** "Add" and "Cancel" buttons at the bottom right.

SNMP Server - The name or IP address of your SNMP manager.

Username - An 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/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 10 to 64 hexadecimal characters only if your SNMP manager requires a fixed Engine ID.

If you are configuring SNMP version 1 or 2, you will see the following:

The screenshot shows a dialog box titled "Add SNMP Server" with a close button (X) in the top right corner. The dialog contains the following fields and controls:

- Version:** A dropdown menu with the value "2" selected.
- SNMP Server:** A text input field with a blue question mark icon to its right.
- SNMP Community:** A text input field.
- Buttons:** "Add" and "Cancel" buttons at the bottom right.

SNMP Server - The name or IP address of your SNMP manager.

SNMP Community - The community name used for SNMP traps (not MIB browsing).

4. Click **Add**.

VirtualTapeLibraryMIB

The falcVtlMonitor MIB has multiple tables that display different capacity, usage, and performance statistics.

falcVtlMonitorMIB - falcVtlMonCapacity

Displays StorSafe server capacity and usage statistics.

falcVtlCapCacheGeneralInfo

Displays StorSafe server 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 StorSafe server performance and statistics.

falcVtlPerfOneDayIntervalDataInfo

Displays StorSafe server performance for the past 24 hours. Note that no data will be available for a failover server that has been taken over.

Object	Description
DataWritten	Amount of data written for all job types (backup, deduplication, import), in MB
CacheSpaceUsed	Amount of space used after compression, in MB
DataRead	Amount of total uncompressed data read, in MB
DataCompressedRead	Amount of compressed data read, in MB
ReplRawDataTx	(StorSafe-S servers only) Raw replication data transferred, in MB
ReplActualDataTx	(StorSafe-S servers only) Actual amount of replication data transferred, in MB
ReplTotalTapesTx	(StorSafe-S servers only) Number of tapes replicated

falcVtlMonPerformanceInfo

Displays StorSafe server performance.

Object	Description
AvgDedupRatio	(StorSafe-S servers only) Average tape total deduplication ratio
CompressRatio	Average tape compression ratio

falcVtlHistoryMIB

The falcVtlHistoryMIB has multiple tables that display historical information about activity on the server, dashboard, tape history, and deduplication status.

Activity

Displays tape activity history, including operations related to client backup and restore, import, export, and deduplication:

Object	Description
StartTime	Start time for the tape operation
EndTime	End time for the tape operation
VlibraryID	Virtual library in which the tape was present when the operation took place
VDriveID	Virtual drive used by the operation
VTapeID	Virtual tape used for the operation
Operation	Operation performed
WriteDataMB	Uncompressed data written, in MB
WriteCompressDataMB	Compressed data written, in MB
ReadDataMB	Uncompressed data read, in MB
ReadCompressDataMB	Compressed data read, in MB
StartEODMB	Location on the tape that marks the end of data at the start of the operation (in MB to denote the location as offset from the beginning of the tape)
EndEODMB	Location on the tape that marks the end of data at the end of the operation (in MB to denote the location as offset from the beginning of the tape)
Barcode	Tape barcode

DashStatistics

Displays the following dashboard activity history:

Object	Description
TimeStamp	Date and time the data was collected
DiskSpaceTotal	Total disk space used, in GB
DiskSpaceAvailable	Available disk space, in GB

Object	Description
PerformanceRead	Read performance, in KB/sec (calculated at the HBA level)
PerformanceWrite	Write performance, in KB/sec (calculated at the HBA level)

TapeHistory

Displays the following statistics history for each tape in a deduplication policy:

Object	Description
TapeID	Tape ID
TapeName	Tape name
TapeBarcode	Tape barcode
PolicyID	Policy ID
DriveSerialNumber	Drive serial number
TapeOperation	Operation conducted on tape: Scan Tape, Write Tape, Resolve Tape, Empty Tape, Upgrade Tape, Verify Tape, Inline Parsing Tape
Parser	Data type/parser used: ARCServe, Atempo, Bacula, Commvault, DataProtector, IBMiSeries, Legato, MicrosoftTapeFormat, NDMP, Netbackup, NetbackupFalconstorOpenStorage, Netvault, OracleSecureBackup, Syncsort, TSM, Virbak, Unknown
VTLServer	StorSafe server
Result	Result of operation: Success, Failed
ErrorCode	Error code
Message	Error message
DedupeStatus	Deduplication status: Total Operation Running, Total Operation Finished, Total Operation Paused, Total Operation Failed, Total Operation Preparing, Total Operation Stopped, Queued, Running, Finished, Paused, Failed, Preparing, Stopped, Unknown
TapeSize	Tape size, in GB
ScannedData	Amount of data processed, in GB
NumberOfTapeBlocks	Physical allocation tape blocks (divided by 1024)
NumberOfFiles	Number of files
Data	Total amount of data on tape, in GB
UniqueData	Unique data, in GB
VVTapeData	Used size of a VIT (the amount of data written to a VIT after deduplication) in GB

Object	Description
WrittenVITData	Same as VVTapeData
ScanStartTime	Scan start time
ScanEndTime	Scan end time
RunTime	Run time, in seconds
Throughput	Throughput, in MB per second
WriteVITTime	VIT write time, in seconds
WriteVITThroughput	VIT write throughput, in MB per second
ReplicationStart	Replication start time
ReplicationEnd	Replication end time
ReplicationThroughput	Replication throughput, in MB per second
DataReplicated	Amount of data replicated, in MB
ReplicationStatus	Replication status: New, Idle, In Progress, Replication Failed, Scanning, Scan Failed
ResolverStartTime	Resolver start time
ResolverEndTime	Resolver end time
ResolverThroughput	Resolver throughput, in MB per second
ResolvedVITData	Amount of resolved VIT data, in GB
ResolveUniqueData	Unique data resolved, in GB
ResolveStatus	Resolver status: Total Operation Running, Total Operation Finished, Total Operation Paused, Total Operation Failed, Total Operation Preparing, Total Operation Stopped, Queued, Running, Finished, Paused, Failed, Preparing, Stopped, Unknown
ResolveErrorCode	Resolver error code
PolicyStartTime	Policy start time
PolicyEndTime	Policy end time
RunType	Run type (manual, scheduled, tape caching)
RunStatus	Status: Complete, Incomplete

falcVtlServer

Displays information about the StorSafe server and the options configured.

Processor

Displays information about all of the processors in the StorSafe server.

Object	Description
falcVtlProcessorInfo	Processor type

NetInterface

Displays information about all of the network interfaces in the StorSafe server.

Object	Description
falcVtlNetInterfaceInfo	Network interface and maximum transfer unit of each IP packet (MTU)

FailoverInfo

Displays failover information for the configured failover node.

Object	Description
Configuration Type	Type of failover (active-passive or mutual)
Failover Partner	Name of failover partner
Failover Configuration Properties	StorSafe servers in failover group Self check interval Heartbeat interval Recovery setting Current failover state

FCInfo

Displays failover information about all of the Fibre Channel HBAs in the StorSafe server.

Object	Description
WWPN	World Wide Port Name
Mode	Mode: Initiator, Target, Standby, Null.

falcVtlServerOptionsInfo

Displays StorSafe server options.

Object	Description
falcVTLServerOptionsInfo	Displays StorSafe server options (Fibre Channel, iSCSI, database, Email Alerts, Hosted Backup, NDMP) and informs whether each is enabled or disabled.

falcVtlServerInfo

Displays StorSafe server information.

Object	Description
ServerName	Server hostname
LoginMachineName	Server IP
ServerVersion	Server version
OsVersion	Server operating system version
KernelVersion	Server kernel version
Memory	Amount of server memory
SwapSpace	Amount of server swap space

falcVtlLibrarySystem

Displays information about virtual and physical libraries, clients, and physical resources. It also contains deduplication policy information, if applicable.

VirtualLibrary / falcVtlVirtualLibsInfo

Displays information about the configuration and properties of each virtual tape library.

Object	Description
falcVtlVirtLibID	Virtual tape library ID
falcVtlVirtLibName	Virtual tape library name
falcVtlVirtLibVendorID	Vendor ID
falcVtlVirtLibProductID	Product ID
falcVtlVirtLibRev	Firmware version
falcVtlVirtLibNumSlots	Number of slots
falcVtlVirtLibNumDrives	Number of drives
falcVtlVirtLibBarcodeBegin	First barcode in range for library
falcVtlVirtLibBarcodeEnd	Last barcode in range for library
falcVtlVirtLibTapeCapacityOnDemand	Indicates if tape capacity on demand (COD) is enabled
falcVtlVirtLibInitAllocSize	Initial tape size, in MB (if tape COD is enabled)
falcVtlVirtLibIncrementSize	Incremental amount, in MB (if tape COD is enabled)
falcVtlVirtLibMaxCapacity	Maximum tape capacity, in MB (if tape COD is enabled)
falcVtlVirtLibAutoArchive	Indicates if auto archive is enabled
falcVtlVirtLibMediaType	Media type
falcVtlVirtLibNumTapes	Number of tapes in library
falcVtlVirtLibSerialNum	Serial number of library
falcVtlVirtLibAutoReplication	Indicates if auto replication is enabled
falcVtlVirtLibAutoTapeCaching	Indicates if tape caching is enabled
falcVtlVirtLibTapeDuplication	Indicates if tape duplication is enabled
falcVtlVirtualLibsInfo	Total number of virtual tape libraries

VirtualDrive / falcVtlVirtualDrivesInfo

Displays information about the configuration and properties of each virtual tape drive (inside a virtual tape library, standalone, and used for deduplication).

Object	Description
falcVtlVirtDriveID	Virtual tape drive ID
falcVtlVirtDriveName	Virtual tape drive name
falcVtlVirtDriveVendorID	Vendor ID ("FALCON" for deduplication tape drives)
falcVtlVirtDriveProductID	Product ID ("SIR" for deduplication tape drives)
falcVtlVirtDriveRevision	Firmware version
falcVtlVirtDriveMediaType	Media type ("SIR001" for deduplication tape drives)
falcVtlVirtDriveLocationType	Virtual tape drive location (virtual tape library or standalone)
falcVtlVirtDriveLocationID	ID of the virtual tape library where the virtual tape drive resides
falcVtlVirtDriveGBRead	GB read from this tape drive
falcVtlVirtDriveGBWritten	GB written to this tape drive
falcVtlVirtDriveCompression	Indicates if compression is enabled
falcVtlVirtDriveStatus	Operational status: unknown, empty, loaded, ejected (but not removed), offline, passthrough (all commands will be sent), becoming ready, unloading, online
falcVtlVirtualDrivesInfo	Total number of virtual tape drives

VirtualTape / falcVtlVirtualTapesInfo

Displays information about the configuration and properties of each virtual tape.

Object	Description
falcVtlVirtTapeID	Virtual tape ID
falcVtlVirtTapeName	Virtual tape name
falcVtlVirtTapeTotalSize	Size, in MB
falcVtlVirtTapeStatus	Status (online, offline)
falcVtlVirtTapeGUID	Globally Unique Identifier (GUID)
falcVtlVirtTapeUsedSize	Used size, in MB
falcVtlVirtTapeBarcode	Barcode
falcVtlVirtTapeMediaType	Media type
falcVtlVirtTapeCapacityOnDemand	Indicates if tape capacity on demand is enabled

Object	Description
falcVtlVirtTapeAutoArchive	Indicates if auto archive is enabled or disabled
falcVtlVirtTapeWriteProtection	Indicates if the tape is write protected
falcVtlVirtTapeLocationType	Tape location (library slot, drive, or vault)
falcVtlVirtTapeLocationID	ID of the virtual resource where the tape resides (-1, not applicable, for vault)
falcVtlVirtTapeLocationSlot	Slot number of virtual tape library that tape resides in (-1, not applicable, for vault and drive)
falcVtlVirtTapeInitAllocSize	Initial tape size, in MB (if tape COD is enabled)
falcVtlVirtTapeIncrementSize	Incremental amount, in MB (if tape COD is enabled)
falcVtlVirtTapeMaxCapacity	Maximum tape capacity, in MB (if tape COD is enabled)
falcVtlVirtTapeRdeTapeType	Type of tape (regular virtual tape, mixed VIT, pure VIT)
falcVtlVirtTapeRdeEndOfVITData	Location on the tape that marks the end of VIT data (in MB to denote the location as offset from the beginning of the tape)
falcVtlVirtTapeRdeUniqueData	Unique data on the tape, in MB
falcVtlVirtTapeRdeDedupedData	Deduplicated data on the tape, in MB
falcVtlVirtualTapesInfo	Total number of virtual tapes

PhysicalLibrary / falcVtlPhysicalLibsInfo

Displays information about the configuration and properties of each physical tape library.

Object	Description
falcVtlPhyLibID	Physical tape library ID
falcVtlPhyLibName	Physical tape library name
falcVtlPhyLibAllocationType	Same as falcVtlPhyLibName
falcVtlPhyLibVendorID	Vendor ID
falcVtlPhyLibProductID	Product ID
falcVtlPhyLibStatus	Status (online, offline)
falcVtlPhyLibGUID	Globally Unique Identifier
falcVtlPhyLibSerialNumber	Serial number
falcVtlPhyLibNumSlots	Number of slots
falcVtlPhysicalLibsInfo	Total number of physical tape libraries

PhysicalDrive / falcVtlPhysicalDrivesInfo

Displays information about the configuration and properties of each physical tape drive (inside a physical tape library and standalone).

Object	Description
falcVtlPhyDriveID	Physical tape ID
falcVtlPhyDriveName	Physical tape drive name
falcVtlPhyDriveVendorID	Vendor ID
falcVtlPhyDriveProductID	Product ID
falcVtlPhyDriveStatus	Status (online, offline)
falcVtlPhyDriveGUID	Globally Unique Identifier
falcVtlPhyDriveSerialNumber	Serial number
falcVtlPhyDriveState	Operational status: empty, loaded, ejected (but not removed from drive), offline, passthrough (all commands will be sent)
falcVtlPhysicalDrivesInfo	Total number of physical tape drives

PhysicalTape / falcVtlPhysicalTapesInfo

Displays information about the configuration and properties of each physical tape.

Object	Description
TapeSlot	Physical slot the tape is in
TapeBarcode	Barcode
falcVtlPhysicalTapesInfo	Total number of physical tapes

VtlJob / falcVtlJobQueueInfo

Displays information about the tape import/export job queue.

Object	Description
falcVTLJobID	Job ID
falcVTLJobType	Job type: • exportToSADriveAndCopy - Export with copy on a standalone physical tape drive • exportToSADriveAndMove - Export with move on a standalone physical tape drive • exportToPhyLibAndCopy - Export with copy to a physical tape library • exportToPhyLibAndMove - Export with move to a physical tape library • importFromSADriveAndCopy - Import in copy mode from a standalone physical tape drive • importFromSADriveAndRecycle - Import in recycle mode from a standalone physical tape drive • importFromPhyLibAndCopy - Import in copy mode from a physical library • importFromPhyLibAndRecycle - Import in recycle mode from a physical library • createFromCacheMetaDataModeAndCopy - Create cache in copy meta data mode • moveTapeIESlot - Moving tape to an IE slot • importFromPhyLibStackAndCopy - Import from physical library with stacking in copy mode • exportToSADriveStackAndCopy - Export to standalone physical drive with stacking in move mode • exportToPhyLibStackAndCopy - Export to physical library with stacking in copy mode • moveTapeIESlotStack - Move tape to an IE slot with stacking • scanPhyTapeStack - Scan physical tape • exportToPhyLibstackAndMove - Export to physical library with stacking in move mode
falcVTLJobPhysicalLibName	Physical library used for import/export
falcVTLJobPhysicalTapeBarcode	Physical tape barcode used for import/export
falcVTLJobPhysicalSlot	Physical slot number used for import/export
falcVTLJobVirtualLibName	Virtual library used for import/export
falcVTLJobTapeName	Virtual tape used for import/export
falcVTLJobTapeBarcode	Virtual tape barcode used for import/export
falcVTLJobVirtualSlot	Library virtual slot number used for import/export
falcVTLJobStatus	Job status: ready, running, completed, cancelled, or failed falcVTLJobDescription - Job description
falcVtlJobQueueInfo	Total number of import/export drives

ReplicaResource

Displays information about replica resources.

Object	Description
VirtualID	Replica resource ID
VirtualName	Resource name
AllocationType	Resource type: Null, Virtual Device, Direct Device, System, Reserved, 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 resource 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 resource ID
Schedule	Current status of the schedule: On Schedule, Suspended, or N/A
Watermark	Watermark set to trigger replication
WatermarkRetry	Retry interval if replication fails
Time	Time when replication is scheduled to occur
Interval	Time interval between replication jobs
falcVtlNumOfReplica	Total number of replica resources

DeduplicationPolicy / falcVtlDeduplicationPoliciesInfo

Displays information about deduplication policies.

Object	Description
PolicyName	Policy name
PolicyID	Policy ID
NumberOfTapes	Number of tapes in the policy
TriggerType	Trigger type: • endOfBackupTapeFull • endOfBackup • noSchedule • scheduleHourly • scheduleDaily • scheduleSunday • scheduleMonday • scheduleTuesday • scheduleWednesday • scheduleThursday • scheduleFriday • scheduleSaturday • tapeCaching

Object	Description
SirCluster	Deduplication cluster to which data will be deduplicated
ReplicationStatus	Indicates if replication is enabled or disabled
TurboDeduplication	Indicates if turbo deduplication is enabled or disabled
ReplicationMode	Replication mode
falcVtlDeduplicationPoliciesInfo	Total number of deduplication policies.

falcVtlPhysicalResources

Displays information about physical resources.

Storage HBAs

Displays information about each storage HBA.

Object	Description
Number	SCSI adapter number
Info	Model/type
WWPN	World wide port name
Mode	Mode: Target or initiator
AliasWWPN	Alias WWPN, if dual mode
AliasMode	Mode of alias: Target or initiator
GBRead	Amount of data read, in GB, for target ports
GBWrite	Amount of data written, in GB, for target ports

StorageDevices

Information about the hardware specifications and characteristics of each SCSI device (physical tape library, physical tape drive, and storage device).

Object	Description
DeviceType	Access type for the attached device (Direct-Access, Sequential-Access, Medium Changer)
VendorID	Product vendor
ProductID	Product model
FirmwareRev	Firmware version
AdapterNo	SCSI adapter number
ChannelNo	SCSI channel
ScsiID	SCSI ID
LUN	SCSI LUN
TotalSectors	Number of sectors or blocks
SectorSize	Number of bytes in each sector or block
TotalSize	Total size of the device, in MB

Object	Description
ConfigStatus	Status (online or offline)
UsedSize	Space used, in MB
FreeSize	Free space, in MB
StorageOwner	Owner of the device. Storage devices will show the server name. Media changers and drives will be displayed as local owner.

StoragePools

Information about the hardware specifications and characteristics of each SCSI device (physical tape library, physical tape drive, and storage device).

Object	Description
StoragePoolName	Storage pool name
StoragePoolID	Storage pool ID
StoragePoolType	Storage pool type (NAS, 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, HostedBackup

FCClientResource

Displays information about virtual resources assigned to each FC client.

Object	Description
ResourceID	Virtual resource ID
ResourceName	Virtual resource name
ClientID	Client ID
ClientName	Client name
ResourceAllocType	Resource type: Direct Device Virtual Library, Direct Device Virtual Drive
LUN	SCSI LUN of client
InitatorWWPN	WWPN of the client's 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 cluster.

Object	Description
Name	Deduplication cluster name
DataDiskAvailable	Repository data storage available, in MB
DataDiskTotal	Total repository data storage, in MB
DataDiskAvailablePercentage	Percentage of total repository data storage that is available
IndexDiskAvailable	Index storage available, in MB
IndexDiskTotal	Total index storage, in MB
IndexDiskAvailablePercentage	Percentage of total index storage that is available
RepositoryObjectRetainedPercentage	Percentage of index cache capacity retained by reclamation
RepositoryObjectUsed	Percentage of index cache capacity that is used
RepositoryObjectAvailablePercentage	Percentage of index cache capacity that is available
FolderDiskAvailable	Folder space available, in MB
FolderDiskTotal	Total folder space, in MB
FolderDiskAvailablePercentage	Percentage of total folder space that is available
ReclamationLastRunTime	Last date and time data reclamation was run
ReclamationDataSaved	Data space saved after reclamation, in MB
IndexReclamationLastRunTime	Last date and time index pruning was run
ReclamationIndexDataSaved	Index space saved after reclamation, in MB

falcSirDedupeRatioRangesInfo

Displays deduplication ratio information for the cluster. The statistics are updated every hour; these are not real-time statistics. By default, the statistics are updated on the hour.

Object	Description
AvgTapeDedupeRatioIn1Hour	Average deduplication ratio (data scanned / data stored) for the last hour
AvgVITReplicaDedupeRatioIn1Hour	Average VIT replication deduplication ratio (data replicated / unique data replicated) for the last hour
AvgSIRDedupeRatioIn1Hour	Average deduplication ratio (total data / total compressed) for the last hour
VITsWith1-2DedupeRatio	The number of tapes with a deduplication ratio between 1.0 and < 2
VITsWith2-4DedupeRatio	The number of tapes with a deduplication ratio between 2.0 and < 4
VITsWith4-8DedupeRatio	The number of tapes with a deduplication ratio between 4.0 and < 8
VITsWith8-16DedupeRatio	The number of tapes with a deduplication ratio between 8.0 and < 16
VITsWithGreaterThan16DedupeRatio	The number of tapes with a deduplication ratio greater than 16

falcSirServerPerformanceInfo

Displays information about total data deduplicated. The statistics are updated every day; these are not real-time statistics. By default, the statistics are updated at midnight.

Object	Description
DataDeduplicatedIn24Hour	Amount of data deduplicated by the SIR server in the past 24 hours, in MB
NumberOfTapesDeduplicatedIn24Hour	Number of tapes deduplicated by the SIR server in the past 24 hours, in MBs
DataReplicatedIn24Hour	Amount of data replicated by the deduplication target replication server in the past 24 hours, in MBs
AverageDedupeRatioIn24Hour	Average deduplication ratio of data processed by the SIR server in the past 24 hours, in MBs
AverageCompressionRatioIn24Hour	Average compression ratio of unique blocks processed by the SIR server in the past 24 hours, in MBs

falcSirMonNodePerformanceInfo

Displays performance information for the cluster. The statistics are updated every day; these are not real-time statistics. By default, the statistics are updated at midnight.

Object	Description
DataDeduplicated	Amount of data deduplicated by the deduplication cluster in the past 24 hours, in MB
OverallDedupeRatio	Overall deduplication ratio by the deduplication cluster
AverageDedupeRatio	Average deduplication ratio for the deduplication cluster in the past 24 hours
AverageCompressionRatio	Average compression ratio for the deduplication cluster of unique blocks in the past 24 hours
ReplicationDataTx	Amount of actual data replicated in the past 24 hours
LastIndexLoadTime	Number of minutes required to load the index into RAM the last time it was loaded

falcSirCluster - falcSirClusterConf

Displays information about the SIR server related to deduplication cluster configuration

falcSirCCGeneralInfo

Displays general cluster information.

Object	Description
ClusterName	Name of the deduplication cluster
FailoverOption	Indicates if failover is enabled or disabled
Guid	Globally Unique Identifier of deduplication cluster
HBIfName	Heartbeat interface name (such as "eth0")
ReplicationOption	Indicates if replication is enabled or disabled
ReplicaCount	Number of replica target deduplication cluster instances assigned to this cluster (this cluster is the primary)
ReplicatorCount	Number of instances where this cluster is a target replica
SirNodeCount	Number of SIR server nodes

VTL

Displays information about associated StorSafe servers.

Object	Description
falcSirCCVtlHostname	Hostname of the StorSafe server
falcSirCCVtlIP	IP address of StorSafe server

SIRReplication

Displays source and target replica information.

Object	Description
Type	Indicates if this is a primary (replicator) or target (replica) server
Name	Replication server name
Guid	Globally Unique Identifier of replication server
Protocol	Protocol of replication server (unknown, iSCSI, fibrechannel, tcp)
SirNodeCount	Number of deduplication nodes on the replication server

Object	Description
NodeIPCount	Number of IP addresses for the replication server
NodePortCount	Number of ports for the replication server

SIRReplicationSirNodeIP

Displays deduplication node IP information.

Object	Description
SirName	Deduplication node name of the replication server
Address	Replica node IP address
RepType	Indicates if this is a primary (replicator) or target (replica) server

SIRReplicationSirNodePort

Displays deduplication node port information.

Object	Description
SirName	Deduplication node name of the replication server
Port	Port information of replication server
RepType	Indicates if this is a primary (replicator) or target (replica) server

falcSirCCReclamationInfo

Displays reclamation information.

Object	Description
Option	Indicates if reclamation is enabled or disabled
Interval	Indicates time interval for reclamation, in seconds

SIRNode

Displays information about the deduplication nodes in each deduplication cluster.

Object	Description
Name	SIR server name
IPAddress	SIR server IP address

Object	Description
Type	SIR server type (active, standby)
PowerFactor	Number of CPUs x 2. This number determines the maximum number of deduplication and resolver processes that can run on each deduplication node.

SIRNodeFailover

Displays failover information for a deduplication node.

Object	Description
Name	Name of the failover SIR server
HeartbeatIP	Heartbeat IP address for failover
HeartbeatNetmask	Heartbeat netmask for failover
HeartbeatGateway	Heartbeat gateway for failover
IpmitIP	IPMI IP address of failover
IpmitNetmask	IPMI netmask of failover
IpmitGateway	IPMI gateway of failover
IPCount	Number of IP addresses for the deduplication failover node

SIRNodeFailoverIP

Displays failover IP address information for a deduplication node.

Object	Description
SIRNodeName	Name of the SIR server
Name	Heartbeat interface
RealIP	Heartbeat IP address
Netmask	Heartbeat netmask
VirtualIP	Virtual IP address

SIRNodeFibreChannel

Displays Fibre Channel information for a deduplication node.

Object	Description
NodeName	Deduplication node name

Object	Description
WWPN	World Wide Port Name

SIRHashIndexStorage

Displays information about the deduplication node hash index list.

Object	Description
SirName	SIR server name
Size	Size of hash index

SIRHashFolderStorage

Displays information about the deduplication node hash folder list.

Object	Description
SirName	SIR server name
Size	Size of hash folder

SIRHashDataStorage

Displays information about the deduplication node hash data list.

Object	Description
SirName	SIR server name
Size	Size of hash data

falcSirCluster - falcSirClusterStats

Displays cluster statistics about folders, repository usage, and deduplication results

Folder

Displays folder information for the deduplication cluster.

Object	Description
Barcode	Barcode of folder
BackupApp	Backup application: ARCServe, Atempo, Bacula, Commvault, DataProtector, IBMiSeries, Legato, MicrosoftTapeFormat, NDMP, Netbackup, NetbackupFalconstorOpenStorage, Netvault, OracleSecureBackup, Syncsort, TSM, Virbak, Unknown
DedupTime	Time of deduplication
DataType	Parser used: arcserve, mtf, netbackup, netbackup2, tsm, legato, generic, legato2, tsm2, commvault, dprotect, ndmp, arcserve2, osb, atempo, netvault, generic2, generic3, ost, IBMi, syncsort, virbak, bacula
MediaType	Media type
Source	Source StorSafe 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

Object	Description
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

Object	Description
UsedSinceReclamation	Percentage of folder disk used since reclamation
Free	Percentage of folder disk that is available
Capacity	Total capacity of folder disk, in MB
SpaceUsed	Total folder disk space used, in MB

Common MIBs

These MIBs monitor both backup and SIR servers.

ServiceEntry

Displays the current state of each StorSafe 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 StorSafe 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.

Command Line

StorSafe 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 StorSafe system, as well as integrate StorSafe 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	StorSafe server name (hostname or IP address)
-u	--server-username	StorSafe server username
-p	--server-password	StorSafe server user password
-c	--client-name	StorSafe client name
-v	--vdevId	StorSafe virtual resource 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 StorSafe 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 removelicense -s <server-name> [-u <username> -p <password>] -k <license-keycode>
[-X <rpc-timeout>]

iscon removelicense --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --license=<license-keycode>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command removes a license keycode.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Register a license keycode

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

```
iscon registerlicense --server-name=<server-name>  
[--server-username=<username> --server-password=<password>]  
--license=<license-keycode> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command registers a specific license key code.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Physical devices

Get physical device information

```
iscon getpdevinfo -s <server-name> [-u <username> -p <password>]
[-F [-M | -C <category>] | [-a] [-A] [-I <ACSL>] ] [-o <output-format>]
[-X <rpc-timeout>]
```

```
iscon getpdevinfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--config [--include-system-info | --category=<category>] |
[--allocated-list] [--available-list] [--scsiaddress=<ACSL>] ]
[--output-format=<output-format>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves information about the physical devices detected by the specified server. By default, the command displays the allocation information for the virtualized disks and the assigned physical libraries and drives owned by the server. The "allocated" and "available" options work as filters for the default execution. The "config" option displays information about all the physical devices that are detected by the server.

-F (--config) is an option to get the physical device configuration information. The default is to exclude the system device information.

-M (--include-system-info) is an option to include the system device information when -F (--config) is used.

-C (--category) is an option to be used as a filter to get the configuration information for the specified category in one of the values, when -F (--config) is used: *virtual* (default), *service-enabled*, or *direct*.

-M (--include-system-info) and -C (--category) options are mutually exclusive.

-o (--output-format) is an option to specify the output format. The <output-format> for the -F (--config) option is one of the following values: *list* or *detail* or *guid* or *scsi*.

-a (--allocated-list) is an option to get the allocated physical device information.

-A (--available-list) is an option to get the available physical device information.

-I (--scsiaddress) is an option to specify the SCSI address as a device filter in the following format:
<ACSL>=#:#:# (adapter:channel: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 StorSafe system with the new configuration.

-a (--adapter-range) is the adapter or adapter range to be rescanned. The default is to rescan all the adapters if it is not specified. For example, -a 5 or -a 5-10 or -a auto.

-i (--scsi-range) is the starting SCSI ID and ending SCSI ID to be rescanned. The default is to rescan all the SCSI IDs if the range is not specified. For example, -i 0-5

-l (--lun-range) is the starting LUN and ending LUN to be rescanned. The default is not to rescan any LUN if it is not specified. For example, -l 0-10

If you want the system to rescan the device sequentially, you can specify the -L (--sequential) option. The default is not to rescan sequentially.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Prepare disk

```
iscon preparedisk -s <server-name> [-u <username> -p <password>]
[-U <target-username> -P <target-password>]
-i <guid> | -I <ACSL> -C <category> [-l <lun-reservation>]
[-X <rpc-timeout>]

iscon preparedisk --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--target-username=<username> --target-password=<password>]
--scsiaddress=<ACSL> | --guid=<guid> --category=<category>
[--lun-reservation=<lun-reservation>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command changes the category for the specified physical device. You must use care with this command; the system will not check the current category and LUN reservation setting before making the specified change.

<guid> is the unique identifier of the physical device.

<ACSL> is the SCSI address of the physical device in the following format: *#:#:#: (adapter:channel:scsi id:lun)*.

Either -i (--guid) or -I (--scsiaddress) has to be specified for the disk to be prepared.

-C (--category) is required to specify the new category for the physical device in one of the following values: *unassigned, virtual, direct*.

-l (--lun-reservation) is needed if the category is set to "virtual". The LUN reservation determines what kind of resources can be created on this device. The accepted values are:

- None
- ConfigurationRepository
- DeduplicationRepository
- Tapes
- NASResources

DeduplicationRepository is only valid for SIR and StorSafe-S servers.

Tapes and *NASResources* are only valid for StorSafe and StorSafe-S servers.

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

If the server is set up for failover, the failover partner has to be rescanned after the disk preparation. <target-username> and <target-password> options specify the user name and password for the failover partner.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Change LUN reservation

```
iscon changelunreservation -s <server-name> [-u <username> -p <password>]
-i <guid> | -I <#:#:#:~> -l <lun-reservation> -f
[-X <rpc-timeout>]
```

```
iscon changelunreservation --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--guid=<guid> | --scsiaddress=<#:#:~>
--lun-reservation=<lun-reservation> --force
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command changes the LUN reservation for a virtualized device. The new reservation type must be compatible with the existing resources on the device or the device must be empty.

Either -i (--guid) or -I (--scsiaddress) must be used in order to identify the device.

-l (--lun-reservation) is required and determines what kind of resources can be created on this device. The accepted values are:

- None
- ConfigurationRepository
- DeduplicationRepository
- Tapes
- NASResources

DeduplicationRepository is only valid for SIR and StorSafe-S servers.

Tapes and *NASResources* are only valid for StorSafe and StorSafe-S servers.

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

-f (--force) is required in order to confirm the operation.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Rename physical device

```
iscon renamephysicaldevice -s <server-name> [-u <username> -p <password>]
-i <guid> | -I <ACSL> -n <new-name> [-X <rpc-timeout>]
```

```
iscon renamephysicaldevice --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--guid=<guid> | --scsiaddress=<ACSL> --name=<new-name> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command renames a physical device.

-i (guid) specifies the unique identifier of the physical device.

-l (--scsiaddress) is the SCSI address of the physical device in the following format: `#:~:~:~` (adapter:channel:scsi id:lun)

Either -i (--guid) or -l (--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> | -l <#:~:~:~> -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 -l (--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:

- NAS (NAS pool type)
- Tape (StorSafe 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 NAS,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:

- NAS (NAS pool type)
- Tape (StorSafe 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 StorSafe 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 StorSafe information type with one of the following values: *VLIBS* or *VDRIVES* or *VAULT* or *PLIBS* or *PDRIVES*.

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

The default is to display all the information.

-L (--tape-library-vid) is an option to specify the virtual tape library when *VLIBS* is specified, or to specify the physical tape library when *PLIBS* is specified.

-F (--vtl-info-filter) is an additional filter that can be combined using the following values separated with commas: *library* or *drive* or *tape*.

- *library* = include physical and/or virtual library information.
- *drive* = include physical and/or virtual drive information.
- *tape* = include physical and/or virtual tape information.

For example: -F "library,drive,tape" or --vtl-info-filter="library,drive,tape"

The default is to display all of the information that applies. There will be an error if <vtl-info-type> is specified and the <vtl-info-filter> specified does not apply. For example, "library" does not apply to "VDRIVES".

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

-A (--long-physical-layout) is an option to display the physical layout associated with the device, if applicable. The argument is ignored if -l (--long-list) is not specified.

-M (--output-delimiter) can be specified when -l is specified to replace the linefeed with the specified delimiter. The maximum length of the delimiter is 8.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Assign virtual library or drive to an FC client

```

iscon assignvdev -s <server-name> [-u <username> -p <password>]
-v <vdev> -c <client-name> [-y]
[-I <initiatorWWPN|*>] [-T <targetWWPN|*>]
[-X <rpc-timeout>]

iscon assignvdev --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --vdev=<vdev>
--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 (--vdev) 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 <vdev> -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>]
--vdev=<vdev> --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 (--vdev) 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 one of the following values: *FC* or *ISCSI*. The client type will only take effect when the client name is *. Be aware that in some platforms you are required to enclose the "*" in double quote to take it as a literal.

-l(--longlist) is an option to display the long format.

-M (--output-delimiter) can be specified when -l is specified to replace the linefeed with the specified delimiter. The maximum length of the delimiter is 8.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Add Fibre Channel client

```
iscon addclient -s <server-name> [-u <username> -p <password>]
-c <client-name>
[-I <initiator-wwpns>] [[-a on] [-A on]] | [-C on]
[-X <rpc-timeout>]

iscon addclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> [--initiator-wwpns=<initiator-wwpns>]
[[-enable-VSA=on] [--enable-iSeries=on]] | [--enable-Celerra=on]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds a Fibre Channel client to the specified server.

-c (--client-name) is a unique client name for the client to be created. The maximum length of the client name is 32. The following characters are invalid for a client name: <>"&\$/'

-I (--initiator-wwpns) is an option to set the initiator WWPNS. An initiator 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 StorSafe 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 StorSafe Administrator)
- S (for StorSafe Read-only User)
- I (for StorSafe 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 StorSafe server from accessing StorSafe. 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>] [-A <auto-archive-mode> [-Y <days>] [-J] | -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]
[-K <data-key-name> -G <data-key-password>] [-k <key-name> -W <key-password>]
[-X <rpc-timeout>]
```

```
iscon createvirtuallibrary --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vlib-type=<vlib-type> [--vlib-name=<vlib-name>] --vdrive-type=<vdrive-type>
[--vdrive-name-prefix=<vdrive-name-prefix>] [--num-of-drives=<num-of-drives>]
[--auto-archive-mode=<auto-archive-mode> [--delay-delete-days=<days>]
[--auto-eject-to-ie] | --auto-replication=<auto-repl-mode> --target-name=<target-name>
[--target-username=<target-username> --target-password=<target-password>]
[--delay-delete-time=<#[D|H|M]>] |
--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]
[--data-key-name=<key-name> --data-key-password=<key-password>]
[--key-name=<key-name> --key-password=<key-password>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a virtual tape library with the specified configuration.

-t (--vlib-type) is required in the following format: "<vendorID>:<productID>"

-n (--vlib-name) is optional. A default name will be provided in the format of <vendorID>-<productID>-<vid> if it is not specified.

-d (--vdrive-type) is required to specify the type of tape drive to be created in the library. The format of <vdrive-type> is as follows: "<vendorID>:<productID>"

-r (--vdrive-name-prefix) is an option to specify the prefix of the virtual drive. The default prefix is in the format of <drive-vendorID>-<drive-productID>-<vid>.

-R (--num-of-drives) is an option to create the specified number of drives, up to the maximum allowed by the library. By default, the library will be created with 1 drive. Use -f (--force) to override the default maximum value for the specified library in order to create up to 256 drives or 500 drives for library IBM 03584L22, IBM 03584L32, VTL 03584L.

-A (--auto-archive-mode) is an option with one of the following values: *copy* or *move*.

-Y (--delay-delete-days) is an option for *move* mode to specify the number of days to wait before deletion. The maximum is 365 days. The default value is 365 days.

-J (--auto-eject-to-ie) is an option to be specified with -A (--auto-archive-mode) to eject the tape to the import/export (IE) slot after the export job.

-N (--auto-replication) is an option with one of the following values: *replication* or *remotemove*.

-S (--target-name) is the target server name for auto-replication. It is required for auto-replication.

-U (--target-username) and -P (--target-password): Target credentials are required if the target server was not already connected to with the login command or if they are not the same as the primary server. If they are not provided, the primary server credentials will be used.

-M (--delay-delete-time) is an option for *remotemove* mode to specify a time to wait before deletion. It can be specified in days(D), hours(H) or minutes(M). For example, 2D, 10H, 150M

-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.

-k (--key-name) and -W (--key-password) are options for tape encryption support to be set in conjunction with Auto-Archive Mode. Specify the key name and key password of the encryption key if you wish to encrypt the data when exporting the virtual tape to the physical tape.

A virtual device ID will be assigned to the virtual library when it is created successfully.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete virtual tape library

```
iscon deletevirtuallibrary -s <server-name> [-u <username> -p <password>]
-v <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.

Set virtual tape library duplication

```
iscon setvirtuallibrarytapeduplication -s <server-name> [-u <username> -p <password>]
-v <vdevvid> -Z <on|off> [-Q <num-of-copies> -q <A:B:C>] [-X <rpc-timeout>]
```

```
iscon setvirtuallibrarytapeduplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevvid=<vdevvid> --tape-duplication=<on|off>
[--num-of-copies=<num-of-copies> --prefix=<A:B:C>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command sets the Tape Duplication property for a virtual tape library.

-v (--vdevvid) is required in order to identify the virtual library.

-Z (--tape-duplication) is required in order to enable or disable the Tape Duplication property: *on* (enable) or *off* (disable).

-Q (--num-of-copies) is an option to specify the number of copies to be made using the same barcode if the tape duplication option is enabled. The maximum value is 5. The default value is 1.

-q (--prefix) is an option to add a single character prefix to the barcode of each copy, in the following format: A:B:C

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Add virtual tape drive

```
iscon addvirtualdrive -s <server-name> [-u <username> -p <password>]
-L <tape-library-vid> [-r <vdrive-name-prefix>] [-R <num-of-drives>] [-X <rpc-timeout>]
```

```
iscon addvirtualdrive --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-library-vid=<tape-library-vid> [--vdrive-name-prefix=<vdrive-name-prefix>]
[--num-of-drives=<num-of-drives>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds a virtual tape drive to a specific virtual tape library.

-L (--tape-library-vid) is required to specify the virtual tape library to add the virtual tape drive(s).

-r (--vdrive-name-prefix) is an option to specify the prefix of the virtual tape drive. The default prefix is in the format of <drive-vendorID>-<drive-productID>-<vid>.

-R (--num-of-drives) is optional, the default is 1 if it is not specified.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create standalone tape drive

```
iscon 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 (--vdev) 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 <vdev> {-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>] --vdev=<vdev>
{--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 (--vdev) 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 StorSafe 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-archive --plib-vid=<plib-vid>
--physical-tape-barcode=<physical-tape-barcode>
--auto-eject-to-ie] | [--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.

-l (--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.

If the parent library has the auto-archive/remotecopy property enabled, use the following options to provide additional information for virtual tape creation:

-A (--enable-auto-archive) is an option when the parent library is enabled with auto-archive option.

-l (--plib-vid) is required when <auto-archive-mode> is specified. It is the physical tape library where the tape will be exported to automatically.

-b (--physical-tape-barcode) is required to specify the list of physical tape barcode(s) when auto-archive option is specified. Separate multiple barcodes with commas. For example, -b 00010001,00010009,0001000A

-J (--auto-eject-to-ie) is optional when <auto-archive-mode> is specified.

-N (--enable-auto-replication) is an option when the parent library is enabled with the auto-replication option.

-S (--target-name) can be specified when 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.

The *count* and *barcode* options cannot be specified when the -A (--enable-auto-archive) option is specified because the number of tapes will be obtained from the list of barcodes specified with -b (--physical-tape-barcode) option.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Set tape properties

```
iscon settapeproperty -s <server-name> [-u <username> -p <password>]
-v <vdev> [-B <barcode>] [-f] [-F] [-w <on|off>] [-A <auto-archive-mode> [-Y <days>]
[-J <on|off>] | -N <auto-repl-mode> -S <target-name>
[-U <target-username> -P <target-password>] [-M <#[D|H|M]>] ]
-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>] --vdev=<vdev>
[--barcode=<barcode>] [--force] [--full-capacity] [--tape-write-protect=<on|off>]
[--auto-archive-mode=<auto-archive-mode> [--delay-delete-days=<days>]
[--auto-eject-to-ie] | --auto-replication=<auto-replication-mode>
--target-name=<target-name>
[--server-username=<username> --server-password=<password>]
[--delay-delete-time=<#[D|H|M]>] ]
--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>] ]
[--key-name=<key-name> --key-password=<key-password> | --disable-key]
[--tape-duplication=<on|off> --num-of-copies=<num-of-copies>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command configures tape properties for the specified virtual tape. The tape must be located in a virtual tape library slot. If the specified virtual tape is in the vault, only the write protection property can be configured.

-v (--vdev) is required to specify the ID of the virtual tape to set the properties.

-B (--barcode) is an option to specify the new barcode for the tape. -f (--force) option is required if the new barcode is not in the barcode range specified for the parent library. Barcode is an alphanumerical value in the length of 4 to 12.

-F (--full-capacity) is an option to expand the tape to the maximum capacity and turn off the <capacity-on-demand> option if it is enabled for the virtual tape.

-w (--tape-write-protect) is an option to turn on and off the tape write protection with the following values: *on* (enable) or *off* (disable).

-A (--auto-archive-mode) is an option with one of the following values: *copy* or *move* or *inherited* or *none*.

- "none" is the value to turn off the auto-archive mode if the virtual tape is enabled with the auto-archive option.
- "inherited" can only be specified when the parent library is enabled with the auto-archive option.

-Y (--delay-delete-days) is an option for auto-archive *move* mode to specify up to 365 days to wait before the deletion. The default value is 365 days.

-J (--auto-eject-to-ie) is an option for auto-archive mode to eject the physical tape to the IE slot after a successful archive job: *on* (enable) or *off* (disable).

-N (--auto-replication) is an option in one of the following values: *localcopy*, *localmove*, *remotecopy*, *remotemove*, or *none*.

-S (--target-name) is the remote server name for auto-replication. It is required for auto-replication.

-U (--target-username) and -P (--target-password) are options to specify a different user ID and password to log in to the remote server.

-M (--delay-delete-time) is an option for auto-replication move mode to specify up to 30 days to wait before deletion. The default value is 1 day. The value can be specified in days(D), hours(H) or minutes(M). For example, 2D, 10H, 150M

-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.

-A (--auto-archive-mode) and -N (--auto-replication) cannot be specified if replication is enabled for the tape.

-k (--key-name), -W (--key-password) and -d (--disable-key) are options for tape encryption support to be set in conjunction with Auto-Archive Mode. Specify the key name and key password of the encryption key if you wish to encrypt the data when exporting the virtual tape to physical tape. Specify -d (--disable-key) if you wish to disable tape encryption for this tape.

-Z (--tape-duplication) is an option to set the Tape Duplication property with one of the following values: *on* (enable), *off* (disable), or *inherited*.

-Q (--num-of-copies) is an option to specify the number of copies to be made using the same barcode if the tape duplication option is enabled. The maximum value is 5. The default value is 1.

At least one of the above properties has to be specified.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete virtual tape

```
iscon deletevirtualtape -s <server-name> [-u <username> -p <password>]
[-v <vdev> ] | [-B <barcode> -l <lib/sa_drive ID | 0 (Vault)>]
[-f] [-X <rpc-timeout>]
```

```
iscon deletevirtualtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
```



```
[--vdevId=<vdevId>] |
[--barcode=<barcode> --from-location-id=<lib/sa_drive ID | 0 (Vault)>]
[--force] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes a virtual tape.

To delete a virtual tape, specify either the `-v (--vdevId)` or the `-B (--barcode)` of the tape, as they are mutually exclusive. You can also specify the `-l (--from-location-id)` option.

`-v (--vdevId)` is an option to specify the tape virtual ID.

`-B (--barcode)` is an option to specify the barcode of the virtual tape. By default, the command queries all libraries, drives, and the vault. The barcode must be unique. If you have duplicate barcodes, use `-l (--from-location-id)` to narrow the search. If the tape's `-v (--vdevId)` is provided, the barcode and location ID options are ignored.

`-l (--from-location-id)` is an option to specify the virtual ID of the library or standalone drive where the virtual tape is located when you use the `-B (--barcode)` option. If the tape is located in the vault, use 0 for the location ID.

`-f (--force)` is an option to force the deletion of a virtual tape configured for replication. The corresponding virtual tape replica will not be deleted or promoted.

`-X (--rpc-timeout)` is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Move virtual tape

```
iscon movevirtualtape -s <server-name> [-u <username> -p <password>]
-v <vdevId> | -B <barcode> [-i]
[-L <tape-library-vid> | -D <tape-drive-vid> | -l <slot-no>] [-X <rpc-timeout>]
```

```
iscon movevirtualtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevId=<vdevId> | --barcode=<barcode> [--include-filter]
[--tape-library-vid=<tape-library-vid> | --tape-drive-vid=<tape-drive-vid> |
--slot-no=<slot-no>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command moves a virtual tape to a different location.

`-v (--vdevId)` or `-B (--barcode)` is required to identify the virtual tape to be moved to a different location.

`-i (--include-filter)` is an optional filter that can be used to uniquely identify a virtual tape when multiple tapes have the same barcode and `-B (--barcode)` is used. This option can be one of the following values:

- `TapeName=***`
- `Location=*`
- `ParentID=#`

"Location" is the current location: the library ID if the tape is in a slot, the drive ID, or Vault. "ParentID" is the ID of the last library that hosted the tape and it is preserved when the tape is moved to the vault. If the tape cannot be uniquely identified, the command will fail.

-L (--tape-library-vid) is the virtual library to move to. It is not required if the virtual tape is moved within the same library.

-D (--tape-drive-vid) is the virtual drive in a library or the standalone drive to move to.

-I (--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 SIR cluster or server. On StorSafe servers, the command always starts a new reclamation job regardless of the reclamation type.

On StorSafe-S servers with deduplication enabled, the command does not start a new space reclamation job unless the *force* argument is used. The *force* argument can be used only on StorSafe-S servers, but cannot be used for index reclamation.

-T (--type) is required to specify the reclamation type. Use *SPACE* for space reclamation or *INDEX* for index pruning.

-f (--force) is an option for StorSafe-S servers with deduplication enabled.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

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>
-TSP <target-storage-pool-name> -TSP2 <target-storage-pool-name>
-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>
--target-storage-pool-name=<target-storage-pool-name>
--target-storage-pool-name2=<target-storage-pool-name>
--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 (--polycname) 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 StorSafe target server or three IP addresses, separated by commas, of the StorSafe source and target servers. When three IP addresses are given, the first IP address is the external IP address of the target server. The second IP address is the IP address of the source server, and the third IP address is the IP address on the target server. When advanced replication is enabled, this information identifies the relationship between server A and server B.

-U (--target-username) and -P (--target-password) are required in order to access StorSafe target server information.

-T2 (--target-server-ip2) is the IP address of server C or three IP addresses, separated by commas, of server B and server C. When three IP addresses are given, the first IP address is the external IP address of server C. The second IP address is the source IP address on server B, and the third IP address is the target IP address on server C. When advanced replication is enabled, this information identifies the relationship between server B and 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.

-TSP (--target-storage-pool-name) is an option for advanced replication to choose the proper physical device in the target storage pool to create the replica resource on server B.

-TSP2 (--target-storage-pool-name2) is an option for advanced replication to choose the proper physical device in the target storage pool to create the replica resource on server C.

-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-vdev> [-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-vdev=<source-vdev> --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-vdev) is required to specify the ID of the virtual tape to be configured for replication.

-S (--target-name) is an option to specify the target server name where the tape replica will be created and replicated to. If the replication is local, use -H (--local) option.

-U (--target-username) and -P (--target-password) are optional for connection and login to the target server if the target server was not logged in with a login command.

-h (--local) is an option to create a local replica. Target server information and credentials are not required when using this option and are ignored if they are specified.

The replication configuration requires a trigger policy to be set. If no trigger policy is specified, the command will automatically apply the appropriate default policy based on the tape caching property of the specified virtual tape.

Any combination of the following two options can be used in order to set up a replication trigger policy for a virtual tape with the tape caching property disabled. The default policy is 1024 MB watermark.

-w (--watermark) is a data size based trigger in MB. The watermark is checked when the tape is unloaded from the tape drive and the replication is triggered if the amount of new data on the tape has reached the specified watermark.

-d (--date) combined with -i (--interval) is a time based trigger. The replication is triggered at the time specified by date and then repeated every interval. -d (--date) format is YYYYMMDDHHMM and -i (--interval) format is a number followed by H for hours or M for minutes (e.g. -i 2H or --interval=120M). The default value for interval is 1H (one hour).

-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).

For virtual tapes with tape caching enabled, replication is triggered based on the tape caching policy:

-r (--repl-first) is an option to replicate the virtual tape before it is migrated. Use *on* in order to enable this policy or *off* to have tape migration executed first. The default policy is to replicate the virtual tape after it is migrated.

Replication is retried based on the timeout policy:

- **-t (--replication-timeout)** in seconds (default 60).
- **-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 <vdev> | -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> --vdev=<vdev> |
--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 (--vdev) 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 <vdev> [-X <rpc-timeout>]
```

```
iscon suspendreplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --vdev=<vdev>
[--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-vdev) 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 with the tape caching property disabled.

-w (--watermark) is a data size based trigger in MB. The watermark is checked when the tape is unloaded from the tape drive and the replication is triggered if the amount of new data on the tape has reached the specified watermark.

-d (--date) combined with -i (--interval) is a time based trigger. The replication is triggered at the time specified by date and then repeated every interval. -d (--date) format is YYYYMMDDHHMM and -i (--interval) format is a number followed by H for hours or M for minutes (e.g. -i 2H or --interval=120M).

-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.

For virtual tapes having the tape caching property enabled, the replication is triggered based on the tape caching policy:

-r (--repl-first) is required to replicate the virtual tape before it is migrated. Use "on" in order to enable this policy or "off" to have tape migration executed first.

The replication retry policy can be changed using the following options:

- -t (--replication-timeout) in seconds (default 60).
- -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-vdev> [-X <rpc-timeout>]
```

```
iscon getreplicationproperties --server-name=<server-name>  
[--server-username=<username> --server-password=<password>]  
--source-vdev=<source-vdev> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command shows the replication configuration for the specified virtual tape.

-v (--source-vdev) 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 <replica-id> [-X <rpc-timeout>]
```

```
iscon getreplicationstatus --target-name=<target-name>
[--target-username=<username> --target-password=<password>]
--replica-id=<replica-id> [--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 (--replica-id) 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 <vdev-id> [-X <rpc-timeout>]
```

```
iscon startreplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdev-id=<vdev-id> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command starts replication on demand for a virtual device.

-v (--source-vdev-id) 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 <vdev-id> [-X <rpc-timeout>]
```

```
iscon stopreplication --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdev-id=<vdev-id> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command stops the replication that is in progress for a virtual device.

-v (--source-vdev-id) 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. If the specified server is part of a cluster, the password will be changed for all nodes in the cluster.

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

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

-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.

Physical libraries and drives

Inventory physical tape library

```
iscon plibinventory -s <server-name> [-u <username> -p <password>]
[-l <physical-tape-library-vid>] [-X <rpc-timeout>]

iscon plibinventory --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--plib-vid=<tape-library-vid>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command refreshes the physical tape information for the specified physical library if the information is out of sync.

-l (--plib-vid) is an option to specify the physical tape library to perform the inventory.

Inventory operation will be performed for all the physical tape libraries if -l (--plib-vid) is not specified.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Mark physical library or drive enabled/disabled

```
iscon markphysicallibdrvstate -s <server-name> [-u <username> -p <password>]
-v <plib-or-pdrive-vid> [-E <on | off>] [-X <rpc-timeout>]

iscon markphysicallibdrvstate --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--plib-or-pdrive-vid=<plib-or-pdrive-vid> [--state=<on | off>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command marks a physical tape library or drive as enabled or disabled.

-v (--plib-or-pdrive-vid) is required to specify the ID of the device.

-E (--state) is required to specify the new state for the device. If the argument is not specified, the command retrieves the current state for the device.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Physical tapes

Move physical tape

```
iscon movephysicaltape -s <server-name> [-u <username> -p <password>]
-m <move-operation> -L <physical-tape-library-vid>
-B <physical-tape-barcode> | -l <from-location-id> -t <to-location-id>
[-X <rpc-timeout>]

iscon movephysicaltape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--move-operation=<move-operation> --tape-library-vid=<physical-tape-library-vid>
--physical-tape-barcode=<barcode> | --from-location-id=<from-location-id>
--to-location-id=<to-location-id> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command moves a physical tape to a new location.

-l (--plib-vid) is the ID of the physical tape library.

-m(--move-operation) is one of the following operations:

- DriveToSlot
- SlotToSlot
- SlotToDrive
- IESlotToSlot
- SlotToIESlot

-L(--tape-library-vid) is the physical library virtual ID where the tape is located.

-B(--physical-tape-barcode) identifies the physical tape to be moved. If barcode is not provided, the current tape location must be provided accordingly to the requested operation.

-l(--from-location-d) is the current slot or import/export (IE) slot number, or the physical drive virtual ID.

-t(--to-location-id) is the destination slot or IE slot number or the physical drive virtual ID.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Eject physical tape

```

iscon ejectphysicaltape -s <server-name> [-u <username> -p <password>]
-L <physical-tape-library-vid> -B <physical-tape-barcode-list>
[-A <acs-lsm-cap>] [-X <rpc-timeout>]

iscon ejectphysicaltape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-library-vid=<physical-tape-library-vid>
--tape-barcode-list=<physical-tape-barcode-list> | [--acs-lsm-cap=<acs-lsm-cap>]
[--rpc-timeout=<rpc-timeout>]

```

Description:

This command ejects physical tapes from the specified library.

-L(--tape-library-vid) is the physical library virtual ID where the tapes are located.

-B(--tape-barcode-list) identifies the physical tapes to be ejected. This argument can be a list of barcodes separated with commas. The list should be enclosed in quotes.

-A <--acs-lsm-cap> is an optional argument representing the Cartridge Access Port for the Automated Cartridge System Library Software libraries. The format of the argument is acs:lsm:cap

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Prepare physical tape for stacking

```

iscon preparestackingtape -s <server-name> [-u <username> -p <password>]
-v <standalone-physical-tape-drive-vid> -b <physical-tape-barcode> -O [-X <rpc-timeout>]

iscon preparestackingtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--pdrive-vid=<physical drive id> --physical-tape-barcode=<barcode>
--overwrite-mode [--rpc-timeout=<rpc-timeout>]

```

Description:

This command prepares a physical tape for stacking. If the tape is not blank, all existing information will be deleted in the process.

-v (--pdrive-vid) is required to provide the ID of the physical tape drive where the physical tape is loaded.

-b (--physical-tape-barcode) is required to provide a barcode for the physical tape. This barcode will replace the old barcode. The barcode must be between 4 and 12 characters.

-O (--overwrite-mode) is required in order to confirm that the existing data will be erased.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Scan stacked physical tape

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

```
iscon retrieveptapebarcode --server-name=<server-name>  
[--server-username=<sername> --server-password=<password>]  
--pdrive-vid=<physical drive id> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command scans a stacked physical tape loaded in the specified standalone physical tape drive in order to retrieve the barcode from the tape header. The command does not wait for the operation to complete. Use "getvdevlist" with "-v" and "-l" to get the barcode displayed.

-v (--pdrive-vid) is required to provide the ID of the physical tape drive where the physical tape is loaded.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Automated tape caching

Set tape caching

```
iscon settapecaching -s <server-name> [-u <username> -p <password>]
-L <library-vid> -t <tape-caching-enable> -m <delete-cache | data-deduplication>
-T <# of days> [-H <hours>] [-S <start-time>] [-W <day-of-the-week>] [-b <and-or>]
[-d <# of hours>] [-c][[-e][-f]] [[-I | -M | -R <# of days> | -N] [-X <rpc-timeout>]
```

```
iscon settapecaching --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-library-vid=<library-vid> --tape-caching-enable=<tape-caching-enable>
[--reclamation-method=<delete-cache|data-deduplication>]
[--turn-into-directlink=<# of days>]
[--hourly <hours>] [--start-time=<start-time>] [--day-of-the-week=<day-of-the-week>]
[--trigger-combine=<and-or>] [--retention-hours=<# of hours>]
[--migration-threshold] [--tape-ejected-to-slot] [--tape-full]]
[--immediately | --reclamation-threshold | --retention-days=<# of days> | --never]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command can be used in order to enable, disable, or change the Automated Tape Caching policy for a virtual tape library. The reclamation method used is "delete cache".

-L (--tape-library-vid) is the virtual device ID of the virtual tape library to be set.

Set -t (--tape-caching-enable) to 1 for enable or 0 to disable. If the *disable* option is used, all other arguments will be ignored. The *enable* option must be used in order to set or change the tape caching policy.

-m (--reclamation-method) is an option to select from the following reclamation methods: delete-cache (default method) and data-deduplication.

Time-based data migration triggers: Time Based

-H (--hourly) is an option to start migration every specified number of hours, up to 23. This option cannot be combined with any other migration trigger.

-S (--start-time) alone can be used to start daily migrations at the time specified. When combined with other data migration triggers, the -S option will delay the migration execution to the specified time.

-W (--day-of-the-week) is the default time based trigger that can be used to start weekly migrations on the specified day at 00:00(am): Sunday: 0, Monday: 1, ..., Saturday: 6. This option is ignored if Policy Based triggers are used.

Policy-based data migration triggers:

-b (--trigger-combine) tells how trigger policies are combined (specified by -d, -c, -e). 1 -- and; 0 -- or. The default value is 1 (and).

-d (--retention-hours) triggers the data migration after the data was retained on the disk for the specified number of hours, up to a year.

-c (--migration-threshold) triggers the data migration when the disk usage percentage is above the global disk space threshold.

-e (--tape-ejected-to-slot) triggers the data migration when unloading a virtual tape from a drive that had data written to it.

-f (--tape-full) applies to the --tape-ejected-to-slot trigger. The data is migrated only if the tape becomes full.

Reclamation triggers:

-I (--immediately) triggers the data reclamation immediately after the data migration completes.

-M (--reclamation-threshold) triggers the data reclamation when the disk usage percentage is above the global disk space threshold.

-R (--retention-days) triggers the data reclamation after the specified number of days, up to 2000.

-N (--never) means that virtual tapes will never be reclaimed.

The reclamation triggers are mutually exclusive.

The command always overwrites the current policy. Therefore, all of the desired properties must be provided each time the command is executed. The following properties are set by default if no triggers are provided:

- Migration trigger: daily at 00:00(am)
- Reclamation trigger: immediately

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get physical tape list

```
iscon getphysicaltapelist -s <server-name> [-u <username> -p <password>]
-l <physical-tape-library-vid> | -v <standalone-physical-tape-drive-vid>
[-F <filter>] [-S] [-X <rpc-timeout>]
```

```
iscon getphysicaltapelist --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--plib-vid=<physical-tape-library-vid> | --pdrive-vid=<standalone-physical-tape-drive-vid>
[--ptape-filter=<filter>] [--stacking-info]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command displays a list of physical tapes located in the specified physical tape library or standalone physical tape drive.

-l (--plib-vid) is the ID of the physical tape library.

-v (--pdrive-vid) is the ID of the standalone physical tape drive where the physical tape is located.

-F (--ptape-filter) is an option to show only the physical tapes having the specified property:

- SYNC (eligible for sync operation, applies to library only) OR
- STACKED (stacked tape).

-S <--stacking-info> is an option to retrieve and display the stacking information of the stacked physical tapes when the stacked tape filter is used.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Sync physical tapes

```
iscon syncphysicaltape -s <server-name> [-u <username> -p <password>]
-l <plib-vid> -b <physical-tape-barcode> -L <virtual-tape-library-id>
-t <virtual-tape-slot-no> [-M <sync-mode>] [-k <key-name> -W <key-password>]
[-I <ACSL list>] [-n <vdevname>] [-g <#(GB)>] [-X <rpc-timeout>]

iscon syncphysicaltape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--plib-vid=<physical-tape-library-vid> --physical-tape-barcode=<physical-tape-barcode>
--tape-library-vid=<virtual-tape-library-id>
--virtual-tape-slot-no=<virtual-tape-slot-no> [--sync-mode=<sync-mode>]
[--key-name=<key-name> --key-password=<key-password>] [--scsiaddress=<ACSL list>]
[--vdevname=<vdevname>] [--size-gb=<#(GB)>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a synchronized virtual tape for each physical tape provided. The physical tapes must be from the specified physical tape library and the virtual tape will be created in the specified virtual tape library. The virtual tape library must have the tape caching feature enabled.

-l (--plib-vid) is the virtual ID of the physical tape library where the physical tapes are located.

-b (--physical-tape-barcode) is the barcode of the physical tape. If the barcode contains leading or trailing spaces, it must be enclosed in double quotes. For batch mode, the argument can be a list of barcodes separated by commas or a file name enclosed in "<>" (i.e., "<file>") containing a barcode on each line. Do not use quotes inside the file. The file must be located in the same folder as the command line utility or a full path is required. The virtual tape(s) will be created with the same barcode as the physical tape(s). The barcode(s) must not be in use by any other virtual tape in the system.

-L (--tape-library-vid) is the ID of the virtual tape library where the virtual tapes will be created.

-t (--virtual-tape-slot-no) is an option to provide an empty destination slot for the virtual tape. Not for "-M cache" mode. For batch mode, virtual tapes will be created starting with the specified slot number.

-M (--sync-mode) is an option to select the synchronization mode from one of the following values (default is "cache"):

- cache (create cache)
- metadata (create cache and copy meta data)
- directlink (create direct link)

-k (--key-name) and -W (--key-password) are options for tape encryption support. If the tape to be synchronized was encrypted through the system, you need to specify the key name and the key password of the encryption key to decrypt the data.

The following three options can be selected for create cache mode only:

-I (--scsiaddress) is an option to specify preferred physical devices for creating the virtual device. It can be a list of ACSLs separated with commas. ACSL=#:#:# (adapter:channel:scsi id:lun)

-n (--vdevname) is an option to specify the virtual tape name or prefix when creating more than one tape. The maximum length of the virtual device name is 64. Leading and trailing spaces will be removed. Enclose the name in double quotes to ensure the proper name is used. The following characters are invalid for the name: <> "&" '\$' '/'

-g (--size-gb) is an option to specify the initial size, in GB, of the virtual tapes, if the capacity-on-demand property for the virtual tape library is enabled. The default is 1 GB.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Migrate virtual tapes

```
iscon migratevirtualtapes -s <server-name> [-u <username> -p <password>]
-T <tape-vid-list> [-f] [-X <rpc-timeout>]

iscon migratevirtualtapes --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-vid-list=<tape-vid-list> [--tape-full] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command migrates the specified virtual tapes to the physical libraries they are synchronized with.

-T (--tape-vid-list) is a list of virtual tape ID(s) separated with commas.

-F (--tape-full) is an option to force full tape migration. By default, the migration operation is incremental.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Reclaim disk space

```
iscon reclaimtapes -s <server-name> [-u <username> -p <password>]
-T <tape-vid-list> [-X <rpc-timeout>]

iscon reclaimtapes --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-vid-list=<tape-vid-list> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command reclaims the disk space occupied by the specified migrated virtual tapes.

-T (--tape-vid-list) is required to specify the ID of the virtual tapes to be reclaimed, separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Renew cache

```
iscon renewcache -s <server-name> [-u <username> -p <password>]
-v <vdevid> [-M <metadata>] [-k <key-name> -W <key-password>] [-I <ACSL>] [-n <vdevname>]
[-g <#(GB)>] [-X <rpc-timeout>]

iscon renewcache --server-name=<server-name>
[--server-username=<username> --server-password=<password>] [--import-mode=<metadata>]
[--key-name=<key-name> --key-password=<key-password>] [--scsiaddress=<ACSL>]
[--vdevname=<vdevname>] [--size-gb=<#(GB)>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command converts a virtual stub tape into a virtual cache tape.

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

-M (--import-mode) is an option to specify that the header area should be copied from the physical tape to the new virtual tape cache. The value of this option must be: *metadata*.

-k (--key-name) and -W (--key-password) are options for tape encryption support. If the tape to be renewed was encrypted through the system, you need to specify the key name and the key password of the encryption key to decrypt the data.

The following properties of the virtual cache tape can be set if the "-M" option is not specified:

-l (--scsiaddress) is an option to specify preferred physical devices for creating the virtual device. It can be a list of ACSLs separated with commas or a file enclosed in <> containing an ACSL on each line.
ACSL=#:#:# (adapter:channel:scsi id:lun)

-n (--vdevname) is an option to specify the virtual tape name or prefix when creating more than one tape. The maximum length of the virtual device name is 64. Leading and trailing spaces will be removed. Enclose the name in double quotes to ensure it is properly parsed and interpreted. The following characters are invalid for the name: <>"&\$\'

-g (--size-gb) is an option to specify the size in GB. The size of the virtual tape will be the size configured in the properties of the virtual tape library or virtual tape drive if it is not specified. This option cannot be specified if the capacity-on-demand option is not enabled at library level.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get tape caching info

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

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

Description:

This command reports a list of virtual tapes that are candidates for tape migration and summarizes the information about the space used by those tapes.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Import/Export

Get import/export job status

```
iscon getimportexportjobstatus -s <server-name> [-u <username> -p <password>]
[-j <job-id-list>] [-T <job-type> -S <job_status> -D <date-range|date> -l]
[-X <rpc-timeout>]

iscon getimportexportjobstatus --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--job-id-list=<job-id-list>] | [--job-type=<job_type> --job_status=<job_status>]
--date-range=<date-range|date> --longlist] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command displays the status of the import/export jobs present in the queue. If no filters are specified, the command displays all the jobs that are in the queue.

-j <--job-id-list> is an optional list of job IDs separated with commas. The command displays the status of specified jobs only. All other filters are ignored.

-T <--job-type> is an optional job type based filter. The command displays those jobs matching the provided type. The accepted job type values are: IMPORT, EXPORT, or OTHER (such as scan).

-S <--job_status> is an optional job status based filter. The command displays those jobs matching the provided status. The accepted job status values are: FAILED, HOLD, READY, or OTHER (such as waiting for tape/drive or cancelled).

-D <--date-range> is an option to specify the date range for the report (future dates are ignored): YYYYMMDD-YYYYMMDD or YYYYMMDD.

-l <--longlist> is an option to display detailed information.

-X <--rpc-timeout> is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Import tape

```
iscon importtape -s <server-name> [-u <username> -p <password>]
[-M <import-mode>] -v <plib-or-pdrive-vid> [-B <barcode> | -l <slot-no>]
-L <tape-library-vid> [-b <virtual-tape-barcode>] -t <virtual-tape-slot-no>
[-j <job-description>] [-k <key-name> -W <key-password>]
[-SP <storage-pool-name>] [-X <rpc-timeout>]

iscon importtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--import-mode=<import-mode>] --plib-or-pdrive-vid=<plib-or-pdrive-vid>
[--barcode=<barcode> | --slot-no=<slot-no>] --tape-library-vid=<tape-library-vid>
--virtual-tape-slot-no=<virtual-tape-slot-no>
[--virtual-tape-barcode=<virtual-tape-barcode>] [--job-description=<job-description>]
[--key-name=<key-name> --key-password=<key-password>]
[--storage-pool-name=<storage-pool-name>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command imports a physical tape to a virtual tape. The virtual tape is created during the import operation.

-M (--import-mode) is an option in one of the following values: *copy* (default) or *direct-access* or *recycle*.

-v (--pdrive-or-pdrive-vid) is required to specify the virtual device ID of the physical tape library or physical tape drive from which the physical tape is to be imported.

If the physical tape is from a physical tape library, either <barcode> or <slot-no> of the physical tape must be specified with -B (--barcode) or -l (--slot-no) to identify the physical tape. If the barcode contains leading or trailing spaces, it must be enclosed in double quotes. Physical tape information is not required if the physical tape is imported from a standalone physical tape drive.

-L (--tape-library-vid) is the virtual device ID of the virtual tape library to which the physical tape is to be imported.

-t (--virtual-tape-slot-no) is required for the virtual tape location.

-b (--virtual-tape-barcode) is optional when the physical tape from a physical tape library contains a barcode. It is required if the physical tape does not have a barcode or when it is from a physical tape drive.

-j (--job-description) is an option to specify a description for the import job.

-k (--key-name) and -W (--key-password) are options for tape encryption support. If the tape to be imported was encrypted through the system, you need to specify the key name and the key password of the encryption key to decrypt the data.

-SP (--storage-pool-name) is an option that will automatically choose the proper physical device in the storage pool to create a 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.

Import stacked tape

```
iscon importstackedtape -s <server-name> [-u <username> -p <password>]
[-l <physical-tape-library-vid> -b | -B <physical-tape-barcode> | -o <slot-no>] |
[-v <standalone-physical-tape-drive-vid> [-b <physical-tape-barcode>]]
-L <virtual-tape-library-vid> [-k <key-name> -W <key-password>]
[-z <originalbarcode:newbarcode:slotnumber>] [-j <job-description>]
[-SP <storage-pool-name>] [-X <rpc-timeout>]
```

```
iscon importstackedtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--plib-vid=<physical-tape-library-vid>
--physical-tape-barcode=<physical-tape-barcode> | --slot-number=<slot-no>] |
[-pdev-vid=<standalone-physical-tape-drive-vid>
[--physical-tape-barcode=<physical-tape-barcode>]]
--tape-library-vid=<virtual-tape-library-vid>
[--key-name=<key-name> --key-password=<key-password>]
[--vtape-info=<originalbarcode:newbarcode:slotnumber>]
[--job-description=<job-description>]
[--storage-pool-name=<storage-pool-name>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command submits an import job in order to copy the content of the specified stacked physical tape to virtual tapes.

If the physical tape is located in a physical tape library:

- -l (--plib-vid) is required to specify the ID of physical tape library where the physical tape is located.
- Either -b | -B (--physical-tape-barcode) or -o (--slot-number) must be used in order to identify the physical tape. If the barcode contains leading or trailing space characters, it must be enclosed in double quotes.

If the physical tape is located in a standalone physical tape drive:

- -v (--pdrive-vid) is required to specify the ID of the standalone physical tape drive where the physical tape is located.
- -b | -B (--physical-tape-barcode) is an option to validate the physical tape by matching this barcode with the barcode in the tape header. If barcode is not provided, the command will import the physical tape without checking. If the barcodes do not match, the operation will fail.

The following arguments are common:

-L (--tape-library-vid) is required to specify the ID of the virtual tape library to which you want to import the stacked tape.

-k (--key-name) and -W (--key-password) are options for tape encryption support. If the tape to be imported was encrypted through the system, you need to specify the key name and the key password of the encryption key to decrypt the data.

-z (--vtape-info) is an option to select the virtual tapes that are to be imported from the specified stacked physical tape and specify how to create those tapes. By default, all virtual tapes found on the physical tape will be imported. This option contains a list of arguments separated with commas, in the following format:
"originalbarcode:newbarcode:slotnumber,..."

Enclose the list in double quotes if the barcodes contain leading or trailing spaces.

In order to create the corresponding virtual tape in the first available slot in the specified virtual tape library, use the keyword "any": originalbarcode:newbarcode:any

In order to preserve the original barcode of the virtual tape, use the keyword "same": originalbarcode:same:any

-j (--job-description) is an option to specify a description for the import job.

-SP (--storage-pool-name) is an option that will automatically choose the proper physical device in the storage pool to create a 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.

Export virtual tape

```
iscon exportvirtualtape -s <server-name> [-u <username> -p <password>]
-v <vdevvid> -L <tape-library-vid> -b | -B <barcode> | -l <slot-no>
[-Z <on> -Q <num-of-copies>] [-J]] | [-pd <standalone-physical-tape-drive-vid>
[-M <export-mode> [-Y <days>] ] [-j <job-description>] [-f]
[-k <key-name> -W <key-password>] [-X <rpc-timeout>]

iscon exportvirtualtape --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdevvid=<vdevvid> --tape-library-vid=<tape-library-vid>
--same-barcode | --barcode=<barcode> | --slot-no=<slot-no>
[--tape-duplication=<on> --num-of-copies=<num-of-copies>]
[--auto-eject-to-ie]] [--pdrive-vid=<standalone-physical-tape-drive-vid>
[--export-mode=<export-mode> [--delay-delete-days=<days>]]
[--job-description=<job-description>] [--force]
[--key-name=<key-name> --key-password=<key-password>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command exports the information from a virtual tape to a physical tape.

-v (--vdevvid) is required to specify the ID of the virtual tape to be exported to the physical tape.

If the physical tape is located in a physical tape library:

-L (--tape-library-vid) is required to specify the ID of the target physical tape library.

You must include one of the following arguments to select the physical tape:

- -b (--same-barcode) lets you select a physical tape with the same barcode of the virtual tape if a physical tape with the same barcode exists.
- -B (--barcode) lets you specify the barcode of an available physical tape in the physical tape library. If the barcode contains leading or trailing spaces, it must be enclosed in double quotes.
- -l (--slot-no) lets you specify the slot number of an available physical tape in the physical tape library.

-Z (--tape-duplication) is an option to enable Tape Duplication for this export job: *on* (enable); default is *off* (disabled).

-Q (--num-of-copies) is an option to specify the number of copies to be made using the same barcode if the tape duplication option is enabled. The maximum value is 5. The default value is 1.

-J (--auto-eject-to-ie) is an option to eject the tape to the IE slot after the export job.

If the physical tape is located in a standalone physical tape drive:

-pd (--pdrive-vid) is required to specify the ID of the standalone physical tape drive where the physical tape is located.

The following arguments are common:

-M (--export-mode) is an option with one of the following values: *copy* (default) or *move*.

-Y (--delay-delete-days) is an option for *move mode* to specify the number of days to wait before deletion. The maximum is 365 days. The default value is 365 days.

-j (--job-description) is an option to specify a description for the tape export job.

-f (--force) is required when the tape is scheduled to be deleted.

-k (--key-name) and -W (--key-password) are options for tape encryption support. Specify the key name and key password of the encryption key if you wish to encrypt the data when exporting the virtual tape to the physical tape.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Export virtual tapes with stacking

```
iscon exportvirtualtapeswithstacking -s <server-name> [-u <username> -p <password>]
-T <tape-vid-list> [-L <physical-tape-library-vid>
-B <physical-tape-barcode> | -l <slot-no>] |
[-v <standalone-physical-tape-drive-vid> -b | -B <physical-tape-barcode>]]
[-M <export-mode> [-Y <days>] [-j <job-description>] [-f] [-J] [-O]
[-k <key-name> -W <key-password>] [-X <rpc-timeout>]

iscon exportvirtualtapeswithstacking --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--tape-vid-list=<tape-vid-list> [--tape-library-vid=<physical-tape-library-vid>
--physical-tape-barcode=<physical-tape-barcode> | --slot-no=<slot-no>] |
[--pdrive-vid=<standalone-physical-tape-drive-vid>
--same-barcode | --physical-tape-barcode=<physical-tape-barcode>]]
[--export-mode=<export-mode> [--delay-delete-days=<days>]]
[--job-description=<job-description>] [--force] [--auto-eject-to-ie] [--overwrite-mode]
[--key-name=<key-name> --key-password=<key-password>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command exports up to one thousand virtual tapes to a physical tape. The virtual tapes must be in the virtual vault.

-T <--tape-vid-list> is required to specify the ID of the virtual tapes to be exported. The argument can be a list of IDs or ranges of IDs, separated with commas, or a file name enclosed in "< >" (i.e., "<file>"). The file must contain a single tape ID or range on each line. The file must be located in the same folder as the command line utility or the full path is required. For example: -T 10000001,10000010-10000020 or -T "</tmp/exporttapes>"

If the physical tape is located in a physical tape library:

-L <--tape-library-vid> is required to specify the ID of the physical tape library hosting the physical tape.

Either -B (--physical-tape-barcode) or -l (--slot-number) must be used in order to identify the physical tape. If the barcode contains leading or trailing space characters, it must be enclosed in double quotes.

If the physical tape is located in a standalone physical tape drive:

-v (--pdrive-vid) is required to specify the ID of the standalone physical tape drive where the physical tape is located.

-b | -B (--physical-tape-barcode) is required to specify the physical tape barcode. The specified barcode must match the the barcode form the physical tape header or the command will fail.

The following arguments are common:

-M (--export-mode) is an option with one of the following values: *copy* (default) or *move*.

-Y (--delay-delete-days) is an option for *move* mode to specify the number of days to wait before deletion. The maximum is 365 days. The default value is 365 days.

-j (--job-description) is an option to specify a description for the tape export job.

-f (--force) is required when the tape is scheduled to be deleted.

-J (--auto-eject-to-ie) is an option to eject the tape to the IE slot after the export job.

-O (--overwrite-mode) is an option to overwrite the physical tape. By default, the data is appended to the physical tape.

-k (--key-name) and -W (--key-password) are options for tape encryption support. Specify the key name and key password of the encryption key if you wish to encrypt the data when exporting the virtual tape to the physical tape.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Scan physical tapes

```
iscon scanphysicaltapes -s <server-name> [-u <username> -p <password>]
-l <physical-tape-library-vid> [-b | -B <physical-tape-barcode-list>] |
[-v <standalone-physical-tape-drive-vid> -b | -B <physical-tape-barcode>]]
[-X <rpc-timeout>]
```

```
iscon scanphysicaltapes --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--plib-vid=<physical-tape-library-vid>
[--tape-barcode-list=<physical-tape-barcode-list>]] |
[--pdrive-vid=<standalone-physical-tape-drive-vid>
[--physical-tape-barcode=<physical-tape-barcode>]] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command submits individual scan jobs for each physical tape found in the specified location. The command does not wait for the scan jobs to finish.

If the location to be scanned is a physical tape library:

-l (--plib-vid) is the ID of the physical tape library.

-b | -B (--tape-barcode-list) is an option to scan only the specified physical tapes from the same library. Multiple barcodes must be separated with commas and the list must be enclosed in double quotes if the barcodes contain leading or trailing space characters.

If the location to be scanned is a standalone physical tape drive:

-v (--pdriveID) is the ID of the standalone physical tape drive.

-b | -B (--physical-tape-barcode) is an option to validate the physical tape by matching this barcode with the barcode in the tape header. If barcode is not provided, the command will scan the physical tape without checking.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Resume import/export jobs

```
iscon resumeimportexportjobs -s <server-name> [-u <username> -p <password>]
-j <job-id-list> [-X <rpc-timeout>]
```

```
iscon resumeimportexportjobs --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --job-id-list=<job-id-list>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command resumes the specified import/export jobs. The jobs must be in the import/export queue in a suspended state.

-j <--job-id-list> is a list of job IDs separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Restart import/export jobs

```
iscon restartimportexportjobs -s <server-name> [-u <username> -p <password>]
-j <job-id-list> [-X <rpc-timeout>]
```

```
iscon restartimportexportjobs --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--job-id-list=<job-id-list> [--rpc-timeout=<rpc-timeout>]
```

Description:

This command restarts the specified import/export jobs. The jobs must be in the import/export queue and they must have either been cancelled or failed.

.j <--job-id-list> is a list of job IDs separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete import/export jobs

```
iscon deleteimportexportjobs -s <server-name> [-u <username> -p <password>]
-j <job-id-list> [-X <rpc-timeout>]
```

```
iscon deleteimportexportjobs --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --job-id-list=<job-id-list>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes the specified import/export jobs. The jobs must be in the import/export queue.

-j <--job-id-list> is a list of job IDs separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Suspend import/export jobs

```
iscon suspendimportexportjobs -s <server-name> [-u <username> -p <password>]
-j <job-id-list> [-X <rpc-timeout>]

iscon suspendimportexportjobs --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --job-id-list=<job-id-list>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command suspends the specified import/export jobs. The jobs must be in the import/export queue and must be idle.

-j <--job-id-list> is a list of job IDs separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Cancel import/export jobs

```
iscon cancelimportexportjobs -s <server-name> [-u <username> -p <password>]
-j <job-id-list> [-X <rpc-timeout>]

iscon cancelimportexportjobs --server-name=<server-name>
[--server-username=<username> --server-password=<password>] --job-id-list=<job-id-list>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command cancels the specified import/export jobs. The jobs must be in the import/export queue and must be running.

-j <--job-id-list> is a list of job IDs separated with commas.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

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.

Hosted Backup

Enable Hosted Backup

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

Description:

This command enables the Hosted Backup option.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Disable Hosted Backup

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

Description:

This command disables the Hosted Backup option.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Failover

Get failover status

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

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

Description:

This command obtains the failover status from either the primary failover server or the secondary failover server.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

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 StorSafe 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.

NAS devices

Create NAS device

```
iscon createnasdevice -s <server-name> [-u <username> -p <password>]
-I <ACSL list> | --SP <storage-pool-name>
[-O] [-F <file-system> [-fo <format-options>]
[-mo <mount-options>] [-n <vdevname>]
[-m <# (MB)>] [-X <rpc-timeout>]

iscon createnasdevice --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--scsiaddress=<ACSL list> | --storage-pool-name=<storage-pool-name>
[--ost] [--file-system=<file-system>
[--format-options=<format-options>] [--mount-options=<mount-options>]
[--vdevname=<vdevname>] [--size-mb=<# (MB)>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a NAS virtual device.

-I (--scsiaddress) is required to specify a physical device to be used to create the virtual device. The argument can be a list of ACSLs separated by commas. `##:##:##:##:##` (adapter:channel:scsi id:lun)

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

-O (--ost) is an option used to create a device for OpenStorage.

-F (--file-system) is an option to choose the file system type (e.g. EXT4). The default file system type will be determined by the NAS configuration.

-fo (--format-options) is an option to specify the format options; e.g. `-F -I 512 -m 0 -v -j -E resize=16383G -J size=128 -b 4096`

-mo (--mount-options) is an option to specify the mount options; e.g. `rw,nosuids,user_attr`

-m (--size-mb) is an option to specify the size in MB for the allocated space. If this option is not specified, the total available space will be allocated. The minimum size for the NAS resource is 1000.

-n (--vdevname) is an option to specify the virtual device name. The maximum length is 64. Leading and trailing spaces will be removed. Enclose the name in double quotes to ensure proper parsing. These characters are invalid for NAS Resource name: `<>"&$\^()'%# :;|'?'`

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Expand NAS device

```
iscon expandnasdevice -s <server-name> [-u <username> -p <password>]
-v <vdev-id> -I <ACSL list> [-m <#(MB)>]
[-W <wait-for-device-status>]
[-X <rpc-timeout>]
```

```
iscon expandnasdevice --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdev-id=<vdev-id>
--scsiaddress=<ACSL list>
[--additional-mb=<#(MB)>]
[--wait-for-device-status=<wait-for-device-status>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command expands a NAS resource. It is recommended that you detach all clients from the resource before executing this command.

-v (--vdev-id) is required to provide the resource ID.

-I (--scsiaddress) is required to specify the SCSI address of the luns to be used for expansion, separated by commas; e.g., `###:###:###:###`

-m (--additional-mb) is an option to specify how much space (in MB) to use for expansion. If this option is not specified, the total available space will be allocated.

-W (--wait-for-device-status) is an option to wait for the operation to complete, between 3 and 3600 seconds. By default, the command does not wait for the expansion result.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

Format NAS resource

```
iscon formatnas -s <server-name> [-u <username> -p <password>]
-v <vdev-id>
[-mo <mount-options>] [-fo <format-options>]
[-X <rpc-timeout>]
```

```
iscon formatnas --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdev-id=<vdev-id>
[--mount-options=<mount-options>]
[--format-options=<format-options>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command formats a NAS resource file system.

-v (--vdev-id) is required to specify the virtual device ID.

-mo (--mount-options) specifies options for mounting the file system. Refer to the Linux mount command man page for all available file system mount options.

-fo (--format-options) specifies options for formatting the file system. Refer to the Linux mke4fs command man page for details. Options must be entered inside double quotes; e.g., "option1 option2 ...".

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete NAS device

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

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

Description:

This command deletes a NAS resource.

-v (--vdev) is required - specify the device ID.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

NAS shares

Create a share

```
iscon createshare -s <server-name> [-u <username> -p <password>]
-v <vdev> [-x <export-path>]
-f <folder-name>
-W <cifs-share-name>
[-C <connection-limit>] [-c <share-comment>]
-A <cifs-account-name> [-a <cifs-account-access>] |
-cn <client-name> | -m <client-machine>
-ca <client-access-mode> [-E] [-Q <squash>]
[-X <rpc-timeout>]

iscon createshare --server-name=<server-name>
[--server-username=<username>--server-password=<password>]
--vdev=<vdev> [--export-path=<export-path>] --folder-name=<folder-name>
--cifs-share-name=<cifs-share-name> [--connection-limit=<connection-limit>]
[--share-comment=<share-comment>]
--cifs-account-name=<cifs-account-name> [--cifs-account-access=<cifs-account-access>] |
--client-name=<client-name> | --client-machine=<client-machine>
--access-mode=<client-access-mode> [--insecure] [--squash=<squash>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a folder and shares it with an existing NFS and/or SMB client.

-v (--vdev) is required to specify the device ID of the NAS resource.

-f (--folder-name) is required to specify the folder name.

-x (--export-path) is an option to specify the parent path of the folder if it is not under the root of the NAS resource. For example, to create a share named 'share-folder-1' under '/nas/NAS-00005/fds/account-depart', use the following arguments:

- --vdev=5
- --folder-name=share-folder-1
- --export-path=account-depart

Share the folder with an existing client.

SMB client:

-W (--cifs-share-name) is required.

-C (--connection-limit) is an option to limit the number of connections to the share. Specify 0 for “unlimited” or a positive number as the maximum number of connections allowed.

-c (--share comment) is an option to add a comment for the share. The maximum length is 255 characters.

-A (--cifs-account-name) is required.

-a (--cifs-account-access) is an option to specify access mode: *ReadOnly* or *ReadWrite* (the default).

Example for an SMB share:

```
iscli createshare --server-name=h41-93 --vdev=10 --folder-name=shrFldTest --cifs-share-name=shrFldTest  
--window-account-name=winTest --window-account-access=ReadOnly
```

NFS client:

Either `-cn` (`--client-name`) or `-m` (`--client-machine`) is required.

`-cn` (`--client-name`) is the name of the NFS client.

`-m` (`--client-machine`) defines the client machine. Enter either the IP address or full hostname. On certain subnetworks, this is defined using network IP/netmask.

- Netmask can be specified using the IP address; e.g., "255.255.0.0" or CIDR notation (a prefix length); e.g., "16". Either specification has the same meaning in the command.
- Whether you use the IP address or CIDR format to specify the netmask, the network IP should have the same number of characters as netmask; e.g., 10.1.0.0/255.255.0.0 or 172.30.123.0/24.

`-ca` `<client-access-mode>` is an option to specify client access. This can be *ReadOnly* or *ReadWrite*.

`-E` (`--insecure`) is a security option. The default is "secure". Enter `-E` (`--insecure`) if the client's operating system does not use a reserved port for NFS.

`-Q` (`--squash`) is an option that maps user IDs. The default is "root", which remaps UID=0 (the root user) to "nfsnobody:nasgrp". Enter `-Q` (`--no`) to prevent mapping UIDs or `-Q` (`--all`) to map all UIDs to "nfsnobody:nasgrp".

Example for an NFS share:

```
iscli createshare --server-name=h41-93 --vdev=10 --folder-name=shrFldTest --client-name=h41-93 --access-  
mode=ReadOnly
```

Delete a share

```
iscon deleteshare -s <server-name> [-u <username> -p <password>]
-v <vdev> -x <export-path> | -W <cifs-share-name> [-X <rpc-timeout>]

iscon deleteshare --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdev=<vdev> --export-path=<export-path> | --cifs-share-name=<cifs-share-name>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes the specified share.

-v (--vdev) is required to specify the device id of the NAS resource.

-x (--export-path) or -W (--cifs-share-name) is required.

-x (--export-path) is required to identify an NFS share and must include the full path; e.g.,
/nas/NAS-00005/fds/account-depart/share-folder-1

Refer to [‘Create a share’](#) for definitions of other parameters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

Assign a share to an SMB client

```
iscon assignshare -s <server-name> [-u <username> -p <password>]
-W <cifs-share-name> -A <cifs-account-name> [-a <cifs-account-access>] [-G]
[-X <rpc-timeout>]

iscon assignshare --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--cifs-share-name=<cifs-share-name>
--cifs-account-name=<cifs-account-name>
[--cifs-account-access=<cifs-account-access>]
[--group]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command assigns an existing SMB share to an existing SMB account.

-W (--cifs-share-name) and -A (cifs-account-name) are required.

-a (--cifs-account-access) is optional. Permitted value is *ReadOnly* or *ReadWrite* (default).

-G (--group) is optional for a group account.

Refer to [‘Create a share’](#) for definitions of other parameters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Unassign a share from an SMB client

```

iscon unassignshare -s <server-name> [-u <username> -p <password>]
-W <cifs-share-name> -A <cifs-account-name> [-X <rpc-timeout>] [-G]

iscon unassignshare --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--cifs-share-name=<cifs-share-name> --cifs-account-name=<cifs-account-name>
[--group]
[--rpc-timeout=<rpc-timeout>]

```

Description:

This command unassigns a share from an SMB account.

-W (--cifs-share-name) and -A (cifs-account-name) are required.

-G (--group) is optional for a group account.

Refer to '[Create a share](#)' for definitions of other parameters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Assign a share to an NFS client

```

iscon assignnfs -s <server-name> [-u <username> -p <password>]
-v <vdev> -x <export-path>
-cn <client-name> | -m <client-machine>
-ca <client-access-mode> [-E] [-Q <squash>]
[-X <rpc-timeout>]

iscon assignnfs --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdev=<vdev> --export-path=<export-path>
--client-name=<client-name> | --client-machine=<client-machine>
[--access-mode=<client-access-mode>] [--insecure] [--squash=<squash>]
[--rpc-timeout=<rpc-timeout>]

```

Description:

This command assigns an existing NFS share to an existing NFS client.

-v (--vdev) is required to specify the device ID.

-x (--export-path) is required.

Either -cn (--client name) or -m (--client-machine) is required in order to assign an NFS share.

Refer to '[Create a share](#)' for definitions of other parameters.

-ca <client-access-mode> is an option to specify client access. This can be *ReadOnly* or *ReadWrite*.

-E (--insecure) is a security option. The default is "secure". Enter -E (--insecure) if the client's operating system does not use a reserved port for NFS.

-Q (--squash) is an option that maps user IDs. The default is “root”, which remaps UID=0 (the root user) to “nfsnobody:nasgrp”. Enter -Q (--no) to prevent mapping UIDs or -Q (--all) to map all UIDs to “nfsnobody:nasgrp”.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Unassign a share from an NFS client

```
iscon unassignnfs -s <server-name> [-u <username> -p <password>]
-v <vdev> -x <export-path> -cn <client-name> | -m <client-machine>
[-X <rpc-timeout>]
```

```
iscon unassignnfs --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--vdev=<vdev> --export-path=<export-path>
--client-name=<client-name> | --client-machine=<client-machine>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command unassigns a share from an existing NFS client.

-v (--vdev) is required to specify the device ID.

-x (--export-path) is required.

Either -cn (--client name) or -m (--client-machine) is required.

Refer to [‘Create a share’](#) for definitions of other parameters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

SMB/NFS clients

Add an SMB client

```
iscon addcifscclient -s <server-name> [-u <username> -p <password>]
-c <client-name> -t user -W <user-password> | -t group
[-d <description>]
[-X <rpc-timeout>]
```

```
iscon addcifscclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> --client-type=user
--password=<user-password> | --client-type=group
[--description=<description>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds an SMB client on a server with authentication mode set to USER.

-c (--client-name) is the user/group name.

-t (--client-type) is required with one of the following values: *user* (default) or *group*.

-W (--user-password) is required if the client type is *user*.

-d (--description) can be added for client type of user or group. The maximum length is 128 characters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete an SMB client

```
iscon deletecifscclient -s <server-name> [-u <username> -p <password>]
-c <client-name> -t <client-type>
[-X <rpc-timeout>]
```

```
iscon deletecifscclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> --client-type=<client-type>
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command deletes an SMB client on a server with authentication mode set to USER.

-c (--client-name) is the user/group name.

-t (--client-type) is required with one of the following values: *user* (default) or *group*.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Add an NFS client

```
iscon addnfsclient -s <server-name> [-u <username> -p <password>]
-cn <client-name> -m <client-machine> [-d <description>]
[-X <rpc-timeout>]
```

```
iscon addnfsclient --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--client-name=<client-name> --client-machine=<client-machine>
[--description=<description>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command adds an NFS client, which can be a single host or a subnetwork.

-cn (--client name) and -m (--client-machine) are required in order to add an NFS client.

-cn (--client-name) is the name of the NFS client.

-m (--client-machine) defines the client machine. Enter either the "IP address" or "Full Hostname". On certain subnetworks, this is defined using "Network IP/Netmask".

- "Netmask" can be specified using the IP address; e.g., "255.255.0.0" or CIDR notation (a prefix length); e.g., "16". Either specification has the same meaning in the command.
- Whether you use the "IP address" or "CIDR" format to specify the Netmask, the Network IP should have the same number of characters as Netmask; e.g., 10.1.0.0/255.255.0.0 or 172.30.123.0/24.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Delete an NFS client

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

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

Description:

This command deletes an NFS client.

-cn (--client-name) is the name of the NFS client.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Set security mode

```

iscon setsecuritymode -s <server-name> [-u <username> -p <password>]
{ -D -S <domain-controller> -N <domain-username> -P <domain-password>
[-B <bind-point>] [-O <OU>,<...>]
[-a <administrative-username> -w <administrative-password>] |
-U [-W <workgroup>] }
[-I <RANGE>,<...>] [-G <RANGE>,<...>] [-f] [-c <comment>]
[-X <rpc-timeout>]

iscon setsecuritymode --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
{ --domain-mode --domain-controller=<domain-controller>
--domain-username=<domain-username> --domain-password=<domain-password>
[--bind-point=<bind-point>] [--organizational-units=<OU>,<...>]
[--administrative-login=<administrative-login> --login-password=<login-password >] |
--user-mode [--workgroup=<workgroup>] }
[--uid-ranges=<RANGE>,<...>] [--gid-ranges=<RANGE>,<...>]
[--force] [--comment=<comment>]
[--rpc-timeout=<rpc-timeout>]

```

Description:

This command is used to set up SMB user security mode in domain mode or user mode.

Domain mode:

-D (--domain-mode) is required to set up domain mode.

-S (--domain-controller) is required to indicate the domain controller(s). Up to two controllers can be specified in the format "-S DC1,DC2" or "--domain-controller=DC1,DC2".

-N (--domain-username) -P (--domain-password) are required to define the account StorSafe will use to log into Active Directory on the domain server.

-B (--bind-point) is optional to direct StorSafe to a starting point or a single OU tree on the domain controller for retrieving organizational units. The default value is "/", which is the root of the OU tree.

-O (--organizational-units) is optional to specify one or more organizational units (OU) to which you will offer NAS shares, in the format "OU1,OU2,OU/OU3,...". Invalid OUs will be omitted. The process will be terminated if no OUs are valid.

- When OU(s) are not specified and the bind point is set to "/", first-level OUs on the domain controller will be used.
- When OU(s) are not specified but a bind point was specified, the bind point will be used.

-a (--administrative-login) -w (--login-password) An account with administrative privileges to allow this NAS server to become a trusted member within the domain realm. If these are empty, the default values will be domain-username and domain-password. These are optional and will not be saved.

User mode:

-U (--user-mode) indicates to set up user mode.

-W (--workgroup) is optional to specify the workgroup name.

Common arguments:

-l (--uid-ranges) -G (--gid-ranges) These are optional to specify id ranges preserved for SMB users or groups. The format of an id range is firstID-lastID, e.g., 12345-23456. It can be a list of id ranges separated by commas, "RANGE1,RANGE2,...". If no specified id range list when changing mode, the default range is "45435-65436". When assigning new id range list without changing mode, the given range value will be validated and insert to list if is valid.

-f (--force) will force to remove all SMB shared assignment(s) when changing security mode.

-c (--comment) is the description of SMB service. The maximum length is up to 128 characters.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

NAS information

Get deduplication statistics

```
iscon getdedupestatsinfo -s <server-name> [-u <username> -p <password>]
[-T <disp-type>] [-l]
[-X <rpc-timeout>]
```

```
iscon getdedupestatsinfo --server-name=<server-name>
[--server-username=<username>
--server-password=<password>]
[--disp-type=<disp-type>] [--longlist]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves repository usage statistics and the global deduplication ratio from the associated deduplication (SIR) cluster or server.

-T (--disp-type) is the option to specify the display type:

- “Usage” displays repository usage statistics.
- “Ratio” displays the global deduplication ratio only.

-l (--longlist) is the option to display the long format.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get realtime performance information

```
iscon getrealtimeperformanceinfo -s <server-name> [-u <username> -p <password>]
-T <disp-type>
[-b [-n <number>] [-d <delay-time>]]
[-S <source-server>] [-v <vdev>] [-l]
[-X <rpc-timeout>]
```

```
iscon getrealtimeperformanceinfo --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--disp-type=<disp-type>
[--batch-mode [--number=<number>] [--delay-time=<delay-time>]]
[--source-server=<source-server>] [--vdev=<vdev>] [--longlist]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves real-time performance statistics for the specified operation.

-T (--disp-type) is required in order to specify the type of performance for which to display statistics:

- “Dedup” - display deduplication performance
- “RepliOut” - display outgoing replication performance
- “RepliIn” - display incoming replication performance
- “FileSystem” - display file system performance

-b (--batch-mode) is the option to include performance of batch mode operations.

-n (--number) is the option to specify the number of iterations of batch mode operations.

-d (--delay-time) is the option to specify the delay in seconds between two iterations of batch mode operations.

-S (--source-server) is the option to specify the name of the source server for incoming replication.

-v (--vdevid) is the option to return performance statistics for a specific virtual device ID.

-l (--longlist) is the option to display the long format.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Get folder details

```
iscon getfolderdetailedinfo -s <server-name> [-u <username> -p <password>]  
-fp <folder-path>  
[-fn <file-name>] [-l]  
[-X <rpc-timeout>]
```

```
iscon getfolderdetailedinfo --server-name=<server-name>  
[--server-username=<username>  
--server-password=<password>]  
--folder-path=<folder-path>  
[--file-name=<file-name>] [--longlist]  
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command returns information on all files in the specified folder.

-fp (--folder-path) is required - specify the folder path.

-fn (--file-name) is the option to only display information only for a specified file on the folder. For example, to display information for a file whose full path is "/nas/NAS-00001/fds/myfolder/myfile.dat", specify "-fp /nas/NAS-00001/fds/myfolder -fn myfile.dat".

-l (--longlist) is the option to display the long format.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

NAS utilities

Start Fast Copy

```
iscon startfastcopy -s <server-name> [-u <username> -p <password>]
-fs <source-path> -ft <target-path> [-f]
[-X <rpc-timeout>]
```

```
iscon startfastcopy --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--source-path=<source-path> --target-path=<target-path> [--force]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command sets up a copy procedure between NAS resources within the same server. If source files have been deduplicated, the stub file will be copied so that the file does not have to be restored.

-fs (--source-path) is required to specify the source directory.

-ft (--target-path) is required to specify the target directory; this can be a new or existing directory.

-f (--force) is an option to overwrite the existing target directory.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

Example: #iscon startfastcopy -s xxx.x.x.x -fs/nas/NAS-xxx/fds/source -ft /nas/NAS-xxx/fds/target/8

The command returns the process identifier (pid) and status:

Fastcopy process started with pid 18313

Command: startfastcopy executed successfully.

Stop Fast Copy

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

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

Description:

This command stops an in-progress copy procedure between NAS resources. If you do not have the process identifier (pid), run *ps aux | grep "fdsccli fastcopy"* on the server to obtain it.

-fp (--pid) is required to specify the process to stop.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

Get Fast Copy info

```
iscon getfastcopystatinfo -s <server-name> [-u <username> -p <password>]  
-fp <pid>  
[-X <rpc-timeout>]  
  
iscon getfastcopystatinfo --server-name=<server-name>  
[--server-username=<username> --server-password=<password>]  
--pid=<pid>  
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command retrieves information about the specified Fast Copy process, including the directory being copied, its size, and percentage of completion. If you do not have the process identifier (pid), run *ps aux | grep "fdscli fastcopy"* on the server to obtain it.

-fp (--pid) is required to specify the process to retrieve.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout for this command is 1,800 seconds.

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 option is not used, 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 createfcacnfreport -s <server-name> [-u <username> -p <password>] [-o <filename>]
[-X <rpc-timeout>]
```

```
iscon createfcacnfreport --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 option is not used, 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 option is not used, 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 option is not used, 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 option is not used, 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 surrounded by single or double quotes and each ID must be separated with semicolons; 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 option is not used, 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 option is not used, the default file name is: DeduplicationTapeUsager-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Deduplication replication status report

```
iscon creatededupereplicationstatusreport -s <server-name> [-u <username> -p <password>]
[-L <server-source-list>] [-d <RNZ>] [-o <filename>] [-X <rpc-timeout>]
```

```
iscon creatededupereplicationstatusreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--source-list=<server-source-list>] [--details=<RNZ>] [--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the replication coverage for all the servers that have deduplication policies set up to replicate to this server. The report can be viewed, printed, emailed or exported to other formats from the console.

-L (--source-list) is an option to include only the specified source servers in the report. This argument can be a list of server names separated by commas or the file name enclosed in "< >" (e.g. "<file>") of a text file containing the list in the first line. The file must be located in the same folder as the command line utility or the full path is required. The server name is case sensitive.

-d (--details) 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 option is not used, 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 option is not used, 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 option is not used, 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 option is not used, 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 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 option is not used, 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 option is not used, the default file name is: DiskSpaceUsageHistory-MM-DD-YYYY-hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Import export job report

```
iscon createimportexportjobreport -s <server-name> [-u <username> -p <password>]
[-z <report period>] | [-D <date-range>] [-i <filter>] [-o <filename>] [-X <rpc-timeout>]
```

```
iscon createimportexportjobreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--include-filter=<filter>] [--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing all import/export and tape caching jobs that were placed in the queue during the specified period of time, regardless of job status. The report can be viewed, printed, emailed, or exported to other formats from the console.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The date option is applied to the local server time. The default value is: "-z t" (today).

-i (--include-filter) is an optional filter to include only the specified jobs. The following values are accepted. Multiple values must be separated with commas.

Job Type:

- ESD[C | M] - Export to standalone drive, Copy or Move
- EPL[C | M] - Export to physical library, Copy or Move
- ISD[C | R] - Import from standalone drive, Copy or Recycle
- IPL[C | R] - Import from physical library, Copy or Recycle
- CC - Create cache with copy meta data
- TS - Tape stacking jobs

The default is to include tapes that were exported with either Copy or Move or were imported with either Copy or Recycle.

Job Status:

- WTD - Waiting for tape/drive
- FAIL - Failed
- COMP - Completed
- CANC - Cancelled
- HOLD - On hold
- WIE - Waiting for IE slot
- RUN - Running

For example: -i EPL,COMP,CANC

This command will include only jobs with export to physical library, *copy* and *move*, *completed* and *cancelled*.

By default all jobs are included.

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the option is not used, 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 RCVC,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 option is not used, 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 option is not used, the default file name is: `LUNReport-MM-DD-YYYY-hh-mm-ss`

`-X (--rpc-timeout)` is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Physical tape usage report

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

```
iscon createphytapeusagereport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--output-file=<filename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing the information queried from the physical tape database. The report can be viewed, printed, emailed, or exported to other formats from the console.

`-o (--output-file)` is an option to specify an output file name for the report. If a report with the same name already exists, a suffix `[#]` will be appended to the report name. If the option is not used, the default file name is: `PhysicalTapeUsage-MM-DD-YYYY-hh-mm-ss`

`-X (--rpc-timeout)` is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Replication status report

```
iscon createreplicationstatusreport -s <server-name> [-u <username> -p <password>]
[-z <report period>] | [-D <date-range>] [-r <repl-resource-type> -R <resourceList>]
[-O <sorting>] [-o <outputFilename>] [-X <rpc-timeout>]
```

```
iscon createreplicationstatusreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--report-period=<report-period>] | [--date-range=<date-range>]
[--repl-resource-type=<repl-resource-type> --resource-list=<resourceList>]
[--options=<sorting>] [--output-file=<outputFilename>] [--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing replication job related information about the specified resources. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD.

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The default value is: "-z t" (today).

-r (--repl-resource-type) is an option to specify a generic resource type to be queried. It can be one of the following: TAPE or TAPEReplica. The default value is TAPE.

-R <--resource-list> 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 option is not used, 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 option is not used, 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 option is not used, 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 option is not used, 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)
- TC (tape caching view)
- RR (replica resources view, includes all replica tapes)
- VV (vault view, includes all tapes from the vault)
- DT (detailed tape view)
- 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 option is not used, 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.

StorSafe 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 StorSafe 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 StorSafe 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 option is not used, 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.

SIR server report

```
iscon getreportdata -s <server-name> [-u <username> -p <password>]
[-d <date>] [-I #[D|W|M]]
[-C <clientList>] [-R <resourceList>]
[-N]
[-o <outputFilename>] [-F <fileFormat>] [-H] [-f]
[-X <rpc-timeout>]
```

```
iscon getreportdata --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--date=<date>] [--units=#[<D|W|M>]]
[--client-list=<clientList>]
[--resource-list=<resourceList>]
[--no-system-info]
[--output-file=<outputFilename>]
[--file-format=<fileFormat>]
[--include-heading] [--force]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command can only be used on a SIR server and generates a report listing SIR server total I/O activity, average system memory, and CPU usage for each interval between data points of the selected unit.

The date format is YYYYMMDD and the default is today's date.

The time unit can be specified in three different ways:

- D = Day
- W = Week
- M = Month

-R (--resourceList) can be any of the following:

- a list of resource ids separated by comma
- "*" for all the resources
- a filename enclosed in <> containing the resource ID on each line.

For example:

- -R 1,3,5,10
- -R "<res_id_file.txt>"

-C (--clientList) can be any of the following:

- a list of client names separated by comma
- "*" for all clients
- a filename enclosed in <> containing client name in each line.

For example:

- -C client1,client2
- -C "<client_file.txt>"

System information, including memory and CPU usage information, is included by default.

-N (--no-system-info) is an option to exclude system memory and CPU usage information.

If an output filename is not specified, the default filename is "ServerThroughputYYYY-MM-DD[.#]"

[.#] is added to the filename automatically if a report with the requested filename already exists.

-F (--fileFormat) can be either "csv" (the default) or "txt".

-H (--include-heading) is the option to include a data header.

-f (--force) overwrites the existing file when the output file already exists.

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create NAS Resource usage report

```
iscon createnasresourceusagereport -s <server-name> [-u <username> -p <password>]
-n <nas-id>
[-z <report period> | -D <date-range>]
[-o <filename>]
[-X <rpc-timeout>]
```

```
iscon createnasresourceusagereport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--nas-id=<nas-id>] [--report-period=<report-period> | --date-range=<date-range>]
[--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing NAS device information. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

The default report is for the current day.

-n (--nas-id) is an option to specify a single device.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD

Either **-z (--report-period)** or **-D (--date-range)** can be specified, but not both. The default value is: "-z t" (today).

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [.#] will be appended to the report name. If the option is not used, the default file name is "IndividualNasResourceUsage-MM-DD-YYYY--hh-mm-ss".

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create NAS SMB share usage report

```
iscon createnascifsshareusagereport -s <server-name> [-u <username> -p <password>]
-S <share-name>
[-z <report period> | -D <date-range>]
[-o <filename>]
[-X <rpc-timeout>]
```

```
iscon createnascifsshareusagereport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
--share-name=<share-name>
[--report-period=<report-period> | --date-range=<date-range>]
[--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing information about the specified NAS SMB share. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-S (--share-name) is required to specify the name of the share.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the date range for the report. The format is: YYYYMMDD-YYYYMMDD or YYYYMMDD

Either **-z (--report-period)** or **-D (--date-range)** can be specified, but not both. The default value is: "-z t" (today).

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [#] will be appended to the report name. If the option is not used, the default file name is: NASShareUsage-MM-DD-YYYY--hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Create NAS statistics summary report

```
iscon createnasstatisticssummaryreport -s <server-name> [-u <username> -p <password>]
[-d] [-v] [-c]
[-z <report period> | -D <date-range>]
[-o <filename>]
[-X <rpc-timeout>]
```

```
iscon createnasstatisticssummaryreport --server-name=<server-name>
[--server-username=<username> --server-password=<password>]
[--nas-isdedup] [--nas-isrepserver] [--nas-isrepclient]
[--report-period=<report-period> | --date-range=<date-range>]
[--output-file=<filename>]
[--rpc-timeout=<rpc-timeout>]
```

Description:

This command creates a report on the server side listing NAS statistics. Data shown in the report is limited to the maximum number of days that database information is retained, which is set in server properties. The report can be viewed, printed, emailed, or exported to other formats from the console.

-d (--nas-isdedup) is an option to include deduplication information in the report.

-v (--nas-isrepserver) is an option to include replication server information in the report.

-c (--nas-isrepclient) is an option to include replication client information in the report.

-z (--report-period) is an option to specify the period of time that the report should cover. The accepted values are:

- t - today
- y - yesterday
- 7 - last seven days
- 30 - last thirty days
- 365 - last 365 days

-D (--date-range) is an option to specify the range of dates the report should cover. There is no limit to the date range, but future dates are ignored. Valid formats are:

- date range: YYYYMMDD-YYYYMMDD
- specific day: YYYYMMDD

Either -z (--report-period) or -D (--date-range) can be specified, but not both. The default value is: "-z t" (today).

-o (--output-file) is an option to specify an output file name for the report. If a report with the same name already exists, a suffix [#] will be appended to the report name. If the option is not used, the default file name is: NASStatisticsSummary-MM-DD-YYYY--hh-mm-ss

-X (--rpc-timeout) is an option to specify a number between 1 and 30000 in seconds for the RPC timeout. The default RPC timeout is 1,800 seconds.

Best Practices

This section provides information about StorSight deployment and deduplication sizing guidelines.

StorSight deployment

StorSight gathers and consolidates information coming from different storage servers into a scalable repository of FalconStor services, users, and historical data. StorSight includes a web service that provides management and monitoring. This section provides procedures for deploying StorSight servers.

Storage considerations

StorSight requires four physical disks, 80 GB minimum each:

- Two are configured as RAID 1 (mirroring) and contain the SQL database in the /DB partition
- The other two are also configured as RAID 1 and contain the OS, StorSight software, and system logs.

Since the database disks are separate from the OS disks, any OS recovery, reimaging of the server, and upgrades will not touch the database partition.

Security considerations (updated October 2024)

- For security reasons, StorSight can reside behind a firewall. In this case, ports 80 and 443 need to be open.
- Since StorSight has public access, the remote shell session `ssh` is disabled by default for the root account. A predefined user account called `fmsadmin` with default password `IPStor101` can remotely connect via `ssh` and then become root with the `su` command.
- The encrypted https protocol can be used instead of the clear text protocol, http. Because the encryption algorithm is based on the browser, customers need to use their own certificate if they want to use https instead of http. For more information, refer to [“HTTPS connection for StorSight with an SSL certificate”](#) in the Appendix.

Protect and recover the StorSight server *(updated January 2025)*

StorSight is a management portal to configure StorSafe servers and to keep track of capacity usage and monitor server status. The StorSafe configuration is saved on StorSafe itself and not on StorSight. The data traffic is between StorSafe and its host clients; StorSight is not on the data path. So, in case StorSight fails, the StorSafe configuration and data traffic do not get impacted.

StorSight management data is kept in two databases that need to be protected, Elasticsearch (statistics for reports) and PostgreSQL (system configuration and settings).

The database disk is normally on a separate partition, /DB. If the StorSight machine break downs while the DB disk is still OK, StorSight can be installed on another machine using the original DB disk. If the disk is not healthy, databases can be recovered on the new StorSight machine from regular backup images that were created at a safe location.

ELASTICSEARCH

Elasticsearch backups are taken using their "snapshot" API. This will take the current state and data in the cluster and save it to a shared repository. This backup process is "smart." The first snapshot will be a complete copy of data; all subsequent snapshots will save the delta between the existing snapshots and the new data. Data is incrementally added, and deleted data is backed up over time. Subsequent backups will be substantially faster since they are transmitting far less data.

There are several repository types that can be used for backups: Shared filesystems (such as a NAS), Amazon S3, HDFS (Hadoop Distributed File System), and Azure Cloud. In our steps below, a local filesystem is used.

Backup

1. Create the path for a safe backup location. In our example, the shared filesystem is mounted to /opt/esbackup. Be sure to replace the path with the correct values for your system with the proper permissions.

```
# mkdir /opt/esbackup
# chmod 777 /opt/esbackup
# chown -R elasticsearch:elasticsearch /opt/esbackup
```

2. Add the following setting to the "/etc/elasticsearch/elasticsearch.yml" file:

```
path.repo: ["/opt/"]
```

3. Restart StorSight for the settings to take effect:

```
# freestor restart
```


4. Create a filesystem repository along with the required metadata for the backup:

```
# curl -XPUT 'http://localhost:9200/_snapshot/my_backup' -H
'Content-Type: application/json' -d
'{"type":"fs", "settings":{"location":"/opt/esbackup",
"compress":true}}'
```

5. Generate a snapshot with a unique name:

```
# curl -XPUT 'http://localhost:9200/_snapshot/my_backup/
snapshot_1?wait_for_completion=true' -H 'Content-Type:
application/json' -d '{"indices": "*", "ignore_unavailable":
true, "include_global_state": false}'
```

The above command generates the snapshot and writes it into the specified directory (/opt/esbackup).

The `wait_for_completion` parameter specifies whether or not the request should return immediately after snapshot initialization (the default) or wait for snapshot completion. This need not be set to true when using backup scripts through cron jobs.

The `snapshot_1` parameter identifies the name of the snapshot. This name must be unique or ES will return an error and the backup will fail. The current time added to the parameter is recommended as a good option for the snapshot name.

6. Verify the snapshot.

```
# curl -XGET 'http://localhost:9200/_snapshot/my_backup/
_all?pretty'
```

7. The simple script below can be used with a cron job at specified intervals to take automatic periodic backups:

```
# tim='date +%Y-%m-%d-%H-%M'
# curl -XPUT "http://localhost:9200/_snapshot/my_backup/
snapshot_${tim}?wait_for_completion=false" -H 'Content-Type:
application/json' -d '{"indices": "*", "ignore_unavailable":
true, "include_global_state": false}'
```

Restore

1. In case you need to recover the database backup on a new StorSight machine, make the database backup location available to this machine. For example, mount the shared filesystem /opt/esbackup on the new machine.
2. Display available snapshots to restore:

```
# curl 'http://localhost:9200/_snapshot/my_backup/
_all?pretty=true'
```

The output will be similar to one shown below. Note the snapshot name that you want to restore from.

```
{
  "snapshots" : [ {
    "snapshot" : "snapshot_1",
    "version_id" : 1070399,
    "version" : "1.7.3",
    "indices" : [ "stats_26099", "replication_sequence", "stats_25457",
"stats_193", "stats_28794", "stats_28642", "stats_19410", "replication_index",
"stats_10713", "stats_400", "stats_25688", "stats_28487", "stats_10855",
"stats_26362", "stats_14096", "stats_14357" ],
    "state" : "SUCCESS",
    "start_time" : "2016-12-23T19:50:34.360Z",
    "start_time_in_millis" : 1482522634360,
    "end_time" : "2016-12-23T19:52:09.077Z",
    "end_time_in_millis" : 1482522729077,
    "duration_in_millis" : 94717,
    "failures" : [ ],
    "shards" : {
      "total" : 16,
      "failed" : 0,
      "successful" : 16
    }
  }, {
    "snapshot" : "snapshot_2",
    "version_id" : 1070399,
    "version" : "1.7.3",
    "indices" : [ "stats_26099", "replication_sequence", "stats_25457",
"stats_193", "stats_28794", "stats_28642", "stats_19410", "replication_index",
"stats_10713", "stats_400", "stats_25688", "stats_28487", "stats_10855",
"stats_26362", "stats_14096", "stats_14357" ],
    "state" : "SUCCESS",
    "start_time" : "2016-12-28T19:15:45.767Z",
    "start_time_in_millis" : 1482952545767,
    "end_time" : "2016-12-28T19:16:04.416Z",
    "end_time_in_millis" : 1482952564416,
    "duration_in_millis" : 18649,
    "failures" : [ ],
    "shards" : {
      "total" : 16,
      "failed" : 0,
      "successful" : 16
    }
  }
]
}
```

3. Close all the indexes to block read/write operations during the restore:

```
# curl -XPOST 'http://localhost:9200/*/_close'
```

4. Restore from the desired snapshot, for example, snapshot_1:

```
# curl -XPOST
'http://localhost:9200/_snapshot/my_backup/snapshot_1/
_restore?wait_for_completion'
```

Reference: <https://www.elastic.co/guide/en/elasticsearch/guide/current/backing-up-your-cluster.html>

POSTGRESQL

Backup

1. Create the path for a safe backup location. In our example, a shared filesystem is mounted to `/opt/pgsqlbackup`.
2. Regularly take a PostgreSQL backup, using the current date for the dump file:

```
# sudo -u postgres /usr/pgsql-$PGVERSION/bin/pg_dump freestor > /opt/pgsqlbackup/YYYY-MM-DD.dump
```

Recover

1. In case you need to recover the database backup on a new StorSight machine, make the database backup location available to this machine. For example, mount the shared filesystem `/opt/pgsqlbackup` on the new machine.

2. Stop the StorSight service:

```
# freestor stop
```

3. Recover the backup dump:

```
# sudo -u postgres /usr/pgsql-$PGVERSION/bin/psql -d freestor < /opt/pgsqlbackup/YYYY-MM-DD.dump
```

4. Start the StorSight service:

```
# freestor start
```

Deduplication repository sizing

A simple formula to estimate how large your deduplication repository needs to be is:

$\text{Repository size} = \text{Amount of data} / \text{deduplication ratio}$

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:

$\text{Repository size} = (\text{Number of VDIs} \times 40 \text{ GB}) / \text{deduplication ratio}$

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 SIR node is 4.3% of deduplication repository capacity.

CPU cores

For StorSafe, 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

For StorSight, you need at least 4 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 SIR Node	Dedupe Repository Size (TB) – 2 SIR Nodes	Dedupe Repository Size (TB) – 4 SIR 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 SIR 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

StorSafe 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 StorSafe 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 total available memory and the load on the system. For example, having hundreds of thousands of virtual tapes, tens of NAS file systems, etc. may result in the need for additional operating system memory. Refer to the following for the system memory requirements.

StorSafe server memory

Each StorSafe server needs the sum of the following values:

- 16 GB (predefined minimum value)
- 150 MB for each inline deduplication stream, which is equivalent to the number of tape drives; maximum 1024 streams or 153.6 GB
- 3 KB for each virtual tape
- 2 GB for each NAS file system
- 2 GB for migration to object storage
- Additional memory for Hosted Backup, as recommended by the backup application vendor

StorSafe-S server memory

Each StorSafe-S server needs:

- 16 GB (predefined minimum value)
- 2 GB memory per 1 TB of SIR capacity for deduplication (based on average hash size 32 KB)
- Additional memory for Hosted Backup, as recommended by the backup application vendor

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
- 2 GB for each NAS file system
- 2 GB for migration to object storage

StorSafe-S with StorSight server memory

Each StorSafe-S with StorSight server needs:

- 32 GB (predefined minimum value)
- 2 GB memory per 1 TB of SIR capacity for deduplication (based on average hash size 32 KB)
- Additional memory for Hosted Backup, as recommended by the backup application vendor

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
- 2 GB for each NAS file system
- 2 GB for migration to object storage

SIR server memory

Each SIR server needs:

- 16 GB (predefined minimum value)
- 2 GB memory per 1 TB of SIR capacity for deduplication (based on average hash size 32 KB)

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 post-deduplication stream; maximum 16 streams or 2.4 GB

Troubleshooting

This section contains general troubleshooting information and a list of error codes generated by StorSafe and SIR servers.

Product registration

If you are unable to complete offline activation successfully, try the following solutions:

1. In order to prevent the possibility of unsuccessful email delivery to the FalconStor activation server, disable Delivery Status Notification (DSN) before you send the activation request email to `Activate.Keycode@falconstor.com`.
2. If you do not receive a reply to your offline activation email from the FalconStor activation server within one hour after sending it, check your email encoding and change it to UNICODE (UTF-8) if set otherwise, then send the email again.
3. If the reply email indicates that the license is successfully registered but the signature file is not attached, you may have set the name of the license information file improperly; you cannot use a single digit before the suffix in the file name. Change the registration file name to a valid alphanumeric string and then try to register again. If the issue persists, contact Technical Support.

Physical resources

The portal does not show physical storage devices correctly

There are several steps to try when physical storage devices have been connected/assigned to the StorSafe server yet they are not showing in the portal.

Rescan devices	Perform a rescan from the portal. 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 (/var/log/messages) for error messages that may correspond to the rescan operation and report failures on SCSI devices. It may be that even though the devices were discovered, they were not accessible due to errors.
Check device type	For external SCSI devices, check the following: • Make sure the system is powered on. Perform a power cycle to make sure. • Physically make sure all the cable connectors are securely plugged in. • Verify SCSI termination. This can be quite involved. If you are not sure, you may have to contact the manufacturer of the devices and have their representatives assist with the troubleshooting.

Once the above conditions are verified, determine the SCSI HBA and the proper driver for it. This can normally be accomplished by going to the website of the HBA manufacturer. From the server portal, 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 StorSafe initiator driver and use persistent binding. Otherwise, StorSafe cannot manage the storage.

An HBA port is missing after rebooting and restarting StorSafe

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 portal

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 portal determine which physical resources comprise this virtual drive. 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 StorSafe 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 portal

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 that it is offline, check the underlying physical device status, device connectivity, and switches.

In a StorSafe-SIR solution, the StorSafe servers associated with SIR servers access physical resources from SIR for successful deduplication. The SIR data disks are presented to the StorSafe servers and the SIR tape drives on the StorSafe servers are assigned to the SIR servers.

During a power outage/recycle of servers, the StorSafe and SIR servers lose access to each other's physical devices. To ensure that the tape drives are online on the SIR servers and the SIR data disks are online on the StorSafe servers, rescanning physical resources on the StorSafe and SIR servers is advised after power outage/recycling the environment.

When you rescan, make sure that the *Discover New Devices* option is selected.

Tape expansion does not work

Highlight the tape in the portal and check that the *Total Size* field shows the correct size of the expanded tape device.

If the portal 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 portal 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 StorSafe server, run the command <code>top</code> or <code>ps -x</code> • 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 StorSafe

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 <i>Medium Changers</i> and tape drives under <i>Tape drives</i>. Usually the tape drive is indicated as <i>\tape<index></i>. • Linux - The tape library is usually indicated by <code>/dev/sg<index></code> (the <i>sg</i> module should be loaded) and the tape drive by <code>/dev/st/<index></code>, <code>/dev/nst/<index></code>, and <code>/dev/sg/<index></code> (The <i>st</i> module should be loaded). • Solaris - The tape library is usually indicated by <code>/dev/sg<index></code> (the <i>sg</i> module should be loaded) and the tape drive by <code>/dev/rmt/<index></code> (the <i>st</i> module should be loaded). • AIX - The tape device is usually indicated by <code>/dev/rmt<index></code> (for LTO1/LTO2) or <code>/dev/mt<index></code> (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 portal, select the virtual device. Check the device status. If the device status is <i>offline</i>, that is the problem as clients cannot see an offline device. Refer to the 'Virtual tapes are displayed as "offline" in the portal' section for more information. If the device status is <i>online</i>, check the client configuration. • Check client assignment - If you do not see virtual devices, assign them to that client. To share a device between several clients the mode should be <i>Read/Write non-exclusive</i>, 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 StorSafe target driver. Otherwise some clients cannot detect more than eight LUNs on StorSafe 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. StorSafe libraries support serialization. Serialization is the conversion of the content of an object into a sequential stream. It identifies the owner of each component, such as robot, slots, and tape drives. If the device appears in the backup software, but it is not attached to the expected component, it may be related to the serialization. Refer to your backup software manual for any patch or upgrade related to serialization on the backup software.

Client sees the tape library/drive but cannot access it

Check device
access by OS

Check if the client’s operating system can access the device or if it is the backup software that cannot access the tape library or drive.

Depending on the OS you can use a raw device utility. Most of these tools work with tape drives; they are not capable of moving tapes into the drives. Even if some can move tapes, you need to know the exact address of the tape and the drive.

We recommend that you use the portal to put a tape in a drive before running these tools. Also, stop the backup software before you use these utilities:

- **Windows** - For IBM Ultrium devices you can use `ntutil`, a command line tool that can check the tape device.
- **Unix systems** - You can use the `mt` or `tar` commands to access the tape device, for example: `mt -f /dev/rmt/0 status`

OS cannot
access device

If the operating system *cannot access* the device, you need to troubleshoot virtual device access.

- Go to the storage to verify that it is not in error or in an abnormal state. The assigned devices have to be in read/write mode.
- Check the Event Log or syslog (`/var/log/messages`) for message indicating IO errors. Such messages usually begin with `log_scsi_error`.

- Check client driver - Go to the client machine and check the adapter driver version. It should be certified for use with StorSafe.

OS can access device If the operating system *can access* the device, you need to troubleshoot the backup software. Verify that you have the correct drivers.

Client can no longer access a virtual device (tape library, drive, or SIR resource)

This can have different causes:

- Client machines may lose device access if you switch between a Multi-ID HBA and a single-ID HBA. If this occurs, you should reboot the client machine.
- If either the StorSafe or the SIR server is shut down for a long period, the devices offered to the clients will time out or be set to *offline*. If this occurs, you will need to perform a rescan from the host machine to regain access.
- ACSLS library users - If you did not select the *Firewall Support* option during configuration, the *portmap* process needs to be running. Otherwise, the system will fail to assign or retrieve the library's status after restarting StorSafe services or rebooting. To enable *portmap*, you will have to run the following command: `chkconfig --add portmap`

VIT tape is marked "Full"

If you see a VIT marked as "*full*", check the log to see if there was enough disk space available during the backup but before the deduplication process started.

If there was not enough space, the tape is marked as "*full*" and this status is preserved after deduplication. If this occurs, you must use a different tape for backups.

Export to direct link tape fails

Writing data to a direct link tape will fail if the media types are not the same between your physical and virtual tapes. When you create a cache for your physical tapes, you must make sure that your physical tapes use the same media type as your virtual tapes.

Fibre Channel connectivity issues

This section provides more detail about FC connectivity issues that cause a client to be unable to see virtual tape libraries/drives. This assumes the devices are properly created and assigned to clients.

- **Check ports** to verify that the QLogic BIOS can see the target port of the 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. To do this, select *Manage* from the menu bar, click the *Physical Resources* tab, click the *Adapters* tab, locate the HBA that matches the recorded target HBA WWPN, and click *SNS Information*. 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.
- **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 portal 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 StorSafe. 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 portal. 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 portal, 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 StorSafe device is still not detected, reboot the system.

NDMP

NDMP backup jobs are failing

If you are using NDMP, you must define the hostname of the StorSafe server in the `/etc/hosts` file in the format "IPAddress Hostname".

For example: 10.7.2.41 Server41

The StorSafe server NDMP daemon fails to start

Because some backup applications use NDMP, if you are running backup software on the StorSafe server, it should be started after StorSafe has started and should be stopped before StorSafe is stopped. Otherwise, the NDMP service that is loaded by the backup software may interfere with StorSafe's NDMP service.

Replication

There are several aspects to replication: replication configuration, replication process, replica resource promotion, replication configuration removal. If a problem occurs, check the Event Log or syslog (/var/log/messages). Look for error messages relating to replication. Some common problems are described below.

Replication configuration

- Virtual tapes Replication configuration fails with a “Failed to add replication target” error. This can occur if the replica server has a device assigned to it from the primary server. You will need to remove the device assignment before you can create your replication configuration.
- Deduplicated tapes Replication configuration fails between two StorSafe-S servers with a message saying that the target server’s WWPN is already assigned to client x. This can occur if the target server is assigned as a client of a virtual library on the primary server. You will need to remove the client assignment before you can create your replication configuration.

Replication process

- Replication fails to start If you enable auto archive/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 auto archive/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 StorSafe89-115 kernel: IOCORE1 [iocore|29557]
OnSANREPRequest, SANREP_STOP_REPLICATION
Jan 13 15:52:54 StorSafe89-115 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 StorSafe89-115 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 and Auto Archive jobs have a higher priority than Post-processing and Turbo deduplication jobs.

If a migration to object storage job or Auto Archive job is running when a deduplication job is triggered, the deduplication job will fail but will attempt to run at the next retry interval.

Inline deduplication jobs fail

With inline deduplication, the backup process is designed to capture a certain amount of unique data in memory before writing it to storage.

When an application requests to write a file mark, the operation must flush all cached data to storage. The flushing process requires time for the write operations to complete.

If the backup process takes too long to flush the cached data to storage, and the default SCSI command timeout is too low, the driver may send an abort request, causing the job to fail.

To avoid this issue, you can adjust the SCSI command timeout value to allow sufficient time for the backup process to complete. This can help ensure that the backup is successful.

For example, to change the SCSI command timeout for an IBM lin_tape, do the following:

1. List all /dev/IBMtape*
2. Set the SCSI timeout. As an example, for IBMtape3, set the SCSI timeout as shown below:

```
echo "30" > /sys/class/lin_tape/IBMtape3/device/timeout
```

Import/Export

Import/Export does not work as expected

Check tape capacity mismatch When you Import/export data between a physical tape device and a virtual device, you must make sure the tape devices are of the same type and the same capacity. If they do not have the same capacity, an end-of-media-hit condition occurs and import/export fails.

If data compression is used, make sure the *actual* capacity matches, not just the *compressed* capacity. Import/export will fail when the destination does not have enough space to hold uncompressed data coming from source.

Check job status Go to the *Import/Export Jobs* page and search for a job related to this operation. If the job is in progress, wait until it is completed. If the job is not there and the import/export operation is not done, look at Event Log to see if there are any job failure messages.

Check barcodes of virtual tapes When you import data from a physical tape, make sure the virtual tapes have different barcodes. Otherwise, the import operation fails. Use the *Inventory* feature in the portal to get the updated bar codes and status from the physical library.

Check physical tape library and device status Make sure the physical tape library does not show any abnormal situation. For example, the tape drives may require cleaning or tapes may need to be moved to the proper location.

Check element address on the physical library When you import data, make sure the assignment of drive in StorSafe follows the element address of the drives in the physical library. Assign the tape drive in the order of their element address.

For example, an import job cannot be executed and the physical library has two DLT 8000 drives with the following configuration:

```
Library SCSI ID: 1
Drive at element address 1200: SCSI ID 10
Drive at element address 1201: SCSI ID 09
```

```
StorSafe Resources:
ABC-00003
DLT8000-0008: SCSI ID 09
DLT8000-0009: SCSI ID 10
```

In this case you need to unassign tape drives, select first DLT8000-0009, assign it, then select DLT8000-0008, and assign it to the physical library ABC-00003.

Check system log for errors Check the Event Log or syslog (/var/log/messages). Look for error messages relating to the physical tape library or drive. If you find error messages but cannot find the cause, contact technical support.

NAS resources/shares

Client cannot access NAS resources or shares

SMB or NFS client If clients are unable to connect to any NAS shares, try to ping the client machine from StorSafe and vice versa.

Also, run `ifconfig` or `ipconfig` and see if there are any dropped packets or transmission errors for the NICs.

SMB client
(Domain/Active Directory mode)

There are several things to try when an SMB client can ping StorSafe but cannot connect to and access NAS resources.

- For Windows clients, try to access the StorSafe server using the IP address of StorSafe instead of the hostname (i.e., `\\13.0.0.44` instead of `\\myserver`).
- If you can connect using the IP address, try to ping the server hostname. If you do not get any response when pinging the server name, the customer's DNS may have a problem.
- Ping the domain controller to see if you get a response. If no response, StorSafe is not communicating with the domain controller. If you try to ping using the short name, the ping should return the long name.
- Check whether the StorSafe server can get user and group information from the domain controller.
 - Run: `getent passwd <user-name>` or `getent group <group-name>` where `<user-name>` is a specific user and `<group-name>` is the name of an existing group on the domain. Look at the results that are returned.
 - Run: `wbinfo --config=$ISHOME/etc/$HOSTNAME/smb.conf -u` and `wbinfo --config=$ISHOME/etc/$HOSTNAME/smb.conf -g` and check results related to users and groups.

SMB client Try the following when an SMB client cannot access a NAS share:

- Check if there is any time difference between the domain controller and the StorSafe server. The difference should not be greater than three minutes.
- We recommend that you set up NTP and have the domain controller and the StorSafe server sync to the same NTP server. If NTP is already set up, use the `ntpd` and `ntpdate` commands to update the time.


```
service ntpd stop;
ntpdate ntp-myNTPserver;
service ntpd start
```
- The domain controller can be used as an NTP server in case the other is not available.
- StorSafe will get tickets from the domain controller. Run the command `klist`, and see if there is any ticket information returned. If not, it means that there is a problem with communication with the domain controller. You can restart the NAS MGTD module to refresh the connection and get new tickets.

If a Windows client is connected to StorSafe through a firewall and cannot connect to its share, make sure the following ports are not blocked: UDP 137, UDP 138, and TCP 139.

NFS client **Issue:** A client must manually reconnect to a NAS share after a server reboot.

Solution: `/etc/fstab` must be modified in order to automatically mount the NFS share at every reboot of the client system in the following way:

```
<hostname or IP address>:/nas/NAS-00006/fds/test /mnt/nfs nfs auto,
user
```

Issue: An NFS client cannot mount NAS resources.

Solution: Verify the NFS client's IP address and netmask and ping the server
`showmount -e <StorSafe IP address>`.

Issue: A Linux or Unix client can ping StorSafe but cannot map NAS shares.

Solution:

- Check the portal and see if the share is mounted. At the command line, use the command `# mount` to check and see if the share is mounted.
- Check the mount path on the server (located in `/nas/<resource>/fds/<share>`).
- Un-assign and re-assign the client.

Issue: A Linux or Unix client can ping StorSafe but cannot connect to and access NAS shares.

Solution:

- Check the mount path (the location where the NAS share is mounted (in the command to mount the file system on an NFS client):
`mount StorSafe-hostname:/nas/nasresourcename/fds/foldername /mnt/share1 /mnt/share1` would be the location for the share.
- Run `Showmount -e <StorSafe server name>` and see if the shares list returns.

Issue: A Linux or Unix client is connected to StorSafe through a firewall and cannot connect to its share.

Solution: Use the `rpcinfo -p` command to verify that the following ports are open:

- 111 - TCP/UDP
- 2049 - TCP/UDP

Issue: An AIX client cannot access an NFS share.

Solution: The client operating system requires connecting to a port less than 1024.

StorSafe includes an insecure CLI option if the client's OS does not use a reserved port for NFS (an Internet port that is less than `IPPORT_RESERVED` -- 1024). AIX is an example of an OS that needs to be flagged as insecure.

NFS client requires remapping user IDs to nobody

Squash can be used to map user IDs to *nobody*. This is set in NFS sharing.

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

Domain controller authentication issue

1. Verify that the StorSafe and Active Directory server clock times are the same.
2. Make sure the Active Directory server is running normally.
3. Verify that the StorSafe NTP setting points to the Active Directory server.

StorSafe server cannot get some users/groups from the domain controller

Below is the command to clean the *winbind* cache. This cache must be purged when the StorSafe server cannot get some users or groups from the domain controller.

Method 1: This method is safe to run on a production server and will not affect client connections. This method does not require a restart of modules or downtime:

```
echo erase | tdbtool $ISHOME/var/tmp/$HOSTNAME/winbindd_idmap.tdb > /dev/null
```

Method 2: This method will require downtime and will need to restart modules.

1. Stop SMB modules by running `vtl stop NAS SMBD`.
2. Go to the directory: `cd $ISHOME/var/tmp/$HOSTNAME/`
3. Remove everything in that specific directory only.
4. Restart SMB modules `vtl start NAS SMBD`.

Problem migrating data with ACLs and extended attributes

A copied file is not accessible or there is a difference in the rights between a copied file and the original.

A simple copy does not copy extended file attributes. Instead, the copied file inherits the file access rights of the owner of the session who did the copy. An additional operation is requested to also copy ACLs and extended attributes as well as data if the targeted file system is managing such attributes and ACLs.

1. Create a new file system with a size no less than the original one. In the following example, the old file system is mounted at /nas/NAS-00001 and the new file system is mounted at /nas/NAS-00002.
2. Log in as root on the StorSafe server, change the current path to the mount point of the old file system, and copy all files from the old file system mount point to the new file system mount point.

```
cd /nas/NAS-00001
cp -rf * /nas/NAS-00002
```

3. Copy the ACLs.

```
getfacl -R . | (cd /nas/NAS-00002; setfacl --restore=- )
```

4. Copy the extended attributes.

```
getfattr -R . | (cd /nas/NAS-00002; setfattr --restore=- )
```

5. Recreate Windows and NFS share information. You can recreate Windows and NFS shares on the new NAS resource from the portal if there are not many of them.

Otherwise, you may need to manually edit: \$ISHOME/etc/\$HOSTNAME/nas.conf/nas.conf and \$ISHOME/etc/\$HOSTNAME/nas.conf/ipstorsmb.conf when StorSafe is not running. Replace the old file system mount point with the new one, when applicable. (It is important to do this under the supervision of Technical Support.)

If a local file copy is not required, the following steps can be used instead of those above. For example, if it is certain that the original file system does not have any corruption; it may be preferable to copy the files instead of copying the raw data as mirroring does.

1. Mirror the file system with the new storage.
2. After synchronization, swap the mirror.
3. Promote the mirror to become a copy.

Truncated folder names

Folder names are truncated to 239 characters in the portal and at the Linux prompt.

This is a limitation of the ext4 file system. Names cannot be longer than 239 characters and must be considered when creating file trees in StorSafe and when defining a naming convention in order to manage backup file organization.

Permission denied to folders

Folder names are using non-US accentuated characters and are not accessible from the host (i.e.: mes-données-sauvegardées).

Use US Code Page 437 characters when creating shares. Non-US Code Page 437 characters like accentuated characters or German ß are translated and authentication may fail as shown below:



The following SMB.conf snip indicates that such accentuated characters are translated according to code page 437, translating the folder name. Do not modify the SMB.conf or even the folder name to manually attempt to fix it.

```
[Test&#228;0&#229;0&#230;0&#231;0&#232;0&#233;0&#234;0]
    path = /nas/FDSDisk-00006/fds/Test&#228;0&#229;0&#230;0&#231;0&#232;0&#233;0&#234;0
    max connections = 0
    read only = yes
    browseable = yes
    valid users = root, "FSTRaining\Administrator",
    write list = "FSTRaining\Administrator",
    available = YES
```

File access issues

The client can access shares, but cannot create, copy, or open files or cannot access a folder.

- Check file permissions on the client side.
For Windows clients, right-click the file and check the permissions and see if they are read-only.
For Linux clients, run `ll` and check the permissions of the file.
- In the portal, make sure the user or group that is assigned to the share has read/write privileges.

System event messages

Information from the `/var/log/messages` file on the StorSafe server can be viewed via the Event Log in the portal. 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.

Error codes

This section contains error messages generated by StorSafe/StorSafe-S and SIR servers.

Number	Type	Text	Probable Cause	Suggested Action
1016	C	Primary device ID %1 failed. The server is switching to the mirror device.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1017	E	The mirror of device ID %1 failed.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1022	E	Replication for virtual tape ID %1 failed; %2.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.
1023	E	Connection to physical device %1 failed; the path was switched to %2.	An adapter or cable might have a problem.	Check for a loose or damaged cable on the affected drive.
1030	E	Replication failed to start because a replication job was already in progress for virtual tape ID %1.	Only one replication job is allowed at a time for a tape.	Try again later. If replication was triggered by a schedule, adjust the schedule in the policy to avoid duplicate sessions.
1031	E	Replication failed to start because the replication control map was missing for virtual tape ID %1.	The configuration may not be valid.	Rescan devices to refresh the configuration.
1032	E	Replication failed to start because access to the replication control map for virtual tape ID %1 failed.	The virtual device may be offline or may have missing segments.	Check the underlying physical device and rescan devices.
1034	W	Replication failed for virtual tape ID %1 due to network transport error %2.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.

Number	Type	Text	Probable Cause	Suggested Action
1035	E	Replication failed for virtual tape ID %1 because the primary disk failed with error %2.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1038	E	Replication failed for virtual tape ID %1 because the local server could not allocate memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
1039	E	Replication could not proceed for virtual tape ID %1 because the replica device failed with error %2.	The replica reports the error specified in the message.	Check the device on the replica server and take necessary actions based on the error.
1046	E	Replica rescan failed for virtual tape ID %1 because the device had error %2.	The primary device reports the error specified in the message.	Check the device and take necessary actions based on the error.
1047	E	Replica rescan failed for virtual tape ID %1 because the replica device had error %2.	The replica device reports the error specified in the message.	Check the device and take necessary actions based on the error.
1048	E	Replica rescan failed for virtual tape ID %1 due to network transport error %2.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.
1049	E	Replica rescan could not proceed because the replication control map was missing for virtual tape ID %1.	The configuration may not be valid.	Rescan devices to refresh the configuration.
1050	E	Replica rescan could not proceed because access to the replication control map for virtual tape ID %1 failed.	The virtual device may be offline or may have missing segments.	Check the underlying physical device and rescan devices.
1052	E	Replica rescan failed for virtual tape ID %1; the replica status is %2.	The replication configuration may not be valid.	Check the configuration on the replica server. Check system logs on both servers for additional information.

Number	Type	Text	Probable Cause	Suggested Action
1053	E	Replica rescan could not proceed because a replication job was already in progress for virtual tape ID %1.	Only one replication job is allowed at a time for a device.	Try again later. If replication was triggered by a schedule, adjust the schedule in the policy to avoid duplicate sessions.
1055	E	Replication failed for virtual tape ID %1; the replica status is %2.	The replication configuration may not be valid.	Check the configuration on the replica server. Check system logs on both servers for additional information.
1056	E	Exchange of the replication control map between servers failed for virtual tape ID %1; error: %2.	This may be due to a connectivity issue.	Check connectivity between the primary and replica. Check system logs on both servers for additional information.
1059	E	Replication failed for virtual tape ID %1; error: %2.	Replication reports the error specified in the message.	Check the replica device and system logs; take necessary actions based on the error.
1060	E	Replica rescan failed for virtual tape ID %1; error: %2.	Replication reports the error specified in the message.	Check the replica device and system logs; take necessary actions based on the error.
1061	W	Storage path with ACSL %1 failed; alternate path %2 will be used.	This may be due to a connectivity issue.	Check the path connectivity between the server and the physical device.
1067	E	Replication could not proceed because connection to replica server %1 failed.	Either the network connection is down or the replica server is down.	Check the state of the replica server. Determine and correct either the network problem or server problem.
1069	E	Replication could not proceed because virtual tape ID %1 no longer has a replica.	The replica device may have been deleted or promoted while the primary server was down.	Reconfigure replication and create a new replica or use the replica that had been promoted.
1071	E	Replication could not proceed because remote device ID %1 does not exist or is not a replica.	The replica device may have been deleted.	Reconfigure replication.
1073	E	Replication could not proceed because the configuration file could not be opened.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.

Number	Type	Text	Probable Cause	Suggested Action
1074	E	Replication could not proceed because memory allocation failed.	The system may have been busy and did not have enough memory.	Check the memory usage of different processes and stop unnecessary processes to free up memory. You may need to restart server modules.
1075	E	Replication could not proceed because unexpected error %1 occurred.	Replication reports the error specified in the message.	Check system logs on both servers and take necessary actions based on the error.
1082	W	Replication for virtual tape ID %1 was cancelled.	This was most probably triggered by a user.	If this was not triggered by a user, check the system log to identify any related errors and take necessary actions based on the error.
1084	W	A SCSI command completed after recovering from an error. The device may have some reliability issues.	This may be due to a temporary error on the physical device.	Check the system log for additional information. Contact the hardware manufacturer for a diagnostic procedure.
1087	W	Replication could not proceed because virtual tape ID %1 is not currently available.	The tape is loaded in a drive and is in use.	Wait for the process to complete before trying again.
1088	E	Replication could not proceed because ID %1 could not 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.

Number	Type	Text	Probable Cause	Suggested Action
1203	E	Storage path failed to be trespassed to %1.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1204	E	Storage path %1 failed to be added in the group.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1206	E	Storage path %1 failed to be activated.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1207	E	A critical storage path failure was detected. Path %1 will be removed.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1208	W	Storage path %1 does not belong to the active path group.	There was an attempt to access the storage via a path that is not part of an active group.	Use only active paths.
1210	W	There is no valid path available for device %1.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1211	W	There is no valid path group available.	Downstream storage path group is not correctly configured.	Check path group configuration on the storage device.
1212	W	There is no active path group for device GUID %1.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1214	E	Storage path %1 is not available.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
1215	W	CLARiiON storage path is trespassing.	A downstream storage path failed or the path was manually trespassed.	If the path was not manually trespassed, check the physical device status, device connectivity and switches, and the storage log.

Number	Type	Text	Probable Cause	Suggested Action
1216	W	T300 storage path is trespassing.	A downstream storage path failed or the path was manually trespassed.	If the path was not manually trespassed, check the physical device status, device connectivity and switches, and the storage log.
1217	W	HSG80 storage path is trespassing.	A downstream storage path failed or the path was manually trespassed.	If the path was not manually trespassed, check the physical device status, device connectivity and switches, and the storage log.
1218	W	MSA1000 storage path is trespassing.	A downstream storage path failed or the path was manually trespassed.	If the path was not manually trespassed, check the physical device status, device connectivity and switches, and the storage log.
7001	E	Patch %1 failed; environment profile is missing in /etc.	Unexpected loss of environment variables defined in /etc/.is.sh occurred on the server.	Check server package installation.
7002	E	Patch %1 failed; it applies only to build %2.	The server is running a different build than the one for which the patch is made.	Get the patch, if any, for your build number or apply the patch on another server that has the expected build number.
7003	E	Patch %1 failed; you must be the root user to apply the patch.	The user account running the patch is not the root user.	Run the patch with the root account.
7004	W	Patch %1 installation failed; it has already been applied.	You tried to apply the same patch again.	No action is needed.
7005	E	Patch %1 installation failed; prerequisite patch %2 has not been applied.	A previous patch is required but has not been applied.	Apply the required patch before applying this one.
7006	E	Patch %1 installation failed because it could not copy new binaries.	Server modules are not in a consistent state or an unexpected error occurred on the binary file name or path in the patch.	Check the patch log file in /usr/local/<product>-archive.

Number	Type	Text	Probable Cause	Suggested Action
7008	W	Patch %1 rollback failed; there is no original file to restore.	You tried to roll back a patch that has not been installed or has already been rolled back.	No action is needed.
7009	E	Patch %1 rollback failed because it could not copy back previous binaries.	Unexpected error on the binary file name or path in the patch.	Check the patch log file in /usr/local/<product>-archive.
7010	E	Patch %1 failed; the file %2 has patch level %3, higher than this patch. You must first roll back %4.	A patch with a higher level patch has been applied that conflicts with this patch.	Roll back the higher level patch, apply this patch, and then reapply the higher level patch.
7011	E	Patch %1 failed; it applies only to kernel %2.	You tried to apply the patch applied to a server that is not running the expected OS kernel.	Apply the patch on a server that has the expected kernel.
7012	E	Patch %1 failed; the available free space is %2 bytes; you need at least %3 bytes to apply the patch.	Patch applied to a server running low on the disk used for server home directory.	Add more storage.
10001	E	User ID %1 has insufficient privileges.	Server modules are not running with root privileges.	Log in to the server with the root account before starting server modules.
10002	W	The server environment is not set properly.	Expected environment variables in /etc/.is.sh are missing.	Restore the file from an X-ray or a backup in /usr/local/<product>-archive.
10003	E	Initialization of configuration %1 failed.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
10004	E	SCSI device information failed to be retrieved.	The device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.

Number	Type	Text	Probable Cause	Suggested Action
10006	E	A write operation to configuration %1 failed.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
10054	E	Server FSID update failed.	This may be due to insufficient system memory or an unhealthy file system.	Check the server memory and file system status. You may need to restart server modules.
10059	E	The server configuration update for FC storage persistent encountered an error.	There is a conflict with the ACSL.	Use a different ACSL for binding.
10100	E	New SCSI devices failed to be scanned.	This may be due to unreliable storage connectivity, hardware failure, or system resources are running low.	Check the storage devices and the connectivity status. Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart the server machine.
10101	E	Update of configuration %1 failed.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
10102	E	SCSI devices failed to be added.	This may be due to unreliable storage connectivity, hardware failure, or system resources are running low.	Check the storage devices and the connectivity status. Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart the server machine.

Number	Type	Text	Probable Cause	Suggested Action
10204	E	Configuration file %1 could not be parsed.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. You may need to restart server modules. If the problem persists, contact Technical Support.
10207	E	Adapter %1 could not be added because there is not enough memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
10209	E	Physical device %1 could not be added because there is not enough memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
10210	W	Physical device %1 was marked as offline because its GUID, %2, did not match SCSI GUID %3.	This may be due to an old device being imported without proper initialization, invalid configuration, or corrupted device header.	Check the physical storage. Replace the device if it is not reliable. Fix any detected issues and rescan devices to refresh the configuration.
10212	W	Physical device %1 was marked as offline because the SCSI status indicated it was offline; GUID: %2.	The device may have been removed, turned off, or is not functioning properly.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
10213	W	Physical device %1 was marked as offline because it did not respond correctly to an inquiry; GUID: %2.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
10214	W	Physical device %1 was marked as offline because its GUID, %2, does not match a valid FSID.	The GUID recorded on the device header does not match the unique ID, called the FSID, which is based on the external properties of the physical device. This may be due to device changes while the server was down.	Make sure devices are not changed directly by 3rd-party applications. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
10215	W	Physical device %1 was marked as offline because its storage capacity has changed; GUID: %2.	The physical device geometry, including the number of sectors, is different from the original record.	Rescan devices to refresh the configuration.
10240	E	SCSI path %1 is missing.	This may be due to a disconnected storage cable, a re-zoned Fibre Channel switch, or a failed storage port.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
10241	E	Physical adapter %1 could not be located in /proc/scsi/.	The adapter driver may not be loaded.	Run 'lsmod' to check loaded drivers and try to load the driver if needed.
10242	E	Physical adapter number %1 is duplicated in /proc/scsi/.	The OS has assigned the same number to two different adapters probably due to continuously loading and unloading drivers.	Do not repeatedly load and unload the Fibre Channel drivers and the server modules individually. You may need to reboot the server.
10244	E	Device %1 LUN in FSID %2 does not match the actual LUN.	The FSID was generated using the physical device LUN but the LUN assignment may have changed on the storage controller.	Do not change the LUN after the device FSID has been generated. Revert back to the original LUN.
10245	E	FSID %1 does not match device ACSL %2, GUID %3.	The FSID was generated using the physical device LUN but the device SCSI path may have changed.	You may need to rescan devices.
10246	E	FSID generation for the device with ACSL %1 failed.	The physical device does not provide reliable data in the SCSI inquiry pages in order to generate a unique ID.	Prepare this device to become a virtual device.
10247	E	GUID for the device with ACSL %1 is blank; FSID validation failed.	The device header may have accidentally been erased by a system command such as fdisk or format.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
10250	W	SCSI alias path %1 was removed because it did not have the same virtualization category as physical device %2.	The hardware configuration may have changed or there is a device or connectivity failure.	No action is necessary if this is due to a configuration change. Otherwise, check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
10251	W	SCSI alias path %1 was removed because it did not have the same GUID as physical device %2.	The hardware configuration may have changed or there is a device or connectivity failure.	No action is necessary if this is due to a configuration change. Otherwise, check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
11000	E	A socket failed to be created.	This may be due to insufficient system resources.	Check your network environment; you may need to restart the system.
11001	E	Setting the socket to re-use address failed.	This may be due to a network configuration error.	Check your network environment; you may need to restart the system.
11002	E	Binding the socket to port %1 failed.	Another process may be using the same port number.	Identify the process using the port and stop it.
11003	E	TCP service failed to be created.	This may be due to insufficient system resources.	Restart the server modules.
11004	E	TCP service failed to be registered; program %1 version %2.	This may be due to a network configuration error.	Restart the server modules.
11006	E	The server communication module failed to start.	This is due to the communication port being unavailable or a network error.	Restart the communication module.

Number	Type	Text	Probable Cause	Suggested Action
11007	W	There is not enough disk space available to successfully complete this operation and maintain the integrity of the configuration file. There is currently %1 MB of disk space available. The server requires %2 MB of disk space to continue.	The available space on the disk holding the configuration file is not enough.	Increase disk space.
11030	E	The Auto-Save option failed to set up a cron job.	The system command to add a cron job returned with an error.	Check the configuration in / etc/crontab and the status of the cron daemon.
11031	E	The Auto-Save option failed to create cron job script %1.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
11032	E	The Auto-Save option failed to connect to FTP server %1 on port %2.	This may be due to an invalid FTP address or port, or the FTP service is not running on the server.	Check that FTP service is enabled on the server; check network connectivity to the FTP site by manually running an FTP session.
11033	E	The Auto-Save option failed to log in to the FTP server with user %1.	The indicated user name is not valid.	Get a valid user name that can connect to the FTP server.
11034	E	The Auto-Save option failed because directory %1 does not exist.	The expected directory to back up the configuration files on the FTP site is missing.	Create the directory on the FTP site.
11035	E	The Auto-Save option failed to copy %1 to the FTP server.	The FTP user account may not have the write access to the directory on the FTP site.	Check the FTP user account and make sure it has valid access rights.
11036	E	The Auto-Save option failed to delete previous file %1 from the FTP server.	The FTP user account may not have the proper access rights for the directory on the FTP site.	Check the FTP user account and make sure it has valid access rights.

Number	Type	Text	Probable Cause	Suggested Action
11101	E	Client %1 failed to be added.	This error is most likely due to a system configuration error, or system resources running low.	Check OS resources running provided utilities such as 'top'.
11104	W	There are too many client connections.	The number of simultaneous client connections exceeded the limit that the current system memory can handle.	This is an unlikely condition as long as the recommended memory is available for the server.
11107	E	Client %1 encountered an illegal access error.	The client host attempted to connect to a session that is no longer available.	Restart the connection and ensure there is no security breach.
11109	E	Client %1 failed to open file %2.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
11112	E	Client %1 failed to parse configuration file %2.	The configuration file is corrupted, or manually tempered to the degree that is no longer recognizable by StorSafe. If corruption is the cause, then it is most likely due to a system drive hardware error.	If there is a valid configuration file saved, it can be restored to the system. Make sure to use reliable storage devices for critical system information.
11113	E	Client %1 failed to restart the authentication module.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11114	E	Client %1 encountered a memory allocation failure.	System resources are running low. This may be due to too little memory installed for the system, or some runaway process that is consuming too much of the memory.	Run 'top' to check the process that is using the most memory. If physical memory is below the server recommendation, install more memory on the system.
11170	E	Virtualization of %1 failed because its size in the configuration file was different from the one on the disk header. Run a rescan and try again.	Attempting to virtualize a LUN that has a different capacity than what was previously seen.	Rescan for new devices and try again.
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.

Number	Type	Text	Probable Cause	Suggested Action
11203	E	User %1 failed to rescan SCSI devices.	This may be due to unreliable storage connectivity, hardware failure, or system resources are running low.	Check the storage devices and the connectivity status. Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart the server machine.
11204	E	An operation by %1 failed to check SCSI devices.	This may be due to unreliable storage connectivity, hardware failure, or system resources are running low.	Check the storage devices and the connectivity status. Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart the server machine.
11205	E	An operation by %1 failed to get information for file %2.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
11206	E	An operation by %1 resulted in a memory allocation failure.	The system does not have enough memory.	Check the memory usage of different processes and stop unnecessary processes to free up memory. You may need to restart server modules.
11207	E	An operation by %1 failed to open file %2.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
11208	E	An operation by %1 failed to read file %2.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
11209	E	User %1 has insufficient access privileges.	Access rights of the specified user are limited.	Retry the operation with a valid user account.

Number	Type	Text	Probable Cause	Suggested Action
11211	E	An operation by %1 failed to save file %2.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
11212	E	An operation by %1 failed to create index file %2 for the Event Log.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
11213	E	An operation by %1 got an illegal time range (%2 - %3) for the Event Log.	The specified time range is not valid.	Retry with a valid range.
11214	E	An operation by %1 failed to get time range (%2 - %3) for the Event Log.	The Event Log does not cover the specified time range.	Retry with a valid range.
11215	E	An operation by %1 failed to open directory %2.	Either the directory is not available or the system is busy and does not have enough resources.	Make sure the directory exists and check the system status. You may need to restart server modules.
11216	E	An operation by %1 failed to fork a process due to insufficient system resources.	The system does not have enough memory.	Check the memory usage of different processes and stop unnecessary processes to free up memory. You may need to restart server modules.
11217	E	An operation by %1 failed to execute program %2.	Either the program is not available or the system is busy and does not have enough resources.	Make sure the program exists and check the system status. You may need to restart server modules.
11218	E	An operation by %1 failed to remove file %2.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.

Number	Type	Text	Probable Cause	Suggested Action
11219	E	User %1 failed to add device %2.	This may be due to insufficient system memory or disk space to update the device information or the underlying physical device may have a failure.	Check system resources, physical device status, device connectivity, switches, and storage log.
11220	E	User %1 failed to remove device %2.	The replica server could not be accessed in order to delete the associated replica device.	Make sure the replica server is accessible and retry.
11221	E	User %1 failed to add client %2 to virtual tape %3.	This may be due to insufficient system memory or disk space to update the device information.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
11222	E	User %1 failed to remove client %2 from virtual tape %3.	This may be due to insufficient system memory or disk space to update the device information.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
11231	E	User %1 failed to get the CPU status.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
11232	E	User %1 failed to get the memory status.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
11233	E	An operation by %1 failed to map the SCSI device name for %2 %3 %4 %5.	The device identified by its SCSI address (Adapter, Channel, SCSI ID, LUN) cannot be located due to a configuration change or storage failure.	Check storage and rescan devices to refresh the configuration.
11234	E	User %1 failed to test device %2.	The system command 'hdparm' to the device failed probably due to insufficient system resources or a storage device failure.	Run the command manually on the system and check storage devices.

Number	Type	Text	Probable Cause	Suggested Action
11237	E	An operation by %1 failed to get file %2.	The file is in use by another process.	The portal automatically retries every 3 seconds. If this occurs repeatedly, close the portal and retry.
11238	E	An operation by %1 failed to restart the authentication module.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11240	E	User %1 failed to start server modules.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11242	E	User %1 failed to stop server modules.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11244	E	User %1 failed to get the user list.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
11257	E	User %1 failed to add client %2.	This error is most likely due to system configuration error, or system resources running low.	Check OS resources using provided utilities such as 'top'.
11261	E	User %1 failed to get the client connection status for virtual tape %2.	Failed to inquire a SAN Client connection status due to system configuration error, storage hardware failure, or system resource access failure. This should rarely happen.	Check the system resource, such as memory, system disk space. Check the system log for specific reason of the failure.
11262	E	An operation by %1 failed to parse configuration file %2.	The configuration file is not readable by the server.	If there is a valid configuration file saved, it can be restored to the system. Make sure to use reliable storage devices for the critical system information.
11263	E	User %1 failed to restore configuration file %2.	Error encountered when writing the server configuration file to the system drive. This can only happen if the system drive ran out of space, is corrupted, or there has a hardware failure.	Check the system drive using OS-provided utilities. Free up space if necessary. Replace drive if not reliable.

Number	Type	Text	Probable Cause	Suggested Action
11265	E	An operation by %1 failed to restart the IO core module.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11266	E	User %1 failed to erase virtual tape %2 partition.	Storage hardware failure occurred.	Check the storage devices, e.g., power status; controller status, etc. Check the connectivity, e.g., cable connectors. With fibre channel switches, even the connection status light indicates the connection is good, it is still not a guarantee. Push the connector in to make sure. Check the specific storage device using OS-provided utilities such as 'hdparm'.
11278	E	User %1 failed to swap device ID %2 with its mirror.	Hardware problem with the repository mirror disk.	Check the repository mirror disk.
11280	E	User %1 failed to create secure communication credentials for server %2.	The remote server cannot be accessed to establish a secure communication channel.	Check that the specified server can be reached, the TCP ports required for communication are open, and both servers are using the same password security rules.
11289	E	User %1 failed to restart server modules.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to restart the server machine.
11291	E	An operation by %1 failed to update metadata of physical device %2.	Storage hardware failure occurred.	Check the storage devices, e.g., power status; controller status, etc. Check the connectivity, e.g., cable connectors. With fibre channel switches, even the connection status light indicates the connection is good, it is still not a guarantee. Push the connector in to make sure. Check the specific storage device using OS-provided utilities such as 'hdparm'.

Number	Type	Text	Probable Cause	Suggested Action
11292	E	User %1 failed to swap IP address %2 with %3.	Storage hardware failure occurred.	Check the storage devices, e.g., power status; controller status, etc. Check the connectivity, e.g., cable connectors. With fibre channel switches, even the connection status light indicates the connection is good, it is still not a guarantee. Push the connector in to make sure. Check the specific storage device using OS-provided utilities such as 'hdparm'.
11295	E	An operation by %1 failed because the configuration file has an invalid format.	The configuration file is not readable by the server.	If there is a valid configuration file saved, it can be restored to the system. Make sure to use reliable storage devices for the critical system information.
11301	E	An invalid password for user %1 was used by a client with IP address %2.	An incorrect username/ password was provided to connect to this server	Provide the correct username/password.
11315	E	You do not have a license for the %1 protocol.	An attempt was made to enable a protocol such as Fibre Channel, iSCSI with no appropriate license keys.	Install appropriate license keys that include support for the needed protocol.
11316	E	You have exceeded the backup cache capacity allowed by your license; used space: %1 MB, requested space %2 MB, licensed capacity: %3 MB.	Physical capacity usage has exceeded the capacity allowed by the license.	Contact FalconStor to purchase additional capacity license.
11406	E	The failover configuration failed to be prepared; %1.	Suspend fail over initiated while packaging fail over configuration information.	No action required. (Operation will be retried automatically)
11407	E	The failover configuration failed to be extracted; %1.	Suspend fail over initiated while packaging fail over configuration information.	No action required. (Operation will be retried automatically)

Number	Type	Text	Probable Cause	Suggested Action
11500	E	Virtual tape %1 expansion failed because there was not enough disk space.	There is no more storage space available for expanding a virtual tape.	Add more storage.
11512	E	User %1 failed to failed to enable replication of virtual tape %2 to server %3; watermark: %4 MB, time: %5, interval: %6, watermark retry: %7, suspended: %8.	It can be connection error while the primary server is synchronizing the initial status of the virtual tape to the replica on the target server, or the virtual tape is loaded to the drive at the moment by backup software.	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 StorSafe 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 StorSafe system has a mechanism to automatically prune the syslog from using up all system disk free space). If enough space is available on system disk, check the integrity of system disk file system using fsck utility.
11516	E	User %1 failed to create replica virtual tape %2.	Could not update the virtual tape partition information on disk.	Check the storage system and make sure that storage is working properly.
11518	E	User %1 failed to start replication of virtual tape %2.	The replication triggered manually by the user failed. It can be due to one of the following reasons. Network problems, Virtual tape is loaded in a drive, replication is in progress or the replica no longer exists.	If the replication is manually triggered by the user, check the replication status at the right panel of the replica before starting another replication. If the replication is triggered by the scheduler, adjust the schedule in the replication policy to avoid replicating too often. Check if the network is working properly as well as the server activity. Remove replication setup from the portal if it no longer has a valid replica.

Number	Type	Text	Probable Cause	Suggested Action
11522	E	User %1 failed to promote virtual tape replica %2 to a virtual tape.	Failed to update the virtual tape partition information.	Check if the physical disk is working properly or the server is busy. Retry the operation when the physical disk is working properly or when the server is not busy.
11524	E	User %1 failed to take an X-ray.	When any server process cannot be started, it is most likely due to insufficient system resources, invalid state left by a server process that may not have been stopped properly, or due to an unexpected OS process failure that left the system in a bad state. This should happen very rarely. If frequent occurrences are encountered, there must be external factors that contributed to the behavior that must be investigated and removed before running the server.	If the system resources are low, run 'top' to check the process that is using the most memory. If the physical memory is below the server recommendation, install more memory on the system. If the OS is suspected to be in a bad state due to an unexpected failure in either hardware or software components, restart the server machine.
11534	E	An operation by %1 failed to reset the replication control map for virtual tape %2.	Storage hardware failure.	Check the storage devices for power status, controller status, etc. Check for proper connectivity. Fibre Channel Switch connection status lights do not guarantee a solid connection. Disconnect/Reconnect the Fibre Channel connector for verification. check the specific storage device using OS provided utility such as 'hdparm'.
11535	E	User %1 failed to update replication parameters for virtual tape %2; server %3, watermark: %4 MB, time: %5, interval: %6, watermark retry: %7, suspended: %8.	Error encountered when writing the updated configuration to the StorSafe configuration file to the system drive. This can only happen if the system drive ran out of space, or is corrupted, or if there is hardware failure in the system drive.	Check if the system drive is out of space or has any corruption.

Number	Type	Text	Probable Cause	Suggested Action
11537	E	User %1 failed to import physical device %2.	This may due a specific version of the StorSafe server limiting the support of the storage.	Check license agreement for the version of StorSafe server.
11539	E	User %1 failed to import physical device %2.	Storage hardware failure.	Check the storage devices for power status, controller status, etc. Check for proper connectivity. Fibre Channel Switch connection status lights do not guarantee a solid connection. Disconnect/Reconnect the Fibre Channel connector for verification. check the specific storage device using OS provided utility such as 'hdparm'.
11542	E	User %1 failed to remove virtual tape replica %2.	Error encountered when writing the updated configuration to the StorSafe configuration file to the system drive. This can only happen if the system drive ran out of space, or is corrupted, or if there is hardware failure in the system drive.	Check if the system drive is out of space or has any corruption.
11554	E	User %1 failed to set the self-checking interval to %2 sec in failover properties.	IP Network communication failure.	Check network connectivity.
11569	E	User %1 failed to set Fibre Channel WWPN %2 to %3 mode.	This is possibly due to the fibre channel driver being improperly loaded, or the wrong version of the driver is loaded. StorSafe FC target mode requires the StorSafe version of the driver to be used.	Run 'lsmod' to check the proper StorSafe 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. StorSafe FC target mode requires the StorSafe version of the driver to be used.	Run 'lsmod' to check the proper StorSafe driver is loaded. If it is, check to make sure it is the correct version. The correct revision should be located in the VTL/lib directory.

Number	Type	Text	Probable Cause	Suggested Action
11581	E	User %1 failed to set the NAS option to %2.	Failed to start the NAS processes. It is most likely due to insufficient system resources, invalid state left by last running server processes that may not have been stopped properly, or due to an unexpected OS processes failure that left the system in a bad state. This should happen very rarely. If frequent occurrences are encountered, there must be external factors that contributed to the behavior that must be investigated and removed before running the server.	If system resources are low, run 'top' to check the process that is using the most memory. If the physical memory is below the server recommendation, install more memory on the system. If the OS is suspected to be in a bad state due to unexpected failure in either hardware or software components, restart the server machine.
11632	E	User %1 failed to set failover options on the partner server; heartbeat interval: %2 sec, auto-recovery interval: %3 sec.	IP Network communication failure.	Check Network connectivity.
11633	E	User %1 failed to set failover options on the partner server; heartbeat interval: %2 sec, auto-recovery interval: disabled.	IP Network communication failure.	Check Network connectivity.
11648	E	The inquiry string on SCSI device %1 failed to be retrieved.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
11649	E	Conversion of the inquiry string on SCSI device %1 failed.	The device SCSI inquiry string contains invalid information.	Check the device configuration and status.
11650	E	The capacity of SCSI device %1 failed to be retrieved.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
11653	W	SCSI device %1 was ignored due to unsupported type %2.	The disk is not from a supported vendor.	Make sure you are using supported devices.
11655	E	SCSI device %1 capacity is not valid.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
11656	W	SCSI device %1 was ignored due to an unsupported Cabinet ID.	The Cabinet ID of the device is not supported.	Make sure you are using supported devices.
11657	W	SCSI device %1 was ignored due to a missing %2 vendor in the inquiry string.	The disk is not from a supported vendor.	Make sure you are using supported devices.
11741	E	An operation by %1 failed to create virtual library with %2 slots because only %3 slots were available.	Specified slot count for the virtual library exceeds the total slot count supported by the system.	Specify appropriate slot count.
11742	E	User %1 created only %2 out of %3 virtual drives due to a memory allocation failure.	System is out of memory, which prevents the creation of the specified number of virtual tape drives.	Increase system memory.
11744	E	An operation by %1 rolled back the configuration update for test mode promotion of %2 %3(s).	Failed to write the StorSafe configuration file.	Check /usr partition for free space. If the partition is out of space, delete unwanted files to create space.
11745	E	An operation by %1 rolled back the disk partition update for test mode promotion of %2 %3(s).	Failed to write update partition information to disks.	Check physical connectivity to the storage system/LUN.
11746	E	An operation by %1 rolled back test mode promotion of %2 %3(s).	Cannot add device to IOCore module.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
11782	E	Barcode [%1] of the source tape ID %2 already exists on target server %3. Auto-replication could not be configured.	A tape with the same barcode exists on the remote server.	Remove the tape with the same barcode from the remote server or change its barcode.
11788	E	Appliance hardware problem %1 was detected.	Detected an error on the StorSafe appliance.	Check the error message for information on the error to determine the solution.
11791	E	Virtual tape %1 failed to shrink to %2 MB; error: %3.	Virtual tape partition information couldn't be updated. This could happen in rare cases when the system is extremely busy and updates to disks take too long.	No user action required since it doesn't cause any problem for backup/restore. If this error happens, the tape will not be resized.
11793	W	Appliance hardware problem %1 was detected.	Detected a hardware problem with the appliance.	Check the error message and take appropriate hardware maintenance.
11796	E	The tape caching data migration trigger policy failed to be set.	Connectivity issues with the storage where StorSafe repository is located.	Check storage connectivity.
11797	E	The tape caching reclamation trigger policy failed to be set; error: %1.	Connectivity issues with the storage where StorSafe repository is located.	Check storage connectivity.
11906	E	File %1 failed to open.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
12002	E	Directory %1 failed to open.	Either the directory is not available or the system is busy and does not have enough resources.	Make sure the directory exists and check the system status. You may need to restart server modules.
12003	E	File %1 failed to open.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
12509	E	The inventory of physical library ID %1 failed.	Physical library inventory process failed.	Check the server log for more information on the cause of HW failure.
12511	E	A tape in library ID %1 failed to move from slot %2 to drive ID %3.	Physical tape move from slot to drive failed.	Check the server log for more information on the cause of HW failure.

Number	Type	Text	Probable Cause	Suggested Action
12514	E	Job %1 failed to restart; command: %2, return code: %3.	Import/Export tape job failed probably due to the tape not being present or the tape drive not being available.	Check the tape exists and the tape drive is available for job to continue, and then restart the job.
12521	E	A tape in library ID %1 failed to be moved from drive ID %2 to slot %3.	Physical library may have hardware issues.	Check the server log for more information on the cause of HW failure.
12525	E	Tape %1 failed to be reclaimed.	The virtual drive is busy with I/O and is not responsive to the upper layer calls.	Try again when the system is less busy, or determine the cause of the extensive I/O, and correct the situation if necessary.
12559	E	Jobs %1 failed to be resumed; error: %2, %3.	The jobs may already be running.	No action is required.
12564	E	Hardware or software compression failed to be set.	The physical device used by the database may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
12565	E	A tape in library ID %1 failed to be moved from import/export slot %2.	The library is in maintenance mode or the destination slot is not available.	Ensure that the library is operational and there is no tape in the slot.
12567	E	A tape in library ID %1 failed to be moved from slot %2 to an import/export slot.	The library is in maintenance mode or the import/export slot is not available.	Ensure that the library is operational and there is no tape in the slot.
12578	E	An import job from physical library ID %1 for tape %2 to virtual library ID %3 for tape %4 failed to be submitted.	This may be due to a memory allocation error, the tape does not have a valid barcode, or the maximum number of tapes in a job has reached.	Check the system log for more details and take necessary actions based on reported errors.
12579	E	An import job from physical drive ID %1 to virtual library ID %2 for tape %3 failed to be submitted.	This may be due to a memory allocation error, the tape does not have a valid barcode, or the maximum number of tapes in a job has reached.	Check the system log for more details and take necessary actions based on reported errors.

Number	Type	Text	Probable Cause	Suggested Action
12581	E	A tape stacking job failed to be submitted for physical tape %1 [%2]; error: %3.	The job reports the error specified in the message.	Take necessary actions based on the error.
12583	E	A tape stacking job failed to be submitted for physical library ID %1.	This may be due to a memory allocation error, the tape does not have a valid barcode, or the maximum number of tapes in a job has reached.	Check the system log for more details and take necessary actions based on reported errors.
12585	E	Importing stacked physical tapes to virtual tapes detected that the virtual and physical tapes did not have the same capacity.	The stacked tape import feature requires the same size for physical and virtual tapes.	Use tapes with the same capacity.
12586	E	Memory allocation failed for stacked tapes.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
12587	E	Memory allocation failed for virtual tapes.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
12588	E	A job to import stacked tapes from physical library ID %1 to virtual library ID %2 failed to be submitted; error: %3.	The job reports the error specified in the message.	Take necessary actions based on the error.
12590	E	Job %1 failed to be purged.	The physical device used by the database may have a failure or the job may have already been purged.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
12592	E	Tape %1 properties failed to be set.	The tape may be in use.	Retry later.

Number	Type	Text	Probable Cause	Suggested Action
12598	E	Physical tape %1 failed to be purged from the database.	The physical device used by the database may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
12600	E	%1 failed to assign %2 to client %3 (%4, %5).	The maximum number of Fibre Channel devices assigned to the client was reached.	Unassign devices that are no longer needed.
12601	E	A tape that was ejected failed to be loaded in virtual drive ID %1.	The eject command may not have completed.	Make sure the tape is completely ejected and try again.
12603	E	The ACSLS/LS physical drive ID %1 failed to be assigned to the server; error: %2.	This may be due to network or FC connection problems.	Check the ACSLS is running properly and is connected to both the server and the physical library. Check the FC connection between the server and the library.
12605	E	Physical tapes failed to be ejected; error: %1.	The server may not be connected to the physical library or the tape is no longer in the library.	Check the FC connection between the server and the library; if the tape is still loaded, try again or eject the tape manually.
12611	E	The ACSLS/LS library failed to be added as physical library ID %1.	This may be due to network connection problems or the maximum number of supported libraries has been reached.	Check the ACSLS is running properly and is connected to both the server and the physical library.
13101	E	Communication with server %1 failed; error: %2	This may be due to a communication error with the partner server.	Check the connection with the partner server.
13102	E	Execution of %1 failed.	The self-monitor module did not start properly.	Try to manually start the module.
13300	E	Authentication to server %1 failed; the failover module stopped.	The credentials used by failover servers are no longer valid.	Reset the root password of both hosts and reconfigure failover.
13301	E	Authentication to the local server failed; the failover module stopped.	The credentials used by failover servers are no longer valid.	Reset the root password of both hosts and reconfigure failover.

Number	Type	Text	Probable Cause	Suggested Action
13302	E	The configuration file of partner server %1 failed to be transferred to server %2; error: %3.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13303	E	Dynamic configuration of server %1 failed to be transferred to partner server %2; error: %3.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13304	E	File %1 failed to be renamed.	The file name already exists or the file system is inconsistent or read-only.	Check the file system.
13305	E	A write operation failed on file %1.	The file does not exist or there may be insufficient disk space, a system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
13306	E	File %1 failed to be opened.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
13307	E	Credential information of server %1 failed to be transferred to server %2; error: %3.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13308	E	Invalid failover configuration was detected. Failover will not occur.	The partner's configuration file is missing.	Check connectivity with the partner to make sure the configuration can be transferred.
13309	E	Server %1 is unable to communicate with server %2.	This may be due to a network error or a connectivity issue with the Configuration Repository.	Restart the network and check the Configuration Repository storage device.
13316	E	A virtual IP address failed to be added; error: %1.	This may be due to a network error.	Restart the network. If the problem persists, you may need to reboot.

Number	Type	Text	Probable Cause	Suggested Action
13317	E	A virtual IP address failed to be released; error: %1.	The server is still holding on to the IP address of the failed server.	Take necessary actions based on the error specified in the message.
13319	E	The failover module failed to stop. You may need to reboot the server.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to reboot.
13320	E	Update of configuration files failed on server %1; %2.	This may be due to a communication error with the partner server.	Check the connection and the filesystem on the partner server.
13700	E	Memory allocation failed; the self-monitoring module stopped.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory.
13701	E	Virtual IP address %1 failed to be released; the operation is being retried.	The server is still holding on to the IP address of the failed server.	No action is needed since this is temporary.
13702	E	Virtual IP address %1 failed to be retrieved; the operation is being retried.	This may be due to a network error.	Restart the network.
13703	E	The self-monitoring module failed to stop.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to reboot.
13710	W	The live trial license expired for %1. Contact 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.
13800	C	The following failure was detected on server %1: %2	The server reported the specified error. The partner server may take over depending on the current condition.	Based on the reported error, take appropriate actions.

Number	Type	Text	Probable Cause	Suggested Action
13807	C	A failure was detected on server %3. Server %1 will take over server %2.	A critical error occurred that triggered failover.	After takeover is complete, check the server log and fix any detected issue prior to failback.
13817	C	Server %1 could not fail back because it failed to update the partner server configuration.	This is related to other error conditions.	Check the system log to identify any related errors and take necessary actions based on the error.
13818	E	Access to the Configuration Repository failed.	The virtual device holding the Configuration Repository is no longer available.	The partner server will take over anyway. Shut down this server in order to avoid any conflicts.
13820	W	Server %1 heartbeat failed to be detected; %2.	This server did not receive a heartbeat from the partner server and is trying to contact other network entities to determine if this server is disconnected from the network or if the partner server is down.	Check related messages in the log for more details.
13821	E	The server could not reach other entities in the network. Takeover was not initiated because it is assumed the failure is on this server.	This server did not receive a heartbeat from the partner server and failed to contact any other network entities in the subnet.	Check network connectivity.
13822	C	Server %1 will not take over because storage devices could not be accessed.	This is due to a storage or connectivity issue that is affecting both servers.	Check the storage and connectivity for both servers and fix any detected issue.
13823	W	Server %1 failed to acknowledge the takeover request in time. Server %2 will forcefully take over.	The partner server is busy or not fully operational.	After takeover is complete, check the partner server and fix any detected issue prior to failback.
13827	E	The Configuration Repository update process with PID %1 failed to be stopped.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
13834	E	Copying files from the Configuration Repository failed.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13836	E	Getting configuration files from the Configuration Repository failed. Check and fix the repository disk.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13841	E	Takeover cannot complete because the partner server status is %1 instead of DOWN or READY.	The self-monitor module may not be running or may be busy on the partner server.	Check that the module is running.
13843	E	This server failed to get the original configuration file from the Configuration Repository of the partner server before failback.	The physical device used by the Configuration Repository may have a failure or there may be a connectivity issue with the partner server.	Check the physical device status, device connectivity and switches, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.
13845	W	The Configuration Repository disk failed. Server %1 is still in takeover mode.	The physical device used by the Configuration Repository may have a failure or there may be a connectivity issue with the partner server.	Check the physical device status, device connectivity and switches, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.
13849	W	The following heartbeat IP addresses are down:%1	This may be due to a network error.	Check connectivity between servers; check network parameters, including jumbo frame configuration, if applicable.
13850	E	Server %1 could not locate the Configuration Repository disk; error: %2.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
13851	E	This server could not take over; reason: %1.	The server reports the error specified in the message.	Check the server and take necessary actions based on the error.

Number	Type	Text	Probable Cause	Suggested Action
13853	E	This server is unable to take over and notified the partner server to remain active.	The server detected a failure on its partner but a local failure prevented takeover.	Check the status of both servers.
13856	E	Server %1 failed to communicate with server %2 through IP %3.	This may be due to a network error.	Check connectivity between servers; check network parameters, including jumbo frame configuration, if applicable.
13859	E	This server will reboot because it detected a storage or network failure while it was taking over its partner.	The server cannot take over if it has any hardware issue.	Check network or storage connectivity.
13860	E	After server recovery, configuration files %1 and %2 failed to be merged.	The configuration files may be inconsistent; this may result in losing some configuration changes made during failover.	Check the configuration and re-apply your changes, if necessary.
13861	E	File %1 failed to be renamed to %2.	The file name already exists or the file system is inconsistent or read-only.	Check the file system.
13863	C	This server was forced to resume its operations; %1.	A user initiated a forceup command even though the server may not be fully functional.	Check the server status.
13864	C	This server is stopping its operations; reason: %1.	The partner server is requesting this server stop and be taken over.	Check the server status.
13878	E	This server has an invalid failover configuration; reason: %1.	The server reports the error specified in the message.	Check the server and take necessary actions based on the error.
13879	C	This server has detected a kernel module failure.	A kernel module on the partner server is not healthy.	Check the partner server; you may need to reboot it.
13882	E	The following error was detected on the Configuration Repository: %1.	The physical device used by the Configuration Repository may have a failure or there may be a connectivity issue with the partner server.	Check the physical device status, device connectivity and switches, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
13900	E	This server failed to take over its partner server due to %1.	This may be due to an issue with the power control device that is needed to turn off the partner or a problem retrieving information from the Configuration Repository.	Check the power control device on the partner server and connectivity with the Configuration Repository storage.
13901	E	Reading %1 from the Configuration Repository failed.	The physical device used by the Configuration Repository may have a failure or there may be a connectivity issue with the partner server.	Check the physical device status, device connectivity and switches, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.
13910	E	The server failover status could not be updated in the Configuration Repository; reason: %1.	The physical device used by the Configuration Repository may have a failure or there may be a connectivity issue with the partner server.	Check the physical device status, device connectivity and switches, the storage log, and the network. Fix any detected issues and rescan devices to refresh the configuration.
13921	C	The following failure was detected: The NAS module terminated abnormally.	A user may have stopped the module or the module had a failure. In a failover configuration, the partner server will try to take over this server.	If the module was not stopped intentionally, check the system log to get more information about the reason for the module failure.
13947	E	Server %1 cannot forcefully take over server %2 since an IPMI/HP iLO power control command failed.	The power control command ipmitool or hponcfg could not query or set the power status of the partner server.	Check the network connectivity with the partner server. Try to run the related command manually to check the power control module status.
14000	E	Virtual tape ID %1 is missing.	The device may have been deleted by another user.	Perform the operation on an existing device.
14001	E	A disk partition checksum mismatch was detected for virtual tape %1.	This is due to some configuration inconsistency issues.	Contact Technical Support.
14003	E	The configuration file failed to be parsed.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. You may need to restart server modules. If the problem persists, Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
14004	E	Virtual tape ID %1 failed to be renamed to %2.	The configuration file could not be updated probably due to insufficient memory, system disk failure, or an unhealthy file system.	Ensure there is enough memory and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
14005	E	Disk partition information could not be updated.	The device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
14006	E	Configuration %1 failed to be written.	The configuration file could not be updated probably due to insufficient memory, system disk failure, or an unhealthy file system.	Ensure there is enough memory and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
14008	E	Physical device %1 failed to be renamed to %2.	The configuration file could not be updated probably due to insufficient memory, system disk failure, or an unhealthy file system.	Ensure there is enough memory and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
14010	E	Client %1 is missing.	The client may have been deleted by another user.	Perform the operation on an existing client.
14011	E	Client %1 failed to be renamed to %2.	The configuration file could not be updated probably due to insufficient memory, system disk failure, or an unhealthy file system.	Ensure there is enough memory and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
14016	E	User %1 failed to renew caching for tape ID %2.	There may not be enough storage available.	Add storage, if applicable.

Number	Type	Text	Probable Cause	Suggested Action
14017	C	The server configuration failed to be saved to the Configuration Repository. Check the storage connectivity.	The Configuration Repository device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
14021	E	The physical device with ACSL %1 failed to be erased.	The device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
14023	E	The physical device with ACSL %1 failed to be claimed.	The device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
16001	E	An operation by %1 failed to convert the file system on %2.	The operation failed due to a system or mount error.	Contact Technical Support.
16100	W	Connection to SMTP server %1 on port %2 failed for email alerts; reason: %3.	The IP address or the port ID of the SMTP server may not be valid.	Check network connectivity, SMTP server IP address and port ID.
16101	W	Email alerts could not be sent via SMTP server %1 on port %2; reason: %3.	The IP address or the port ID of the SMTP server may not be valid.	Check network connectivity, SMTP server IP address and port ID.
18000	E	The Deduplication Repository disk %1 with GUID %2 is missing.	The server could not detect a data disk from another node.	Confirm services are running on the node that owns the missing device. Check the underlying physical storage, FC connectivity, and switch zoning.
18001	E	The repository module failed to start because the Deduplication Repository disk is missing. Check that each member of the cluster is operating properly.	The server could not connect to a physical device used by the deduplication repository.	Check the underlying physical storage, FC connectivity, and switch zoning. Ensure that each cluster node can connect to the storage.

Number	Type	Text	Probable Cause	Suggested Action
18050	E	The reclamation triggering module stopped because index pruning has failed.	The index disk may be in an invalid state.	Contact Technical Support.
18055	E	Deduplication index reclamation failed; reason: %1.	The operation failed due to the specified reason.	Take necessary actions based on the reported error.
18060	E	Deduplication Repository space reclamation did not complete; reason %1	The operation failed due to the specified reason.	Take necessary actions based on the reported error.
18066	E	Deduplication index pruning did not complete; reason: %1.	The operation failed due to the specified reason.	Take necessary actions based on the reported error.
18067	E	Deduplication index pruning did not start; reason: %1.	The operation failed due to the specified reason.	Take necessary actions based on the reported error.
18068	E	Space reclamation did not start; reason: %1.	The operation failed due to the specified reason.	Take necessary actions based on the reported error.
18069	E	Space reclamation will fail because folder %1 does not exist.	The folder referenced by the VIT is not present in the deduplication repository.	Contact Technical Support.
18070	E	A hardware failure was detected on the data disk. The repository module has been shut down.	For consistency and safety, the repository has been shut down.	Contact Technical Support immediately.
18071	C	A data consistency error was detected. The repository module has been shut down.	For consistency and safety, the repository has been shut down.	Contact Technical Support immediately.
18072	C	A hardware failure was detected on the metadata disk used for the index. The repository module has been shut down.	For consistency and safety, the repository has been shut down.	Contact Technical Support immediately.
18073	C	An invalid block was detected while reading the metadata disk for the index. The repository module has been shut down.	For consistency and safety, the repository has been shut down.	Contact Technical 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 switches, and the storage log.
18087	W	The Deduplication Repository usage threshold of %1 was reached. %2 of repository capacity has been used.	The specified threshold has been reached.	Run reclamation, add storage, or change the threshold.
18090	E	Deduplication index reclamation failed because there is no folder.	You may have deleted, erased, or relabeled all tapes when you detected the SIR repository was full. At least one VIT tape is required for reclamation to proceed.	Contact Technical Support.
18101	E	The power control module status for node %1 at IP address %2 failed to be retrieved.	The power control module may not be configured properly or may not be functioning correctly.	Make sure a power control module such as IPMI or HP iLO is present on the server, is enabled, and is configured correctly.
18104	E	This standby server failed to take over node %1 because that node reported its hostname as %2, which is unknown to this server.	The hostname is not in the cluster configuration file probably because the server was not registered correctly.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
18105	E	This standby server failed to take over node %1 because it cannot retrieve required parameters in the cluster configuration.	The standby server needs cluster parameters in order to take over.	Contact Technical Support.
18107	E	This standby server failed to take over node %1 because it could not add virtual IP %2 of the failed node.	The heartbeat IP address may not have been properly set.	Contact Technical Support.
18108	E	This standby server failed to take over node %1 because it could not spoof the WWPN of the failed node.	The standby server cannot assume the FC identity of the failed node.	Contact Technical Support.
18109	E	This standby server failed to take over node %1 because it could not change its hostname to the name of the failed node.	The standby server cannot assume the identity of the failed node.	Contact Technical Support.
18110	E	This standby server failed to take over node %1 because it could not load physical resources of the failed node.	One or more physical devices may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
18111	E	This standby server failed to take over node %1 because it could not retrieve the failed node configuration from the Configuration Repository.	The standby server needs the failed node's configuration in order to assume its identity.	Contact Technical Support.
18112	E	This standby server failed to take over node %1 because it could not start up the deduplication module.	The standby server needs to start certain services to become active and take over the failed node.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
18113	E	This standby server failed to take over node %1; reason: %2.	The standby server reports the error specified in the message.	Take necessary actions based on the specified reason. If failover is desired, it should be performed manually.
18114	E	This standby server could not see physical device %1 [%2], used for the Deduplication Repository.	The standby server could not connect to a physical device used by the deduplication repository.	Check the underlying physical storage, FC connectivity, and switch zoning. Rescan devices on the standby server.
18116	E	This standby server detected physical device %1 [%2], used for the Deduplication Repository, is offline.	The standby server could not connect to a physical device used by the deduplication repository.	Check the underlying physical storage, FC connectivity, and switch zoning.
18119	E	This standby server failed to check the health of node %1 at IP address %2 because of an authentication error.	The remote server cannot be accessed via a secure communication channel.	Check that the server can be reached. You may need to restart services on this node.
18120	E	This server cannot be selected as a standby server in a failover configuration because it is running as an active node.	This server is running the deduplication module as an active node.	Select an appropriate server for the standby server. If the module has been started manually, restart services on this server and try again.
18133	W	This standby server is being forced to take over node %1, which may not be fully shut down.	The standby server cannot take over if it is unable to confirm that the failed node is shut down.	Make sure that the node that is to be taken over has actually been shut down, and if not, shut it down.
18135	E	User %1 failed to set a reclamation schedule for deduplication.	An I/O operation to the deduplication configuration file failed probably due to insufficient system resources.	Ensure there is enough memory and disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.

Number	Type	Text	Probable Cause	Suggested Action
19004	C	Storage usage has reached the threshold; %1 of total storage is being used; total storage capacity is %2; available storage space is %3.	Storage utilization has reached the limit specified by the user.	Check storage utilization and delete unused virtual tapes to free up space or add more storage.
19056	E	Remote copy failed; operation: %1; error: %2; server: %3, virtual tape ID: %4, target server: %5, replica tape ID: %6.	The operation reports the error specified in the message.	Check the log for related errors and take necessary actions.
19057	E	Remote copy of virtual tape failed because target server %1 could not be connected.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.
19058	E	Remote copy of virtual tape ID %1 to the target server failed because the replica tape was missing.	This is due to some configuration inconsistency issues.	Contact Technical Support.
19059	E	Remote copy of virtual tape ID %1 to the target server failed because the replica tape was invalid.	This may be due to a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable.
19060	E	Remote copy of the virtual tape to the target server failed because the configuration file could not be opened.	The system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
19061	E	Remote copy of the virtual tape to the target server failed because there was not enough memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory.
19063	W	Remote copy of virtual tape ID %1 to the target server was cancelled.	This was most probably triggered by a user.	If this was not triggered by a user, check the system log to identify any related errors and take necessary actions based on the error.

Number	Type	Text	Probable Cause	Suggested Action
19064	E	Remote copy of virtual tape ID %1 to the target server failed because the tape could not be loaded.	This may be due to insufficient system resources or storage issues.	Check the physical device status, device connectivity and switches, and the storage log.
19065	W	Remote copy of virtual tape ID %1 to the target server failed because the tape was in a drive.	The virtual tape cannot be in a tape drive when remote copy is in progress.	Run remote copy when the virtual tape is moved back to the vault or a slot.
19066	E	Remote copy of virtual tape ID %1 to the target server failed because the tape could not be located by the tape emulation module.	This is due to some configuration inconsistency issues.	Contact Technical Support.
19068	E	Replica ID %1 of virtual tape ID %2 failed to be promoted because the target server was busy or there was a communication failure; error: %3.	This may be due to too many pending requests or a network error.	Check connectivity between the primary and replica; check network parameters, including jumbo frame configuration, if applicable, and try again.
19073	E	Remote copy failed because target server %1 was busy.	This may be due to too many pending requests.	Try again later.
19074	E	Replication failed because replica server %1 was busy.	This may be due to too many pending requests.	Try again later.
19075	E	Remote copy of virtual tape ID %1 from %2 to %3 failed due to a version mismatch.	The target server does not have the required version for replication.	Upgrade the target server or select a different target.
19076	E	Replication of virtual tape ID %1 from %2 to %3 failed due to a version mismatch.	The replica server does not have the required version for replication.	Upgrade the target server, select a different target, or stop replication.
19077	E	Remote copy failed because target server %1 did not have enough space.	The target server is running low on disk space.	Check the disk usage and try to reclaim disk space or add more storage.

Number	Type	Text	Probable Cause	Suggested Action
19078	E	Replication failed because replica server %1 did not have enough space.	The replica server is running low on disk space.	Check the disk usage and try to reclaim disk space or add more storage.
19200	E	User %1 failed to get the encryption keys.	This may be due to insufficient system resources.	Ensure there is enough memory and the system disk is healthy. Check the event log messages for more information.
19201	E	User %1 failed to get the encryption key.	This may be due to insufficient system resources.	Ensure there is enough memory and the system disk is healthy. Check the event log messages for more information.
19202	E	User %1 failed to create encryption key %2.	A key with the same name already exists.	Retry the operation using a different name.
19204	E	User %1 failed to delete encryption key %2.	The key may have been deleted.	Wait for the portal to refresh and retry, if necessary.
19206	E	User %1 failed to update information for encryption key %2.	The key may have been deleted.	Retry the operation.
19208	E	User %1 failed to export encryption keys to package %2.	This may be due to insufficient system resources.	Ensure there is enough disk space and the system disk is healthy and the file system is not read-only. Check the event log messages for more information.
19210	E	User %1 failed to get encryption key package information.	The format or contents of the key package are not valid.	Use a valid key package.
19211	E	User %1 failed to import encryption keys from a package.	The format or contents of the key package are not valid.	Use a valid key package.
19381	E	The Deduplication Repository could not be reconfigured using the Flexible method.	The configuration file could not be updated for the repository storage allocation mode.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.

Number	Type	Text	Probable Cause	Suggested Action
31003	E	File %1 failed to be opened.	Either the file is not available or the system is busy and does not have enough resources.	Make sure the file exists and check the system status. You may need to restart server modules.
31004	E	Adding user %1 to the NAS server failed.	The user name already exists in the /etc/passwd system file, nasgrp is missing from /etc/group, or the password file cannot be updated.	Make sure the user name does not already exist. If nasgrp does not exist, add it with the 'groupadd nasgrp' command. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31005	E	Memory allocation failed.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory.
31011	E	%1 failed to unmount.	The NAS resource may still be in use or system resources are insufficient to run the system command to unmount the file system.	Check the system status and retry later. If the problem persists, you may need to reboot.
31012	E	The predefined group for Samba users, nasgrp, is missing.	The predefined NAS group does not exist in /etc/group probably due to a configuration error.	Add the group with the 'groupadd nasgrp' command.
31013	E	%1 mount failed.	This may be due to an invalid mount point, obsolete NAS configuration, or insufficient system resources.	Ensure there is enough disk space and the system disk is healthy and the file system is not read-only. Check the event log messages for more information. If necessary, try to mount the resource manually to see the actual error.
31017	E	A write operation failed on file %1.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
31020	E	File %1 failed to be renamed to file %2.	The file name already exists or the file system is inconsistent or read-only.	Check the file system.

Number	Type	Text	Probable Cause	Suggested Action
31028	E	File %1 failed to be locked.	The file is still in use by another process.	Retry later. If the error persists, you may need to restart the server modules.
31029	E	File %1 failed to be created.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
31030	E	Directory %1 failed to be created.	This may be due to insufficient disk space, system disk failure, or an unhealthy file system.	Ensure there is enough disk space and the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
31031	E	Directory %1 failed to be removed.	Another process may be accessing the directory or the system may have been busy and did not have enough resources.	Check the system status. You may need to restart server modules.
31032	E	Execution of program %1 failed.	This may be due to insufficient system resources or an invalid process state.	Check the system status. You may need to reboot.
31035	E	Group %1 failed to be added to the NAS server.	All the group IDs in the reserved range have been used.	Add a new GID range for NAS Windows clients.
31036	E	User %1 failed to be deleted from the NAS server.	The user is currently logged in and cannot be deleted.	Wait until the user logs off or stop running processes that belong to this user account.
31039	E	File %1 failed to be renamed to file %2.	The file name already exists or the file system is inconsistent or read-only.	Check the file system.
31041	E	A NAS user cannot be created because the maximum number of reserved UIDs was reached.	All the user IDs in the reserved range have been used.	Add a new UID range for NAS Windows clients.

Number	Type	Text	Probable Cause	Suggested Action
31042	E	A NAS group cannot be created because the maximum number of reserved GIDs was reached.	All the group IDs in the reserved range have been used.	Add a new GID range for NAS Windows clients.
31045	E	File /etc/passwd failed to be updated.	This may be due to insufficient system resources or an unhealthy file system.	Check the server memory and file system status. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31047	E	The NAS management module is not running.	A user may have stopped the module or the module had a failure. In a failover configuration, the partner server will try to take over this server.	If the module was not stopped intentionally, check the system log to get more information about the reason for the module failure.
31048	E	NAS resource %1 failed to mount.	This may be due to an invalid mount point, obsolete NAS configuration, or insufficient system resources.	Ensure there is enough disk space and the system disk is healthy and the file system is not read-only. Check the event log messages for more information. If necessary, try to mount the resource manually to see the actual error.
31050	E	NAS resource ID %1 failed to unload during failover.	The device information could not be retrieved from the configuration file probably due to a storage error.	Check the underlying physical storage.
31051	E	NAS resource ID %1 failed to load during failover.	The device information could not be retrieved from the configuration file probably due to a storage error.	Check the underlying physical storage.
31054	E	The server hostname failed to be retrieved.	This is due to a network configuration issue.	Check that the host name returned by the 'uname -a' command corresponds to definitions in the DNS database or the /etc/hosts file.
31056	W	The mount operation was delayed because the file system on NAS resource ID %1 is in an inconsistent state.	This may be due to an ungraceful shutdown or the file system was mounted many times.	If the issue persists, run 'fsck' to check the file system.

Number	Type	Text	Probable Cause	Suggested Action
31058	W	Some NAS resources could not be unmounted.	The resources may still be in use.	You may need to restart the machine.
31060	W	Some NAS resources could not be mounted.	This may be due to insufficient system resources or storage issues.	Check the physical device status, device connectivity and switches, and the storage log.
31061	E	NAS process %1 failed.	This may be due to insufficient system resources.	Check the event log messages for more information and take necessary actions.
31062	E	A read operation failed on file %1.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. You may need to run 'fsck' to check the file system.
31063	E	A write operation failed on file %1.	This may be due to an unhealthy file system or a storage issue.	Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Check the underlying physical storage.
31064	E	An invalid file system type of %1 was retrieved from the configuration file.	The configuration file may be corrupted.	Contact Technical Support.
31066	E	The configuration file cannot be read.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. You may need to run 'fsck' to check the file system.
31067	E	Dynamic configuration file %1 failed to be parsed.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31068	E	Device information is missing from dynamic configuration file %1.	This is due to some inconsistency issues.	Contact Technical Support.
31069	E	The file system format operation on NAS resource %1 failed.	The file system superblock could not be erased probably due to an unhealthy file system or a storage issue.	Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Check the underlying physical storage.
31071	E	The status of file %1 failed to be retrieved.	The 'stat' system command failed to run for the file.	Run the command manually to display errors.

Number	Type	Text	Probable Cause	Suggested Action
31072	E	The SMB native configuration failed to be updated.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31073	E	The NFS native configuration failed to be updated.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31074	E	Configuration file %1 failed to be parsed.	This may be due to insufficient system memory, an unhealthy file system, or the file is corrupted.	Check the server memory and file system status. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system.
31075	E	NAS resource %1 failed to unmount during failback.	The resource may still be in use.	You may need to restart the machine.
31076	C	Due to a storage failure, the partner server had to reboot to resume NAS services.	There was a storage error preventing recovery without a reboot.	No action is required.
31078	E	NAS resources failed to load during failback when processing %1.	The device information could not be retrieved from the configuration file probably due to a storage error.	Check the underlying physical storage.
31079	E	On NAS resource %1, parameter %2 is missing from a file system command.	The file system command requires the parameter, which is not specified.	Specify the parameter in the command associated with the file system.
31086	E	Replication throttling for VITs failed to be set to %1 KB/s.	The configuration file could not be updated probably due to a disk failure or an unhealthy file system.	Ensure the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.

Number	Type	Text	Probable Cause	Suggested Action
32087	E	File system properties failed to be set on directory %1 and its subdirectories.	This may be due to a disk failure or an unhealthy file system.	Ensure the system disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
32090	E	The description of user %1 failed to be changed.	The user may not exist or the NAS user configuration file is missing.	Check the validity of the user name and make sure NAS users and groups are still available.
32091	E	Group %1 failed to be deleted.	The group may not exist or the NAS user configuration file is missing.	Check the validity of the group name and make sure NAS users and groups are still available.
32092	E	The description of group %1 failed to be changed.	The group may not exist or the NAS user configuration file is missing.	Check the validity of the group name and make sure NAS users and groups are still available.
32093	E	Members of group %1 failed to be changed.	The group or the group members may not exist.	Check the validity of the group name and members.
32094	E	Connection to Domain Controller %1 failed.	This may be due to an invalid user/password or a network error.	Validate the domain account and check connectivity.
40003	C	The virtual tape library database on device ID %1 is inconsistent.	The database consistency check failed, probably due to a storage error.	Check the underlying physical storage.
40004	E	The following device is not supported: [%1][%2][%3].	The device could not be found in the list of supported devices.	Contact Technical Support.
40007	E	Unloading a tape from virtual drive ID %1 failed; error code: %2.	The virtual drive is busy and is not responsive.	Take necessary actions based on the error stated in the message.
40009	E	The Move Medium command failed on virtual library ID %1 from source element address %2 to %3.	The tape library may be unavailable, the source element may be missing, or the destination may be occupied.	Check the library status, make sure the source element is available, and the destination element is not occupied.
40018	E	The tape in physical library ID %1, in drive ID %2, with barcode [%3], failed to be loaded.	There is a hardware problem with the physical tape drive.	Check with the tape drive vendor.

Number	Type	Text	Probable Cause	Suggested Action
40020	W	Job %1 was cancelled.	The physical library may have been reset or the related tape was requested by a backup software.	No action is required.
40029	E	There is not enough memory to complete the operation.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
40032	W	Physical tape [%1] is not available to start an auto archive job. The job will be started when the tape becomes available.	The physical tape may have been removed from the slot.	Insert a tape with an appropriate barcode in the physical library and perform an inventory operation.
40036	E	The connection to tape database %1 failed.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
40039	E	The status of physical library ID %1 could not be reported because the 'Read Element Status' SCSI command failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40040	E	Initialization of device %1 failed; error code: %2.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
40043	E	The 'Move Medium' command for physical library ID %1 failed; source address: %2, destination address: %3, error code: %4.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.

Number	Type	Text	Probable Cause	Suggested Action
40044	E	The 'Unload' command for physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40046	E	Writing to physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40047	E	Writing a filemark to physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40048	E	The 'Mode Sense' command for physical drive ID %1 failed; page code: %2, error code: %3.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40049	E	The 'Mode Select' command for physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40050	E	The 'Rewind' command for physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.

Number	Type	Text	Probable Cause	Suggested Action
40051	E	The 'Inquiry' command for physical drive ID %1 failed; error code: %2.	There is a hardware problem with the physical tape library.	Check the event log messages prior to this message for more information and take necessary actions. You may need to check with the library vendor.
40067	E	Physical drive ID %1 failed to be added to tape database %2.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
40074	E	Export job %1 failed; error code: %2, source tape: [%3], destination tape: [%4].	The job reports the error specified in the message.	Check the event log messages prior to this message for more information and take necessary actions.
40076	E	Import job %1 failed; error code: %2, source tape: [%3], destination tape: [%4].	The job reports the error specified in the message.	Check the event log messages prior to this message for more information and take necessary actions.
40077	E	Import job %1 failed due to a duplicate virtual tape barcode; destination tape: [%2].	A virtual tape with the same barcode already exists in the virtual library.	Delete the existing tape or use a different barcode.
40078	E	Import job %1 for a standalone drive failed due to a duplicate virtual tape barcode; destination tape: [%2].	A virtual tape with the same barcode already exists in the virtual library.	Delete the existing tape or use a different barcode.
40084	W	Tape [%1] is blank and could not be exported.	Import/export jobs do not allow blank tapes.	Make sure that cleaning tapes are not used for import/export jobs.
40089	E	A process failed to be created to communicate with an ACSLS/LS server because of insufficient system resources; error: %1.	System resources are running low.	Check the system status, log, and the error stated in the message. You may need to reboot the server.

Number	Type	Text	Probable Cause	Suggested Action
40090	E	The ACSLS/LS communication module failed to be executed; error: %1.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details.
40091	E	%1 failed to establish ACSLS/LS communication; error: %2.	This may be due to a network error or the ACSLS server is not responsive.	Check the ACSLS server status and the connectivity between servers. Take necessary actions based on the error stated in the message.
40092	E	The ACSLS/LS server could not resolve the StorSafe server name; DNS or the /etc/hosts file returned IP address %1 for %2.	This is due to a network configuration issue.	Check that the StorSafe host name returned by the 'uname -a' command corresponds to definitions in the DNS database or the /etc/hosts file.
40093	E	The connection to the ACSLS/LS server using IP address %1 and ACS %2 failed.	This may be due to a network error or the IP returned by host name resolution is not valid.	Check the ACSLS server status and the connectivity between servers. Check that either a DNS server is running or the '/etc/hosts' file contains the correct IP address.
40094	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 (%3) ACS %4.	The ACSLS server did not provide the requested information in time.	Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details and take necessary actions.
40095	E	The ACSLS/LS server failed to mount %1 on physical drive ID %2; error from %3 (%4): %5.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details.
40096	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 (%3) when attempting to mount %4 on drive ID %5.	The ACSLS server did not provide the requested information in time.	Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details and take necessary actions.

Number	Type	Text	Probable Cause	Suggested Action
40097	E	The ACSLS/LS server failed to dismount %1 from drive ID %2; error from %3 (%4): %5.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details.
40098	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 (%3) when attempting to dismount %4 from drive ID %5.	The ACSLS server did not provide the requested information in time.	Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details and take necessary actions.
40099	E	The ACSLS/LS server failed to retrieve drive information in ACS %1; error from %2 (%3): %4.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details.
40100	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 (%3) when attempting to retrieve drive information in ACS %4.	The ACSLS server did not provide the requested information in time.	Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details and take necessary actions.
40101	E	The ACSLS/LS server failed to retrieve volume information in ACS %1 and pool %2; error from %3 (%4): %5.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details.
40102	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 when attempting to retrieve volume information in ACS %3 and pool %4.	The ACSLS server did not provide the requested information in time.	Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details and take necessary actions.
40103	E	The ACSLS/LS server failed to retrieve LSM information in ACS %1; error from %2 (%3): %4.	The ACSLS server reports the error specified in the message.	Take necessary actions based on the error stated in the message. Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details.

Number	Type	Text	Probable Cause	Suggested Action
40104	E	A timeout occurred after %1 minute(s) while waiting for a response from %2 (%3) when attempting to retrieve LSM information in ACS %4.	The ACSLS server did not provide the list of tape drives in time.	Check the ACSLS server's event log '\$ASCSLSHOME/log/acsss_event.log' for details and take necessary actions.
40112	W	Physical tape [%1] is not available in physical library ID %2 to start the tape caching job. The job will be started when the tape becomes available.	The tape is missing from the physical library.	Make sure the tape is in the library.
40114	W	Export jobs are not allowed because tape [%1] in physical library [%2][%3] is used for tape caching.	A wrong physical tape was selected.	Retry with a valid tape.
40118	E	A block compressed by hardware compression failed to be decompressed using software decompression; error code: %1.	The compressed data may not be valid or the compression software had a failure.	You may need to use a hardware compression adapter.
40120	W	Tape [%1] had no data so no export job was submitted for tape stacking.	The tape is empty.	Retry with a valid tape.
40123	E	The tape in library ID %1, drive ID %2, barcode [%3], failed to load because it is a cleaning tape.	A wrong slot was used for the tape operation.	Retry with a valid tape and slot.
40124	E	A write command to the Configuration Repository failed.	The physical device used by the Configuration Repository may have a failure or the system may have been busy and did not have enough resources.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.

Number	Type	Text	Probable Cause	Suggested Action
40130	W	The cache could not be renewed for tape ID %1. Data will be redirected to physical tape [%2].	The space reserved for tapes may not be sufficient.	Increase the space available for tapes.
40132	E	The tape shredding job failed for tape [%1]; error code: %2.	A write operation has failed probably due to a storage device issue.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues, rescan devices to refresh the configuration, and try again.
40136	E	A tape position could not be set by the SCSI SPACE command on physical drive ID %1; error code: %2.	The physical tape or drive may have a failure.	Fix any detected issue and try again.
40137	E	An import/export job failed to be added to the queue because the maximum of %1 jobs was reached; job ID: %2, physical tape: [%3].	The maximum number of jobs has been reached.	Delete some jobs.
40141	E	Tape stacking scan job %1 failed; error code: %2, physical tape: [%3].	The job reports the error specified in the message.	Take necessary actions based on the error stated in the message.
40142	E	Tape stacking import job %1 failed; error code: %2, physical tape: [%3].	The job reports the error specified in the message.	Take necessary actions based on the error stated in the message.
40143	E	Tape stacking export job %1 failed; error code: %2, virtual tape: [%3], physical tape: [%4].	The job reports the error specified in the message.	Take necessary actions based on the error stated in the message.
40148	E	Tape stacking export job %1 failed because there is not enough memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.

Number	Type	Text	Probable Cause	Suggested Action
40152	E	Tape stacking import job %1 failed because there is not enough memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
40153	E	A stacked tape could not be imported because slot %1 in virtual library ID %2 is occupied.	Another virtual tape is in the import destination slot.	Choose another slot.
40163	E	The ACSLS/LS server failed to find a compatible drive to mount tape [%1]; error from %2 (%3): %4.	The ACSLS/LS server reports the error specified in the message.	Take necessary actions based on the error stated in the message.
40164	E	After %1 second(s), %2 server (%3) did not respond to the mount request for physical tape [%4]; the operation timed out.	The ACSLS/LS server may have been busy.	Check the server status and try again.
40165	W	The door was opened on physical library ID %1 %2 %3, which can cause tape operations to stop.	A user opened the library door.	Close the library door.
40167	W	Encryption key %1 does not exist. Decryption and writing operations were disabled for tape ID %2.	The key may have been deleted.	Recreate the key.
40168	W	There is no license for encryption. Decryption and writing operations were disabled for encrypted tape ID %1.	There is no valid license for this option.	Obtain a valid license keycode.
40172	W	The tape shredding job was cancelled for tape [%1], ID %2.	A user requested to stop the operation.	No action is required.

Number	Type	Text	Probable Cause	Suggested Action
40173	E	Physical drive ID %1 [%2][%3] became offline due to a hardware issue; error: %4.	The physical device or the connection to the device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues and rescan devices to refresh the configuration.
40175	W	Virtual tape ID %1 [%2] became read-only due to metadata inconsistencies. %3	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
40178	E	User %1 failed to run patch %2.	Execution of a patch could not be launched.	Check to see if there are further errors reported by the patch that describe the reason.
40183	W	Tape duplication job %1 could not create copy job %2 because there were no more physical tapes with the same barcode [%3] in any library.	Either tapes are missing from the library or they have a different barcode.	Try again using physical tapes with the correct barcodes in the library.
40184	W	Tape duplication job %1 running copy job %2 detected that tape [%3] would not be able to be deleted at the end of export because all related duplication jobs did not complete successfully.	To export tape data to a physical tape, multiple jobs were created but one of them did not complete.	Normally, the tape will be moved during the next export job. If this error persists, check related messages and take necessary actions.
40185	W	Tape duplication job %1 running copy job %2 detected that tape [%3] could not be reclaimed because all physical copies were not created.	To export tape data to a physical tape, multiple jobs were created but some of them did not complete.	Wait until all jobs complete.
40186	W	Tape duplication job %1 detected that tape [%2] could not be reclaimed because a related tape copy job did not complete successfully.	To export tape data to a physical tape, multiple jobs were created but one of them did not complete.	Normally, the tape will be moved during the next export job. If this error persists, check related messages and take necessary actions.

Number	Type	Text	Probable Cause	Suggested Action
40187	W	Tape duplication job %1 detected that tape [%2] could not be moved to the vault because a related tape copy job did not complete successfully.	To export tape data to a physical tape, multiple jobs were created but one of them did not complete.	Normally, the tape will be moved during the next export job. If this error persists, check related messages and take necessary actions.
40188	W	Tape stacking import job %1 completed; virtual tape library ID: %2; source tape: [%3], throughput: %4 MB/min. [%5] tapes were not imported due to duplicate barcodes.	Some tapes are using the same barcodes as other tapes in this library probably because they were previously used in another library.	Try again with different barcodes for virtual tapes to eliminate duplicates.
40189	W	Tape stacking scan job %1 detected that physical tape [%2] is not a stacked tape.	The physical tape header does not indicate that the tape is a stacked tape.	Make sure a stacked tape is used for this operation.
40191	E	The hardware compression card failed; error code: %1, complete code: %2, hardware complete code: %3, tape: [%4], ID [%5].	The compression card reported a failure.	Check the error details and take necessary actions to fix the hardware; you may need to contact the compression card vendor.
40193	E	A query to the %1 server with IP %2 and ACS %3 failed because ACS ID %4 was invalid.	The ACS ID specified in the ACSLS configuration is incorrect.	Use a correct ACS ID in the ACSLS configuration.
40194	E	The portmap/rpcbind daemon, required for ACSLS/LS servers, is not running. Start the daemon and then rescan physical devices.	The daemon may have been inadvertently stopped.	Check and make sure the daemon is running.
40195	W	An export job could not be submitted for tape [%1] because it is being deduplicated.	An attempt was made to manually export a tape while a deduplication operation was in progress.	Run export jobs after deduplication is completed.

Number	Type	Text	Probable Cause	Suggested Action
40197	W	The tape database consistency check detected non-critical errors. Check the system log for more information.	The underlying physical device may have some issues that compromise the integrity of the tape database.	Check the physical device status, device connectivity and switches, and the storage log.
40198	E	The tape database consistency check detected critical errors. Check the system log for more information.	The underlying physical device may have some issues that compromise the integrity of the tape database.	Check the physical device status, device connectivity and switches, and the storage log.
40199	E	The consistency check for compressed data detected an invalid data signature on tape [%1], ID %2.	The data integrity signature in compressed blocks does not match the original signature probably due to some storage issues.	Check the physical device status, device connectivity and switches, and the storage log.
40201	E	Tape ID %1 could not be created because the ID already exists.	This may be due to some inconsistency issues with the tape database.	Try to create the tape again; if the issue persists, contact Technical Support.
40202	W	Tape [%1], ID %2 could not be found.	This may be due to some inconsistency issues with the tape database.	Contact Technical Support.
40203	W	Tape [%1], ID %2 became read-only due to a maximum capacity mismatch.	The tape does not have the expected value for the maximum capacity.	Contact Technical Support.
40204	W	Tape [%1], ID %2 became read-only due to an allocated size mismatch.	The tape does not have the expected allocation size.	Contact Technical Support.
40205	W	Tape [%1], ID %2 became read-only due to a property mismatch for stub tape %3.	The stub tape does not have an expected property.	Contact Technical Support.
40207	W	Virtual tape ID %1 [%2] became read-only due to tape header inconsistencies.	This is due to some inconsistency issues.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
40208	E	ACSLS library ID %1 is offline; ACS: %2, pool: %3. Check the physical library and then perform an inventory to bring the library back online.	The ACSLS library might have some issues.	Check ACSLS logs and take necessary actions.
40209	E	ACSLS drive ID %1 (%2,%3,%4,%5) is offline. Check the physical library and then perform an inventory to bring the drive back online.	The ACSLS library might have some issues.	Check ACSLS logs and take necessary actions.
40210	E	Physical library ID %1 is disabled for maintenance. Enable the library when maintenance completes.	A user has disabled the library for maintenance purposes.	Wait until maintenance is complete.
40211	W	Metadata was rolled back on virtual tape %1 because of a write failure; details: %2.	The write error may be due to a storage issue.	Check the physical device status, device connectivity and switches, and the storage log.
40214	W	Physical tape [%1] was not available to start a tape caching job. The job will be started when the tape becomes available.	The current physical tape position got changed by third party application during tape stacking job running.	Make sure physical drives are not shared with other applications.
40216	E	Tape stacking job %1 did not complete because the current physical tape position (%2) did not match with the number of blocks (%3) transferred to physical tape [%4] in physical drive ID %5. The start position is %6; the expected current position is %7.	This is due to some inconsistency issues.	Contact Technical Support.

Number	Type	Text	Probable Cause	Suggested Action
40220	W	Turbo deduplication was stopped for tape [%1], ID %2, in virtual drive ID %3; this tape will be deduplicated using the post-processing method.	The tape scanner encountered an error.	Check the log files to identify the actual reason and take necessary actions.
40222	W	Inline deduplication was stopped for tape [%1], ID %2, in virtual drive ID %3.	The tape scanner encountered an error.	Check the log files to identify the actual reason and take necessary actions.
40223	W	Turbo deduplication was stopped for tape [%1], ID %2, in virtual drive ID %3.	The tape scanner encountered an error.	Check the log files to identify the actual reason and take necessary actions.
40224	W	Inline deduplication was not used for tape [%1], ID %2, in virtual drive ID %3; this tape will be deduplicated using the post-processing method.	The tape scanner does not support this tape format for inline deduplication or encountered an error.	Check the log files to identify the actual reason and take appropriate actions, if necessary.
40226	W	Tape [%1], ID %2, could not be loaded into virtual drive ID %3.	This is due to some other errors.	Check the system log to identify any related errors and take necessary actions.
40227	E	Tape [%1] could not be imported into virtual library ID %2 from stacked physical tape [%3] by job ID %4.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.
40228	E	The tape import/export job queue could not be initialized on tape database %1.	This may be due to insufficient system resources.	Restart the server modules.
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.
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 and the job will be automatically retried.

Number	Type	Text	Probable Cause	Suggested Action
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 and the job will be automatically retried.
40404	E	Physical tape [%1] cannot be labeled as a stacked tape in physical drive ID %2.	The physical tape has some data and cannot be labeled without overwrite mode.	Use an empty tape or use overwrite mode if you do not want to keep existing data on this tape.
40405	W	Encryption is not activated, resulting in I/O failures for encrypted tape [%1], ID %2.	The encryption activation password was not entered after server startup.	Activate encryption in order to access encrypted virtual tapes.
50000	E	iSCSI target name is missing in login session from initiator %1.	The iSCSI initiator may not be compatible.	Check the iSCSI initiator on the client side.
50002	E	iSCSI login to non-existent target %1 was requested by initiator %2	The iSCSI target does not exist any longer.	Check the iSCSI initiator on the client side and the iSCSI configuration on the server. Remove targets from the configuration if they do not exist.
50003	E	iSCSI CHAP authentication method was rejected in the login request to target %1 from initiator %2.	The CHAP settings are not valid.	Check the iSCSI CHAP secret settings on the server and the client sides.
50115	E	Deduplication cluster %1 could not detect Fibre Channel targets of %2.	This may be due to invalid Fibre Channel zoning or connectivity.	Check the Fibre Channel zoning and make sure the cluster nodes can see each other's target ports.
50204	E	A virtual drive with serial number %1 could not be found on deduplication cluster node %2.	This may be due to a previous redundant node configuration that has not been removed completely.	You need to remove the obsolete configuration. Contact Technical Support for help.
50206	E	Deduplication data disk %1 is not available.	The underlying physical device may have a failure.	Check the physical device status, device connectivity and switches, and the storage log.

Number	Type	Text	Probable Cause	Suggested Action
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.
50514	E	Deduplication Policy %1: The scan process failed to start for tape [%2] with ID %3; reason: %4, server: %5, cluster: %6.	Either the deduplication server is not healthy or it cannot be connected.	Make sure services are up and running on both the local server and the deduplication server. Refer to the error specified in the message and check the log for additional information and take necessary actions. You may need to restart services on this server to release tapes for deduplication.
50516	E	Deduplication Policy %1: The scan process failed for tape [%2] with ID %3; error: %4, server: %5, cluster: %6.	Either the deduplication server is not healthy or it cannot be connected.	Make sure services are up and running on both the local server and the deduplication server. Refer to the error specified in the message and check the log for additional information and take necessary actions.
50538	W	Space reclamation failed to start on node %1; reason: %2.	Space reclamation encountered a network connection issue and reported the error specified in the message.	Check the log for related errors and take necessary actions.
50546	E	Space reclamation failed for tape %1 [%2].	The tape cannot be read due to an issue with the underlying physical device or the connection to the device.	Check the physical device status, device connectivity and switches, and the storage log.
50549	W	Deduplication Policy %1: The policy was not started because the repository on deduplication server %2 is still being loaded.	The deduplication repository is in the startup phase and is not ready yet.	Wait until the server startup phase is complete.

Number	Type	Text	Probable Cause	Suggested Action
50550	E	Deduplication Policy %1: The policy was not started because deduplication server %2 could not be accessed.	This may be due to a network connectivity issue with the deduplication server or the server is not healthy.	Check that the deduplication server can be reached and all modules are up and running.
50551	W	The deduplication policy was not started because the repository on deduplication server %1 is still being loaded.	The deduplication repository is in the startup phase and is not ready yet.	Wait until the server startup phase is complete.
50552	E	The deduplication policy was not started because deduplication server %1 could not be accessed.	This may be due to a network connectivity issue with the deduplication server or the server is not healthy.	Check that the deduplication server can be reached and all modules are up and running.
50553	W	Deduplication Policy %1: Tape [%2] scan process was cancelled by a backup client or a user.	A user requested to stop the operation or the tape was requested by backup software.	No action is required.
50554	E	Space reclamation failed because VIT %1 [%2] has an old format.	The VIT was created by a previous version.	Run the deduplication policy to convert the VIT to the new format.
50557	W	Deduplication Policy %1: Tape [%2] unique data replication was cancelled by a backup client or a user.	A user requested to stop replication or the tape was requested by backup software.	No action is required.
50561	W	Deduplication Policy %1: Tape [%2] unique data replication was cancelled because it was outside of the specified time range for the policy. Replication will be triggered during the next specified time range.	The replication duration was longer than the policy time range.	Extend the policy time range or wait for the next time range.

Number	Type	Text	Probable Cause	Suggested Action
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.
50568	E	The replica tape ID %1 could not be resolved to an LVIT due to the following reason: %2	The replication process for a VIT reports the error specified in the message.	Check the log for related errors and take necessary actions.
50569	W	Deduplication Policy %1: Processing will be stopped for tape [%2]; reason: %3	The deduplication process reports the error specified in the message.	Check the log for related errors and take necessary actions.
50570	W	Deduplication Policy %1: Processing is not allowed at this moment for tape [%2]; reason: %3. Post-processing deduplication will be performed for the tape.	The deduplication process reports the error specified in the message.	Check the log for related errors and take necessary actions.
50573	E	Deduplication Policy %1: Processing failed for tape [%2] in drive ID %3; reason: %4	The deduplication process reports the error specified in the message.	Check the log for related errors and take necessary actions.
50575	E	Deduplication Policy %1: Processing for tape [%2] with ID %3 failed %4 times. The job will not be retried.	Too many errors occurred during tape data deduplication.	Check the system log for related errors and take necessary actions.
50602	E	NAS replication failed due to a network connectivity error.	This may be due to a network connectivity issue.	Check connectivity between the primary and replica; check network parameters.
50603	E	A NAS resource being replicated is full.	The file system is full.	Remove some unnecessary files or expand the resource.
50609	E	The Deduplication Repository is full on the replica server.	The replica server reported that its repository was full.	Run reclamation on the replica server.

Number	Type	Text	Probable Cause	Suggested Action
50612	E	A NAS operation failed because information could not be exchanged between the server and the portal.	The system location that is used temporarily for information exchange cannot be accessed probably due to low system resources.	If this issue persists, you may need to restart the server machine.
50613	E	The NAS server encountered a memory allocation failure.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory.
50615	E	NAS replication failed.	An unexpected error occurred during replication.	Wait for the next scheduled start time or run replication manually.
50617	E	Replication cannot proceed because the replica server is running an older version.	Replication is only allowed from an older version to a newer version.	Make sure you are using compatible versions on both servers.
50631	E	The NAS deduplication configuration failed to be saved; error: %1.	This may be due to insufficient disk space, an unhealthy file system, or an issue with the underlying physical device or device connectivity.	Ensure there is enough disk space and the disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
50653	E	NAS outgoing replication session ID %1 terminated.	Replication was stopped unexpectedly, probably due to services being restarted.	Wait for the next scheduled start time or run replication manually.
50662	E	NAS outgoing replication session failed for file %1.	The stub file could not be read. This may be due to an unhealthy file system or an issue with the underlying physical device or device connectivity.	Check the system log to determine the nature of the error. Fix any file system, hardware, or connectivity issues.
50669	E	NAS reclamation session ID %1 failed.	Due to previous file system errors, the reclamation process stopped.	Check that file systems are mounted and can be browsed. Check related errors and take necessary actions.
50675	E	NAS incoming replication stopped; user: %1, host: %2, session ID: %3.	The server did not receive data from the source probably due to a network issue or an error on the source server.	Check connectivity between the primary and replica; check network parameters and the source server.

Number	Type	Text	Probable Cause	Suggested Action
50682	E	NAS incoming replication failed for file %1.	The server could not reach the source to acknowledge the replication probably due to a network issue.	Check connectivity between the source and replica; check network parameters.
50688	E	NAS integrity check detected a checksum mismatch on file %1; type: %2, recorded: %3, expected: %4.	The deduplicated file did not contain expected data.	Check the system log to see if there are any file system, storage, or connectivity issues. If there are no issues with your environment, Contact Technical Support.
50693	E	NAS integrity check failed; start time: %1, end time: %2, files processed: %3, policy ID: %4.	A stub file could not be read because the related target file was deleted or was incomplete.	Run the integrity check again.
50695	E	NAS integrity check detected some problems; total files: %1, passed: %2, failed: %3, mismatched: %4, skipped: %5, ineligible: %6.	Some deduplicated files did not contain expected data.	Check the system log to see if there are any file system, storage, or connectivity issues. If there are no issues with your environment, Contact Technical Support.
50699	E	NAS integrity check failed to read file %1.	This may be due to an unhealthy file system or an issue with the underlying physical device or device connectivity.	Check the system log to determine the nature of the error. Fix any file system, hardware, or connectivity issues.
50702	E	Scan of tape [%1] in drive %2 on server %3 failed to start; reason: %4.	The deduplication server may have lost access to the StorSafe resources.	Check FC connectivity to ensure that the deduplication cluster can access tape drives and tapes.
50706	E	Scan of tape [%1] in drive %2 failed to create a folder resource; reason: %3.	The folder disk may not be accessible.	Make sure the storage that contains the folder disk is accessible by the deduplication cluster.
50707	E	Scan of tape [%1] in drive %2 failed to store data in the repository; reason: %3.	This may be due to a storage or a connectivity issue.	Check the physical device status, device connectivity and switches, and the storage log. Fix any detected issues.

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, the deduplication server may be unavailable or busy due to heavy I/O, or a storage issue.	Take necessary actions based on the specified reason in the message, which may include checking FC connectivity to ensure that the deduplication cluster can access tape drives and repository storage.
50710	W	Scan of tape [%1] in drive %2 was cancelled.	Scanning was cancelled manually or the tape was requested by backup software.	No action is required.
50714	E	Scan of tape [%1] in drive %2 failed to convert the tape to a VIT after running for %3.	Due to previous errors, the scan process stopped.	Check related errors and take necessary actions.
50719	E	Scan of tape [%1] in drive %2 on server %3 failed to generate the index information.	The deduplication server may have lost access to the StorSafe resources used for the index information or there is not enough storage on the StorSafe server.	Check FC connectivity to ensure that the deduplication cluster can access tape drives and tapes; check that there is enough storage on the StorSafe 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 StorSafe server %4; reason: %5.	The resolver process reports the error specified in the message.	Check the log for related errors and take necessary actions.

Number	Type	Text	Probable Cause	Suggested Action
50803	E	Replication failed to process a data block from VIT %1 using source drive %2 and destination drive %3; reason: %4.	A block of data could not be replicated to the target due to the specified reason in the message.	Take necessary actions based on the specified error.
50805	E	Replication of VIT %1 failed; source drive: %2, destination drive: %3, StorSafe 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.
50907	E	The deduplication client module failed to reach the deduplication server; error: %1.	The deduplication server may not be running or there may be a connectivity issue.	Take necessary actions based on the specified error.
50908	E	The deduplication client module failed to initialize the deduplication server; error: %1.	The deduplication server is not configured or it has already been started.	Check configuration and take necessary actions based on the specified error.
50909	E	The deduplication client module failed to get the deduplication server configuration; error: %1.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.

Number	Type	Text	Probable Cause	Suggested Action
50910	E	The deduplication client module failed to set the deduplication server configuration; error: %1.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50911	E	The deduplication client module failed to get hashed data; error: %1.	This may be due to data inconsistencies.	Contact Technical Support.
50912	E	The deduplication client module failed to get the location of hashed data; error: %1.	This may be due to data inconsistencies.	Contact Technical Support.
50913	E	The deduplication client module failed to create a folder; error: %1.	Folder disks may not be large enough or cannot be accessed.	If this is due to insufficient capacity, run reclamation and add folder disks if necessary. Otherwise, check storage connectivity and the status of the folder disks.
50915	E	The deduplication client module failed to close folder %1; reason: %2.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50916	W	The deduplication client module failed to discard folder %1; reason: %2.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50917	E	The deduplication client module failed to flush data for folder %1; reason: %2.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50919	W	The deduplication client module failed to initialize; error %1.	The deduplication client module reports the error specified in the message.	Take necessary actions based on the error.
50920	E	Data could not be written to the deduplication index; reason: %1.	The index disk may be full or a SCSI write error occurred.	If the index disk is full, run pruning. Otherwise, check storage connectivity and the status of the index disk.
50950	W	The SSD index cache is being rebuilt from the index drive due to an ungraceful shutdown.	This is due to a server failure, an ungraceful shutdown, or failover to the standby server.	Wait for the index cache to complete loading; this may take a while.

Number	Type	Text	Probable Cause	Suggested Action
50952	E	The deduplication server initialization failed; reason: %1.	The deduplication client module could not start successfully due to the specified reason.	Ignore this messages if it occurs when the system is initially configured. Otherwise, take necessary actions based on the error.
50956	E	Deduplication data disk %1 is full.	There is not enough space on the data disk for backup.	Run space 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.
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.

Number	Type	Text	Probable Cause	Suggested Action
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.
51801	E	The NAS option failed to be enabled; error: %1.	The system may have been busy and did not have enough resources.	Check the system status. Retry later. If this error persists, you may need to restart server modules.
51803	E	The NAS option failed to be disabled; error: %1.	The system may have been busy and did not have enough resources.	Check the system status. Retry later. If this error persists, you may need to restart server modules.
52009	W	The file system on device ID %1 could not be detected; the reclamation session will be aborted.	The NAS file system is not mounted; reclamation requires all file systems to be mounted.	Mount the file system or restart NAS services.
52093	E	NAS Fast Copy process with ID %1 stopped; reason: %2.	The process reports the error specified in the message; this may be due to a storage issue.	Check the system log for additional information and take necessary actions based on the error.
52102	W	In network bonding configuration %1, slave interface %2 is down.	This may be due to an issue with the network port or adapter.	Check the network port or adapter related to the device and fix any problems.
52104	W	In network bonding configuration %1, slave interface %2 was removed.	This was most probably triggered by a user.	If this was not triggered by a user, check the system log to identify any related errors and take necessary actions based on the error.

Number	Type	Text	Probable Cause	Suggested Action
52106	E	Network bonding device %1 is down.	This may be due to an issue with the network port, adapter, or connectivity.	Check the network port or adapter related to the device and fix any hardware or connectivity problems.
53000	W	User account %1 used an incorrect password too many times when attempting to query SNMP information. The SNMP module has been shut down temporarily and will automatically restart in %2 minutes.	For security purposes, invalid passwords result in module interruption.	Try again later.
60003	E	The OST server module failed to allocate memory.	The system memory is low.	Check the memory usage of different processes and stop unnecessary processes to free up memory or add memory to the system. You may need to restart server modules.
60004	E	The OST server module failed to write to image ID %1; error: %2.	This may be due to insufficient disk space, an unhealthy file system, or an issue with the underlying physical device or device connectivity.	Ensure there is enough disk space and the disk is healthy. Make sure the file system is not read-only. You may need to run 'fsck' to check the file system. Replace the drive if it is not reliable.
60005	E	The OST server module failed to read from image ID %1; error: %2.	This may be due to an unhealthy file system or an issue with the underlying physical device or device connectivity.	Check the system log to determine the nature of the error. Fix any file system, hardware, or connectivity issues.
61002	E	User %1 failed to enable the OST option.	The system may have been busy and did not have enough resources.	Check the system status. Retry later. If this error persists, you may need to restart server modules.
61004	E	User %1 failed to disable the OST option.	The system may have been busy and did not have enough resources.	Check the system status. Retry later. If this error persists, you may need to restart server modules.

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.

Appendix

This appendix contains information about:

- [Port usage \(updated October 2024\)](#)
- [IP address and netmask update](#)
- [Storage LUN migration](#)
- [FIPS security](#)
- [Shared library environment variable](#)
- [Linux auditing](#)
- [Block device support](#)
- [HTTPS connection for StorSight with an SSL certificate](#)

Port usage (updated October 2024)

The following ports are used for incoming requests. The communication direction of each port listed below is a one-way communication from a source to a destination. The reply to the request is sent back to a dynamic port number. Network firewalls should allow access through these ports for successful communications.

To maintain a high level of security, it is recommended that you disable all unnecessary ports. Although you may temporarily open some ports during initial setup of an appliance, such as the telnet port (23), you should shut them down after your work is complete.

Note: Make sure there are no blocked ports and the loopback device access is open.

The ports are not used unless the associated option is enabled. For FalconStor appliances, the ports marked ● are enabled by default.

Port	Protocol	Open On	Used By	Description
20	TCP/UDP	StorSafe and SIR servers	FTP clients	Standard FTP port used for file data transfer.
21	TCP/UDP	StorSafe and SIR servers	FTP clients	Standard FTP port used for sending commands.
22 ●	TCP	StorSafe and SIR servers	Host clients	Standard Secure Shell (SSH) port used for remote connection to the server.
25	TCP/UDP	StorSight servers	Mail servers	Standard SMTP port used for Email Alerts.

Port	Protocol	Open On	Used By	Description
80 ●	TCP	- StorSight, StorSafe and SIR servers - Clients	Web license server	Standard Internet port used for online registration of license keycodes. License registration material is sent back using HTTP protocol, where a local random port number is used on the server, just like a typical Web-based page. The firewall does not block the reply if the 'established bit' is set to let established traffic in.
80 ●	HTTP	StorSight, StorSafe and SIR servers	Web portal or REST API clients	Standard HTTP port used to access FalconStor servers via web or REST APIs.
111	TCP/UDP	StorSafe servers	NFS v3 clients	NFS and ACSLS* port used for RPC program number mapper. The port is assigned via the SUNRPC protocol. The ports vary, so it is not feasible or convenient to keep checking them and reprogramming a firewall. Most firewalls have a setting to "Enable NFS", upon which they will change the settings if the ports change.
123	UDP	StorSafe and SIR servers	NTP servers	Standard Network Time Protocol (NTP) transport layer used to access external time servers.
137	UDP	StorSafe servers	SMB clients	ipstornmbd NETBIOS Name Service used for SMB protocol.
138	UDP	StorSafe servers	SMB clients	ipstornmbd NETBIOS Datagram Service used for SMB protocol.
139	TCP	StorSafe servers	SMB clients	ipstorsmbd NETBIOS Session Service used for SMB protocol.
161	UDP	StorSafe and SIR servers	SNMP servers	Standard Simple Network Management Protocol (SNMP) port used to query StorSafe, and SIR MIBs
199	TCP	StorSafe and SIR servers	SNMP clients	Standard SNMP multiplexing (SMUX) protocol port used to query Dell OpenManage system MIBs
389	LDAP	StorSight servers	Web portal	Used for Active Directory and LDAP authentication.
443	HTTPS	StorSight, StorSafe, and SIR servers	Web portal or REST API clients	Secure HTTP port used to access FalconStor servers via web or REST APIs.
443 ●	TCP	StorSafe and SIR servers	Failover partner server	HP iLO power control secure port used to power off the failed server in a failover configuration.
445	TCP	StorSafe servers	SMB clients	Microsoft Directory Services for Windows shares using Active Directory.
465	TCP/UDP	StorSight servers	Mail servers	Secure SMTP port used for Email Alerts.

Port	Protocol	Open On	Used By	Description
623 ●	UDP	StorSafe and SIR servers	Failover partner servers	HP iLO and IPMI power control port used for Alert Standard Format (ASF) Remote Management and also used to power off the failed server in a failover configuration.
636	LDAP	StorSight servers	Web portal	Used for Active Directory authentication when TLS is enabled and LDAP authentication when SSL is enabled.
647 ●	UDP/TCP	StorSafe and SIR servers	DHCP servers	Used for DHCP server failover.
705	UDP	StorSafe and SIR servers	SNMP clients	Standard SNMP AgentX port used to query agents, such as Fujitsu ServerView
847 ●	UDP/TCP	StorSafe and SIR servers	DHCP servers	Used for DHCP server failover.
1311	HTTPS	StorSafe and SIR servers	Dell OpenManage servers	HTTPS port used for hardware configuration of DELL servers
2049	TCP/UDP	StorSafe servers	NFS clients	nfsd NFS server (NFS).
3260	TCP	StorSafe and SIR servers	iSCSI client or storages	iSCSI port used for communication between iSCSI clients and the server.
10000	TCP	StorSafe servers	Veritas NetBackup servers	Communication port between Veritas NetBackup and StorSafe for NDMP backup/restore and for the FalconStor OpenStorage Option (FSOST) Copy to Tape feature.
10161	TCP/UDP	StorSafe and SIR servers	SNMP	SNMP communication port.
11576 ●	TCP	StorSafe and SIR servers	• CLI from client • StorSight portal • REST API client • FSOST • Server management console	Secure RPC communication port used for sending requests to the configuration management module on the server.
11577 ●	TCP/UDP	StorSafe servers	Replica servers	Communication port for incoming data replication. This port is only open while replication is being performed.
11579	TCP/UDP	StorSafe and SIR servers	Replication servers	Replication authentication.
11580 ●	TCP	StorSafe servers	Failover servers	Communication port used between a pair of failover servers It is not required for standalone servers.
11581 ●	TCP	StorSafe and SIR servers	StorSight server	Statistics gathering
11582 ●	TCP	StorSafe and SIR servers	CLI from client	Communication port used to send CLI commands to the server.

Port	Protocol	Open On	Used By	Description
11584	TCP	StorSafe servers	Replication servers	Communication port between replication servers for data replication of deduplicated data.
11676	TCP	StorSafe and SIR servers	StorSafe and SIR servers	Communication port for deduplication management and repository health check.
11676 ●	TCP	SIR servers	SIR cluster nodes	Communication port between SIR cluster nodes for deduplication management and repository health check.
11753 ●	TCP	StorSight servers	StorSight portal	Secure RPC communication port used for retrieving server configuration.
11781 ●	TCP	SIR servers	Replication servers	Unencryption port for replication.
11782 ●	TCP	SIR servers	Replication servers	Encryption port for replication.
18651	TCP	StorSafe servers	Replication servers	Communication port between servers for non-encrypted replication.
18652	TCP	StorSafe servers	Replication servers	Communication port between servers for encrypted replication.
18720	TCP	StorSafe servers	FSOST clients	OST-related operations
18750	TCP	StorSafe servers	FSOST	Communication port between FSOST and the StorSafe server.
46373	TCP	StorSafe servers	FSOST and NDMP	Used to transfer an FSOST image to physical tape.

IP address and netmask update

When failover, replication, or deduplication are configured, you cannot change the server IP parameters of the StorSafe server from the StorSight portal. However, you can use the procedure below to make these changes when failover, replication, or deduplication are configured. It also explains how to update an FSOST configuration based on the StorSafe server changes.

This procedure allows you to change the IP address/netmask, the failover heartbeat IP address/netmask, IPMI address/netmask, and replica IP address/netmask.

Environment	This procedure is for StorSafe/StorSafe-S servers with IPMI power control, including failover and replica servers. It is not supported for SIR cluster nodes, for StorSafe servers with a NAS configuration, or for servers with HP iLO.
Pre-operation checklist	If you are changing service or heartbeat IP addresses for failover servers, confirm that the new IP addresses meet the requirements in Failover requirements.
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. Stop all client IO. 2. Stop all jobs, such as replication, deduplication, import, export, migration, recovery, and space reclamation. 3. Suspend failover. 4. Open a command prompt and go to the utility directory. <pre># cd \$ISHOME/ci</pre> 5. Run the <code>ci-cli</code> command with the following syntax to generate the configuration files: <pre># ./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 <code>server</code> 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, <code>ServerIP:PortNum</code>. You can specify up to four servers; all servers must use the same username/password.</pre> The <code>fromip</code> and <code>toip</code> parameters specify the IP address/netmask to change; up to 24 <code>fromip/toip</code> pairs can be specified. The <code>ci-cli</code> command generates a <code>ci-<hostname>.txt</code> file for each StorSafe server and a <code>ci-Localcluster-<hostname>.txt</code> file for each StorSafe-S 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. Reboot the servers after running `ci-updateinfo`.
8. Re-add the servers with the new IP addresses into the StorSight portal.
9. Resume failover.
10. Suspend each deduplication policy with replication enabled and then resume it.

Examples There are three servers, a failover pair and a replica server:

- Local failover pair:
 - SUMMER21-14 (StorSafe-S)
 - WINTER21-15 (StorSafe)
- Remote replication target server:
 - H21-70 (StorSafe-S)

Scenario 1: IP addresses are to be changed for all three servers:

Servers	Current IP Parameters	New IP Parameters
SUMMER21-14	Service IP (eth0:0): 172.22.21.14 Heartbeat IP (eth0): 172.22.20.214 IPMI IP: 172.22.21.114 IPMI Netmask: 255.255.252.0	Service IP (eth0:0): 172.22. 20 .14 Heartbeat IP (eth0): 172.22. 22 .214 IPMI IP: 172.22.21.214 IPMI Netmask: 255.255. 255 .0
WINTER21-15	Service IP (eth0:0): 172.22.21.15 Heartbeat IP (eth0): 172.22.20.215 IPMI IP: 172.22.21.115 IPMI Netmask: 255.255.252.0	Service IP (eth0:0): 172.22. 20 .15 Heartbeat IP (eth0): 172.22. 22 .215 IPMI IP: 172.22.21.215 IPMI Netmask: 255.255. 255 .0
H21-70	IP: 172.22.21.70	IP: 172.22. 20 .70

Start the procedure on SUMMER21-14:

```
# cd $ISHOME/ci
# ./ci-cli --username=root --password=abc123
--server=172.22.21.14
--server=172.22.21.15
--server=172.22.21.70
--fromip=172.22.21.14 --toip=172.22.20.14
--fromip=172.22.21.15 --toip=172.22.20.15
--fromip=172.22.20.214 --toip=172.22.22.214
--fromip=172.22.20.215 --toip=172.22.22.215
--fromip=172.22.21.114/255.255.252.0 --toip=172.22.21.114/255.255.255.0
--fromip=172.22.21.115/255.255.252.0 --toip=172.22.21.115/255.255.255.0
--fromip=172.22.21.70 --toip=172.22.20.70
```

This command generates the following configuration files for each server:

- ci_Localcluster-SUMMER21-14.txt
- ci_WINTER21-15.txt
- ci_Localcluster-H21-70.txt:

Run the following commands on SUMMER21-14 to apply changes to each server:

```
# ./ci-updateinfo ci-Localcluster-SUMMER21-14.txt
# ./ci-updateinfo ci-WINTER21-15.txt
# ./ci-updateinfo ci-Localcluster-H21-70.txt
```

Reboot all servers.

Scenario 2: Only the IP address of the replica server is to be changed:

Server	Current IP Parameters	New IP Parameters
H21-70	IP: 172.22.21.70	IP: 172.22.20.70

Start the procedure on H21-70:

```
# cd $ISHOME/ci
# ./ci-cli --username=root --password=abc123
--server=172.22.21.70
--fromip=172.22.21.70 --toip=172.22.20.70
```

This command generates the following configuration file:

- ci_Localcluster-H21-70.txt

Run the following command on H21-70 to apply the change:

```
# ./ci-updateinfo ci-Localcluster-H21-70.txt
```

Run the following command on SUMMER21-14 to update the replication target IP address on failover servers:

```
# cd $ISHOME/ci
# ./ci-cli --username=root --password=abc123
--server=172.22.21.14
--server=172.22.21.15
--fromip=172.22.21.70 --toip=172.22.20.70
```

This command generates the following configuration files for each server:

- ci_Localcluster-SUMMER21-14.txt
- ci_WINTER21-15.txt

Run the following commands on SUMMER21-14 to apply changes to each server:

```
# ./ci-updateinfo ci-Localcluster-SUMMER21-14.txt
# ./ci-updateinfo ci-WINTER21-15.txt
```

Reboot all servers.

FSOST configuration

Follow the procedure below to update an FSOST configuration based on the StorSafe server IP address changes:

1. Update the hosts file for all NetBackup media servers with the new IP parameters.

The default location for the host file is:

- **Windows:** c:\WINDOWS\system32\driver\etc\hosts
- **Linux:** /etc/hosts

2. Update the FSOST configuration file for all NetBackup media servers.

The file can be found in the following locations:

- **NetBackup media servers running on Windows:**
C:\Program Files\VERITAS\NetBackup\bin\ost-plugins\fsost\fsost.cfg
- **NetBackup media servers running on Linux:**
/etc/fsost.conf
- **NetBackup media servers running on a NetBackup appliance:**
/usr/open/ostconf/fsost.conf

3. Rescan the LSU via the FSOST command on the source and target NetBackup media servers.

```
#fsost rescanlsu -storage_server VTL-Source -u root -p *****  
#fsost rescanlsu -storage_server VTL-Target -u root -p *****
```


Storage LUN migration

StorSafe 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 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.
 - StorSafe 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, make sure servers are not in a failover state and then run the following clean-up command before re-running the migration job; in case of failover, fix the error condition to bring back servers to a normal state before running the clean-up command:

```
lmclean.sh
```

If errors occur, rescan physical devices from the StorSight portal and re-run the migration job.

FIPS security

StorSafe and SIR can be configured to be Federal Information Processing Standard (FIPS) compliant. FIPS is a Federal government standard for security and interoperability and is useful for government organizations that have FIPS certification as a product requirement.

The option is enabled when StorSafe/SIR 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 StorSafe servers.

Note that because block devices are not shareable, there is no failover support for downstream block devices. In addition, block devices cannot be used on a deduplication server.

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 portal.

Note: You cannot change the hostname if you are using block devices. If you do, all block devices claimed by StorSafe will be marked offline and seen as foreign devices.

HTTPS connection for StorSight with an SSL certificate

A Secure Sockets Layer (SSL) certificate is required to use the encrypted HTTPS protocol instead of the clear text protocol HTTP on StorSight. You can get an SSL certificate from a Certificate Authority (CA) entity or use a self-signed key.

To generate a self-signed SSL certificate, run the `keytool` command and follow the prompts to enter information for the certificate. For example:

```
# keytool -genkeypair -alias storsight -keyalg RSA -keysize 2048 -
keystore /usr/local/apache-tomcat/conf/keystore.jks -validity 365
```

Follow steps below on the StorSight server to configure HTTPS:

1. Edit the Apache Tomcat server configuration file:

```
# vim /usr/local/apache-tomcat/conf/server.xml
```

2. Add the following configuration section:

```
<Connector port="443"
protocol="org.apache.coyote.http11.Http11NioProtocol"
maxThreads="150" SSLEnabled="true">
  <SSLHostConfig>
    <Certificate certificateKeystoreFile="/usr/local/apache-tomcat/
conf/keystore.jks" type="RSA"
certificateKeystorePassword="keystore_password"/>
  </SSLHostConfig>
</Connector>

<Connector port="80"
protocol="org.apache.coyote.http11.Http11NioProtocol"
redirectPort="443" />
```

3. Restart the Tomcat Service to apply the new configuration:

```
# freestor restart tomcat
```

4. Adjust your firewall settings to ensure that it allows traffic on port 443, which is used for HTTPS:

```
# sudo firewall-cmd --permanent --add-port=443/tcp
# sudo firewall-cmd --reload
```

Index

A

- ACSLS 136
 - Add/remove tapes 139
 - Configure 137
 - Configure with failover 138
 - Eject policy 138
 - Filter tapes 139
 - Hardware configuration 136
 - Replace ACSLS physical drive 140
- Active Directory 67, 68
 - Self-signed certificate 74
- Activity Log
 - Configure purge policy 370
- AD/LDAP settings 35
- Adapters
 - Adapters page 97
 - Enable/disable FC target mode 99
 - Icons 98
 - Monitor 448
 - Rename 99
 - Rescan 98
- Administration 39
 - AD/LDAP 67
 - Alert notifications 75
 - Customers 54
 - License usage prediction 80
 - Servers 48
 - StorSight 39
 - Automatic logout 46
 - HTTP proxy 45
 - License 40
 - Multi-factor authentication 42
 - Patches 44
 - Restart/reboot 46
 - Server message 46
 - X-ray 44
 - Users 62
- Administrator
 - Definition 62
 - StorSight 62
- Alarm policies
 - Command line 584
 - Get 584
 - Remove 585
 - Set 584
- Alerts 34, 418
 - Acknowledge 421
 - All 421
 - Custom scripts 78
 - Delete 421
 - Descriptions 419
 - Email settings 75
 - Global settings 79
 - Manage 421
 - Notification 75
 - SMS settings 77
 - Unacknowledge 421
- Alerts widget 27
- Analytics 34, 380
- Architecture 18
- Associate StorSafe server with an SIR server 335
- Attention required
 - Command line
 - Get 603
- Audit log 34, 424
- Auditing 743
- Authentication methods 35, 67
 - Reset 61
- Auto Archive 320
- Auto archive 109
- Auto replication 109, 241
- Auto Save Config 360
- Automated Tape Caching
 - Command line
 - Get tape caching info 561
 - Migrate virtual tapes 560
 - Reclaim disk space 560
 - Renew cache 560
 - Set 557
 - Sync physical tapes 559
- Autopathing 93

B

- Backups
 - Troubleshooting 643
- Best practices 625
 - Deduplication repository sizing 630
 - Deduplication system sizing 631
 - Index/folder disk sizing 630
 - Memory sizing 632
 - Protect/recover StorSight server 626
 - StorSight deployment 625
 - StorSight security considerations 625
 - StorSight storage considerations 625

Block device 744
Configuration 744

C

Capacity

Servers 49
Trend report 396

Capacity prediction 34, 422

Chargeback 426

Report 427
Settings 426

Clients

Add Fibre Channel 151
Add iSCSI 149
Agents 19
Assign virtual tape libraries and drives 154

Command line

Add FC client 511
Add iSCSI client 513
Add NFS 595
Add SMB 594
Assign iSCSI target 514
Assign share to NFS client 592
Assign share to SMB client 591
Create iSCSI target 514
Delete 512
Delete iSCSI client initiators 515
Delete iSCSI target 515
Delete NFS 595
Delete SMB 594
Get iSCSI client initiators 515
Properties 513
Rename 512
Set Security Mode 596
Unassign share from NFS client 593
Unassign share from SMB client 592
Virtual device list 511

Inventory 383

Manage 154

Multipathing 156

Troubleshooting 638, 639, 640

Unassign virtual tape libraries and drives 157

Update properties 157

Cloud Tier page 259

Icons 260

Cloud Tiering 111

Add object storage account 249
Configuration 247
Jobs 245

Object storage provider page 248

Object storage providers 246

Cloud tiering 244

Policies 244

Command line 491

Alarm policies 584
Automated tape caching 557
Clients 511
Commands 493
Common arguments 492
Data encryption 550
Deduplication 535
Failover 583
Hosted Backup 582
Import/Export 562
Information 598
Licensing 495
Login/logout 493
Mirroring 579
NAS devices 586
NAS shares 589
Object Storage 570
Physical devices 497
Physical libraries and drives 553
Physical tapes 554
Replication 542
Reports 604
Server info 494
Server licensing 495
SMB/NFS clients 594
Support utilities 602
Usage 491
Users 516
Utilities 600
Virtual devices 507
Virtual library and drives 519
Virtual tapes 527

Configuration

StorSafe 81

Configuration repository 346

Create 346
Manage 346
Reconfigure 346

CSV

Export 37

Customers 54

Add 60
Assign resources 60
Assign server 60

- Customers page 60
- Delete 61
- Exclusive mode 56
- Inventory 385
- Reset authentication 61
- Shared mode 56
- Synchronize users 61
- Unassign resources 61
- Unassign server 61
- Customize
 - Dashboard 26
- D**
- Dashboard 25
 - Customize 26
- Data encryption
 - Command line
 - Activate data encryption 551
 - Change encryption password 551
 - Enable 550
 - Enable virtual tape encryption 550
 - Get data encryption information 551
- Deduplication
 - Activity page 223
 - Classic storage configuration 351
 - Command line
 - Add policy 536
 - Add tape to policy 539
 - Delete policy 539
 - Get NAS statistics 598
 - Get tape activity 541
 - List policies 535
 - Remove tape from policy 540
 - Start policy 536
 - Stop policy 536
 - Deduplication jobs page 373
 - Deduplication methods 210
 - Disable 336
 - Discover data disks 336
 - Enable 335
 - Flexible storage configuration 350, 357
 - Index pruning 354
 - Inline 210
 - Memory 393
 - Methods 210
 - NAS
 - Deduplication overview 209
 - Post-processing 210
 - Ratio 220, 436
 - Replication 224, 438
 - Replication targets 229
 - Add 229
 - Run History page 225
 - Space reclamation 354
 - Tape History page 227
 - Tapes page 219
 - Turbo 210
 - Virtual tapes 208
- Deduplication Activity
 - Page 223
- Deduplication cluster
 - Start nodes 206
- Deduplication cluster failover
 - Behavior 201
 - Change power control password 206
 - Configure 204
 - Remove 207
 - Requirements 202
 - Status 206
 - Suspend/resume 207
 - Terminology 200
 - Triggers 201
- Deduplication jobs
 - Deduplication jobs page 373
 - Icons 374
 - Manage 374
 - Page 373
- Deduplication policies 208
 - Add tapes 221
 - Create 212
 - Deduplication Policies page 211
 - Delete 218
 - Icons 212
 - Manage 217
 - Priority 213
 - Remove tapes 221
 - Replicate 222
 - Replication modes 215
 - Replication options 216
 - Run tapes 222
 - Triggers 214
 - Update 217
- Deduplication repository 350
 - Add deduplication data storage 353
 - Add deduplication index/folder disks 354
 - Delete 357
 - Encryption 328
 - Expand 353

- Manage 353
- Mirror 357
- Reclaim space 354
- Set space reclamation policy 355
- Start index pruning 356
- Start space reclamation 355
- Switch to flexible mode 357
- View device physical layout 357
- Deduplication Usage widget 431
- Direct access tapes 134, 144
- Direct link tape 317
- Disks
 - Troubleshooting 637
- Dissociate StorSafe server from a deduplication cluster 336
- Domain 54
- Drives
 - Drives page 113, 141
- Drives page
 - Icons 114

E

- Email
 - Format 76
 - Settings 75, 410, 414
 - Test settings 75
- Email/SMS settings
 - Administration 36
- Encryption 328, 332, 371
 - Activate 331
 - Change encryption activation password 331
 - Deduplication repository 328
 - Enable virtual tape 330
 - Encryption Keys page 332
 - Encryption page 329
 - Icons 330
 - Keys
 - Create 333
 - Manage encryption keys 333
 - Virtual tape 328
- Encryption keys 332
 - Create 333
 - Delete 334
 - Icons 332
 - Manage 333
 - Update 334
- Encryption support 371
- Error codes 654
- Event Log

- Command line 603
- Event log 338
- Export 37
 - Troubleshooting 640
- Export History
 - Icons 146
- Export stacked tape
 - Command line 566
- Export to tape
 - Command line 565

F

- Failover
 - Auto recovery 197
 - Backup application server configuration 189
 - Behavior 182, 201
 - Change power control 197, 206
 - Command line
 - Get status 583
 - Configure 190, 204
 - Heartbeat interval 197
 - Port swapping 194
 - Power control IP address 188
 - Recovery 197, 198, 199
 - Remove 199, 207
 - Report 391
 - Requirements 187, 202
 - Fibre Channel requirements 188
 - General requirements 187
 - iSCSI requirements 188
 - Network requirements 188
 - Storage requirements 189
 - Resume backups after failover/failback 185
 - Self-checking interval 197
 - SIR cluster 200
 - Status 196, 206
 - Status report 391
 - StorSafe server 178
 - Suspend/resume 198, 207
 - Takeover 198
 - Terminology 181, 200
 - Triggers 179, 201
 - Update configuration 197
- FalconStor Deduplication Repository 19
- FalconStor Management Server 18
- FalconStor StorGuard Server 18
- FalconStor StorSafe Server 19
- FalconStor summary widget 27
- Fast copy 267

- Command line 600
- FC/iSCSI targets
 - Performance 449
- Fibre Channel
 - Add clients 151
 - Fabric 158
 - Point-to-Point 158
- Fibre Channel client
 - Add 151
- Fibre Channel Target Mode 158
 - Data rate 164
 - fshba.conf 162
 - Hardware configuration
 - Server 160
 - Link speed 162
 - Persistent binding 162
 - Ports 160
 - QLogic driver 160
 - RSCN 159
 - Switches 159
 - Target port binding 162
 - WWNN 165
 - Zoning 161
- Fibre Channel Target mode
 - Enable 359
- Filter 38
- FIPS compliance 741
- Firewall 731
- Folders 279
- Folders page
 - Icons 280
- FSOST
 - IP address change 738

G

- General rules 411
 - Create 412
 - Delete 415, 416
 - Enable 416
 - Global time exclusions 416
 - Network error 412
 - System Log errors 413
 - Update 415, 416
- Glossary 20

H

- HBA
 - Multi-ID
 - Troubleshooting 640

- Troubleshooting 636
- Host
 - Agents 19
- Host clients
 - Icons 148
 - Monitor 445
 - Performance 445
 - Performance drill down 446
- Hosted Backup 367
- Hosted backup
 - Command line
 - Disable 582
 - Enable 582
- Hostname 345
- HP iLO 192
- HTTP 732
- HTTPS 732

I

- IBM System i configuration 451
- Icons
 - Adapters page 98
 - Administrators 22
 - Cloud Tier page 260
 - Deduplication jobs page 374
 - Deduplication policies page 212
 - Deduplication tapes page 221
 - Drives page 114
 - Encryption Keys page 332
 - Encryption page 330
 - Export History page 146
 - Folders page 280
 - Host Clients page 148
 - Import/export jobs page 378
 - Integrity Check page 309
 - Libraries page 103
 - NAS Replication Configuration page 298
 - NAS Replication Policies page 305
 - NAS Resources page 263
 - NFS Clients page 295
 - Object storage page 248
 - Physical devices page 91
 - Physical tape drives page 142
 - Physical tape libraries page 132
 - Replication jobs page 376
 - Server 33
 - Shares page 269
 - SMB Clients page 283
 - SMB Groups page 292

- Storage pools page 87
- Storage web links 101
- Tapes page 124, 221
- User access 22
- Vault page 118
- Import stacked tape
 - Command line 563
- Import tape
 - Command line 562
 - Copy 134, 144
 - Direct access 134, 144
 - Recycle 134, 144
 - Stacking 135, 144
- Import/export
 - Command line
 - Cancel jobs 569
 - Delete jobs 568
 - Job status 562
 - Restart jobs 568
 - Resume jobs 568
 - Suspend jobs 569
 - Troubleshooting 647
- Import/export jobs
 - Icons 378
 - Import/export jobs page 377
 - Manage 378
 - Report 379
 - Retry parameters 378
- Index pruning 354
- Inline deduplication 210
- Integrity Check
 - Page icons 309
- Integrity checking
 - Configure 310
 - Disable 312
 - Refresh status 312
 - Start/suspend 312
 - Update 312
- Introduction 17
- Inventory
 - Command line 553
- IP address
 - Change 735
 - FSOST 738
- IPMI 192
- iSCSI
 - Add clients 149
- iSCSI client
 - Add 149
- iSCSI initiator
 - Rescan 149
- iSCSI target
 - Create 150
- iSCSI Target Mode
 - AIX
 - Configuration 176
 - Requirements 176
 - Run a backup 177
 - IBM i
 - Configuration 172
 - Requirements 172
 - Initiators 167
 - Linux
 - Configuration 170
 - Requirements 170
 - Targets 167
 - Users 167
 - Windows
 - Configuration 168
 - Requirements 168
- iSCSI target mode
 - Enable 359
- iSCSI users 358
- J**
 - Jobs
 - Monitor 440
- K**
 - Keycodes
 - Command line
 - Add 495
 - Get information 495
 - Register 496
 - Remove 495
 - Offline registration 635
 - Troubleshooting 635
- L**
 - LD_LIBRARY_PATH
 - Set 742
 - LDAP 67, 71
 - Self-signed certificate 74
 - Libraries
 - Cloud Tiering 111
 - Libraries page
 - Icons 103
 - License

- Renewal 41
- StorSight 40
- License keycodes
 - Offline registration 635
 - Troubleshooting 635
- Licensing
 - Command line
 - Add keycode 495
 - Get keycode information 495
 - Register keycode 496
 - Remove keycode 495
- Linux
 - Auditing 743
- Local users 358
- Lock user access 47
- Logical resources
 - Troubleshooting 637
- LUN reservation
 - Command line
 - Change reservation 500

M

- Manage page 32
- Manage patches 339
- Map share
 - Windows/SMB client 275
- Message
 - Lock user access 47
- Message of the day 46
- Migration
 - Storage devices 739
- Mirroring
 - Command line
 - Create 579
 - Remove 580
 - Status 579
 - Swap 580
 - Sync 581
- Monitor 428
 - Client view page 450
 - Overview page 428
 - Server view page 429
 - Adapters 448
 - Deduplication 436
 - FC/iSCSI Targets 449
 - Host clients 445
 - Jobs 440
 - NAS 442
 - Overview 429

- Overview tab
 - Server overview 429
- Physical devices 447
- Physical resources 447
- Physical tape drives 448
- Physical tape libraries 447
- Storage pools 447
- Summary tab
 - Storage usage 430
 - Usage history 430
- Virtual tape libraries 432
- Monitor page 31
 - Client view 31
 - Overview 31
 - Server view 31
- Mount share
 - NFS client 276
- Multi-ID
 - HBA
 - Troubleshooting 640
- Multipathing 156
- Multi-tenancy 54
 - Configuration 56
 - Example 57
 - Exclusive mode 56
 - Rules and restrictions 59
 - Shared mode 56

N

- NAS 262
 - ACLs
 - Troubleshooting 650
 - Authentication 284
 - Client mount options 276
 - Command line
 - Create device 586
 - Create share 589
 - Delete device 588
 - Delete share 591
 - Expand device 587
 - Format resource 587
 - Get deduplication statistics 598
 - Get fast copy information 601
 - Get folder details 599
 - Get realtime performance information 598
 - Start fast copy 600
 - Stop fast copy 600
 - Configure integrity checking 310
 - Create NFS client 295

- Create SMB group 292
- Create SMB user 283
- Delete NFS client 296
- Delete SMB group 293
- Delete SMB user 284
- Disable integrity checking 312
- Domain controller authentication issue 650
- Domain mode 284
- Enable 371
- File access issues 652
- Folders 279
- Integrity check 309
- Map/mount shares 275, 276
- NFS clients 294
- Permission denied to folders 652
- Refresh status of integrity checking 312
- Replication 297
- Replication policies 304
- Resources 262
- Security mode 284
- Set authentication mode 284
- Set SMB options 288
- Set Windows ACLs 288
- Shares 268
- SMB clients 282
- SMB connection status 290
- SMB groups 291
- Start/suspend integrity checking 312
- Sync SMB Clients 287
- Troubleshooting 648
- Truncated folder names 652
- Update domain controller 287
- Update integrity checking 312
- Update NFS client 296
- Update share access rights 290, 293, 296
- Update SMB group 293
- Update SMB user 284
- User mode 284
- Users/groups
 - Troubleshooting 650
- NAS authentication
 - Domain mode 284
 - User mode 284
- NAS folder
 - Display statistics 281
 - Exclude from deduplication 280
 - Include for deduplication 280
 - Manage 280
 - Set properties 281
- NAS replication
 - Configure incoming 302
 - Create policy 305
 - Delete policy 307
 - Disable outgoing replication 301
 - Edit outgoing replication 300
 - Enable outgoing 298
 - Purge replica files 302
 - Refresh session status 303
 - Set outgoing replication throttle 301
 - Start replication 307
 - Suspend/resume replication schedule 308
 - Synchronize files 307
 - Update policy 307
 - View purge results 303
- NAS resources
 - Check file system 266
 - Create 264
 - Delete 265
 - Display info 267
 - Exclude from deduplication 266
 - Exclude from replication 266
 - Expand 265
 - Fast copy 267
 - Format 266
 - Include for deduplication 266
 - Include for replication 266
 - Inventory 385
 - Mount 265
 - Troubleshooting 648
 - Unmount 265
- NAS resources page 262
 - Icons 263
- NAS share
 - Create 270
 - Delete 272
 - Display statistics 275
 - Exclude from deduplication 272
 - Exclude from replication 273
 - Include for deduplication 272
 - Include for replication 273
 - Set properties 274
 - Show files 275
 - Update 272
- NAS utilities
 - Command line 600
- NDMP 368
- NDMP backup
 - Configuration 368

- Troubleshooting 643
- Netmask
 - Change 735
- Network bonding 342
- Network Settings 340
- NFS client 294
 - Create 295
 - Delete 296
 - Mount share 276
 - Update 296
- NFS Clients
 - Page icons 295
- NFS user
 - Update share access rights 296
- Notification
 - Portal 46
- Notification message
 - Global 46
 - Server specific 47
- NPIV 160
- NTP
 - Time service 340
- O**
- Object storage 248
 - Add account 249
 - Command line 570
 - Add account 570
 - Convert migrated tape to stub tape 577
 - Delete 575
 - Display information 575
 - Get job status 577
 - Manage jobs 578
 - Modify information 573
 - Recover tape from object storage 576
 - Start object migration 576
 - Delete accounts 258
 - Icons 248
 - Update accounts 258
- Object storage account
 - Amazon Web Services (AWS) S3 249
 - Generic S3 256
 - Google 255
 - Hitachi Content Platform (HCP) 251
 - IBM Cloud Object Storage (COS) 253
 - Microsoft Azure 250
 - Wasabi 254
- Offline tapes
 - Troubleshooting 637
- OST
 - Settings 370
- P**
- Passwords
 - Strong 372
 - StorSight 43
- Patches
 - Apply
 - StorSight 44
 - Manage 339
 - Roll back
 - StorSight 44
 - StorSafe 339
- PDF
 - Export 37
- Performance
 - Clients 396, 445
 - FC/iSCSI targets 449
 - NAS resource 396
 - Physical devices 396, 447
 - Virtual tape library 396
- Persistent binding 162
- Physical device
 - Command line
 - Delete 501
 - Get adapter information 498
 - Get device information 497
 - Prepare disk 499
 - Rename 500
 - Rescan 498
 - Restore path 501
 - Storage allocation 506
 - Storage pool
 - Add physical devices 503
 - Assign ACL 505
 - Create 502
 - Delete 503
 - Get list 504
 - Get users 517
 - Remove physical devices 504
 - Rename 505
 - Unassign ACL 505
- Physical devices 90
 - Autopathing 93
 - Change device category 95, 96
 - Change reservation 96
 - Icons 91
 - Inventory 382

- Monitor 447
 - Performance 447
 - Physical devices page 90
 - Prepare 95
 - Rename 94
 - Rescan 92
 - Reserve 96
 - Test 96
 - Unassign 95
 - Physical drive
 - Command line
 - Enable/disable 553
 - Physical library
 - Command line
 - Enable/disable 553
 - Inventory 553
 - Physical resources 84
 - Troubleshooting 636
 - Physical tape
 - Command line
 - Eject 555
 - Move 554
 - Physical tape list 558
 - Prepare for stacking 555
 - Scan 567
 - Scan stacked tape 556
 - Physical tape drives
 - Disable 142
 - Enable 142
 - Inventory 384
 - Monitor 448
 - Physical tape drives page 141
 - Prepare for stacking 142
 - Retrieve barcode 143
 - Scan for stacked tapes 143
 - Physical tape drives page
 - Icons 142
 - Physical tape libraries 131
 - ACSL 136
 - Assign 133
 - Disable 132
 - Enable 132
 - Import 143
 - Import a tape in another virtual tape library 135
 - Import tapes 134
 - Inventory 132, 384
 - Libraries page 131
 - Monitor 447
 - Move tapes 133
 - Reset 133
 - Scan for stacked tapes 133
 - Unassign 134
 - Physical tape libraries page
 - Icons 132
 - Physical tapes 145
 - Export history 146
 - Tapes page 145
 - Port
 - Status 33
 - Port swapping 194
 - Port usage 731
 - Portal 24
 - Arrange columns 38
 - Filter columns 38
 - Global functions 37
 - Lock user access 47
 - Notification messages 46
 - Save/reset layout 37, 38
 - Search 38
 - Sort columns 37
 - Power control management 192
 - Power control options 196
 - Protect
 - Server configuration 360
- ## R
- Reboot 337
 - Reclamation
 - Command line 535
 - Recover server configuration 362
 - Redundant node failover
 - WWNN 165
 - Refresh 37
 - Remote export 324
 - Replica device 123, 236
 - Replication
 - Command line
 - Create replica 542
 - Demote in test mode 548
 - Get properties 546
 - Get status 547
 - Network throttling 549
 - Promote in test mode 548
 - Promote replica 543
 - Remove 544
 - Resume 545
 - Set properties 545
 - Start 547

- Stop 547
- Suspend 544
- Deduplicated tapes 232
 - Access data on tape 235
 - Configure 235
 - FVIT 233
 - LVIT 233
 - Requirements 234
- Global throttle speed 365
- Jobs page 375
- NAS 297
- NAS policies 304
- Promote replica 129, 239
- Test mode 239
- Troubleshooting 644
- Virtual tapes 232, 236
 - Auto replication 241
 - Configure 237
 - Delete 240
 - Local 236
 - Primary tape 236
 - Remote 236
 - Replica device 236
 - Requirements 236
 - Setup 237
 - Troubleshooting 644
 - Update 240
- Replication Configuration
 - Page icons 298
- Replication jobs
 - Icons 376
 - Manage 376
 - Replication jobs page 375
- Replication Policies page
 - Icons 305
- Replication targets
 - Deduplication 229
 - Add 229
- Reports 381
 - Capacity Management 385
 - Deduplication Resource 386
 - NAS SMB Shares 387
 - Server Inventory 386
 - Storage Pool 386
- Command line 604
- Create your own 399
- Custom 399
 - Capacity 400
 - Fields 400
 - Load 402
 - Performance 403
- Deduplication Activity 224
- Deduplication Policy Tapes 222
- Deduplication Run History 226
- Deduplication Tape History 228
- Import/Export Jobs 379
- Inventory 382
 - Clients 383
 - Customer Inventory 385
 - NAS Resources 385
 - Physical Devices 382
 - Physical Tape Drives 384
 - Physical Tape Libraries 384
 - Server 382
 - Virtual Tape Drives 384
 - Virtual Tape Libraries 383
- Performance Monitoring 388
 - Client Performance 389
 - Fibre Channel Target Performance 389
 - iSCSI Target Performance 390
 - NAS Resource Performance 389
 - Physical Device Performance 388
 - Server Throughput 390
 - Virtual Tape Drive Performance 389
 - Virtual Tape Library Performance 388
- Schedule 407
- Server Load 393
 - CPU Load 393
 - Memory Usage 393
 - Network Throughput 394
- Service Status 391
 - Deduplicated Tape Replication History 391
 - Failover 391
 - NAS Deduplication and Replication 391
 - NAS Replication Sessions 392
 - Tape Deduplication and Replication 391
- Top N Components 398
- Trends 395
 - Client performance 396
 - Deduplication Replication 397
 - Deduplication Repository 397
 - Modify 397
 - NAS resource performance 396
 - Physical device 396
 - Storage capacity 396
 - Virtual tape library performance 396
- Reports page 381
- Rescan

- Physical devices 92
- Reset layout 37, 38
- Restart network 336
- Restart services 337
- Root password
 - Change 358
- Rules
 - General rules 411
 - Apply to multiple servers 417
 - Create 412
 - Delete 415, 416
 - Enable 416
 - Global time exclusions 416
 - Network error 412
 - System Log errors 413
 - Update 415, 416
 - Smart rules 408
 - Alerts 418
 - Create 408
 - Delete 410
 - Modify 410
- Run History
 - Run History page 225

S

- Save layout 37, 38
- Save server configuration 361
- Scan stacked tape
 - Command line 556
- Scheduled reports 406
- Search 38
- Server
 - Capacity information 49
 - Icons 33
 - Info 33
 - Module status 33
- Server configuration
 - Recover 362
 - Save 361
- Servers
 - Add 51
 - Delete 53
 - Disconnect 52
 - Force rediscovery 53
 - Inventory 382
 - Lock 53
 - Reset password 52
 - Servers page 48
 - StorSafe server settings 335
- Services
 - Restart 337
- Share 37
- Shared library environment variable
 - Set 742
- Shares
 - Command line
 - Assign share to NFS client 592
 - Assign share to SMB client 591
 - Create 589
 - Delete 591
 - Unassign share from NFS client 593
 - Unassign share from SMB client 592
 - Shares page 268
- Shares page
 - Icons 269
- Shut down server 337
- Simple Name Server (SNS) 152
- SIR 19
- SIR cluster failover 200
- SIR server failover
 - Auto recovery 197
 - Backup application server configuration 189
 - Behavior 182
 - Change power control 197
 - Configure 190
 - Fibre Channel/iSCSI requirements 188
 - General requirements 187
 - Network requirements 188
 - Port swapping 194
 - Recovery 198
 - Requirements 187
 - Resume backups after failover/failback 185
 - Status 196
 - Storage requirements 189
 - Suspend/resume 198
 - Takeover 198
 - Terminology 181
 - Triggers 179
 - Update configuration 197
- Sizing guidelines 625
- Smart rules 408
 - Alerts 418
 - Behavior 408
 - Create 408
 - Delete 410
 - Modify 410
- SMB client
 - Map share 275

- SMB Clients
 - Page icons 283
 - SMB clients 282
 - SMB group
 - Create 292
 - Delete 293
 - Update 293
 - Update share access rights 293
 - SMB Groups
 - Page icons 292
 - SMB groups 291
 - SMB user
 - Create 283
 - Delete 284
 - Set authentication mode 284
 - Set SMB options 288
 - Set Windows ACLs 288
 - SMB connection status 290
 - Sync SMB Clients 287
 - Update 284
 - Update domain controller 287
 - Update share access rights 290
 - SMS settings 75, 77, 410, 414
 - Test 77
 - SNMP 359
 - Deduplication MIBs 478
 - Event Log messages 455
 - Implementing 456
 - Integration 455
 - MIBs 455
 - Traps 455
 - VTL MIBs 458
 - Sort 37
 - Space reclamation 354
 - Stacking
 - Command line
 - Prepare physical tape 555
 - Scan stacked physical tape 556
 - Standalone tape drive
 - Command line
 - Create 524
 - Storage
 - Monitoring 370
 - Storage LUN migration 739
 - Storage pools 84
 - Access rights 85
 - Create
 - StorSafe server 87
 - Delete 88
 - Icons 87
 - Migrate disks to pool 89
 - Monitor 447
 - Storage pools page 86
 - Update 88
 - Storage usage
 - Threshold 370
 - Storage usage widget 28
 - Storage web links
 - Icons 101
 - Page 101
 - StorSafe
 - Configuration 81
 - StorSafe info
 - Command line 508
 - StorSight 18
 - Administration 39
 - Administrators 62
 - HTTP proxy 45
 - License 40
 - Load average 39
 - Modules 46
 - Multi-factor authentication 42
 - Patches 44
 - Reboot 46
 - Restart services 46
 - Server summary 39
 - System disk usage 39
 - Tape Data Load 39
 - Used memory 39
 - User roles 62
 - Users 62
 - X-ray 44
 - Strong passwords
 - Enable 372
 - StorSight 43
 - Super Administrator
 - Definition 62
 - System i configuration 451
 - System time 358
- ## T
- Tape Caching 313
 - Create cache for physical tapes 318
 - Create policy 315
 - Create virtual tapes 319
 - Direct link tape 317
 - Policies 314
 - Space reclamation policy 317

- Space reclamation triggers 318
- Triggers 315
- Update policy 315
- Tape caching 120, 129
- Tape database 348
 - Manage 349
 - Mirror 349
- Tape drives
 - Tape drives page 113
- Tape Duplication 322
- Tape duplication 107
 - Command line 523
- Tape expansion
 - Troubleshooting 637
- Tape History
 - Tape History page 227
- Tape history 146
- Tape migration
 - Convert tape to stub 120, 260
 - Migrate to object storage 120
 - Recover from object storage 121, 260
- Tapes
 - Copy 119, 129
 - Deduplication tapes page 219
 - Export to physical 313
 - Filter 125
 - Manage 119, 128
 - Move 119, 128
 - Page icons 124, 221
 - Rename 119, 128
 - Replica 123
 - Shred 119
 - Tapes page 123
 - Write protect 119
- Target port binding 162
- Tenant 54
- Terminology 20
- Throttle speed
 - Replication
 - Global setting 365
- Top physical devices widget 30
- Troubleshooting 635
 - Clients 638, 639, 640
 - Direct link tape 640
 - Export 640
 - HBA missing 636
 - Import/export 647
 - NAS clients 648, 650
 - NAS resources 648

- NAS users/groups 650
- NDMP 643
- Offline disks 637
- Offline tapes 637
- Physical resources 636
- Replication 644
- Tape expansion 637
- VIT 640
- Turbo deduplication 210

U

- Usage history widget 29
- User
 - Definition 62
- User roles 62
- User types 62
- Users
 - Add 64
 - Command line
 - Add 516
 - Delete 516
 - List 517
 - Reset password 518
 - Set password 517
 - Delete 66
 - Local 65, 73
 - StorSight 62
 - Synchronize users 65
 - Update 65
 - Users page 64

V

- Vault page
 - Convert tape to stub 120, 260
 - Icons 118
 - Migrate to object storage 120
 - Recover from object storage 121, 260
- Viewer
 - Definition 62
- Virtual devices
 - Command line
 - Assign 509
 - Assign to iSCSI client 509
 - Create NAS device 586
 - Delete NAS device 588
 - Expand NAS device 587
 - Format NAS resource 587
 - List 507
 - Unassign 510

Virtual tape drives

- Load tapes 116
- Add to library 110
- Command line
 - Add 523
 - Get supported 519
- Delete 116
- Inventory 384
- Update 116

Virtual tape libraries

- Add virtual tape drives 110
- Assign to clients
 - Client multipathing 156
- Auto Archive 320
- Barcode 105
- Capacity On Demand 106
- Command line
 - Create 520
 - Delete 522
 - Get supported 519
 - Tape duplication 523
- Create 104
- Delete 107
- Inventory 383
- Monitor 432
- Reset 110
- Tape Caching 313
- Tape Duplication 322
- Tape duplication 107
- Update 107

Virtual tape libraries page 102

Virtual Tape Library 102

Virtual tapes 123

- Auto archive 109
- Auto replication 109, 241
- Capacity on Demand 106
- Command line
 - Copy 533
 - Create 527
 - Delete 530
 - Get tape info 527
 - Move 531
 - Set properties 529
 - Shred 534
- Copy 119, 129, 243
- Create 107, 116
- Delete 129
- Dynamic LUN Allocation 109

- Encryption 328
- Export to physical 313
- Filter 125
- How they are allocated 109
- Information 122, 130
- Load 110
- Manage 119, 128
- Move 119, 128
- Rename 119, 128
- Replica 123
- Replication 232
- Round Robin Logic 109
- Unload 110
- Update 128
- WORM 20, 108
- Write protect 119
- Virtual vault 117
- VIT
 - Troubleshooting 640
- VLAN tagging 344

W

Widgets 27

- Alerts 27
- Deduplication Usage 431
- FalconStor Summary 27
- Storage Usage 28
- Top Physical Devices 30
- Top StorSafe Servers 29
- Top Virtual Tape Libraries 30
- Usage History 29

Windows client

- Map share 275

World Wide Port Names 152

WORM 20, 108

WWPN 152

X

X-ray 337

- Command line 602
- StorSight 44

Z

Zoning 161