

Veritas Access Appliance Administrator's Guide

8.0 Linux

Access Appliance Administrator's Guide

Last updated: 2023-10-29

Legal Notice

Copyright © 2023 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

https://www.veritas.com/content/support/en_US/dpp.Appliances.html

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

APPL.docs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1	Introducing Access Appliance	15
Chapter 1	Introducing Access Appliance	16
	About Access Appliance	16
	About the Access Appliance Dashboard	20
	Getting started with the Access CLISH	21
	Accessing the Access CLISH	21
	Navigating the Access CLISH	22
	Getting help using the Access CLISH	22
	Displaying the command history	24
	Using the more command	24
	Getting started with Access Appliance APIs	25
Section 2	Configuring Access Appliance	32
Chapter 2	Managing users	33
	About user management	33
	About the naming requirements for adding new users	35
	About the naming requirements when assigning roles to domain users	36
	Adding and removing user roles using GUI	37
	Performing user management using CLISH	37
Chapter 3	Configuring the network	42
	About configuring the Access Appliance network	42
	About bonding Ethernet interfaces	43
	Bonding Ethernet interfaces	44
	Configuring DNS settings	46
	About Ethernet interfaces	48
	Displaying current Ethernet interfaces and states	48
	Configuring IP addresses	49
	Configuring VLAN interfaces	52
	Configuring NIC devices	53

	About configuring routing tables	53
	Configuring routing tables	53
	Changing the firewall settings	56
	Configuring Access Appliance in IPv4 and IPv6 mixed mode	59
	Support for multiple data subnets	62
Chapter 4	Configuring authentication services	64
	About configuring LDAP settings	64
	Configuring LDAP server settings	66
	Administering the Access Appliance cluster's LDAP client	68
	About Active Directory (AD)	69
	Configuring AD server settings	70
	Configuring entries for Access Appliance DNS for authenticating to Active Directory (AD)	71
	Configuring AD/LDAP using the GUI	73
	Configuring the NIS-related settings	73
	Configuring NSS lookup order	75
	Sign-in options for the Access Appliance UI	76
	Configuring user authentication using digital certificates or smart cards	76
	Adding CA certificates for smart card authentication	77
	Deleting CA certificates	78
Section 3	Managing Access Appliance storage	79
Chapter 5	Configuring storage	80
	About storage provisioning and management	80
	About configuring disks	81
	About configuring storage pools	81
	Configuring storage pools	82
	About quotas for usage	85
	Enabling, disabling, and displaying the status of file system quotas	86
	Setting and displaying file system quotas	87
	Setting user quotas for users of specified groups	89
	About quotas for CIFS home directories	89
	Workflow for configuring and managing storage using the Access Appliance CLI	91
	Displaying information for all disk devices associated with the nodes in a cluster	92

	Displaying WWN information	93
	Importing new LUNs forcefully for new or existing pools	94
	Initiating host discovery of LUNs	94
Chapter 6	Managing disks	96
	Adding disks to a storage pool	96
	Removing disks from a storage pool	97
	Viewing information about disks	97
	Accessing disk details	98
	Discovering disks	99
	Formatting a disk	99
Chapter 7	Access Appliance as an iSCSI target	100
	About Access Appliance as an iSCSI target	100
	Managing the iSCSI target service	101
	Managing the iSCSI targets	102
	Managing the LUNs	104
	Managing the mappings with iSCSI initiators	108
	Managing the users	109
	Creating an iSCSI target and provisioning LUNs	109
Section 4	Managing Access Appliance file access services	111
Chapter 8	Configuring the NFS server	112
	About using the NFS server with Access Appliance	112
	Using the kernel-based NFS server	113
	Accessing the NFS server	113
	Displaying and resetting NFS statistics	114
	Configuring Access Appliance for ID mapping for NFS version 4	114
	Configuring the NFS client for ID mapping for NFS version 4	115
	About authenticating NFS clients	115
	Setting up Kerberos authentication for NFS clients	116
	Adding and configuring Access Appliance to the Kerberos realm	117
Chapter 9	Using Access Appliance as a CIFS server	119
	About configuring Access Appliance for CIFS	120
	About configuring CIFS for standalone mode	121
	Configuring CIFS server status for standalone mode	122

Changing security settings	123
About configuring CIFS for Active Directory (AD) domain mode	123
Joining Access Appliance to Active Directory (AD)	123
Verifying that Access Appliance has joined Active Directory (AD)	
successfully	124
Using multi-domain controller support in CIFS	124
About leaving an AD domain	125
Changing domain settings for AD domain mode	125
Removing the AD interface	126
Setting NTLM	126
About setting trusted domains	127
Specifying trusted domains that are allowed access to the CIFS	
server	128
Allowing trusted domains access to CIFS when setting an IDMAP	
backend to rid	129
Allowing trusted domains access to CIFS when setting an IDMAP	
backend to ldap	130
Allowing trusted domains access to CIFS when setting an IDMAP	
backend to hash	131
Allowing trusted domains access to CIFS when setting an IDMAP	
backend to ad	131
About configuring Windows Active Directory as an IDMAP backend	
for CIFS	132
Configuring the Active Directory schema with CIFS-schema	
extensions	133
Configuring the LDAP client for authentication using the CLI	137
Setting Active Directory trusted domains	137
About storing account information	138
Storing user and group accounts	140
Reconfiguring the CIFS service	142
About mapping user names for CIFS/NFS sharing	143
About the mapuser commands	145
Adding, removing, or displaying the mapping between CIFS and NFS	
users	145
Automatically mapping UNIX users from LDAP to Windows users	
.....	146
About managing home directories	146
Setting the home directory file systems	147
Setting up home directories	148
Displaying home directory usage information	150
Deleting home directories and disabling creation of home	
directories	150
About CIFS clustering modes	151

	About switching the clustering mode	152
	About migrating CIFS shares and home directories	152
	Migrating CIFS shares and home directories from normal to ctdb clustering mode	152
	Migrating CIFS shares and home directories from ctdb to normal clustering mode	153
	Setting the CIFS aio_fork option	154
	About managing local users and groups	154
	Creating a local CIFS user	154
	Configuring a local group	155
	Enabling CIFS data migration	156
Chapter 10	Using Access Appliance as an Object Store server	157
	About the Object Store server	157
	Use cases for configuring the Object Store server	158
	Configuring the Object Store server	158
	About buckets and objects	163
	File systems used for objectstore buckets	166
	Enabling WORM on buckets	167
	S3 with NFS use case	169
	S3 with NSP use case	171
	Configuring the S3 server using GUI	171
	Setting Object Access server default parameters	171
	Setting up the Object Access server group-specific parameters	172
	Viewing information about S3	173
	Provisioning the S3 bucket using GUI	173
Section 5	Managing Access Appliance security	179
Chapter 11	Managing security	180
	Security overview	180
Chapter 12	Setting up FIPS mode	181
	FIPS 140-2 conformance for Access Appliance	181
	Viewing FIPS status for Access Appliance	182
	Enabling FIPS for Access Appliance	183
	To enable FIPS using the command-line interface	184

Chapter 13	Configuring STIG	186
	STIG overview for Access Appliance	186
	Enabling OS STIG hardening for Access Appliance	186
	Viewing STIG status of an Access Appliance	189
	Enabling STIG using the command-line interface	190
Chapter 14	Setting the banner	192
	Managing the login banner using the UI	192
	Managing the banner from the command-line interface	193
Chapter 15	Setting the password policy	195
	Managing the password policy using the UI	195
	Managing the password policy from the command-line interface	197
Chapter 16	Immutability in Access Appliance	203
	Support for immutability in Access Appliance	203
	About lockdown modes	204
	Selecting or changing the lockdown mode	205
	Accessing the root shell in lockdown mode	206
	Configuring immutability using GUI	207
Section 6	Monitoring and troubleshooting	210
Chapter 17	Configuring event notifications and audit logs	211
	About troubleshooting	211
	Monitoring command activity	212
	Monitoring alerts	212
	About alert management	213
	Viewing information about alert management	213
	Managing alerts	213
	Monitoring events	214
	Viewing reports	214
	Viewing cluster storage usage	214
	Viewing file system usage	215
	About event notifications	215
	About severity levels and filters	215
	About SNMP notifications	217
	Configuring a syslog server	217

	Displaying events on the console	218
Chapter 18	Appliance log files	219
	About appliance log files	219
	Gathering device logs with the DataCollect command	220
	Collecting default and additional diagnostic logs	223
	Downloading logs using the Log Transfer Console	224
	Forwarding logs to an external server	225
Section 7	Provisioning and managing Access	
	Appliance file systems	227
Chapter 19	Creating and maintaining file systems	228
	About creating and maintaining file systems	228
	About encryption at rest	229
	Considerations for creating a file system	231
	Best practices for creating file systems	231
	Choosing a file system layout type	234
	Determining the initial extent size for a file system	234
	About striping file systems	235
	About FastResync	238
	About fsck operation	239
	Enabling WORM on a file system	239
	Setting retention in files	240
	Setting WORM over NFS	241
	Manually setting WORM-retention on a file over CIFS	242
	About managing application I/O workloads using maximum IOPS	
	settings	249
	Creating a file system	250
	Bringing the file system online or offline	255
	Listing all file systems and associated information	256
	Modifying a file system	256
	Adding or removing a mirror from a file system	256
	Adding or removing a column from a file system	258
	Increasing the size of a file system	259
	Decreasing the size of a file system	261
	Managing a file system	263
	Defragmenting a file system	263
	Checking and repairing a file system	265
	Configuring FastResync for a file system	265
	Disabling the FastResync option for a file system	266

	Checking and resynchronizing stale mirrors	267
	Setting file system alerts	268
	Displaying file system alert values	269
	Removing file system alerts	269
	Destroying a file system	270
	Upgrading disk layout versions	270
Section 8	Provisioning and managing Access	
	Appliance shares	272
Chapter 20	Creating shares for applications	273
	About file sharing protocols	273
	About concurrent access	274
	Sharing directories using CIFS and NFS protocols	275
	Sharing a file system as a CIFS home directory	277
	About concurrent access with NFS and S3	277
Chapter 21	Creating and maintaining NFS shares	278
	About NFS file sharing	278
	About the NFS shares	279
	Creating an NFS share with continuous replication	279
	Creating an NFS share with episodic replication	280
	Creating an NFS share with encryption	281
	Creating an NFS share with episodic replication and encryption	282
	Creating an NFS share without replication and encryption	283
	Displaying file systems and snapshots that can be exported	284
	Exporting an NFS share	284
	Displaying exported directories	288
	About managing NFS shares using netgroups	288
	Unexporting a directory or deleting NFS options	289
	Exporting an NFS share for Kerberos authentication	290
	Mounting an NFS share with Kerberos security from the NFS client	291
	Exporting an NFS snapshot	293
Chapter 22	Creating and maintaining CIFS shares	294
	About managing CIFS shares	295
	About the CIFS shares	295
	Creating a CIFS share with continuous replication	295

	Creating a CIFS share with episodic replication	296
	Creating a CIFS share with encryption	297
	Creating a CIFS share with episodic replication and encryption	298
	Creating a CIFS share without replication and encryption	299
	Exporting a directory as a CIFS share	300
	Configuring a CIFS share as secondary storage for an Enterprise Vault store	300
	Exporting the same file system/directory as a different CIFS share	301
	About the CIFS export options	302
	Setting share properties	306
	Displaying CIFS share properties	307
	Hiding system files when adding a CIFS normal share	308
	Allowing specified users and groups access to the CIFS share	309
	Denying specified users and groups access to the CIFS share	310
	Exporting a CIFS snapshot	311
	Deleting a CIFS share	311
	Modifying a CIFS share	312
	Making a CIFS share shadow copy aware	313
	About managing CIFS shares for Enterprise Vault	313
	Creating a CIFS share for Enterprise Vault with replication	313
	Creating a CIFS share for Enterprise Vault without replication	314
Chapter 23	Integrating Access Appliance with Data Insight	316
	Access Appliance integration with Data Insight	316
Section 9	Managing Access Appliance storage services	322
Chapter 24	Configuring episodic replication	323
	About Access Appliance episodic replication	324
	How Access Appliance Replication works	325
	Starting Access Appliance episodic replication	326
	Setting up communication between the source and the destination clusters	328
	Setting up the file systems to replicate	332
	Setting up files to exclude from an episodic replication unit	334
	Scheduling the episodic replication	336

Defining what to replicate	338
About the maximum number of parallel episodic replication jobs	340
Managing an episodic replication job	340
Replicating compressed data	345
Displaying episodic replication job information and status	345
Synchronizing an episodic replication job	346
Behavior of the file systems on the episodic replication destination target	347
Accessing file systems configured as episodic replication destinations	347
Episodic replication job failover and fallback	347
Process summary	349
Overview of the planned failover process	349
Overview of the planned fallback process	350
Overview of the unplanned failover process	351
Overview of the unplanned fallback process	351

Chapter 25	Configuring continuous replication	353
	About Access Appliance continuous replication	353
	How Access Appliance continuous replication works	354
	How data flows in continuous replication synchronous mode	356
	How data flows in continuous replication asynchronous mode	357
	Starting Access Appliance continuous replication	359
	Setting up communication between the source and the destination clusters	361
	Setting up the file system to replicate	365
	Managing continuous replication	367
	Displaying continuous replication information and status	368
	Unconfiguring continuous replication	374
	Preserving the file system on the destination cluster	376
	Continuous replication failover and fallback	376
	Process summary	377
	Overview of the planned failover process	377
	Overview of the planned fallback process	378
	Overview of the unplanned failover process	379
	Overview of the unplanned fallback process	379
	Addition of multiple file systems to a Replicated Volume Group	380

Chapter 26	Using snapshots	384
	About snapshots	384
	Enabling WORM on storage snapshots	385
	Creating snapshots	386
	Displaying snapshots	387
	Managing disk space used by snapshots	388
	Bringing snapshots online or taking snapshots offline	390
	Restoring a snapshot	391
	About snapshot schedules	391
	Configuring snapshot schedules	392
	Managing automated snapshots	395
Chapter 27	Using instant rollbacks	398
	About instant rollbacks	398
	Creating a space-optimized rollback	400
	Creating a full-sized rollback	400
	Listing Access Appliance instant rollbacks	401
	Restoring a file system from an instant rollback	401
	Refreshing an instant rollback from a file system	402
	Bringing an instant rollback online	402
	Taking an instant rollback offline	402
	Destroying an instant rollback	403
	Creating a shared cache object for Access Appliance instant rollbacks	403
	Listing cache objects	405
	Destroying a cache object of a Access Appliance instant rollback	407
Section 10	Reference	408
Appendix A	Access Appliance documentation	409
	Using the Access Appliance product documentation	409
	About accessing the online man pages	410
Index		412

Introducing Access Appliance

- [Chapter 1. Introducing Access Appliance](#)

Introducing Access Appliance

This chapter includes the following topics:

- [About Access Appliance](#)
- [About the Access Appliance Dashboard](#)
- [Getting started with the Access CLISH](#)
- [Getting started with Access Appliance APIs](#)

About Access Appliance

You can use Access Appliance in any of the following ways.

Table 1-1 Interfaces for using Access Appliance

Interface	Description
GUI	Getting Started wizard with operations for managing the Access Appliance. Centralized dashboard and Quick Actions with operations for managing your storage. See the GUI and the Online Help for more information.
RESTful APIs	Enables automation using scripts, which run storage administration commands against the Access Appliance cluster. See the <i>Access Appliance RESTful API Guide</i> for more information.

Table 1-1 Interfaces for using Access Appliance (*continued*)

Interface	Description
Command-line interface (CLI)	Single point of administration for the entire cluster. See the manual pages for more information.

Table 1-2 Access Appliance key features

Feature	Description
Supported protocols	Access Appliance includes support for the following protocols: ■ Amazon-compatible S3 See “About the Object Store server” on page 157. ■ CIFS See “About configuring Access Appliance for CIFS” on page 120. ■ iSCSI target ■ NFS See “About using the NFS server with Access Appliance” on page 112.
Creation of Partition Secure Notification (PSN) file for Enterprise Vault Archiving	A Partition Secure Notification (PSN) file is created at a source partition after the successful backup of the partition at the remote site. For more information, see the <i>Access Appliance Solutions Guide for Enterprise Vault</i> .
Managing application I/O workloads using maximum IOPS settings	The MAXIOPS limit determines the maximum number of I/Os processed per second collectively by the storage underlying the file system. See “About managing application I/O workloads using maximum IOPS settings” on page 249.
Snapshot	Access Appliance supports snapshots for recovering from data corruption. If files, or an entire file system, are deleted or become corrupted, you can replace them from the latest uncorrupted snapshot. See “About snapshots” on page 384.
Access Appliance as an iSCSI target for RHEL 7.x	Access Appliance as an iSCSI target can be configured to serve block storage. An iSCSI target as service is hosted in an active/active mode in the Access Appliance cluster. See “About Access Appliance as an iSCSI target” on page 100.

Table 1-2 Access Appliance key features (*continued*)

Feature	Description
Configuring Access Appliance in IPv4 and IPv6 mixed mode	Support for configuring the Access Appliance cluster in an IPv4 environment, or an IPV6 environment, or in a mixed mode environment where you have both IPv4 and IPv6 addresses. See “Configuring Access Appliance in IPv4 and IPv6 mixed mode” on page 59.
NetBackup integration	Built-in NetBackup client for backing up your file systems to a NetBackup primary or media server. Once data is backed up, a storage administrator can delete unwanted data from Access Appliance to free up expensive storage for more data. See the <i>Access Appliance Solutions Guide for NetBackup</i> for more information.
Quotas	Support for setting file system quotas, user quotas, and hard quotas. See “About quotas for usage” on page 85.
Replication	Periodic replication of data over IP networks. See “About Access Appliance episodic replication” on page 324. See the <code>episodic(1)</code> man page for more information. Synchronous replication of data over IP networks See “About Access Appliance continuous replication” on page 353. See the <code>continuous(1)</code> man page for more information.
Support for LDAP, NIS, and AD	You can configure LDAP, NIS, and AD authentication services with Access Appliance. See “About configuring LDAP settings” on page 64. See “Configuring the NIS-related settings” on page 73. See “About Active Directory (AD)” on page 69.

Table 1-2 Access Appliance key features (*continued*)


Feature	Description
Partition Directory	With support for partitioned directories, directory entries are redistributed into various hash directories. These hash directories are not visible in the namespace view of the user or operating system. For every new create, delete, or lookup, this feature performs a lookup for the respective hashed directory and performs the operation in that directory. This leaves the parent directory inode and its other hash directories unobstructed for access, which vastly improves file system performance. By default this feature is not enabled. See the <code>storage_fs(1)</code> manual page to enable this feature.
Veritas Data Deduplication	Veritas Data Deduplication technology is installed on top of Access Appliance and integrates with NetBackup. It catalogs and organizes incoming deduplicated backup data and stores it on Access Appliance storage. For more information, see the <i>Access Appliance Solutions Guide for NetBackup</i>.
FIPS	FIPS 140-2 standard is enabled by default for the Veritas Operating System (VxOS). See “Enabling FIPS for Access Appliance” on page 183.
STIG	You can enable OS STIG hardening rules for increased security. These rules are based on the following profile from the Defense Information Systems Agency (DISA). See “Enabling OS STIG hardening for Access Appliance” on page 186. For more information, see the <i>Appliance security</i> chapter in the <i>Veritas Access Appliance Initial Configuration Guide</i>.
Support for Cloud tiering	The cloud as a tier feature for a file system lets you move data to different cloud services. The data is always written to the on-premises storage tier and then data can be moved to the cloud tier using a tiering mechanism. For more information, see the <i>Access Appliance Cloud Storage Tiering Guide</i>.
Separation of management and data network	Ability to configure a separate management and data network during cluster configuration. For more information, see the <i>Veritas Access Appliance Initial Configuration Guide</i>.

Table 1-2 Access Appliance key features (*continued*)

Feature	Description
Support for multiple data subnets	Access Appliance supports multiple data subnets. This is applicable to all the protocols that the Access Appliance supports. See “Support for multiple data subnets” on page 62.
Support for immutability in Access Appliance	Access Appliance supports lockdown modes that protects your cluster data from internal and external threats by securing all the external endpoints from unauthorized access. See “About lockdown modes” on page 204.

About the Access Appliance Dashboard

The **Dashboard** provides an overview of your cluster storage, configured storage, and information about alerts and performance. You can monitor the overall health of your cluster from the **Dashboard** and navigate to more detailed information.



The **Dashboard** displays the following information:

- The **Cluster storage** area displays an aggregated view of the storage utilization on the appliance. It provides information about the total storage available, the configured storage, and the utilized disk space for provisioning the storage.
- The **Alerts** area displays information about the issues that you may need to take action on. To view a more detailed list of all the alerts, click **View details** in the Alerts area.
- The **Provisioned storage** area displays the storage that is configured for Veritas Data Deduplication, buckets, and NFS and CIFS shares.
- The **Infrastructure** area displays the infrastructure components of the appliance. This includes the cluster nodes and their health, the total number of storage disks and their health, and the status of the public and private network that is configured on the appliance.
- The **Performance** area displays the number of I/O requests processed per second, the amount of data read and written per second, and the latency for the read and write operations at any given time.
- The **Resource monitoring** area displays the aggregate data about the CPU, memory, and network usage in the cluster.
- The **License summary** area displays whether a valid trialware or perpetual license is installed.
- The **Cloud storage tiers** area displays the configured cloud services where you can move the on-premise data to.

Getting started with the Access CLISH

Access Appliance provides many features to help you when you enter commands on the command line.

Accessing the Access CLISH

To access the Access CLISH

- 1 Connect to the management console using the console IP address you assigned during the configuration.
- 2 Log on to the management console using the password that you set during initial configuration.

Navigating the Access CLISH

All of the Access CLISH commands are organized in different command modes depending on the operation you want to perform. You can get a list of the different command modes with descriptions of all the available modes by typing a question mark (?) at the CLI prompt.

To navigate the Access CLISH

- ◆ After logging on to the Access CLISH, type a question mark (?) to see the available command modes.

For example, enter the `Storage` mode by typing `storage` for example.

Getting help using the Access CLISH

Access Appliance provides the following features to help you when you enter commands on the command line:

- Auto-completion

The following keys both perform auto-completion for the current command line. If the command prefix is not unique, then the bell rings and a subsequent repeat of the key displays possible completions.

 - [enter] - Auto-completes, syntax-checks then executes a command. If there is a syntax error, then the offending part of the command line is highlighted and explained.
 - [space] - Auto-completes, or if the command is already resolved inserts a space.
- Command-line help

Type a question mark at the command line to display context-sensitive Help. This is either a list of possible command completions with summaries, or the full syntax of the current command. A subsequent repeat of this key, when a command has been resolved, displays a detailed reference.
- Keyboard shortcut keys

Move the cursor within the command line or delete text from the command line.
- Command-line manual pages

Type `man` and the name of the command.
- Error reporting

The ^ (caret) indicates a syntax error occurred in the preceding command statement. The location of a caret in the command statement indicates the location of the syntax error.
- Escape sequences

Substitute the command line for a previous entry.

Table 1-3 Conventions used in the Access online command-line man pages

Symbol	Description
(pipe)	Indicates you must choose one of elements on either side of the pipe.
[] (brackets)	Indicates that the element inside the brackets is optional.
{ } (braces)	Indicates that the element inside the braces is part of a group.
< >	Indicates a variable for which you need to supply a value.

Table 1-4 Access command-line keyboard shortcut keys for deletions

Shortcut key	Description
[CTRL-C]	Delete the whole line.
[CTRL-U]	Delete up to the start of the line from the current position.
[CTRL-W]	Delete one word to the left from the current position.
[ALT-D]	Delete one word to the right from the current position.
[CTRL-D]	Delete the character to the right on the insertion point.
[CTRL-K]	Delete all the characters to the right of the insertion point.
[CTRL-T]	Swap the last two characters.
[backspace]	Delete the character to the left of the insertion point.
[Del]	Delete one character from the current position.

Table 1-5 Escape sequences

Escape sequence	Description
!!	Substitute the last command line.
!N	Substitute the Nth command line (you can find the Nth command from using the <code>history</code> command).
!-N	Substitute the command line entered N lines before (the number is relative to the command you are entering).

Note: Most of the Access Appliance commands are executed asynchronously, so control may be returned to the command prompt before the operation is fully completed. For critical commands, you should verify the status of the command before proceeding. For example, after starting a CIFS service, verify that the service is online.

Displaying the command history

The `history` command displays the commands that you have executed. You can also view commands executed by another user.

In addition to the commands that users execute with the Access Appliance command-line interface (CLI), the `history` command displays internal commands that were executed by Access Appliance.

You must be logged in to the system to view the command history.

To display command history

- ◆ To display the command history, enter the following:

```
CLISH> history [username] [number_of_lines]
```

username Displays the command history for a particular user.

number_of_lines Displays the number of lines of history you want to view.

The information displayed from using the `history` command is:

Time	Displays the time stamp as MM-DD-YYYY HH:MM
Status	Displays the status of the command as Success, Error, or Warning.
Message	Displays the command description.
Command	Displays the actual commands that were executed by you or another user.

Using the more command

The `System> more` command enables, disables, or checks the status of the `more` filter. The default setting is enable, which lets you page through the text one screen at a time.

To modify and view the more filter setting

- ◆ To modify and view the `more` filter setting, enter the following:

```
System> more enable|disable|status
```

<code>enable</code>	Enables the more filter on all of the nodes in the cluster.
<code>disable</code>	Disables the more filter on all of the nodes in the cluster.
<code>status</code>	Displays the status of the <code>more</code> filter.

Getting started with Access Appliance APIs

The Access Appliance API provides a web-service based interface to manage and administer the Veritas Access Appliance server. The Access Appliance API is built on the Representational State Transfer (REST) architecture, which is the most widely used style for building APIs. The Access Appliance API uses the HTTP protocol to communicate with the Access Appliance server. It is easy to use in cloud-based applications, as well as across multiple platforms and programming languages.

The Access Appliance API uses JavaScript Object Notation (JSON) as the message format for the request and response messages. It employs client-server communication in the form of HTTP requests and responses. The API client (your program) uses the HTTP protocol to make an API request to the Access Appliance server. The server processes the request and responds to the client with an appropriate HTTP status code indicating either success or failure. The client then extracts the required information from the server's response.

The Access Appliance server authenticates the incoming API requests based on a JSON Web Token (JWT) that needs to be provided in the Authorization HTTP header when making the API requests. A JSON Web Token (JWT) is acquired by executing a login API request. The login API accepts a username and password. The port that is used to access the Access Appliance API is the standard Access Appliance PBX port, 14161.

Example of generating and using JWT for authentication

The following procedure provides a sample workflow to retrieve task information from Access Appliance server. This procedure involves logging in to Access Appliance server to receive a JWT and then requesting task information for a specific task (task ID f42ac680-acba-11ea-b342-1f3ab58ec019 in this scenario).

1. Step 1:

Use the Access Appliance API endpoint `POST`

`/api/appliance/v1.0/authentication/login` to create a login request:

```
curl --cookie-jar /tmp/cookies.txt -g -k -X POST
https://hostname:14161/api/appliance/v1.0/authentication/login\
--header 'Accept: application/json'\
--header 'Content-Type: application/x-www-form-urlencoded'\
--data 'userName=myuser&password=mypassword'
```

The following response to the login request contains the JSON Web Token (JWT):

The `userRole` attribute indicates the Role of the user.

```
{
  "token": "eyJ0eXAiOiJKV1QiLCJhbGciOiJSUzI1NiIsInpccCI6IkdFRiJ9.
eNp0VF1v2jAU_S9-nHiArXQrb15yaT2cOLIdEJ2mKGWZmnXARMKEVvW_13EciO3weu65H
-d-vaIyr9FsMp18Gt_dTsefb6bjESqrCs3Q7umYbfOqLg5ohH7XpYIWqVfzNM1k6tHvCK
JmHsnv63_5MvHO-V_cztC1fFJkQ_b3bGaHar6UGxeitm_k2Lnx_o5q_cvxW44W111-c9t2V
jrw7FokW2-eS53hcJ-5X-qwoT5n232u7o4qfJfEQchmZfZN_YVzb6jD-jHSGMQOYifA9nd
Elg5dJAZg9MyDMUcMASsmUkgC-96jGYrFDoKCwgFEKGSQsPlsoETILgEsyJ4HyE2dLxEIyX
9txdJYwSoK1Xa_F01EFZcLnZJgS3MOBNxEhDhNGYi9Tr7KLdJZ4JXBYsgUM0nUtnFFfmirQ
7ZBUUZzm6EwEHClDmbTBKczEtacZK55kkXJuwrjcf1m6gq45Tgne7I1_o9XF8CNPqdMTnIZE
UnYvXLIjQZPnEI0QS2yjAwPVuJLrIFioFfYy4VQ-6JY70hLgERGiv6gad2SYVR9Ya822h6m
OKI3sSXRnwSGEWEJmTLsszTzhpbfarnh9cfI9_saN08DbzZF87m6RcSxyqAKoK63o9Rsr
JWnPd2zdzq-5A8MRqUjUN11N6I4BgpQTuU44S7zzHOqQyeJdbodjAKiVor2YK8_nMutL7DQ
Jux-m_KIIYukvnrthumPCbaS9BCbwbk5j8vW0cYzGfPvXZHQrHrvU7tfsWmY9wf0Onakt
7d3AAAA__8.VVE25rQqbrC-isGOqbRTqMPoK4ts5-9_6zSgz0fUg11m9GCClq10PS9u1D
laXye-S2MYyYHVEHSVs6uKcPVvN2WGBHkv7t-c4Hixc908zrJYhJaP979wF_gn08YnRlX
7_o4Qj6muc1IWHjK0hPMIgg0X-sBU2Git9uppVW1jbLA",
  "userRole": "Appliance Administrator"
}
```

2. Step 2:

Get the task information using the Access Appliance API endpoint `GET`

`/api/appliance/v1.0/tasks/{taskId}`. In this example, the information for the task ID `f42ac680-acba-11ea-b342-1f3ab58ec019` is requested.

Note: The Authorization header uses the value of the token attribute from the response to the login request made in the previous step.

```
curl --cookie /tmp/cookies.txt -k -X GET
https://hostname:14161/api/appliance/v1.0/tasks/f42ac680-acba-11ea-
```

```

b342-1f3ab58ec019 \
--header 'Accept: application/json'
--header 'Authorization: Bearer eyJ0eXAiOiJKV1QiLCJhbGciOiJSUz
I1NiIsInppcCI6IkpRFRiJ9.eNp0VF1v2jAU_S9-nHiArXQrb15yaT2cOLIdEJ2mKGWZmnX
ARMKEVvW_13EciO3weu65H-d-vaIyr9FsMp18Gt_dTsefb6bjESqrCs3Q7umYbfOqLg5oh
H7XpYIWqwVfzNM1k6tHvCKJmHsnv63_5MvHO-V_cztC1fFJkQ_b3bGaHar6UGxeitm_k2
LnX_o5q_cvxW44W1l1-c9t2Vjrw7FokW2-eS53hcJ-5X-qwoT5n232u7o4qfJfEQchMZfz
N_YVzb6jD-jHSGOMQ0Yifa9ndElg5dJAZg9MyDMUCMASsmUkgC-96jGYrFD0KCwgFEKgs
QsPlsoETILgEsyJ4HyE2dLxEIyX9txDJYwSoK1Xa_F01EFZcLnZJgS3MOBNxEhDhNGYi
9Tr7KLdJZ4JXBYsgUM0nUtnFFfmirQ7ZBUUZzm6EwEHClDmbTBKczEtacZK55kkXJuw
rjcf1m6gq45Tgne7I1_o9XF8CNPqdMTnIZEUnYvXLIjQZPnEIoQS2yJAwPVuJLrIFi
oFfYy4VQ-6JY70hLgERGiv6gad2SYVR9Ya822h6mOKI3sSXRnwsGEWBJmTlsszTZh
pbfarnh9cfI9_saN08DbzZF87m6RcSxyqAKoK63o9RsrJWnPd2zdzq-5A8MRQuJUN1
lN6I4BgpQTuU44S7zzHOqQyeJdbodjAKiVor2YK8_nMutL7DQJux-m_KIIYukvnr
thumPCbaS9BCbwbKb5j8vW0cYZGfPvrxZHqrHrvU7tfsWmY9wf0Onakt7d3AAAA__
8.VVE25rQqbrC-isGOqbRTqMPoK4ts5-9_6zSgz0fUg11m9GCC1q10PS9u1Dla
Xye-S2MYyYHVEHSVs6uKcPVvN2WGBHkv7t-c4Hixc9O8zrJYhJaP979wF_gn
08YnR1X7_o4Qj6muc1IWHjK0hPMIqg0X-sBU2Git9uppVW1jbLA'

```

The response to the GET request contains the task information for task ID **f42ac680-acba-11ea-b342-1f3ab58ec019**:

```

{
  "data": {
    "links": {
      "self": {
        "href": "/api/appliance/v1.0/tasks/f42ac680-acba-
11ea-b342-1f3ab58ec019"
      }
    },
    "type": "tasks",
    "id": "1",
    "attributes": {
      "isParent": true,
      "numCompleted": 1,
      "state": "SUCCESS",
      "taskId": "f42ac680-acba-11ea-b342-1f3ab58ec019",
      "source": "MASTER_FS",
      "numChildren": 2,
      "parentTaskId": "Unknown",
      "progress": "Unknown",
      "userName": "Unknown",
      "startTime": "2020-06-12 14:42:35",
      "sourceType": "vrts_raidgroup",

```

```

        "taskName": "Create File System MASTER_FS",
        "endTime": 2020-06-12 14:43:39,
        "output": ["ACCESS fs SUCCESS V-493-10-2110 Created
                    mirrored file system MASTER_FS"],
        "childrenTaskIdList": ["f7d750fa-acba-11ea-a760-000c29
                               1f5977", "035d1158-acbb-11ea-9ea1-000c291f5977"]
    }
}
}

```

Executing the REST APIs through Swagger UI

Swagger UI link: <https://hostname:14161/swagger>

1. Step 1: Authentication

- Click on POST `/api/appliance/v1.0/authentication/login`
- Click on **Try it out**. Enter username and password. Click on **Execute**.
- This will return a token, userRole and warning in Response body 200.
- Copy the token from Response body.

2. Step 2: Authorization

Click on **Authorize** in the upper right corner of Swagger UI. This will open a popup modal. In the **Value** field, type Bearer (space) and paste the token copied in the previous step 1.

Example: Bearer

```

eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCJ9.eyJzdWIiOiI5OWU5ZGRlMzkyNWl1LCJpYXQiOiJlE2NDk5MzQyMDZ9.Eyv2lmRz7pM__PL
UdgkyevuV8dM78-ie65rr5xzVgcshWS25HuN7bmELscNIIC13n1EKILb953ne
BfW_o9R9RA8IalmL_m6NsBvNWb6AfXbAFxVXHRAPAZnDpkkt6hls6koONkc0Ld
I1707w7wrceZMgWMe-F8cGmffN3dSiCvG89JkSLOWL0K6gQncwqkkzJnI5S7by
R_YMrB5hbgWV2zJMbWIIwAJogChWE2uhhGobKqztxo0Y1ZN-XrCk4E_AftLrbZ
4GjO_AddvbpYcVlQ7DdCZsKgSMpHG5N8FFwDtk3n_8zriVV0XPLnMGMadbsM
9qm5YmO-MkuydSI07KlQ

```

Click on **Authorize** and close the popup modal.

3. Step 3: Execute REST APIs

To get all the file systems using the Access Appliance API, click on GET `/api/appliance/v1.0/storage/filesystems`. Click on **Try it out** and then click **Execute**. This will return a list of file systems in Response body 200.

```

{
  "links": {
    "self": {
      "href": "/api/appliance/v1.0/storage/filesystems"
    }
  },
  "data": [
    {
      "links": {
        "self": {
          "href": "/api/appliance/v1.0/storage/filesystems/fs1"
        }
      },
      "type": "storage",
      "id": "1"
    },
    {
      "links": {
        "self": {
          "href": "/api/appliance/v1.0/storage/filesystems/fs2"
        }
      },
      "type": "storage",
      "id": "1"
    }
  ]
}

```

To get the specific file system information using the Access Appliance API, click on `GET /api/appliance/v1.0/tasks/{fsName}`. In this example, the information for file system name is requested.

Click on **Try it out**. Enter the file system name in the input box and then **Execute**. This will return details of the requested file system in Response body 200.

```

{
  "links": {
    "self": {
      "href": "/api/appliance/v1.0/storage/filesystems/fs1"
    }
  },
  "data": {
    "links": {

```

```
    "self": {
      "href": "/api/appliance/v1.0/storage/filesystems/fs1"
    }
  },
  "type": "storage",
  "id": "1",
  "attributes": {
    "fsBlockSize": 8192,
    "size": "1.00 GB",
    "usedSize": "49.30 MB",
    "fullFsckStatus": "Not Running",
    "fullFsck": 0,
    "isDedupeEnabled": "No",
    "storageLayout": "striped",
    "numColumn": "5",
    "resiliency": "-",
    "fsName": "fs1",
    "offlineNodeNameList": [],
    "onlineNodeNameList": [
      "r6515-003v011",
      "r6515-003v012"
    ],
    "status": "online",
    "stripeUnit": "64 K",
    "isWormEnabled": "No",
    "reconfigState": "Unknown",
    "isReconfigRunning": false,
    "minimumRetention": "",
    "maximumRetention": "",
    "diskPoolName": "pool_default",
    "diskNameList": [
      "vrts_appliances0_4",
      "vrts_appliances0_6",
      "vrts_appliances0_7",
      "vrts_appliances0_8",
      "vrts_appliances0_9"
    ],
    "fsId": "{2d2c9fed-7620-9bd8-d298-c74bcc50a140}_fs1",
    "usage": []
  }
}
```

Possible response codes an Access Appliance API can return:

- 200 - Successful
- 400 - Error: Bad Request
- 401 - Error: Unauthorized
- 500 - Error: Internal Server Error

Configuring Access Appliance

- [Chapter 2. Managing users](#)
- [Chapter 3. Configuring the network](#)
- [Chapter 4. Configuring authentication services](#)

Managing users

This chapter includes the following topics:

- [About user management](#)
- [About the naming requirements for adding new users](#)
- [About the naming requirements when assigning roles to domain users](#)
- [Adding and removing user roles using GUI](#)
- [Performing user management using CLISH](#)

About user management

Access Appliance enables you to add users for administering your appliance. Your privileges within Access Appliance are based on what user role you have been assigned.

The following table provides an overview of the user roles within Access Appliance.

Table 2-1 User roles within Access Appliance

User role	Description
Primary	Primary user has all of the permissions including adding and deleting users, displaying users, and managing passwords. Only the primary user can add or delete other administrators.
Vxdefault	A vxdefault user role does not have any privilege. If the user role is not specified while creating the user, then the user is created without any privileges and is termed as a vxdefault user.

You can add local users as well as users from an Active Directory (AD) server and a Lightweight Directory Access Protocol (LDAP) server. The LDAP server should

be an OpenLDAP Linux/UNIX server. Access Appliance does not support LDAP users from the AD server. Registering remote users lets you leverage your existing directory service for user management and authentication.

Each user account must authenticate itself with a user name and password to access the appliance. For a local user, the user name and password are managed on the appliance. For a registered remote user, the user name and password are managed by the remote directory service.

To add the different administrator roles, you must have `primary` privilege. After you assign the `primary` role to AD and LDAP users, they can also perform cluster management using the GUI, Access CLISH, Appliance CLISH and Restful APIs.

When you add a new user role, remember that:

- Any user who has been assigned the primary role can assign the primary role to an AD/LDAP user and groups for managing the configuration.
- You can remove the primary role for a given AD/LDAP user and groups.
- AD/LDAP users without a primary role can generate S3 keys.
- AD/LDAP users can login even after a console failover or HA scenarios.
- Domain is an optional argument with the default value as local.
- For the local user, only the username is required.
- You can assign only the primary role to the LDAP and AD user.
- The LDAP username should not be the same as a local username.
- For AD, enter `domain\\username`.
- For LDAP, enter `username`.
- Veritas recommends that LDAP, AD, and NIS UIDs start from 10000. Otherwise, when you assign a role to the AD/LDAP user, the UID of the local user may conflict with the UID of a user from the directory server.
- Nested LDAP group for role assignment is not supported.

The `Support` account is reserved for Technical Support use only, and it cannot be created by administrators.

To perform user role management using GUI, see [“Adding and removing user roles using GUI”](#) on page 37.

To perform user management using CLISH, see [“Performing user management using CLISH”](#) on page 37.

About the naming requirements for adding new users

The following table provides the naming requirements for adding new Access Appliance users.

Table 2-2 Naming requirements for adding new users

Guideline	Description
Starts with	Must begin with an alphabetic character and the rest of the string should be from the following character set: [a-zA-Z0-9_-] Can start with an underscore (_).
Length	Can be up to 31 characters. If user names are greater than 31 characters, you get the error, "Invalid user name".
Case	Access Appliance CLI commands are case-insensitive (for example, the user command is the same as the USER command). However, user-provided variables are case-sensitive (for example, the username Primary1 is not the same as the username PRIMARY1).
Can contain	Hyphens (-) and underscores (_) are allowed.

Note: Each user account needs a password for authentication. The following special characters are not supported in the user password: ` (backtick), : (colon), ; (semicolon), " (double quote), ' (single quote), ((opening parenthesis),) (closing parenthesis), ? (question mark), * (asterisk), **white space** (blank), \ (back slash), / (forward slash).

When the STIG option is not enabled, the following password requirements must be met:

- The password must contain at least eight characters.
- The password must contain at least one lowercase character (a-z), uppercase character (A-Z), numeric character (0-9), and special character (!#\$%&+,-.<=>@[]^_{}~).
- The password cannot be a Dictionary word.
- The last seven passwords cannot be reused, and the new password cannot be similar to the previous passwords.

When the STIG option is enabled, the following password requirements must be met:

- The password must be at least 15 characters long and must include at least one character from the following categories:
 - Numbers
 - Lowercase characters
 - Uppercase characters
 - Special characters (!#\$%&+,-.<=>@[]^_{|}~)
- The password can contain a maximum of two consecutive repeating characters.
- The password can contain a maximum of four consecutive repeating characters from the same category.
- The password must contain a minimum number of eight different characters.
- The password must be retained for a minimum of one day before you change it.
- The password must be changed after a maximum of 60 days.
- The password cannot be a Dictionary word.
- The last seven passwords cannot be reused.
- The new password must have at least eight characters that are different from the old password.

About the naming requirements when assigning roles to domain users

The following table provides the naming requirements when assigning roles to AD/LDAP users.

Table 2-3 Naming requirements when assigning roles to domain users

Guideline	Description
Starts with	Must begin with an alphabetic character and the rest of the string should be from the following character set: A-Za-z0-9_- Must not start with hyphen (-) or underscore (_).

Table 2-3 Naming requirements when assigning roles to domain users
(continued)

Guideline	Description
Length	LDAP: Length can be up to 74 characters. AD: Length can be up to 20 characters.
Case	AD user names are case-insensitive. LDAP user names are case-sensitive.
Cannot contain	Space (blank)

Adding and removing user roles using GUI

You can use the Access Appliance GUI to add local, AD and LDAP users.

To add a new user

- 1 Navigate to **Settings > User management** and click **Add** to add a new user.
- 2 The **Add user** form appears.
 - To add a local user, select the user type as **Local**.
 - To add an AD/LDAP user, select the user type as **Directory** and add a valid user name or group name.
 - For AD, enter **domain\username**.
 - For LDAP, enter **username**.

Click **Add**.

- 3 The list of users gets updated after the operation is complete.

To remove a user

- 1 You can remove a user by clicking the menu button next to the user and selecting **Remove**.
- 2 A confirmation window appears. Click **Ok** to remove the user. The list of users gets updated after the operation is complete.

Performing user management using CLISH

The following administrator role is included with Access Appliance:

- Primary

You can add additional users with these roles. To add the different administrator roles, you must have `primary` privilege.

Note: When adding a new user, you must assign a password.

To add a user with specific privileges

- ◆ Enter the following to add a user with specific privileges:

```
Admin> user add username role
```

where *role* is the role that you want to assign to the user.

Note: The AD username cannot contain space.

To delete a user

- ◆ Enter the following to delete a specified user from the current system.

```
Admin> user delete username
```

To change a user's password

- 1 Enter the following command to change the password for the current user:

```
Admin> passwd
```

You are prompted to enter the old password first. If the password matches, then you are prompted to enter the new password for the current user.

Note: A user is not required to have primary privileges to change the password. Users with **vxdefault** or with no assigned role can also log in to CLISH and change their own password by using `Admin> passwd` command.

- 2 Enter the following command to change the password for a user other than the current user:

```
Admin> passwd [username]
```

You are prompted to enter the old password first. If the password matches, then you are prompted to enter the new password for the user.

- 3 To change the default password of the sysadmin user (IPMI user account), enter the following command:

```
Admin> ipmi passwd username old_password new_password
```

where *username* is sysadmin

To display a list of current users

- 1 Enter the following to display the current user:

```
Admin> show [username]
```

- 2 Enter the following to display a list of all the current users:

```
Admin> show
```

Enter the following to display the details of the administrator with the user name `primary`:

```
Admin> show primary
```

To add primary role to a local or domain user

- ◆ Enter the following to add a user role to Access Appliance:

```
Admin> user modify role add username role domain
```

where *domain* can be local, ad or ldap.

- Domain is an optional argument with the default value as local.
- For local user, only the username is required.
- You can assign only a primary role to LDAP and AD users.
- For AD, enter `domain\\username`.
- For LDAP, enter `username`.

Note: If you want to add a CIFS local user as an appliance administrator, you have to first add the user using the `admin user add` command and then add the user as a CIFS local user using the `cifs local user add` command. Users added directly using the `cifs local user add` command are not allowed to do appliance administration.

To delete a primary role from a local or domain user

- ◆ Enter the following to delete the primary role from a user account:

```
Admin> user modify role delete username role domain
```

To add a role to a domain group

- ◆ Enter the following to add a role to a specific domain group:

```
Admin> group modify role add groupname role domain
```

where *domain* can be ad or ldap.

Note: The AD group name cannot contain a space.

- For AD, enter `domain\\groupname`.
- For LDAP, enter `groupname`.

To remove a role from a domain group

- ◆ Enter the following to delete the specified role from a specific domain group:

```
Admin> group modify role delete groupname role domain
```

where *domain* can be ad or ldap.

Note: If the domain user is AD or LDAP, then you must configure NSS for that user using the `Network nsswitch conf` command. For an LDAP user, the NSS lookup should be set to *ldap* and for an AD user, the NSS lookup should be set to *winbind*. For more information, See [“Configuring NSS lookup order”](#) on page 75.

Configuring the network

This chapter includes the following topics:

- [About configuring the Access Appliance network](#)
- [About bonding Ethernet interfaces](#)
- [Bonding Ethernet interfaces](#)
- [Configuring DNS settings](#)
- [About Ethernet interfaces](#)
- [Displaying current Ethernet interfaces and states](#)
- [Configuring IP addresses](#)
- [Configuring VLAN interfaces](#)
- [Configuring NIC devices](#)
- [About configuring routing tables](#)
- [Configuring routing tables](#)
- [Changing the firewall settings](#)
- [Configuring Access Appliance in IPv4 and IPv6 mixed mode](#)
- [Support for multiple data subnets](#)

About configuring the Access Appliance network

Access Appliance has the following types of networks:

- Private network

The network between the nodes of the cluster itself. The private network is not accessible to Access Appliance client nodes.

- Public network

The public network is visible to all clients. Access Appliance uses static IP address for its public interface networking. Access Appliance does not support DHCP for public network configuration.

About bonding Ethernet interfaces

Bonding associates a set of two or more Ethernet interfaces with one IP address. The association improves network performance on each Access Appliance cluster node by increasing the potential bandwidth available on an IP address beyond the limits of a single Ethernet interface. Bonding also provides redundancy for higher availability.

For example, you can bond two 1-gigabit Ethernet interfaces together to provide up to 2 gigabits per second of throughput to a single IP address. Moreover, if one of the interfaces fails, communication continues using the single Ethernet interface.

When you create a bond, you need to specify a bonding mode. In addition, for the following bonding modes: `802.3ad`, `balance-rr`, `balance-xor`, `broadcast`, `balance-tlb`, and `balance-alb`, make sure that the base network interface driver is configured correctly for the bond type. For type `802.3ad`, the switch must be configured for link aggregation.

Consult your vendor-specific documentation for port aggregation and switch set up. You can use the `-s` option in the Linux `ethtool` command to check if the base driver supports the link speed retrieval option. The `balance-alb` bond mode type works only if the underlying interface network driver enables you to set a link address.

Note: An added IPv6 address may go into a TENTATIVE state while bonding Ethernet interfaces with `balance-rr`, `balance-xor`, or `broadcast` bond modes. While bonding with those modes, Access Appliance requires the switch to balance incoming traffic across the ports, and not deliver looped back packets or duplicates. To work around this issue, enable EtherChannel on your switch, or avoid using these bond modes.

Table 3-1 Bonding mode

Index	Bonding mode	Fault tolerance	Load balancing	Switch setup	Ethtool/base driver support
0	balance-rr	yes	yes	yes	no
1	active-backup	yes	no	no	no
2	balance-xor	yes	yes	yes	no
3	broadcast	yes	no	yes	no
4	802.3ad	yes	yes	yes	yes (to retrieve speed)
5	balance-tlb	yes	yes	no	yes (to retrieve speed)
6	balance-alb	yes	yes	no	yes (to retrieve speed)

Note: When you create or remove a bond, SSH connections with Ethernet interfaces involved in that bond may be dropped. When the operation is complete, you must restore the SSH connections.

Bonding Ethernet interfaces

The `Network> bond create` and `Network> bond remove` operations involve bringing down the interface first and then bringing them back up. This may cause the SSH connections that are hosted over those interfaces to terminate. Use the physical console of the client rather than SSH when performing `Network> bond create` and `Network> bond remove` operations.

To display a bond

- ◆ To display a bond and the algorithm that is used to distribute traffic among the bonded interfaces, enter the following:

```
Network> bond show
```

To create a bond

- ◆ To create a bond between sets of two or more Ethernet interfaces on all Access Appliance cluster nodes, enter the following:

```
Network> bond create interfacelist mode option
```

interfacelist	Specifies a comma-separated list of public Ethernet interfaces to bond.
mode	Specifies how the bonded Ethernet interfaces divide the traffic.
option	Specifies a comma-separated option string. Available only when the bond mode is 2 (balance-xor) or 4 (802.3ad) xmit_hash_policy - specifies the transmit hash policy to use for slave selection in balance-xor and 802.3ad modes.

If the option is not specified correctly, you get an error.

You can specify a mode either as a number or a character string, as follows:

0	balance-rr	This mode provides fault tolerance and load balancing. It transmits packets in order from the first available slave through the last.
1	active-backup	Only one slave in the bond is active. If the active slave fails, a different slave becomes active. To avoid confusing the switch, the bond's MAC address is externally visible on only one port (network adapter).
2	balance-xor	Transmits based on the selected transmit hash policy. The default policy is a simple. This mode provides load balancing and fault tolerance.
3	broadcast	Transmits everything on all slave interfaces and provides fault tolerance.
4	802.3ad	Creates aggregation groups with the same speed and duplex settings. It uses all slaves in the active aggregator based on the 802.3ad specification.
5	balance-tlb	Provides channel bonding that does not require special switch support. The outgoing traffic is distributed according to the current load (computed relative to the speed) on each slave. The current slave receives incoming traffic. If the receiving slave fails, another slave takes over its MAC address.

6	balance-alb	Includes balance-tlb plus Receive Load Balancing (RLB) for IPV4 traffic. This mode does not require any special switch support. ARP negotiation load balances the receive.
---	-------------	--

To remove a bond

- ◆ To remove a bond from all of the nodes in a cluster, enter the following:

```
Network> bond remove bondname
```

where *bondname* is the name of the bond configuration.

Configuring DNS settings

The Domain Name System (DNS) service resolves names to IP addresses. You can configure Access Appliance to use DNS to look up domain names and IP addresses. You enable the DNS service for the cluster, then specify up to three DNS servers.

Note:

- DNS names can contain only alphabetical characters (A-Z), numeric characters (0-9), the minus sign (-), and the period (.). Period characters are allowed only when they are used to delimit the components of domain style names.
- The following characters are not supported : comma (,), tilde (~), colon (:), exclamation point (!), at sign (@), number sign (#), dollar sign (\$), percent (%), caret (^), ampersand (&), apostrophe ('), period (.), parentheses (), braces {}, underscore (_), white space (blank). Please refer to the DNS nomenclature in your environment for more details.
- If you change the DNS domain name using the Access Appliance command-line interface while the NFS server is running, the NFS server continues to use the old domain name for ID mapping. You are required to restart the NFS server for the change in the domain name to take effect.

To display DNS settings

- ◆ To display DNS settings, enter the following:

```
Network> dns show
```

To enable DNS service

- ◆ To enable Access Appliance hosts to do DNS lookups, enter the following commands:

```
Network> dns enable
```

You can verify using the `dns show` command.

To disable DNS settings

- ◆ To disable DNS settings, enter the following:

```
Network> dns disable
```

You can verify using the `dns show` command.

To specify the IP addresses of the DNS name servers

- ◆ To specify the IP addresses of the DNS name servers used by the Access Appliance DNS service, enter the following commands:

```
Network> dns set nameservers nameserver1 [nameserver2] [nameserver3]
```

You can verify using the `dns show` command.

To remove the name servers list used by DNS

- ◆ To remove the name servers list used by DNS, enter the following commands:

```
Network> dns clear nameservers
```

You can verify using the `dns show` command.

To set the domain name for the DNS server

- ◆ To set the domain name for the DNS server, enter the following:

```
Network> dns set domainname domainname
```

where *domainname* is the domain name for the DNS server.

You can verify using the `dns show` command.

To allow multiple DNS search domains

- ◆ To allow multiple DNS search domains, enter the following:

```
Network> dns set searchdomains searchdomain1[,searchdomain2]  
[,searchdomain3]
```

where *searchdomain1* is the first DNS search domain to be searched. Specify the search domains in the order in which the search domains should be used.

To configure multiple DNS search domains that have already been entered

- ◆ To configure multiple DNS search domains that have already been entered, add the existing domain name with the new domain name as comma-separated entries.

```
Network> dns set searchdomains domain1.access.com,  
domain2.access.com.
```

You can verify using the `dns show` command.

To remove the domain name used by DNS

- ◆ To remove the domain name used by DNS, enter the following:

```
Network> dns clear domainname
```

You can verify using the `dns show` command.

About Ethernet interfaces

Internet Protocol (IP) commands configure your routing tables, Ethernet interfaces, and IP addresses, and display the settings.

Each Ethernet interface can be configured with a virtual IP address for clustering purposes in Access Appliance. This does not imply that each interface must have a virtual IP to communicate with the network.

Displaying current Ethernet interfaces and states

To display current Ethernet interfaces and states

- ◆ To display current configurations, enter the following:

```
Network> ip link show [nodename] [device]
```

nodename Specifies which node of the cluster to display the attributes.
Enter `all` to display all the IP links.

device Specifies which Ethernet interface on the node to display the attributes.

To display all configurations, enter the following:

```
Network> ip link show
```

Configuring IP addresses

During installation, you specified a range of public IP addresses to be used for physical interfaces. You also specified a range for virtual interfaces. You can see which of these addresses are assigned to each node. You can use this procedure to verify the IP addresses in your configuration. You can add additional IP addresses if you want to add additional nodes and no other IP addresses are available.

To display all the IP addresses for the cluster

- ◆ To display all of a cluster's IP addresses, enter the following:

```
Network> ip addr show
```

The output headings are:

IP	Displays the IP addresses for the cluster.
Netmask	Displays the netmask for the IP address. Netmask is used for IPv4 addresses. Specify an IPv4 address in the format AAA.BBB.CCC.DDD, where each number ranges from 0 to 255.
Prefix	Displays the prefix used for IPv6 addresses. The value is an integer in the range 0-128.
Device	Displays the name of the Ethernet interface for the IP address.
Node	Displays the node name associated with the interface.
Type	Displays the type of the IP address: physical or virtual.
Status	Displays the status of the IP addresses: ■ ONLINE ■ ONLINE (console IP) ■ OFFLINE ■ FAULTED A virtual IP can be in the FAULTED state if it is already being used. It can also be in the FAULTED state if the corresponding device is not working on all nodes in the cluster (for example, a disconnected cable).

To add an IP address to a cluster

- ◆ To add an IP address to a cluster, enter the following:

```
Network> ip addr add ipaddr netmask | prefix type
[device] [nodename] [fqdns]
```

ipaddr	Specifies the IP address to add to the cluster. Do not use physical IP addresses to access the Access Appliance cluster. In case of failure, the IP addresses cannot move between nodes. A failure could be either a node failure, an Ethernet interface failure, or a storage failure. You can specify either an IPv4 address or an IPv6 address.
netmask	Specifies the netmask for the IP address. Netmask is used for IPv4 addresses.
prefix	Specifies the prefix for the IPv6 address. The accepted range is 0-128 integers.
type	Specifies the IP address type, either virtual or physical. If type is virtual, the device is used to add new IP address on that device. If type is physical, the IP address gets assigned to given node on given device. In this case, you have to specify the nodename.
device	Only use this option if you entered <i>virtual</i> for the <i>type</i>.
nodename	Any node of the cluster
fqdns	Specifies a comma-separated list of Fully Qualified Domain Name (FQDN) of the IP address. The <i>fqdn</i> can include the characters: a-z A-Z 0-9 or a hyphen (-). Each level of the FQDN should be between 1 and 63 characters long and should not start or end with a hyphen (-). The last Top Level Domain (TLD) must be at least two characters and have a maximum of six characters.

To change an IP address to online on a specified node

- ◆ To change an IP address to online on a specified node, enter the following:

```
Network> ip addr online ipaddr nodename
```

ipaddr	Specifies the IP address that needs to be brought online. You can specify either an IPv4 address or an IPv6 address.
nodename	Specifies the nodename on which the IP address needs to be brought online. If you do not want to enter a specific nodename, enter <i>any</i> with the IP address.

You can also modify an IP address. But note that you cannot use the `ipr addr modify` command to modify the IP addresses in the following scenarios:

- IP is used by the NetBackup client
- Replication IP
- IP is used by deduplication service
- IP is used by Veritas Data Deduplication server
- IP is used by the NetBackup primary server
- IP is used by the NetBackup media server

To modify an IP address

- ◆ To modify an IP address, enter the following:

```
Network> ip addr modify oldipaddr newipaddr netmask | prefix fqdns
```

oldipaddr	Specifies the old IP address to be modified, as either an IPv4 address or an IPv6 address. The specified <i>oldipaddr</i> must be assigned to the cluster.
newipaddr	Specifies the new IP address, as either an IPv4 address or an IPv6 address. The new IP address must be available.
netmask	Specifies the netmask for the new IP address. Netmask is used for IPv4 addresses.
prefix	Specifies the prefix for the IPv6 address. The value is an integer in the range 0-128.
fqdns	Specifies a comma-separated list of Fully Qualified Domain Name (FQDN) of the IP address. The <i>fqdn</i> can include the characters: a-z A-Z 0-9 or a hyphen (-). Each level of the FQDN should be between 1 and 63 characters long and should not start or end with a hyphen (-). The last Top Level Domain (TLD) must be at least two characters and have a maximum of six characters. <i>fqdns</i> can have NONE value. NONE is used to remove the existing FQDN entry from the <code>/etc/hosts</code> file.

To remove an IP address from the cluster

- ◆ To remove an IP address from the cluster, enter the following:

```
Network> ip addr del ipaddr
```

where *ipaddr* is either an IPv4 address or an IPv6 address.

Configuring VLAN interfaces

The virtual LAN (VLAN) feature lets you create VLAN interfaces on the Access Appliance nodes and administer them as any other VLAN interfaces. The VLAN interfaces are created using Linux support for VLAN interfaces.

Use the `Network> vlan` commands to view, add, or delete VLAN interfaces.

Note: To use VLAN, your network must have VLAN-supported switches.

To display the VLAN interfaces

- ◆ To display the VLAN interfaces, enter the following:

```
Network> vlan show
```

To add a VLAN interface

- ◆ To add a VLAN interface, enter the following:

```
Network> vlan add device vlan_id
```

device	Specifies the VLAN interface on which the VLAN interfaces will be added.
--------	--

vlan_id	Specifies the VLAN ID which the new VLAN interface uses. Valid values range from 1 to 4095.
---------	---

Note: Only one VLAN interface can be tagged on eth1 interface.

To delete a VLAN interface

- ◆ To delete a VLAN interface, enter the following:

```
Network> vlan del vlan_device
```

where the `vlan_device` is the VLAN name from the `Network> vlan show` command.

Configuring NIC devices

To list NIC devices on a specified node

- ◆ To list NIC devices on a specified node, enter the following:

```
Network> device list nodename
```

where *nodename* is the specified node for which bus id, mac addresses, device info, and device type for all devices are listed.

To monitor a NIC device

- ◆ To configure a remote IP address to monitor the NIC device, enter the following:

```
Network> device ipmonitor show
```

About configuring routing tables

Sometimes a Access Appliance cluster must communicate with network services (for example, LDAP) using specific gateways in the public network. In these cases, you must define routing table entries.

These entries consist of the following:

- The target network node's IP address and accompanying netmask.
- Gateway's IP address.
- Optionally, a specific Ethernet interface via which to communicate with the target. This is useful, for example, if the demands of multiple remote clients are likely to exceed a single gateway's throughput capacity.

Configuring routing tables

To display the routing tables of the nodes in the cluster

- ◆ To display the routing tables of the nodes in the cluster, enter the following:

```
Network> ip route show [nodename]
```

where *nodename* is the node whose routing tables you want to display. To see the routing table for all of the nodes in the cluster, enter `all`.

For example:

```
Network> ip route show all
```

Destination	Displays the destination network or destination host for which the route is defined.
Gateway	Displays a network node equipped for interfacing with another network.
Genmask	Displays the netmask.
Flags	The flags are as follows: U - Route is up H - Target is a host G - Use gateway
MSS	Displays maximum segment size. The default is 0. You cannot modify this attribute.
Window	Displays the maximum amount of data the system accepts in a single burst from the remote host. The default is 0. You cannot modify this attribute.
irtt	Displays the initial round trip time with which TCP connections start. The default is 0. You cannot modify this attribute.
Iface	Displays the interface. On UNIX systems, the device name <code>lo</code> refers to the loopback interface.

To add to the route table

- ◆ To add a route entry to the routing table of nodes in the cluster, enter the following:

```
Network> ip route add nodename ipaddr netmask  

| prefix via gateway [dev device]
```

<i>nodename</i>	Specifies the node to whose routing table the route is to be added. To add a route path to all the nodes, use <code>all</code> in the <i>nodename</i> field. If you enter a node that is not a part of the cluster, an error message is displayed.
<i>ipaddr</i>	Specifies the destination of the IP address. You can specify either an IPv4 address or an IPv6 address. If you enter an invalid IP address, then a message notifies you before you fill in other fields.

netmask	Specifies the netmask associated with the IP address that is entered for the <code>ipaddr</code> field. Use a netmask value of 255.255.255.255 for the netmask to add a host route to <code>ipaddr</code>.
prefix	Specifies the prefix for the IPv6 address. Accepted ranges are 0-128 integers.
via	This is a required field. You must type in the word.
gateway	Specifies the gateway IP address used for the route. If you enter an invalid gateway IP address, then an error message is displayed. To add a route that does not use a gateway, enter a value of 0.0.0.0.
dev	Specifies the route device option. You must type in the word.
device	Specifies which Ethernet interface on the node the route path is added to. This variable is optional. You can specify the following values: ■ any - Default ■ pubeth0 - Public Ethernet interface ■ pubeth1 - Public Ethernet interface The Ethernet interface field is required only when you specify <code>dev</code> in the <code>dev</code> field. If you omit the <code>dev</code> and <code>device</code> fields, Access Appliance uses a default Ethernet interface.

To delete route entries from the routing tables of nodes in the cluster

- ◆ To delete route entries from the routing tables of nodes in the cluster, enter the following:

```
Network> ip route del nodename ipaddr
netmask | prefix
```

<i>nodename</i>	Specify the node from which the route entry has to be deleted. To delete the route entry from all nodes, use the <code>all</code> option in this field.
<i>ipaddr</i>	Specifies the destination IP address of the route entry to be deleted. You can specify either an IPv4 address or an IPv6 address. If you enter an invalid IP address, a message notifies you before you enter other fields.
<i>netmask</i>	Specifies the IP address to be used. Netmask is used for IPv4 addresses.
<i>prefix</i>	Specifies the prefix for the IPv6 address. Accepted ranges are 0-128 integers.

Changing the firewall settings

The `network firewall` commands are used to view or change the firewall settings.

To display the current firewall status

- ◆ To display whether the current firewall status is enabled or disabled, enter the following:

```
Network> firewall status
Firewall status : DISABLED
```

To enable the firewall setting

- ◆ To enable the firewall setting to allow specific IPs to connect to the ports while blocking the other connections, enter the following:

```
Network> firewall enable
ACCESS net INFO V-288-0 Firewall successfully enabled
```

To disable the firewall setting

- ◆ To disable the firewall setting and allow connections on any port from any IP, enter the following:

```
Network> firewall disable
It is not advisable to disable firewall. Do you want you continue (y/n):
ACCESS net INFO V-288-0 Firewall successfully disabled
```

The applied rules do not work when the firewall setting is disabled.

To display the list of firewall rules

- ◆ To display the list the firewall rules set on the cluster nodes by the user, enter the following:

```
Network> firewall rule list
```

iptype	filter	interface	client	protocols	sport	dport	match_state
=====	=====	=====	=====	=====	=====	=====	=====
ipv4	INPUT	pubeth0	ALL	tcp	101	102	NEW
ipv6	INPUT	pubeth0	ALL	tcp	101	102	NEW
ipv4	INPUT	pubeth0	ALL	udp	101	102	NEW
ipv6	INPUT	pubeth0	ALL	udp	101	102	NEW
ipv4	INPUT	pubeth1	ALL	tcp	101	102	NEW
ipv6	INPUT	pubeth1	ALL	tcp	101	102	NEW
ipv4	INPUT	pubeth1	ALL	udp	101	102	NEW
ipv6	INPUT	pubeth1	ALL	udp	101	102	NEW
ipv4	INPUT	pubeth2	ALL	tcp	101	102	NEW
ipv6	INPUT	pubeth2	ALL	tcp	101	102	NEW
ipv4	INPUT	pubeth2	ALL	udp	101	102	NEW
ipv6	INPUT	pubeth2	ALL	udp	101	102	NEW
ipv4	INPUT	pubeth3	ALL	tcp	101	102	NEW
ipv6	INPUT	pubeth3	ALL	tcp	101	102	NEW
ipv4	INPUT	pubeth3	ALL	udp	101	102	NEW
ipv6	INPUT	pubeth3	ALL	udp	101	102	NEW

To add a firewall rule

- ◆ To add the iptable rule in the current iptable configuration, enter the following:

```
Network> firewall rule add iptype filter interface
client protocols [sport] [dport] [match_state]
```

iptype	Specifies the network IP type. Allowed values are ipv4/ ipv6/ ipv4,ipv6.
filter	Specifies the iptable chain filter type. Allowed values are INPUT or OUTPUT
interface	Specifies the name of the network interface by which the packet will be received. If you enter ALL , an iptable entry is added for all public interfaces which are in control of the product.
client	Specifies the source IP from which the packet is received. ALL should be entered to apply the rule to all sources.
protocols	Specifies the protocol. Allowed values are tcp, udp, icmp and tcp,udp. When (tcp,udp) is given, two separate rules will be added, one for each protocol.
sport	Specifies the port through which the packet leaves the machine. The ALL option applies the rule to all the ports. NONE option is used to unspecify a port or enter specific port number. Note: <i>sport</i> and <i>dport</i> cannot both have NONE value at the same time.
dport	Specifies the port through which the packet is received. The ALL option applies the rule to all the ports. NONE option is used to unspecify a port or enter specific port number. Note: <i>sport</i> and <i>dport</i> cannot both have NONE value at the same time.
match_state	Specifies the match state for the connection. Enter NONE if you do not want to specify any state. Match state can be NEW/ ESTABLISHED/ RELATED.

To remove a firewall rule

- ◆ To remove the rule from the current iptable configuration, enter the following:

```
Network> firewall rule remove iptype filter interface
client protocols [sport] [dport] [match_state]
```

Configuring Access Appliance in IPv4 and IPv6 mixed mode

Access Appliance supports both IPv4 and IPv6 addresses. You can configure the Access Appliance cluster in an IPv4 environment, or an IPV6 environment, or in a mixed mode environment where you have both IPv4 and IPv6 addresses. You can use the Access Appliance services from the client by using IPv4 as well as IPv6 addresses but there are some limitations for specific services. When you configure Access Appliance in mixed mode, you are required to add both IPv4 and IPv6 physical IPs to the NIC(s) of all the cluster nodes.

Note: If you want to use the Access Appliance service over an IPv4 address, then you must configure the IPv4 address on both Access Appliance and the client. Similarly, if you want to use the Access Appliance service over an IPv6 address, then you must configure the IPv6 address on both Access Appliance and the client. Access Appliance does not convert IPV4 addresses to IPV6 addresses and vice versa for communication.

Converting an IPv4 Access Appliance cluster to mixed mode

You can use the below mentioned steps to convert an IPv4 Access Appliance cluster to mixed mode.

To convert an IPv4 Access Appliance cluster to mixed mode

- 1
- Install a Access Appliance cluster with IPv4 addresses.
- 2
- Connect to the cluster using the primary user credentials using the physical or primary IP of the cluster.

```
network> ip addr show
```

IP	Netmask/Prefix	Device	Node	Type	Status
---	-----	-----	----	----	-----
192.168.30.10	255.255.252.0	pubeth0	vascan_01	Physical	
192.168.30.11	255.255.252.0	pubeth0	vascan_02	Physical	
192.168.30.12	255.255.252.0	pubeth0	vascan_02	Virtual	ONLINE (Con IP)
192.168.30.13	255.255.252.0	pubeth0	vascan_02	Virtual	ONLINE
192.168.30.14	255.255.252.0	pubeth0	vascan_01	Virtual	ONLINE

3 Add the physical IPv6 IP.

```
network> ip addr add 2001:0:0:0::11 64 physical pubeth0 vascan_01
ACCESS ip addr WARNING V-493-10-1027 Default gateway for IPv6 is not
configured.
ACCESS ip addr SUCCESS V-493-10-1381 ip addr add successful.
network> ip addr add 2001:0:0:0::12 64 physical pubeth0 vascan_02
ACCESS ip addr WARNING V-493-10-1027 Default gateway for IPv6 is not
configured.
ACCESS ip addr SUCCESS V-493-10-1381 ip addr add successful.
```

4 Add the virtual IPV6 IP.

```
network> ip addr add 2001:0:0:0::13 64 virtual pubeth0
ACCESS ip addr WARNING V-493-10-1027 Default gateway for IPv6 is not
configured.
ACCESS ip addr SUCCESS V-493-10-1381 ip addr add successful.
network> ip addr add 2001:0:0:0::14 64 virtual pubeth0
ACCESS ip addr WARNING V-493-10-1027 Default gateway for IPv6 is not
configured.
ACCESS ip addr SUCCESS V-493-10-1381 ip addr add successful.
```

5 Add the IPv6 default gateway.

```
network> ip route add all :: 0 via 2001:0:0:0::1
```

6 Check the cluster configuration in mixed mode.

IP	Netmask/Prefix	Device	Node	Type	Status
--	-----	-----	----	----	-----
192.168.30.10	255.255.252.0	pubeth0	vascan_01	Physical	
192.168.30.11	255.255.252.0	pubeth0	vascan_02	Physical	
2001:0:0:0:0:0:0:11	64	pubeth0	vascan_01	Physical	
2001:0:0:0:0:0:0:12	64	pubeth0	vascan_02	Physical	
192.168.30.12	255.255.252.0	pubeth0	vascan_02	Virtual	ONLINE (Con IP)
192.168.30.13	255.255.252.0	pubeth0	vascan_02	Virtual	ONLINE
192.168.30.14	255.255.252.0	pubeth0	vascan_01	Virtual	ONLINE
2001:0:0:0:0:0:0:13	64	pubeth0	vascan_01	Virtual	ONLINE
2001:0:0:0:0:0:0:14	64	pubeth0	vascan_02	Virtual	ONLINE

Converting an IPv6 Access Appliance cluster to mixed mode

You can use the below mentioned steps to convert an IPv6 Access Appliance cluster to mixed mode.

To convert an IPv6 Access Appliance cluster to mixed mode

- 1 Install a Access Appliance cluster with IPv6 addresses.
- 2 Connect to the cluster using the primary user credentials using the physical or primary IP of the cluster.

```
network> ip addr show
```

IP	Netmask/Prefix	Device	Node	Type	Status
--	-----	-----	----	----	-----
2001:0:0:0:0:0:0:112	64	pubeth0	vascan_01	Physical	
2001:0:0:0:0:0:0:113	64	pubeth1	vascan_02	Physical	
2001:0:0:0:0:0:0:114	64	pubeth0	vascan_01	Virtual	ONLINE (Con IP)
2001:0:0:0:0:0:0:116	64	pubeth1	vascan_01	Virtual	ONLINE

- 3 Add the physical IPv4 IP.

```
network> ip addr add 192.168.30.10 255.255.252.0 physical pubeth0
vascan_01
ACCESS ip addr SUCCESS V-493-10-1381 ip addr add successful.
network> ip addr add 192.168.30.11 255.255.252.0 physical pubeth1
vascan_02
ACCESS ip addr SUCCESS V-493-10-1381 ip addr add successful.
```

- 4 Add the virtual IPV4 IP.

```
network> ip addr add 192.168.30.12 255.255.252.0 virtual pubeth0
ACCESS ip addr SUCCESS V-493-10-1381 ip addr add successful.1381 ip addr
add successful.
network> ip addr add 192.168.30.13 255.255.252.0 virtual pubeth1
ACCESS ip addr SUCCESS V-493-10-1381 ip addr add successful.
```

5 Add the IPv4 default gateway.

```
network> ip route add all 0.0.0.0 0.0.0.0 via 192.168.30.1
```

6 Check the cluster configuration in mixed mode.

```
network> ip addr show
```

IP	Netmask/Prefix	Device	Node	Type	Status
--	-----	-----	----	----	-----
192.168.30.10	255.255.252.0	pubeth0	vascan_01	Physical	
192.168.30.11	255.255.252.0	pubeth1	vascan_02	Physical	
2001:0:0:0:0:0:0:112	64	pubeth0	vascan_01	Physical	
2001:0:0:0:0:0:0:113	64	pubeth1	vascan_02	Physical	
2001:0:0:0:0:0:0:114	64	pubeth0	vascan_01	Virtual	ONLINE (Con IP)
192.168.30.12	255.255.252.0	pubeth0	vascan_01	Virtual	ONLINE
192.168.30.13	255.255.252.0	pubeth1	vascan_02	Virtual	ONLINE
2001:0:0:0:0:0:0:116	64	pubeth1	vascan_01	Virtual	ONLINE

Support for multiple data subnets

Access Appliance now supports multiple data subnets. This is applicable to all the protocols that the Access Appliance supports. If you have different subnets, you can access the services of the Access Appliance cluster via those subnets. The protocol specific entities and features decide how the subnets are utilized by the individual protocol. Multiple subnets can be configured with or without multiple VLANs.

Note: Data network isolation depends on the customer network configuration. The application configuration in Access Appliance does not restrict the customer from using the same network for multiple protocols.

If you have multiple subnets and you want to create multiple VLANs for each subnet, then Veritas recommends the following:

- Create a VLAN using the `network vlan add nic VLAN_ID` command, which creates a new device called `nic.VLAN_ID`.
- Add new virtual IP address belonging to the VLAN using the `network ip addr add` command.
- Add a route for the newly created subnet, if required.

If you want to create another VLAN, repeat the above steps.

```
va800a> network vlan add eth5 1979
Access Appliance INFO V-493-10-4071 VLAN Device will not get Physical
IP by default use following command to add IP
Access Appliance INFO V-493-10-4072 network ip addr add <ipaddr>
<netmask> <type> [device] [nodename]
100% [#] Configuring VLAN device
```

```
va800a> network vlan show
VLAN          DEVICE      VLAN id
----          -
eth5.1979     eth5          1979
```

```
va800a> network ip addr add 10.11.12.95 255.255.255.0 virtual
eth5.1979
Access Appliance ip addr SUCCESS V-493-10-1381 ip addr add successful.
```

```
va800a> network ip addr show
```

IP	Netmask/Prefix	Device	Node	Type	Status	FQDNs
192.168.10.130	255.255.255.0	eth4	va800a-01	Physical		
192.168.10.120	255.255.255.0	eth1	va800a-01	Physical		
192.168.10.131	255.255.255.0	eth4	va800a-02	Physical		
192.168.10.121	255.255.255.0	eth1	va800a-02	Physical		
192.168.10.124	255.255.255.0	eth1	va800a-01	Virtual	ONLINE (Con IP)	
192.168.10.140	255.255.255.0	eth4	va800a-01	Virtual	ONLINE	
192.168.10.141	255.255.255.0	eth4	va800a-02	Virtual	ONLINE	
192.168.10.142	255.255.255.0	eth5	va800a-02	Virtual	ONLINE	
192.168.10.143	255.255.255.0	eth5	va800a-01	Virtual	ONLINE	
10.11.12.95	255.255.255.0	eth5.1979	va800a-01	Virtual	ONLINE	

Configuring authentication services

This chapter includes the following topics:

- [About configuring LDAP settings](#)
- [Configuring LDAP server settings](#)
- [Administering the Access Appliance cluster's LDAP client](#)
- [About Active Directory \(AD\)](#)
- [Configuring AD server settings](#)
- [Configuring entries for Access Appliance DNS for authenticating to Active Directory \(AD\)](#)
- [Configuring AD/LDAP using the GUI](#)
- [Configuring the NIS-related settings](#)
- [Configuring NSS lookup order](#)
- [Sign-in options for the Access Appliance UI](#)
- [Configuring user authentication using digital certificates or smart cards](#)

About configuring LDAP settings

The Lightweight Directory Access Protocol (LDAP) is the protocol used to communicate with LDAP servers. The LDAP servers are the entities that perform the service. In Access Appliance, the most common use of LDAP is for user authentication.

For sites that use an LDAP server for access or authentication, Access Appliance provides a simple LDAP client configuration interface.

Before you configure Access Appliance LDAP settings, obtain the following LDAP configuration information from your system administrator:

- IP address or host name of the LDAP server. You also need the port number of the LDAP server.
- Base (or root) distinguished name (DN), for example:

```
cn=employees,c=us
```

LDAP database searches start here.

- Bind distinguished name (DN) and password, for example:

```
ou=engineering,c=us
```

This allows read access to portions of the LDAP database to search for information.

- Base DN for users, for example:

```
ou=users,dc=com
```

This allows access to the LDAP directory to search for and authenticate users.

- Base DN for groups, for example:

```
ou=groups,dc=com
```

This allows access to the LDAP database, to search for groups.

- Base DN for Netgroups, for example:

```
ou=netgroups,dc=com
```

This allows access to the LDAP database, to search for Netgroups.

- Root bind DN and password. This allows write access to the LDAP database, to modify information, such as changing a user's password.
- Secure Sockets Layer (SSL). Configures a cluster to use the Secure Sockets Layer (SSL) protocol to communicate with the LDAP server.
- Password hash algorithm, for example, `md5`, if a specific password encryption method is used with your LDAP server.

See [“Configuring LDAP server settings”](#) on page 66.

See [“Administering the Access Appliance cluster's LDAP client”](#) on page 68.

Configuring LDAP server settings

You can set the LDAP base Distinguished Name (base DN). LDAP records are structured in a hierarchical tree. You access records through a particular path, in this case, a Distinguished Name, or DN. The base DN indicates where in the LDAP directory hierarchy you want to start your search.

Note: For Access Appliance to access an LDAP directory service, you must specify the LDAP server DNS name or IP address.

To set the base DN for the LDAP server

- ◆ To set the base DN for the LDAP server, enter the following:

```
Network> ldap set basedn value
```

where *value* is the LDAP base DN in the following format:

```
dc=yourorg,dc=com
```

To set the LDAP server hostname or IP address

- ◆ To set the LDAP server hostname or IP address, enter the following:

```
Network> ldap set server value
```

where *value* is the LDAP server hostname or IP address.

To set the LDAP server port number

- ◆ To set the LDAP server port number, enter the following:

```
Network> ldap set port value
```

where *value* is the LDAP server port number.

To set Access Appliance to use LDAP over SSL

- ◆ To set Access Appliance to use LDAP over SSL, enter the following:

```
Network> ldap set ssl {on|off}
```

To set the bind DN for the LDAP server

- ◆ To set the bind DN for the LDAP server, enter the following:

```
Network> ldap set binddn value
```

where *value* is the LDAP bind DN in the following format:

```
cn=binduser,dc=yourorg,dc=com
```

The *value* setting is mandatory.

You are prompted to supply a password. You must use the password used to connect to the LDAP service on the specified LDAP server.

To set the root bind DN for the LDAP server

- ◆ To set the root bind DN for the LDAP server, enter the following:

```
Network> ldap set rootbinddn value
```

where *value* is the LDAP root bind DN in the following format:

```
cn=admin,dc=yourorg,dc=com
```

You are prompted to supply a password. You must use the password used to connect to the LDAP service on the specified LDAP server.

To set the LDAP users, groups, or netgroups base DN

- ◆ To set the LDAP users, groups, or netgroups base DN, enter the following:

```
Network> ldap set users-basedn value
```

```
Network> ldap set groups-basedn value
```

```
Network> ldap set netgroups-basedn value
```

users-basedn Specifies the value for the users-basedn. For example:

value

```
ou=users,dc=example,dc=com (default)
```

groups-basedn Specifies the value for the groups-basedn. For example:

value

```
ou=groups,dc=example,dc=com (default)
```

netgroups-basedn Specifies the value for the netgroups-basedn. For example:

value

```
ou=netgroups,dc=example,dc=com (default)
```

To set the password hash algorithm

- ◆ To set the password hash algorithm, enter the following:

```
Network> ldap set password-hash {clear|crypt|md5}
```

To display the LDAP configured settings

- ◆ To display the LDAP configured settings, enter the following:

```
Network> ldap get {server|port|basedn|binddn|ssl|rootbinddn|  
users-basedn|groups-basedn|netgroups-basedn|password-hash}
```

To clear the LDAP settings

- ◆ To clear the previously configured LDAP settings, enter the following:

```
Network> ldap clear {server|port|basedn|binddn|ssl|rootbinddn|  
users-basedn|groups-basedn|netgroups-basedn|password-hash}
```

To clear all the LDAP settings

- ◆ To clear the LDAP client configuration settings for all parameters, enter the following:

```
Network> ldap clearall
```

Administering the Access Appliance cluster's LDAP client

You can display the Lightweight Directory Access Protocol (LDAP) client configurations. LDAP clients use the LDAPv3 protocol to communicate with the server.

To display the LDAP client configuration

- ◆ To display the LDAP client configuration, enter the following:

```
Network> ldap show [users|groups|netgroups]
```

users	Displays the LDAP users that are available in the Name Service Switch (NSS) database.
groups	Displays the LDAP groups that are available in the NSS database.
netgroups	Displays the LDAP netgroups that are available in the NSS database.

If you do not include one of the optional variables, the command displays all the configured settings for the LDAP client.

To enable the LDAP client configuration

- ◆ To enable the LDAP client configuration, enter the following:

```
Network> ldap enable
```

LDAP clients use the LDAPv3 protocol for communicating with the server. Enabling the LDAP client configures the Pluggable Authentication Module (PAM) files to use LDAP. PAM is the standard authentication framework for Linux.

To disable the LDAP client configuration

- ◆ To disable the LDAP client configuration, enter the following:

```
Network> ldap disable
```

LDAP clients use the LDAPv3 protocol for communicating with the server. This command configures the PAM configuration files so that they do not use LDAP.

About Active Directory (AD)

In order to provide CIFS services, Access Appliance must be able to authenticate within the Windows environment.

Active Directory (AD) is a technology created by Microsoft that provides a variety of network services including LDAP directory services, Kerberos-based authentication, Domain Name System (DNS) naming, secure access to resources, and more.

You can configure Active Directory by navigating to **Settings > Directory Services management**.

Access Appliance will not join the AD domain if its clock is excessively out-of-sync with the clock on the AD domain controller. Ensure that Network Time Protocol (NTP) is configured on Access Appliance, preferably using the same NTP server as the AD domain controller.

Configuring AD server settings

The `network ad` commands are used to configure the Active Directory (AD) client for authentication. These commands configure the Access Appliance to use AD users and groups when logging into the Access Appliance.

If the AD client's domain controller is set to refuse NTLM authentication, run the following command to disable NTLM prior to configuring the Active Directory client:

```
CLFS> set ntlm_auth no
```

To set the AD client configuration

- ◆ You can set the AD client's domain, domain controller, workgroup and domain user. To set the AD client configuration details, enter the following:

```
Network> ad set domain domaincontroller workgroup  
domainuser
```

domain	Active Directory domain name or Windows NT domain name
domaincontroller	Primary[,backup] domain-controller names
workgroup	Windows WORKGROUP name or NetBIOS domain name
domainuser	domain user name which is used for authentication in the domain join operation

Access Appliance displays the cluster time as well as the time on the Active Directory Domain Controller.

If NTP has been configured correctly, there will be no time skew. Otherwise, you will need to reconfigure NTP correctly.

You will be prompted to enter the password of `domainuser`.

To enable the AD client

- ◆ To enable the AD client to use Active Directory for authentication, enter the following:

```
Network> ad enable
```

To display the AD client configuration

- ◆ To display the AD client configuration, enter the following:

```
Network> ad show
```

To disable the AD client

- ◆ To disable the AD client so that Active Directory is not used for authentication, enter the following:

```
Network> ad disable
```

To clear the AD client configuration

- ◆ To clear the AD client configuration, enter the following:

```
Network> ad unset
```

Configuring entries for Access Appliance DNS for authenticating to Active Directory (AD)

Name resolution must be configured correctly on Access Appliance. Domain Name System (DNS) is usually used for name resolution.

To configure entries for Access Appliance DNS for authenticating to Active Directory

- 1 Create an entry for the Access Appliance cluster name.

The cluster name is chosen at the time of installation, and it cannot be reset afterwards. It is also the NetBios name of the cluster, hence it must resolve to an IP address.

- 2 Configure the Access Appliance cluster name in DNS so that queries to it return the Virtual IP Addresses (VIPs) associated with the Access Appliance cluster in a round-robin fashion.

This is done by creating separate A records that map the cluster name to each VIP. So, if there are four VIPs associated with the Access Appliance cluster (not including special VIPs for backup, replication for Access Appliance, and so on), then there must be four A records mapping the cluster name to the four VIPs.

- 3 Verify that the DNS server has correct entries for Access Appliance by querying from a client:

```
# nslookup cluster name
```

After configuring the DNS server correctly, Access Appliance must be configured as a DNS client.

This is done during installation, but may be modified by using the following commands:

```
Network> dns set domainname domain_name
```

```
Network> dns set nameservers IP address
```

```
Network> dns enable
```

- 4 Verify that DNS client parameters are set correctly by entering the following command:

```
Network> dns show
```

- 5 Ensure host resolution is querying DNS by checking nsswitch:

```
Network> nsswitch show
```

In the above scenario, host resolution first looks at files, and then DNS.

Configuring name resolution correctly is critical in order to successfully join Access Appliance to Active Directory.

Configuring AD/LDAP using the GUI

You can use the Access Appliance GUI to configure an AD or LDAP user.

Prerequisites:

- Configure DNS to resolve AD and LDAP server domains on the Access Appliance cluster using the `network dns set` commands.
- Ensure host resolution is querying DNS by checking nsswitch configurations using the `network nsswitch show` command.
- Ensure that the password and groups are authenticated by winbind.
- Ensure that the time difference between AD/LDAP server and the Access Appliance cluster is less than 5 minutes.

To configure AD or LDAP user

- 1 Go to **Settings > Directory Services management**. Click **Configure**.
- 2 In the **Configure Directory Services** form, select the **Directory Type**. Select **Active Directory** for an AD user and **Open LDAP** for an LDAP user. Enter all the details on the form. Click **Configure**.
- 3 Once configuration is complete, the available configurations appear on the same page. Click on the individual hostname or the server IP to see the details of the configuration.

Configuring the NIS-related settings

Access Appliance supports Network Information Service (NIS), implemented in a NIS server, as an authentication authority. You can use NIS to authenticate computers.

If your environment uses NIS, enable the NIS-based authentication on the Access Appliance cluster.

Note: IPv6 addresses are not supported for NIS.

To display NIS-related settings

- ◆ To display NIS-related settings, enter the following:

```
Network> nis show [users|groups|netgroups]
```

users Displays the NIS users that are available in the Access Appliance cluster's NIS database.

groups Displays the NIS groups that are available in the Access Appliance cluster's NIS database.

netgroups Displays the NIS netgroups that are available in the Access Appliance cluster's NIS database.

To set the NIS domain name on all nodes in the cluster

- ◆ To set the NIS domain name on the cluster nodes, enter the following:

```
Network> nis set domainname [domainname]
```

where *domainname* is the domain name.

To set NIS server name on all nodes in the cluster

- ◆ To set the NIS server name on all cluster nodes, enter the following:

```
Network> nis set servername servername
```

where *servername* is the NIS server name. You can use the server's name or IP address.

To enable NIS clients

- ◆ To enable NIS clients, enter the following:

```
Network> nis enable
```

To view the new settings, enter the following:

```
Network> nis show
```

To disable NIS clients

- ◆ To disable NIS clients, enter the following:

```
Network> nis disable
```

Configuring NSS lookup order

Name Service Switch (NSS) is a cluster service that provides a single configuration location to identify the services (such as NIS or LDAP) for network information such as hosts, groups, netgroups, passwords, and shadow files.

For example, host information may be on an NIS server. Group information may be in an LDAP database.

The NSS configuration specifies which network services the Access Appliance cluster should use to authenticate hosts, users, groups, and netgroups. The configuration also specifies the order in which multiple services should be queried.

To display the current value set on NSS for all groups, hosts, netgroups, passwd, and shadow files

- ◆ To display the current value set on nsswitch for all groups, hosts, netgroups, passwd, and shadow files

```
Network> nsswitch show
```

To change the order of group items

- ◆ To configure the NSS lookup order, enter the following:

```
Network> nsswitch conf {group|hosts|netgroups|passwd|shadow}  
value1 [[value2]] [[value3]] [[value4]]
```

group	Selects the group file.
hosts	Selects the hosts file.
netgroups	Selects the netgroups file.
passwd	Selects the password.
shadow	Selects the shadow file.
value	Specifies the following NSS lookup order with the following values: ■ <i>value1</i> (required)- { files/nis/winbind/ldap } ■ <i>value2</i> (optional) - { files/nis/winbind/ldap } ■ <i>value3</i> (optional) - { files/nis/winbind/ldap } ■ <i>value4</i> (optional) - { files/nis/winbind/ldap }

For example:

```
Network> nsswitch conf group nis files  
Network> nsswitch show
```

To select DNS, you must use the following command:

```
Network> nsswitch conf hosts
```

Sign-in options for the Access Appliance UI

The following user authentication methods are supported for accessing the UI:

- Local user name and password
- LDAP domain users
- Digital certificate or smart card, including Common Access Card (CAC) for LDAP users

Configuring user authentication using digital certificates or smart cards

You can configure Access Appliance to authenticate users with a smart card or a digital certificate. After configuration, the users can use the **Sign in with certificate or smart card** option to sign in to Access Appliance UI using smart cards or digital certificates.

Before you configure user authentication using smart cards or digital certificates, note the following:

- Ensure that LDAP is configured. Go to **Settings > Directory Services management** and click **Configure** to configure LDAP, if not already done so. For details about how to configure LDAP:
See [“About configuring LDAP settings”](#) on page 64.
See [“Configuring LDAP server settings”](#) on page 66.
- Smart card authentication can be configured only for LDAP users.
- Smart card authentication requires a list of trusted root or intermediate CA certificates. You must add the CA certificates that are associated with the user digital certificates or the user smart cards.

To configure Access Appliance to authenticate users with a certificate or smart card:

- 1 Log in to Access Appliance UI.
- 2 In the left navigation pane, click **Settings > Security management**, and then click **Smart card authentication**.
- 3 Use the slider to turn on smart card authentication.

- 4 In the **Configure smart card authentication** dialog box, specify the following options:
 - In the **User authentication domain** list, select the configured LDAP server.
 - Click **Common name** to select the common mapping attribute.
 - Optionally, enter the Online Certificate Status Protocol (OCSP) URI. OCSP is used for checking the validity of the certificate.
If you do not provide the OCSP URI, the URI in the user certificate is used.
- 5 Click **Save**.
- 6 To the right of **CA certificates** click **Add**.
- 7 Click **Browse** to select the CA certificate or drag and drop the CA certificate and click **Add**.

Certificates must be in PEM format, with certificate file type as `.pem`. Only one certificate can be added at a time. The web server is restarted after adding the certificate.

The selected CA certificate is displayed under **CA certificates**.
- 8 Upload the client certificate to the browser's certificate store. See the browser documentation for importing client certificates.

The users can now use the **Sign in with certificate or smart card** option to sign in to the Access Appliance UI. LDAP users with Appliance administrator role have access to all settings in the UI. LDAP users that do not have this role, can only create S3 keys.

Adding CA certificates for smart card authentication

Smart card authentication requires a list of trusted root or intermediate CA certificates. Add the CA certificates that are associated with the user digital certificates or the user smart cards.

To add a CA certificate:

- 1 Log in to Access Appliance UI.
- 2 In the left navigation pane, click **Settings > Security management**, and then click **Smart card authentication**.
- 3 To the right of **CA certificates** click **Add**.
- 4 Click **Browse** to select the certificate or drag and drop the CA certificate and click **Add**.

The web server is restarted after adding the certificate.

The selected certificate is displayed under **CA certificates**.

Deleting CA certificates

You can delete a CA certificate if it is no longer used for smart card authentication. If a user attempts to use the associated digital certificate or smart card certificate, they are not able to sign in to Access Appliance.

To delete a CA certificate:

- 1 Log in to Access Appliance UI.
- 2 In the left navigation pane, click **Settings > Security management**, and then click **Smart card authentication**.
- 3 Select the CA certificate that you want to delete and click **Delete**.

The certificate is deleted and no longer displayed under **CA certificates**. The web server is restarted after the certificate is deleted.

Managing Access Appliance storage

- [Chapter 5. Configuring storage](#)
- [Chapter 6. Managing disks](#)
- [Chapter 7. Access Appliance as an iSCSI target](#)

Configuring storage

This chapter includes the following topics:

- [About storage provisioning and management](#)
- [About configuring disks](#)
- [About configuring storage pools](#)
- [Configuring storage pools](#)
- [About quotas for usage](#)
- [Enabling, disabling, and displaying the status of file system quotas](#)
- [Setting and displaying file system quotas](#)
- [Setting user quotas for users of specified groups](#)
- [About quotas for CIFS home directories](#)
- [Workflow for configuring and managing storage using the Access Appliance CLI](#)
- [Displaying information for all disk devices associated with the nodes in a cluster](#)
- [Displaying WWN information](#)
- [Importing new LUNs forcefully for new or existing pools](#)
- [Initiating host discovery of LUNs](#)

About storage provisioning and management

When you provision storage, you want to be able to assign the appropriate storage for the particular application. Access Appliance supports a variety of storage types. To help the users that provision the storage to select the appropriate storage, you

classify the storage into groups called storage pools. A storage pool is a user-defined way to group the disks that have similar characteristics.

Access Appliance supports a wide variety of storage arrays, direct attached storage as well as in-server SSDs and HDDs. During the initial configuration, you add the disks to the Access Appliance nodes. For a storage array, a disk is a LUN from the storage array. For best performance and resiliency, each LUN should be provisioned to all Access Appliance nodes. Local disks and fully shared disks have unique names, but partially shared disks across nodes may have the same name. Make sure that you do not assign LUNs from the same enclosure to different nodes partially.

Before you can provision storage to Access Appliance, the physical LUNs must be set up and zoned for use with the Access Appliance cluster. The storage array administrator normally allocates and zones the physical storage.

After the disks are correctly discovered by Access Appliance, you assign the disks to storage pools. You create a file system on one or more storage pools. You can mirror across different pools.

You can also use local disks that are shared over the network. Both DAS disks and SAN disks (LUNs) can be used by the same cluster, and you can have a mix of DAS and SAN disks in the same storage pool.

About configuring disks

Disks and pools can be specified in the same command provided the disks are part of an existing storage pool.

The pool and disk that are specified first are allocated space before other pools and disks.

If the specified disk is larger than the space allocated, the remainder of the space is still utilized when another file system is created spanning the same disk.

About configuring storage pools

Access Appliance uses storage pools to provision storage. Pools are more a logical construct rather than an architectural component. Pools are loosely collections of disks.

In the Access Appliance context, a disk is a LUN provisioned from a storage array. Each LUN should be provisioned to all Access Appliance nodes. Disks must be added to pools before you use them.

During the initial configuration, you create storage pools, to discover disks, and to assign them to pools. Disk discovery and pool assignment are done once. Access Appliance propagates the disk information to all the cluster nodes.

You must first create storage pools that can be used to build file systems on.

By default, all of the storage pools in Access Appliance share the same configuration. Copies of the configuration reside on disks in the storage pools. The first storage pool you create uses the default configuration. You can create additional storage pools to be part of that default configuration or to be isolated. An isolated storage pool protects the pool from losing the associated metadata even if all configuration disks in the main storage pool fail. If isolated storage pools exist, you cannot remove the disks from a non-isolated pool which has an internal file system. You also cannot destroy a non-isolated pool if it is the last remaining non-isolated pool.

Configuring storage pools

A storage pool is a group of disks that Access Appliance uses for allocation. Before creating a file system, you must create a storage pool.

To create the storage pool used to create a file system

- 1 List all of the available disks, and identify which ones you want to assign to which pools.

```
Storage> disk list
```

- 2 To create a storage pool, enter the following:

```
Storage> pool create pool_name disk1[,disk2,...] [isolated=yes|no]
```

pool_name	Specifies what the created storage pool will be named. The storage pool name should be a string.
disk1, disk2,...	Specifies the disks to include in the storage pool. If the specified disk does not exist, an error message is displayed. Use the <code>Storage> disk list</code> command to view the available disks. Each disk can only belong to one storage pool. If you try to add a disk that is already in use, an error message is displayed. To specify additional disks to be part of the storage pool, use a comma with no space in between.
isolated=yes no	Optional. Specifies whether or not the storage pool is isolated from other storage pools. Isolating the storage pool means that the configuration information is not shared. By default, storage pools are not isolated.

To add a set of disks to a logical pool

- ◆ To add a set of disks to a logical pool, enter the following:

```
Storage> pool adddisk pool_name disk1 [,disk2,...]
```

Where *pool_name* specifies the name of the storage pool to which the disks have to be added.

disk1, disk2,.. specifies the disks to be added to the pool.

To list your pools

- ◆ To list your pools, enter the following:

```
Storage> pool list
```

If a node is down, the `Storage> pool list` command shows local disks of that node.

To rename a pool

- ◆ To rename a pool, enter the following:

```
Storage> pool rename old_name new_name
```

<i>old_name</i>	Specifies the name for the existing pool that will be changed. If the old name is not the name of an existing pool, an error message is displayed.
<i>new_name</i>	Specifies the new name for the pool. If the specified new name for the pool is already being used by another pool, an error message is displayed.

To destroy a storage pool

- 1 Because you cannot destroy an unallocated storage pool, you need to remove the disk from the storage pool using the `Storage> pool rmdisk` command prior to trying to destroy the storage pool.

See [“About configuring disks”](#) on page 81.

If you want to move the disk from the unallocated pool to another existing pool, you can use the `Storage> pool mvdisk` command.

- 2 To destroy a storage pool, enter the following:

```
Storage> pool destroy pool_name
```

Where *pool_name* specifies the storage pool to delete. If the specified *pool_name* is not an existing storage pool, an error message is displayed.

If a node is down temporarily, it is not a good practice to destroy a storage pool that contains local disks of that node.

Note: You cannot destroy the last non-isolated pool if isolated pools exist.

To mark a disk as a spare disk

- ◆ To mark a disk as a spare disk and add it to a pool which can be later used for hot relocation, enter the following:

```
Storage> pool markdiskspare pool_name disk1 [, disk2,...]
```

Where *pool_name* specifies the name of the storage pool to which the disks have to be added.

disk1, disk2,.. specifies the disks to be marked as spare.

In case of failure of a disk or a plex, the affected sub disks are relocated to disks designated as spare disks.

To remove the spare disk flag set on a disk

- ◆ To remove the spare disk flag set on a disk by the `Storage> pool markdiskspare` command, enter the following:

```
Storage> pool removediskspare pool_name  
disk1 [, disk2,...]
```

Where *pool_name* specifies the name of the storage pool to which the disks belong.

disk1, disk2,.. specifies the disks from which the spare disk flag has to be removed.

To list free space for pools

- ◆ To list free space for your pool, enter the following:

```
Storage> pool free [pool_name]
```

Where *pool_name* specifies the pool for which you want to display free space information.

If a specified pool does not exist, an error message is displayed.

If *pool_name* is omitted, the free space for every pool is displayed, but information for specific disks is not displayed.

About quotas for usage

Disk quotas limit the usage for users or user groups. You can configure disk quotas for file systems or for CIFS home directories.

Note: Quotas work over NFS, but quota reporting and quota details are not visible over NFS.

Users and groups visible through different sources of name service lookup (nsswitch), local users, LDAP, NIS, and Windows users can be configured for file systems or CIFS home directory quotas.

There are two types of disk quotas:

- Usage quota (numspace) - limits the amount of disk space that can be used on a file system.
The numspace quota value must be an integer with a unit. The minimum unit is KB. VxFS calculates numspace quotas based on the number of KBs. The range for numspace is from 1K to 9007199254740991(2⁵³ - 1)K.
- Inode quota (numinodes) - limits the number of inodes that can be created on a file system.
An inode is a data structure in a UNIX or UNIX-like file system that describes the location of some or all of the disk blocks allocated to the file.
The numinodes quota value must be an integer without a unit, and the range is from 1 to 999999999999999999(19bit).
0 is valid for numspace and numinodes, which means the quota is infinite.

Access Appliance supports disk quota limits greater than 2 TB.

In addition to setting a limit on disk quotas, you can also define a warning level, or soft quota, whereby the Access Appliance administrator is informed that they are nearing their limit, which is less than the effective limit, or hard quota. Hard quota limits can be set so that a user is strictly not allowed to cross quota limits. A soft quota limit must be less than a hard quota limit for any type of quota.

Note: The alert for when a hard limit quota or a soft limit quota is reached in Access Appliance is not sent out immediately. The hard limit quota or soft limit quota alert is generated by a cron job scheduled to run daily at midnight.

Enabling, disabling, and displaying the status of file system quotas

To configure file system quotas, you must enable the file system quotas. You can enable file system quotas for all file systems or specify a file system name. You can enable quotas per user (**userquota**), quotas per group (**groupquota**), or both.

Once the quotas are enabled, you set the values for the number of blocks or the number of inodes that can be created. Quotas can be hard limits or soft limits.

See [“Setting and displaying file system quotas”](#) on page 87.

To enable a file system quota

- ◆ To enable a file system quota, enter the following:

```
Storage> quota fs enable [fs_name] [userquota | groupquota]
```

To disable a file system quota

- ◆ To disable a file system quota, enter the following:

```
Storage> quota fs disable [fs_name] [userquota | groupquota]
```

To display the status of a file system quota

- ◆ To display the status of a file system quota, enter the following:

```
Storage> quota fs status [fs_name] [userquota | groupquota]
```

Note: If the LDAP client is disabled, then the quota information may not be displayed using the `Storage> quota show` command for LDAP users and groups.

Setting and displaying file system quotas

You can set usage quotas for users or for groups.

Before the file system quotas take effect, you must enable the quotas for the file system.

See [“Enabling, disabling, and displaying the status of file system quotas”](#) on page 86.

To set the quota value

- ◆ To set the quota value for a file system, enter the following:

```
Storage> quota fs set {userquota | groupquota} user_or_group_names  
domain_name [hardlimit | softlimit] [numinodes | numspace]  
[value] [fs_name]
```

Domain name is the first section of the domain, for example:

```
veritas.example.com
```

The domain name is `veritas` in the example above.

If a value is not provided, the default value is used.

To set all quota values

- ◆ To set all of the quota values, enter the following:

```
Storage> quota fs setall {userquota | groupquota}
[hardlimit | softlimit] [numinodes | numspace]
[value] [fs_name]
```

For example, to set all existing user quotas to default values for the file system fs1:

```
Storage> quota fs show fs1

Storage> quota fs setall userquota
```

To display the file system settings

- ◆ To display the file system settings, enter the following:

```
Storage> quota fs show [fs_name] [userquota | groupquota]
[user_or_group_names]
```

For example, to display quota values for the file systemx:

```
Storage> quota fs show
```

Note: If the LDAP client is disabled, then the quota information may not be displayed using the `Storage> quota show` command.

fs_name	File system name you want to set the quota for.
userquota	User quota can be set with hard or soft limits on usage. Usage is dictated by the number of blocks and number of inodes that are created by the user.
groupquota	Group quota can be set with hard or soft limits on usage. Usage is dictated by the number of blocks and number of inodes that are created by all the users in the group.

<code>user_or_group_names</code>	Name of the user or the name of the group for which a quota value is set. You can specify a comma-separated list of user or group names. To delete quota values for a user, you have to set all the user quota entries to 0. A user with a UID of 0 is not allowed in a <code>Storage> quota fs set</code> command.
----------------------------------	---

To set the default quota values

- ◆ To set the default quota values, enter the following:

```
Storage> quota fs setdefault {userquota | groupquota}
{hardlimit | softlimit} {numinodes | numspace} [value] [fs_name]
```

To display the default values

- ◆ To display the default values, enter the following:

```
Storage> quota fs showdefault [fs_name] [userquota | groupquota]
```

For example, to display the default quota values:

```
Storage> quota fs showdefault
```

Setting user quotas for users of specified groups

You can set the same quota for each user in a group with a single command. As with the other quota commands, you can specify hard or soft limits for the number of inodes or the usage space.

To set user quotas for users of specified groups

- ◆ To set user quotas on users of specified groups, enter the following:

```
Storage> quota fs setbygroup group_names domain_name
[hardlimit | softlimit] [numinodes | numspace]
[value] [fs_name]
```

About quotas for CIFS home directories

You use `Storage> quota cifshomedir` commands to configure quotas for CIFS home directories. Users and groups visible through different sources of name service

lookup (nsswitch), local users, LDAP, NIS, and Windows users can be configured for CIFS home directory quotas.

Default values are entered in a configuration file only. The actual application of the quota is done with the `set` and `setall` commands using the default values provided.

When a CIFS home directory file system is changed, quota information for a user's home directory is migrated from the existing home directory file system to the new home directory file system.

Quota migration results are based on the following logic:

- Case 1:

In the case where the existing home directory file system is NULL, you can set the new home directory file system to be multiple file systems (for example, fs1, fs2). If the multiple file systems previously had different quota values, the quota status and values from the first file system are migrated to other file systems in the new home directory. The first file system is the template. Only the user/group quota values that existed on the first file system are migrated. Other user/group quota values remain the same on the other file system.

For example, assume the following:

- The new home directory file systems are fs1 and fs2.
- user1, user2, and user3 have quota values on fs1.
- user2, user3, and user4 have quota values on fs2.

For the migration, user/group quota values for user1, user2, and user3 are migrated from fs1 to fs2. Quota values for user4 are kept the same on fs2, and user4 has no quota values on fs1.

- Case 2:

When the existing home directory file systems are already set, and you change the file systems for the home directory, the quota status and values need to be migrated from the existing home directory file systems to the new file systems. For this migration, the first file system in the existing home directory acts as the template for migrating quota status and values.


For example, if the existing home directory file systems are fs1 and fs2, and the file systems are changed to fs2, fs3, and fs4, then the user/group quota values on fs1 are migrated to fs3 and fs4. Other user/group values on fs3 and fs4 remain the same.

Workflow for configuring and managing storage using the Access Appliance CLI

Figure 5-1 describes configuring and managing storage using the Access Appliance CLI.

See the Access Appliance manual pages for the detailed syntax for completing the operations.

Figure 5-1 Workflow for configuring and managing Access Appliance storage using the CLI



Displaying information for all disk devices associated with the nodes in a cluster

You can display disk information for the disk devices associated with the nodes in the Access Appliance cluster. If local disks are present, the information includes entries for the local disks.

See the `storage_disk(1)` man page for the detailed examples.

The information displayed depends on the form of the command that you use. The following information is available:

Disk	Indicates the disk name.
Serial Number	Indicates the serial number for the disk.
Enclosure	Indicates the type of storage enclosure.
Size	Indicates the size of the disk.
Use%	Indicates the percentage of the disk that is being used.
Transport	Indicates transport protocol values like SCSI, FC, and other values.
ID	ID column consists of the following four fields. A ":" separates these fields. ■ VendorID - Specifies the name of the storage vendor, for example, HITACHI, IBM, EMC, and so on. ■ ProductID - Specifies the ProductID based on vendor. Each vendor manufactures different products. For example, HITACHI has HDS5700, HDS5800, and HDS9200 products. These products have ProductIDs such as DF350, DF400, and DF500. ■ TargetID - Specifies the TargetID. Each port of an array is a target. Two different arrays or two ports of the same array have different TargetIDs. TargetIDs start from 0. ■ LunID - Specifies the ID of the LUN. This should not be confused with the LUN serial number. LUN serial numbers uniquely identify a LUN in a target. Whereas a LunID uniquely identifies a LUN in an initiator group (or host group). Two LUNS in the same initiator group cannot have the same LunID. For example, if a LUN is assigned to two clusters, then the LunID of that LUN can be different in different clusters, but the serial number is the same.
Array Type	Indicates the type of storage array and can contain any one of the three values: Disk for JBODs, Active-Active, and Active-Passive.

To display a list of disks and nodes

- ◆ To display a list of disks and nodes, enter the following:

```
Storage> disk list
```

This form of the command displays local disk information for all nodes in the cluster.

To display the disk information

- ◆ To display the disk information, enter the following:

```
disk list detail
```

This form of the command displays local disk information from all the nodes in the cluster.

To display the disk list paths

- ◆ To display the disks multiple paths, enter the following:

```
Storage> disk list paths
```

This form of the command displays local disk information from all the nodes in the cluster.

Displaying WWN information

The `Storage> hba` (Host Bus Adapter) command displays World Wide Name (WWN) information for all of the nodes in the cluster. If you want to find the WWN information for a particular node, specify the node name (host name).

To display WWN information

- ◆ To display the WWN information, enter the following:

```
Storage> hba [host_name]
```

where you can use the `host_name` variable if you want to find WWN information for a particular node.

To display WWN information for all the running nodes in the cluster, enter the following:

To display WWN information for a particular node, enter the following:

```
Storage> hba [host_name]
```

HBA_Node_Name	Displays the node name for the Host Bus Adapter (HBA).
WWN	Displays World Wide Name (WWN) information.
State	Available values include: ■ online ■ offline
Speed	Displays the speed per second.
Support_Classes	Displays the class value from <code>/sys/class/fc_host/\${host}/supported_classes</code> .
Transmitted_FC_Frames	Displays a value equal to the number of total transmitted serial attached SCSI frames across all protocols.
Received_FC_frames	Displays a value equal to the number of total received serial attached SCSI frames across all protocols.
Link_Failure_Count	Displays a value equal to the value of the LINK FAILURE COUNT field of the Link Error Status.

Importing new LUNs forcefully for new or existing pools

The `Storage> scanbus force` command tries to import Logical Unit Numbers (LUNs) forcefully. This may help when using `Storage> scanbus` alone does not work.

To import LUNs forcefully

- ◆ To import LUNs forcefully, enter the following:

```
Storage> scanbus [force]
```

Initiating host discovery of LUNs

The `Storage> scanbus` command scans all of the SCSI devices connected to all of the nodes in the cluster. When you add new storage to your devices, you must scan for new SCSI devices. You only need to issue the command once and all of the nodes discover the newly added disks. The `scanbus` command updates the device configurations without interrupting the existing I/O activity. The scan does

not inform you if there is a change in the storage configuration. You can see the latest storage configuration using the `Storage> disk list` command.

You do not need to reboot after scanbus has completed.

To scan SCSI devices

- ◆ To scan the SCSI devices connected to all of the nodes in the cluster, enter the following:

```
Storage> scanbus
```

Managing disks

This chapter includes the following topics:

- [Adding disks to a storage pool](#)
- [Removing disks from a storage pool](#)
- [Viewing information about disks](#)
- [Accessing disk details](#)
- [Discovering disks](#)
- [Formatting a disk](#)

Adding disks to a storage pool

Storage pools are used to classify storage using user-defined criteria. The best practice is to group disks with similar characteristics into a storage pool. A single storage pool can contain both SSD and HDD storage.

To add disks to storage pool

1 Prerequisites:

Disks must be initialized and discovered by the cluster before you can add them to a storage pool.

2 To add the selected disks to an existing pool or specify a new storage pool, select the disks and click **Add to storage pool**.

If a disk already belongs to a pool, you cannot add it to a new pool.

If you select an existing pool, the Storage Pool Capacity shows the existing capacity for the selected pool.

3 View the **Recent Activity** panel for the status of the task to create the storage pool.

See [“Removing disks from a storage pool”](#) on page 97.

See [“About storage provisioning and management ”](#) on page 80.

See [“Viewing information about disks”](#) on page 97.

See [“Monitoring command activity”](#) on page 212.

Removing disks from a storage pool

Storage pools are used to classify storage using user-defined criteria. The best practice is to group disks with similar characteristics into a storage pool. A single storage pool can contain both SSD and HDD storage.

To remove disks from a storage pool

1 Prerequisites:

Disks must not be in a used state.

2 Go to **Infrastructure > Disks**. Select the disks and click **Remove from storage pool**.

3 View the **Recent Activity** panel for the status of the task.

See [“Adding disks to a storage pool”](#) on page 96.

See [“About storage provisioning and management ”](#) on page 80.

See [“Monitoring command activity”](#) on page 212.

Viewing information about disks

From the **Infrastructure > Disks** page, you can view the following details:

The following details are displayed for the disks:

- Name of the disk
- Status of the disk
- Percentage of disk space that is provisioned
- Total disk space
- Storage pool the disk is assigned to
- Vendor name

Click a disk to see more details.

See [“About storage provisioning and management ”](#) on page 80.

See [“Adding disks to a storage pool”](#) on page 96.

See [“Removing disks from a storage pool”](#) on page 97.

Accessing disk details

Click on the name of the disk to view detailed information for the specified disk.

To access disk details

- 1 Navigate:
- On the **Infrastructure > Disks** page, click the disk name to see its details.
- 2 You can view the following information for the disk:

Disk Name	Name of the disk.
Serial No	Serial number for the disk.
ID	Vendor ID for the disk.
Version	Firmware version for the disk.
Storage Pool	Name of the storage pool that is associated with the disk.
Node(s)	The nodes on which the disk is visible.
File System	File systems that are associated with the disk.
Usage	The capacity consumption for the disk, including the following values: ■ Used Space - currently allocated disk capacity. ■ Free Space - unallocated disk capacity. ■ Total Space - total space available on the disk.

In the **Disk Paths** table, you can view the following information for the nodes that are associated with the specified disk:

Path Name	The name of the path that is associated with the specified disk.
Status	The status of the path.

See [“About storage provisioning and management ”](#) on page 80.

See [“Viewing information about disks”](#) on page 97.

Discovering disks

The discover disks operation scans all of the SCSI devices connected to both the nodes in the cluster. When you add new storage to your devices, you must scan for new SCSI devices. You only need to issue the command once and all of the nodes discover the newly added disks. The discover disks operation updates the device configuration without interrupting the existing I/O activity.

To discover the new disks:

1 Prerequisites:

Disks must be initialized and discovered by the cluster before you can add them to a storage pool.

Add the selected disks to an existing pool or specify a new storage pool.

2 On the **Infrastructure** page, on the **Disks** tab click **Discover disks**.

3 View the **Recent Activity** panel for the status of the task.

See [“About storage provisioning and management”](#) on page 80.

See [“Viewing information about disks”](#) on page 97.

Formatting a disk

You can format or reinitialize a disk. If the disk does not belong to any disk group, the format disk operation erases the first 100-MB space on the disk(s).

To format a disk

1 Prerequisites:

Disks must be initialized and discovered by the cluster.

2 **Infrastructure > Disks**: Select the disk that you want to format.

3 Click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Format**.

4 Follow the wizard prompts.

5 View the **Recent Activity** panel for the status of the task.

See [“About storage provisioning and management”](#) on page 80.

See [“Viewing information about disks”](#) on page 97.

Access Appliance as an iSCSI target

This chapter includes the following topics:

- [About Access Appliance as an iSCSI target](#)
- [Managing the iSCSI target service](#)
- [Managing the iSCSI targets](#)
- [Managing the LUNs](#)
- [Managing the mappings with iSCSI initiators](#)
- [Managing the users](#)
- [Creating an iSCSI target and provisioning LUNs](#)

About Access Appliance as an iSCSI target

You can create an iSCSI target and provisions LUNs for storage. Access Appliance as an iSCSI target feature enables a Access Appliance cluster to serve block storage. Through the use of multiple portal IPs, an iSCSI target can be served in active-active fashion.

This feature enables the block storage to be capable of supporting multipathing at the initiator end. Access Appliance eases provisioning of block storage, with the functionality to resize, clone, and snapshot the LUNs, ACL controls such as initiator mapping and user management.

Note: Access Appliance as an iSCSI target supports VMware version 5.5.0 as an initiator.

You can perform the following functions on an iSCSI target:

- Start, stop, and check status of the iSCSI target service.
- Create, destroy, check status, and list iSCSI targets and add and delete multiple portal addresses.
- Add or delete, resize, manage, grow or shrink LUNs, and clone LUNs snapshots.
- Map or remove mapping of iSCSI initiators.
- Add or delete users to set up CHAP authentication.
- Support for multiple portal IPs per target makes the targets active-active.

Managing the iSCSI target service

To start the iSCSI target service

- ◆ To start the iSCSI target service, enter the following:

```
Target> iscsi service start
```

To stop the iSCSI target service

- ◆ To stop the iSCSI target service, enter the following:

```
Target> iscsi service stop
```

To check the status of the iSCSI target service

- ◆ To check the status of the iSCSI target service, enter the following:

```
Target> iscsi service status
```

Managing the iSCSI targets

To create an iSCSI target

- ◆ To create an iSCSI target, enter the following:

```
Target> iscsi target create target-name
```

target-name Name of the iSCSI target at which SCSI LUNs are available.
 target-name should conform to the naming rules defined in RFC3721.

To destroy an iSCSI target

- ◆ To destroy an iSCSI target, enter the following:

```
Target> iscsi target destroy target-name
```

To list all iSCSI targets

- ◆ To list all iSCSI targets, enter the following:

```
Target> iscsi target list
```

To list the specific iSCSI target

- ◆ To list the specific iSCSI target, enter the following:

```
Target> iscsi target list target-name
```

To check the status of a specific iSCSI target

- ◆ To check the status of a specific iSCSI target, enter the following:

```
Target> iscsi target status target-name
```

To add multiple portal addresses to an iSCSI target

- ◆ To add multiple portal addresses to an iSCSI target, enter the following:

```
Target> iscsi target portal add target-name portal-ip
```

target-name	Name of the iSCSI target at which SCSI LUNs are available. <code>target-name</code> should conform to the naming rules defined in RFC3721.
-------------	--

portal-ip	The virtual IP through which the target is accessible. You can specify multiple portal addresses that you want to add.
-----------	--

To delete multiple portal addresses from an iSCSI target

- ◆ To delete multiple portal addresses from an iSCSI target, enter the following:

```
Target> iscsi target portal del target-name portal-ip
```

target-name	Name of the iSCSI target at which SCSI LUNs are available. <code>target-name</code> should conform to the naming rules defined in RFC3721.
-------------	--

portal-ip	The virtual IP through which the target is accessible. You can specify multiple portal addresses that you want to delete.
-----------	---

To map a file system with a specified iSCSI target

- ◆ To map a file system with a specified iSCSI target, enter the following:

```
Target> iscsi target store add fs-name target-name
```

fs-name	Name of the file system which is to be mapped to an iSCSI target. LUNs are created as files on this file system.
---------	--

target-name	Name of the iSCSI target at which SCSI LUNs are available. <code>target-name</code> should conform to the naming rules defined in RFC3721.
-------------	--

To remove the file system from a specified iSCSI target

- ◆ To remove the file system from a specified iSCSI target, enter the following:

```
Target> iscsi target store delete fs-name target-name
```

fs-name	Name of the file system which is to be removed from an iSCSI target.
target-name	Name of the iSCSI target at which SCSI LUNs are available. <code>target-name</code> should conform to the naming rules defined in RFC3721.

Managing the LUNs

To create a LUN with the specified name and size

- ◆ To create a LUN with the specified name and size on the specified file system, enter the following:

```
Target> iscsi lun create lun-name target-name size option=dense|sparse
```

lun-name	Name of the LUN which is to be created on the store that is associated with an iSCSI target.
target-name	Name of the iSCSI target at which SCSI LUNs are available. <code>target-name</code> should conform to the naming rules defined in RFC3721.
size	Size of the LUN that you want to create.
option	Name of the LUN type such as Dense or Sparse. By default, Sparse is selected.

To destroy a specific LUN

- ◆ To destroy a specific LUN, enter the following:

```
Target> iscsi lun destroy lun-name target-name [force]
```

lun-name Name of the LUN which is to be destroyed from an iSCSI target.

target-name Name of the iSCSI target at which SCSI LUNs are available.
 target-name should conform to the naming rules defined in
 RFC3721.

Note: You can use this command to destroy a clone of a specified LUN as well.

To list the details of all the LUNs present in all targets

- ◆ To list the details of all the LUNs present in all targets, enter the following:

```
Target> iscsi lun list
```

To list the details of all the LUNs present in a specific target

- ◆ To list the details of all the LUNs present in a specific target, enter the following:

```
Target> iscsi lun list target-name
```

target-name Name of the iSCSI target at which SCSI LUNs are available.
 target-name should conform to the naming rules defined in
 RFC3721.

To grow LUN size to specified size

- ◆ To grow LUN size to specified size, enter the following:

```
Target> target iscsi lun growto lun-name target-name size
```

lun-name	Name of the LUN for which size is to be increased on the store that is associated with an iSCSI target.
target-name	Name of the iSCSI target at which SCSI LUNs are available. <code>target-name</code> should conform to the naming rules defined in RFC3721.
size	A new size that is to be specified for the LUN.

To shrink LUN size to specified size

- ◆ To shrink LUN size to specified size, enter the following:

```
Target> target iscsi lun shrinkto lun-name target-name size
```

lun-name	Name of the LUN for which, size is to be decreased on the store that is associated with an iSCSI target.
target-name	Name of the iSCSI target at which SCSI LUNs are available. <code>target-name</code> should conform to the naming rules defined in RFC3721.
size	A new size that is to be specified for the LUN.

To create a clone of a specified LUN

- ◆ To create a clone of LUN, enter the following:

```
Target> iscsi lun clone create lun-name clone-name
```

lun-name	Name of the LUN for which a clone is to be created on the store that is associated with an iSCSI target.
clone-name	Name of the clone of a LUN.

To list details of all clones

- ◆ To list details of all clones, enter the following:

```
Target> iscsi lun clone list
```

To list details of the specified clone

- ◆ To list details of the specified clone, enter the following:

```
Target> iscsi lun clone list clone-name
```

To create a snapshot of LUN

- ◆ To create a snapshot of LUN, enter the following:

```
Target> iscsi lun snapshot create lun-name snapshot-name
```

lun-name Name of the LUN for which a snapshot is to be created on the store that is associated with an iSCSI target.

snapshot-name Name of the snapshot of a LUN.

To destroy a snapshot of LUN

- ◆ To destroy a snapshot of LUN, enter the following:

```
Target> iscsi lun snapshot destroy lun-name snapshot-name
```

lun-name Name of the LUN for which a snapshot is to be destroyed from an iSCSI target.

snapshot-name Name of the snapshot of a LUN.

To restore a snapshot of specified LUN

- ◆ To restore a snapshot of specified LUN, enter the following:

```
Target> iscsi lun snapshot restore lun-name snapshot-name
```

lun-name Name of the LUN for which a snapshot is to be restored on the store that is associated with an iSCSI target.

snapshot-name Name of the snapshot of a LUN.

To list details of all the snapshots

- ◆ To list details of all the snapshots, enter the following:

```
Target> iscsi lun snapshot list
```

To list details of a specified snapshot

- ◆ To list details of a specified snapshot, enter the following:

```
Target> iscsi lun snapshot list snapshot-name
```

where *snapshot-name* is the snapshot name for which details are displayed.

Managing the mappings with iSCSI initiators

To map an iSCSI initiator to a specific iSCSI target

- ◆ To map an iSCSI initiator to a specific iSCSI target, enter the following:

```
Target> iscsi target map add target-name initiator-name
```

target-name	Name of the iSCSI target at which SCSI LUNs are available. <i>target-name</i> should conform to the naming rules defined in RFC3721.
-------------	---

initiator-name	where <i>initiator-name</i> is a name that conforms to the naming rules for initiator and target names as specified in RFC3721.
----------------	---

To remove the mapping of iSCSI initiator from specific iSCSI target

- ◆ To remove the mapping of iSCSI initiator from specific iSCSI target, enter the following:

```
Target> iscsi target map delete target-name initiator-name
```

target-name	Name of the iSCSI target at which SCSI LUNs are available. <i>target-name</i> should conform to the naming rules defined in RFC3721.
-------------	---

initiator-name	where <i>initiator-name</i> is a name that conforms to the naming rules for initiator and target names as specified in RFC3721.
----------------	---

Managing the users

To create an incoming user and bind the account to a specified, existing iSCSI target

- ◆ To create an incoming user and bind the account to a specified, existing iSCSI target, enter the following:

```
Target> iscsi target auth incominguser add target-name user-name
```

target-name	Name of the iSCSI target at which SCSI LUNs are available. target-name should conform to the naming rules defined in RFC3721.
-------------	--

user-name	Name of the incoming user that is to be added to an iSCSI target.
-----------	---

To remove an incoming user and unbind the account from its corresponding iSCSI target

- ◆ To remove an incoming user and unbind the account from its corresponding iSCSI target, enter the following:

```
Target> iscsi target auth incominguser delete target-name user-name
```

target-name	Name of the iSCSI target at which SCSI LUNs are available. target-name should conform to the naming rules defined in RFC3721.
-------------	--

user-name	Name of the incoming user that is to be removed from an iSCSI target.
-----------	---

Creating an iSCSI target and provisioning LUNs

You can create an iSCSI target and provision LUNs.

To create an iSCSI target and provision LUNs

1 Prerequisites:

Configure storage.

Start the iSCSI service.

2 Quick Actions: Go to **Provision Storage > iSCSI Block storage > Provision iSCSI LUNs**, and then click **Continue**.

- 3** Click the **Configure Storage** tab and specify the storage options.
- 4** Click the **Start iSCSI Service** tab, and then click **Start iSCSI Service** to start the iSCSI service.
- 5** Click the **Provision Storage** tab, and then click **Select storage pool** to select a storage pool for the file system.
- 6** Click **Target** and choose the required options for the target and LUNs.
- 7** Click **Provision iSCSI**.
- 8** View the **Recent Activity** panel for the status of the task.

Managing Access Appliance file access services

- [Chapter 8. Configuring the NFS server](#)
- [Chapter 9. Using Access Appliance as a CIFS server](#)
- [Chapter 10. Using Access Appliance as an Object Store server](#)

Configuring the NFS server

This chapter includes the following topics:

- [About using the NFS server with Access Appliance](#)
- [Using the kernel-based NFS server](#)
- [Accessing the NFS server](#)
- [Displaying and resetting NFS statistics](#)
- [Configuring Access Appliance for ID mapping for NFS version 4](#)
- [Configuring the NFS client for ID mapping for NFS version 4](#)
- [About authenticating NFS clients](#)
- [Setting up Kerberos authentication for NFS clients](#)

About using the NFS server with Access Appliance

Access Appliance provides file access services to UNIX and Linux client computers using the Network File System (NFS) protocol. Access Appliance supports NFSv3 and NFSv4. Access Appliance provides the following NFS server support:

- Kernel-based NFS server
See [“Using the kernel-based NFS server”](#) on page 113.

The kernel NFS server is enabled by default.

Using the kernel-based NFS server

The kernel-based NFS server supports NFS version 3 and version 4. The kernel NFS server is enabled by default. Kernel NFS supports Active-Active mode serving NFS version 3 and 4. Veritas recommends that you use the default kernel-based NFS server.

Accessing the NFS server

To check on the NFS server status

- ◆ Prior to starting the NFS server, check on the status of the server by entering:

```
NFS> server status
```

The output shows the status. The output also indicates whether the NFS server being used.

The states (ONLINE, OFFLINE, and FAULTED) correspond to each Access Appliance node identified by the node name. The states of the node may vary depending on the situation for that particular node.

The possible states of the `NFS> server status` command are:

ONLINE	Indicates that the node can serve NFS protocols to the client.
OFFLINE	Indicates the NFS services on that node are down.
FAULTED	Indicates something is wrong with the NFS service on the node.

You can run the `NFS> server start` command to restart the NFS services, and only the nodes where NFS services have problems, are restarted.

To start the NFS server

- ◆ To start the NFS server, enter the following:

```
NFS> server start
```

You can use the `NFS> server start` command to clear an OFFLINE state from the `NFS> server status` output by only restarting the services that are offline. You can run the `NFS> server start` command multiple times without it affecting the already-started NFS server.

Run the `NFS> server status` command again to confirm the change.

To stop the NFS server

- ◆ To stop the NFS server, enter the following:

```
NFS> server stop
```

Displaying and resetting NFS statistics

You can display the statistics for a specific node or for all the nodes in the cluster

To display NFS statistics for a specific node in the cluster, enter the following:

```
NFS> stat show [nodename]
```

where *nodename* specifies the node name for which you are trying to obtain the statistical information. If the *nodename* is not specified, statistics for all the nodes in the cluster are displayed.

To display the NFS statistics for all the nodes in the cluster for the NFS server, enter the following:

```
NFS> stat show all
```

To reset NFS statistics for a specific node or for all the nodes in the cluster to zero

- ◆ To reset NFS statistics for the kernel NFS server, enter the following:

```
NFS> stat reset [nodename]
```

where *nodename* specifies the node name for which you want to reset the NFS statistics to zero. If *nodename* is not specified, NFS statistics for all the nodes in the cluster are reset to zero. Statistics are automatically reset to zero after a reboot of a node.

Configuring Access Appliance for ID mapping for NFS version 4

If you plan to use NFS version 4, you must configure Access Appliance to map the user IDs to the required format. In NFS version 3, each user is identified by a number, the user ID (uid). A UNIX file also identifies the owner of the file by a uid number. NFS version 4 has a different way of identifying users than that used by NFS version 3. In NFS version 4, each user is identified by a string, such as

```
user1@example.com.
```

Access Appliance requires a mechanism to map the user strings from NFS version 4 to uids on the server and the client. This process, called ID mapping, uses a file `/etc/idmapd.conf`.

NFS version 4 uses the `/etc/idmapd.conf` file to map the IDs. The Domain field needs to be set to the DNS domain of the Access Appliance server. If the DNS domain is not set, the ID mapping maps all of the users on the client to the user 'nobody'.

To configure Access Appliance for ID mapping

- ◆ Configure the DNS domain of Access Appliance using the following command:

```
Network> dns set domainname domainname
```

When the NFS server is started, the `/etc/idmapd.conf` file is updated with the domain information of the Access Appliance server.

You must also configure the NFS client.

Configuring the NFS client for ID mapping for NFS version 4

For NFS version 4, you must configure the NFS client so that the NFS version 4 user strings can be mapped to the uids. You must also configure the NFS server.

To configure the NFS client for ID mapping

- 1 For proper ID mapping, set the `Domain` field in the `/etc/idmapd.conf` file as the DNS domain name of the NFS client. Make sure that the DNS domain is the same for the NFS client and the Access Appliance server.

This setting in the `/etc/idmapd.conf` file should be updated on the NFS client.

- 2 Clear the ID mapping cache on the NFS client using the command `nfsidmap -c` and restart the ID mapping service.

About authenticating NFS clients

See [“About managing NFS shares using netgroups”](#) on page 288.

The kernel NFS server support Kerberos authentication.

Setting up Kerberos authentication for NFS clients

Kerberos provides a secure way of authenticating NFS clients. In this configuration, the Access Appliance server behaves as a Kerberos client. The Kerberos KDC (Key Distribution Center) server must already be set up and running outside of Access Appliance. For NFS version 3, when a Access Appliance share is exported with the `krb5` security option, the NFS clients have to mount the Access Appliance share with the `krb5` mount option. Otherwise the mount fails with an authentication error. For NFS version 4, the NFS clients automatically find the security type and mount the Access Appliance share with the same mount option.

Note: When CIFS security is configured with `ads`, Kerberos for NFS cannot be configured. When NFS is configured for Kerberos authentication, CIFS security cannot be configured with `ads`.

To configure Access Appliance for authenticating NFS clients using Kerberos, perform the tasks in the order that is listed in [Table 8-1](#).

Table 8-1 Tasks for configuring Access Appliance for authenticating NFS clients using Kerberos

Task	Where to find more information
Add and configure Access Appliance to the Kerberos realm	See “Adding and configuring Access Appliance to the Kerberos realm” on page 117.
Configure the NFS server for ID mapping	See “Configuring Access Appliance for ID mapping for NFS version 4” on page 114.
Configure the NFS client for ID mapping	See “Configuring the NFS client for ID mapping for NFS version 4” on page 115.
Exporting an NFS share for Kerberos authentication	See “Exporting an NFS share for Kerberos authentication” on page 290.
Mount the NFS share from the NFS client	See “Mounting an NFS share with Kerberos security from the NFS client” on page 291.

Adding and configuring Access Appliance to the Kerberos realm

Kerberos authentication support on Access Appliance is available only if the Key Distribution Center (KDC) server is running on a standalone computer (in a non-AD (Active Directory) environment), and there is a single KDC server. Before Access Appliance can be used as a Kerberos client, the NFS service principal of Access Appliance has to be added to the KDC server. Use the Access Appliance cluster name (either the short name or the fully qualified domain name) in small letters as the host name when creating the NFS service principal.

For example, if `access_ga_01` and `access_ga_02` are two nodes in the Access Appliance cluster, then `access_ga` (or the fully qualified domain name `access_ga.example.com`) should be used for adding the NFS service principal. The Domain Name System (DNS) or `/etc/hosts` is then set up to resolve `access_ga` to all the virtual IPs of the Access Appliance cluster.

To configure the KDC server

- 1 Create the NFS service principal on the KDC server using the `kadmin.local` command.

```
addprinc -randkey nfs/access_ga
```

- 2 Create a `keytab` file for the NFS service principal on KDC.

```
ktadd -k /etc/access.keytab nfs/access_ga
```

- 3 Copy the created `keytab` file (`/etc/access.keytab`) to the Access Appliance console node.

- 4 Use the `network krb standalone set` command to set the Kerberos configuration on Access Appliance.

The `network krb standalone set` command takes the KDC server name, Kerberos realm, and the location of the `keytab` that is located on the Access Appliance console node. This command sets up the Kerberos configuration file `/etc/krb5.conf` with the KDC server name and realm on all the nodes of the Access Appliance cluster. The command then copies the `keytab` file to `/etc/krb5.keytab` on all the nodes of the Access Appliance cluster.

```
Network> krb standalone set kdc_server TESTKDC.COM  
/home/support/krb5.keytab
```

The `network krb standalone set` command checks for the correct domain in the `/etc/idmapd.conf` file. If the domain is not set, the command gives a warning message saying that the DNS domain name needs to be set.

See [“Configuring Access Appliance for ID mapping for NFS version 4”](#) on page 114.

- 5 Use the `network krb standalone show` command to show the Kerberos configuration.
- 6 Use the `network krb standalone unset` command to reset the Kerberos configuration.

After the KDC server is configured, you can export the NFS shares with Kerberos authentication options.

Using Access Appliance as a CIFS server

This chapter includes the following topics:

- [About configuring Access Appliance for CIFS](#)
- [About configuring CIFS for standalone mode](#)
- [Configuring CIFS server status for standalone mode](#)
- [Changing security settings](#)
- [About configuring CIFS for Active Directory \(AD\) domain mode](#)
- [Setting NTLM](#)
- [About setting trusted domains](#)
- [About storing account information](#)
- [Storing user and group accounts](#)
- [Reconfiguring the CIFS service](#)
- [About mapping user names for CIFS/NFS sharing](#)
- [About the mapuser commands](#)
- [Adding, removing, or displaying the mapping between CIFS and NFS users](#)
- [Automatically mapping UNIX users from LDAP to Windows users](#)
- [About managing home directories](#)
- [About CIFS clustering modes](#)

- [About migrating CIFS shares and home directories](#)
- [Setting the CIFS aio_fork option](#)
- [About managing local users and groups](#)
- [Enabling CIFS data migration](#)

About configuring Access Appliance for CIFS

The Common Internet File System (CIFS), also known as the Server Message Block (SMB), is a network file sharing protocol that is widely used on Microsoft and other operating systems. Access Appliance supports the SMB3 protocol.

Access Appliance supports the following clustering modes:

- Normal
- Clustered Trivial Database (CTDB) - a cluster implementation of the TDB (Trivial database) based on the Berkeley database API

Note: In case of network or node failover, the application which performs the I/O operation on the CIFS share needs to have a retry logic for a failed I/O to survive from an I/O failure.

Access Appliance supports the following CIFS security modes:

- User
- ADS

Each clustering mode supports both of the CIFS security modes.

See [“About CIFS clustering modes”](#) on page 151.

Access Appliance can be integrated into a network that consists of machines running Microsoft Windows. You can control and manage the network resources by using Active Directory (AD) domain controllers.

Before you use Access Appliance with CIFS, you must have administrator-level knowledge of the Microsoft operating systems, Microsoft services, and Microsoft protocols (including AD and NT services and protocols).

You can find more information about them at: www.microsoft.com.

When serving the CIFS clients, Access Appliance can be configured to operate in one of the security mode environments described in [Table 9-1](#).

Table 9-1 CIFS security mode environments

Mode	Definition
Standalone	Information about the user and group accounts is stored locally on Access Appliance. Access Appliance also authenticates users locally using the Linux password and group files. This mode of operation is provided for Access Appliance testing and may be appropriate in other cases, for example, when Access Appliance is used in a small network and is not a member of a Windows security domain. In this mode of operation, you must create the local users and groups; they can access the shared resources subject to authorization control.
Active Directory (AD)	Access Appliance becomes a member of an AD security domain and is configured to use the services of the AD domain controller, such as DNS, LDAP, and NTP. Kerberos, NTLMv2, or NTLM authenticate users. When Access Appliance operates in the AD domain mode, it acts as a domain member server and not as the domain controller.

About configuring CIFS for standalone mode

If you do not have an AD server, you can use Access Appliance as a standalone server. Access Appliance is used in standalone mode when testing Access Appliance functionality and when it is not a member of a domain.

Before you configure the CIFS service for the standalone mode, do the following:

- Make sure that the CIFS server is not running.
- Set security to user.
- Start the CIFS server.

To make sure that the configuration has changed, do the following:

- Check the server status.
- Display the server settings.

Configuring CIFS server status for standalone mode

To check the CIFS server status

- 1 To check the status of the CIFS server, enter the following:

```
CIFS> server status
```

By default, `security` is set to `user`, the required setting for standalone mode. The following example shows that `security` was previously set to `ads`.

- 2 If the server is running, enter the following:

```
CIFS> server stop.
```

To check the security setting

- 1 To check the current settings before setting security, enter the following:

```
CIFS> show
```

- 2 To set security to `user`, enter the following:

```
CIFS> set security user
```

To start the CIFS service in standalone mode

- 1 To start the service in standalone mode, enter the following:

```
CIFS> server start
```

- 2 To display the new settings, enter the following:

```
CIFS> show
```

- 3 To make sure that the server is running in standalone mode, enter the following:

```
CIFS> server status
```

The CIFS service is now running in standalone mode.

See [“About managing local users and groups”](#) on page 154.

See [“About managing CIFS shares”](#) on page 295.

Changing security settings

To change security settings

- ◆ To set the security to user, enter the following:

```
CIFS> set security user
```

To stop the CIFS server:

```
CIFS> server stop
```

About configuring CIFS for Active Directory (AD) domain mode

This section assumes that an Active Directory (AD) domain has already been configured and that Access Appliance can communicate with the AD domain controller (DC) over the network. The AD domain controller is also referred to as the AD server.

Before you configure CIFS, open the `nsswitch.conf` file and check the entry in the line which has the password. Ensure that the line has **winbind** as the first entry.

For example:

```
passwd: winbind files sss
```

Joining Access Appliance to Active Directory (AD)

To join Access Appliance to Active Directory (AD)

- 1 To stop the CIFS server, enter the following command.

```
CIFS> server stop
```

- 2 Ensure that AD is configured.

See [“Configuring AD server settings”](#) on page 70.

- 3 To set the CIFS security mode, enter the following command:

```
CIFS> set security ads
```

The other CIFS security mode is `user` for local users. For authenticating to Active Directory, use the `ads` CIFS security mode.

- 4 To start the CIFS server, enter the following command:

```
CIFS> server start
```

Verifying that Access Appliance has joined Active Directory (AD) successfully

To verify that Access Appliance has joined Active Directory (AD) successfully

- ◆ To verify that Access Appliance has joined Active Directory successfully, enter the following command:

```
CIFS> server status
```

Refer to the `Domain membership status` line of the output to verify that the Access Appliance cluster has joined the domain (displays as `Enabled`) if the join is successful.

If the cluster did not join the domain, an informative error message is provided indicating why the Access Appliance cluster cannot join the domain.

Using multi-domain controller support in CIFS

Access Appliance allows you to set a comma-separated list of primary and backup domain controllers for the given domain.

Note: You need to set `dns nameserver` for other domain controller (i.e. backup domain controller) using the `network dns set nameserver` command.

You will need to stop and start the CIFS server.

See [“Reconfiguring the CIFS service”](#) on page 142.

To display the list of domain controllers

- ◆ To display the list of domain controllers, enter the following:

```
CIFS> show
```

If the primary domain controller goes down, the CIFS server tries the next domain controller in the list until it receives a response. You should always point Access Appliance to the trusted domain controllers to avoid any security issues. Access Appliance does not perform list reduction or reordering, instead it uses the list as it is. So, avoid entering the redundant name for the same domain controller.

About leaving an AD domain

Changing domain settings for AD domain mode

Each case assumes that the Access Appliance cluster is part of an AD domain.

To verify the cluster is part of an AD domain

- ◆ To verify that the cluster is part of an AD domain, enter the following:

```
CIFS> server status
```

To change domain settings for AD domain mode

- 1 To stop the CIFS server, enter the following:

```
CIFS> server stop
```

- 2 To change the domain, enter the following:

```
Network> ad set domain domaincontroller workgroup domainuser
```

When you start the CIFS server, it tries to leave the existing domain. This requires the old domainuser to enter its password. After the password is supplied, and the domain leave operation succeeds, the CIFS server joins an AD domain with the new settings.

- 3 To start the CIFS server, enter the following:

```
CIFS> server start
```

To change the security settings for the AD domain mode

- ◆ To set the security to user, enter the following:

```
CIFS> set security user
```

To stop the CIFS server:

```
CIFS> server stop
```

Changing security settings with stopped server on the AD domain mode

- ◆ To set security to a value other than ads, enter the following:

```
CIFS> set security user
```

Removing the AD interface

You can remove the Access Appliance cluster from the AD domain by using the Active Directory interface.

To remove the Access Appliance cluster from the Active Directory

- 1 Open the interface **Active Directory Users and Computers**.
- 2 In the domain hierarchy tree, click on **Computers**.
- 3 In the details pane, right-click the computer entry corresponding to Access Appliance (this can be identified by the Access Appliance cluster name) and click **Delete**.

Setting NTLM

When you use Access Appliance in AD domain mode, there is an optional configuration step that can be done. You can disable the use of Microsoft NTLM (NT LAN Manager) protocol for authenticating users.

When the Access Appliance CIFS service is running in the standalone mode (with security set to user) some versions of the Windows clients require NTLM authentication to be enabled. You can do this by setting the value of `ntlm_auth` to `yes` by using the `CIFS> set ntlm_auth yes` command.

When NTLM is disabled and you use Access Appliance in AD domain mode, the available authentication protocols are Kerberos and NTLMv2. The one used depends on the capabilities of both the Access Appliance clients, and domain controller. If no special action is taken, Access Appliance allows the NTLM protocol to be used.

For any specific CIFS connection, all the participants, that is the client machine, Access Appliance and the domain controller select the protocol that they all support and that provides the highest security. In the AD domain mode, Kerberos provides the highest security.

To disable NTLM

- 1 If the server is running, enter the following:

```
CIFS> server stop
```

- 2 To disable NTLM, enter the following:

```
CIFS> set ntlm_auth no
```

- 3 To start the CIFS service, enter the following:

```
CIFS> server start
```

To enable NTLM

- 1 If the server is running, enter the following:

```
CIFS> server stop
```

- 2 To enable the NTLM protocol, enter the following:

```
CIFS> set ntlm_auth yes
```

- 3 To start the CIFS service, enter the following:

```
CIFS> server start
```

About setting trusted domains

The Microsoft Active Directory supports the concept of trusted domains. When you authenticate users, you can configure domain controllers in one domain to trust the domain controllers in another domain. This establishes the trust relation between the two domains. When Access Appliance is a member in an AD domain, both Access Appliance and the domain controller are involved in authenticating the clients. You can configure Access Appliance to support or not support trusted domains.

You can obtain unique user IDs (UIDs) or group IDs (GIDs) from domains by reading ID mappings from an Active Directory server that uses RFC2307/SFU schema extensions. This is a read-only idmap backend..

A valid user from a domain or trusted domain should have a UID as well as a GID for the user's primary group.

By default, the `uid_range` is set to 10000-1000000. Change it in cases where there are more than 1,000,000 users existing on a local Access Appliance cluster where there are joined Active Directory domains or trusted domains.

Note: The `uid_range` is adjusted automatically according to the search results of the defined UNIX IDs from the domain after a CIFS server restart.

Table 9-2 Set trusted domains commands

Command	Definition
<code>set allow_trusted_domains yes</code>	Enables the use of trusted domains in the AD domain mode. Note: If the security mode is <code>user</code>, it is not possible to enable AD trusted domains. All the IDMAP backend methods (<code>rid</code>, <code>ldap</code>, and <code>hash</code>) are able to support trusted domains. See “Setting Active Directory trusted domains” on page 137.
<code>set allow_trusted_domains no</code>	Disables the use of trusted domains in the AD domain mode. See “Setting Active Directory trusted domains” on page 137.

Specifying trusted domains that are allowed access to the CIFS server

You can specify the trusted domains that are allowed access to a CIFS server when the `CIFS> set allow_trusted_domains` option is set to `yes` and `idmap_backend` is set to `rid` or `ad`.

See [“Allowing trusted domains access to CIFS when setting an IDMAP backend to rid”](#) on page 129.

By default, all the trusted domains of the joined active directory domain are included in the CIFS settings and configuration if `allow_trusted_domains` is set to `yes`.

By default, `CIFS> set allow_trusted_domains` is set to `no`.

To specify the trusted domains that are allowed access to the CIFS server

- ◆ To specify the trusted domains that are allowed access to the CIFS server, enter the following:

```
CIFS> set allow_trusted_domains yes|no [trusted_domains]
```

where *trusted_domains* are the trusted domains that you want to allow access to the CIFS server.

Allowing trusted domains access to CIFS when setting an IDMAP backend to rid

To allow trusted domains access to CIFS when setting IDMAP backend to rid

- 1 If the CIFS server is running, enter the following:

```
CIFS> server stop
```

- 2 To set the *idmap_backend* to *rid*, enter the following:

```
CIFS> set idmap_backend rid [uid_range]
```

where *uid_range* represents the range of identifiers that are used by Access Appliance when mapping domain users and groups to local users and groups.

```
CIFS> set idmap_backend rid
```

- 3 To set *allow_trusted_domains* to *yes*, enter the following:

```
CIFS> set allow_trusted_domains yes
```

- 4 To start the CIFS server again, enter the following:

```
CIFS> server start
```

- 5 To verify the CIFS server status when there are trusted domains, enter the following:

```
CIFS> server status
```

Domain names containing square brackets indicate that the domain used to be a trusted domain, but the domain is currently obsolete.

Allowing trusted domains access to CIFS when setting an IDMAP backend to ldap

To allow trusted domains access to CIFS when setting an IDMAP backend to ldap

- 1 To configure AD as an IDMAP backend, follow the steps provided at:

See [“About configuring Windows Active Directory as an IDMAP backend for CIFS”](#) on page 132.

- 2 To set `idmap_backend` to `ldap`, enter the following:

```
CIFS> set idmap_backend ldap [idmap_ou] [uid_range]
```

`idmap_ou` Specifies the CIFS idmap Organizational Unit Name (OU) configured on the LDAP server, which is used by Access Appliance when mapping users and groups to local users and groups. The default value is `cifsidmap`.

`uid_range` Specifies the range of identifiers that are used by Access Appliance when mapping domain users and groups to local users and groups.

```
CIFS> set idmap_backend ldap
```

- 3 To set `allow_trusted_domains` to `yes`, enter the following:

```
CIFS> set allow_trusted_domains yes
```

- 4 To restart the CIFS server again, enter the following:

```
CIFS> server start
```

- 5 To verify the CIFS server status when there are trusted domains, enter the following:

```
CIFS> server status
```

Allowing trusted domains access to CIFS when setting an IDMAP backend to hash

To allow trusted domains access to CIFS when setting an IDMAP backend to hash

- 1 If the CIFS server is running, enter the following:

```
CIFS> server stop
```

- 2 To set `idmap_backend` to `hash`, enter the following:

```
CIFS> set idmap_backend hash
```

- 3 To set `allow_trusted_domains` to `yes`, enter the following:

```
CIFS> set allow_trusted_domains yes
```

- 4 To verify the CIFS server status when there are trusted domains, enter the following:

```
CIFS> server status
```

Allowing trusted domains access to CIFS when setting an IDMAP backend to ad

To allow trusted domains access to CIFS when setting IDMAP backend to ad

- 1 If the CIFS server is running, enter the following:

```
CIFS> server stop.
```

- 2 To set the `idmap_backend` to `ad`, enter the following:

```
CIFS> set idmap_backend ad [uid_range]
```

where `uid_range` represents the range of identifiers that are used by Access Appliance when mapping domain users and groups to local users and groups.

- 3 To set `allow_trusted_domains` to `yes`, enter the following:

```
CIFS> set allow_trusted_domains yes
```

- 4 To start the CIFS server again, enter the following:

```
CIFS> server start
```

- 5 To verify the CIFS server status when there are trusted domains, enter the following:

```
CIFS> server status
```

Domain names containing square brackets indicate that the domain used to be a trusted domain, but the domain is currently obsolete.

About configuring Windows Active Directory as an IDMAP backend for CIFS

The CIFS server requires equivalent UNIX identities for Windows accounts to service requests from Windows clients. In the case of trusted domains, Access Appliance has to store the mapped UNIX identities (IDMAP) in a centralized database that is accessible from each of the cluster nodes.

Active Directory (AD), as with any LDAP V3 compliant directory service, can function as the backend for CIFS IDMAP backend storage. When the CIFS server joins a Windows Active Directory Domain as a member server, and you want to use LDAP as an IDMAP backend, then it is necessary to create an Active Directory application partition for the IDMAP database. To support the creation of an Active Directory application partition, Windows 2003 R2 and above version is required.

Active Directory application partition provides the ability to control the scope of replication and allow the placement of replicas in a manner more suitable for dynamic data. As a result, the application directory partition provides the capability of hosting dynamic data in the Active Directory server, thus allowing ADSI/LDAP access to it.

By extending the AD schema with the necessary CIFS-schema extensions, and creating an AD application partition, it is possible to store CIFS IDMAP data entries in AD, using one or more domain controllers as IDMAP LDAP backend servers. Also, it is possible to replicate this information in a simple and controlled manner to a subset of AD domain controllers located either in the same domain or in different domains in the AD forest.

Note: A single domain user account is used, for example, **cifsuser** for setting application partition Access Control List (ACL) settings. Make sure the selected user naming context has no space key inside (for example, **CN=cifsuser1,CN=Users,DC=example,DC=com**). Here, a sample AD server is used, for example, **adserver.example.com**. Use relevant values when configuring your AD server.

Configuring the Active Directory schema with CIFS-schema extensions

To extend the Active Directory schema with the necessary CIFS-schema extensions

- 1 Login with **Schema Admins** privileges on the Active Directory Forest Schema Master domain controller.
- 2 Download `ADCIFSSchema.zip` from the Access Appliance server (`/opt/VRTSnas/tools/cifs/ADCIFSSchema.zip`) with software such as `WinSCP.exe`.
- 3 Unzip the file and open each `.ldf` file to perform a search and replace of the string `dc=example,dc=com`, replacing the string with the top-level domain component (that is, `dc=yourdomain,dc=com`) values for the AD forest.
- 4 Install the schema extensions by executing the `schemaupdate.bat` file from the command prompt.

To validate the schema extensions

- 1 Execute `regsvr32 schmmgmt.dll` in a command prompt window to install the Active Directory Schema Snap-In on the AD server.
- 2 Enter `mmc` in **Run**.
- 3 On the **File** menu, click **Add/Remove Snapin**.
- 4 In **Available snap-ins**, click **Active Directory Schema**, and then click **Add**.
- 5 Click **OK**.
- 6 Click **Attributes** in the left frame, and try to find **uidNumber** and **gidNumber** in the right frame.

Validate that the **uidNumber** and **gidNumber** attributes have no minimum or maximum value setting by viewing the properties of the attribute objects.

To create an application partition

- 1 Open a command prompt window on the domain controller that will hold the first replica of the application partition.
- 2 Enter `ntdsutil` in the command prompt window.
- 3 At the `ntdsutil` command prompt, enter the following:

```
domain management
```

If you are using Windows 2008, change this command to the following:

```
partition management
```

- 4 At the domain management command prompt, enter the following:

```
connection
```

- 5 At the connection command prompt, enter the following:

```
connect to server adserver.example.com
```

- 6 At the connection command prompt, enter the following:

```
quit
```

- 7 At the domain management command prompt, enter the following such as:

```
create nc dc=idmap,dc=example,dc=com null
```

Example settings:

```
C:\>ntdsutil
ntdsutil: domain management
domain management: connection
server connections: connect to server adserver.example.com
Binding to adserver.example.com ...
Connected to adserver.si2m.com using credentials of locally logged
on user.
server connections: quit
domain management: create nc dc=idmap,dc=example,dc=com NULL
adding object dc=idmap,dc=example,dc=com
domain management: quit
ntdsutil: quit
Disconnecting from adserver.example.com...
```

- 8 Once the application partition has been created, open **ADSIedit.msc** from **Run**, then right-click on **ADSI Edit** in the left frame, and click **connect to ...** to connect to the application partition using the settings as indicated:

Name Enter **Domain**.

Connection Point Select or enter a **Distinguished Name** or **Naming Context**, as in:

`dc=idmap,dc=example,dc=com`

Computer Select or enter a domain or server, as in:

`adserver.example.com`

Connection Settings

Name:

Path:

Connection Point:

☒ Select or type a Distinguished Name or Naming Context:

☐ Select a well known Naming Context:

Computer:


☒ Select or type a domain or server: (Server | Domain [:port])

☐ Default (Domain or server that you logged in to)

☐ Use SSL-based Encryption

Advanced... OK Cancel

- 9 Once connected, select the top-level application partition (for example, **dc=idmap,dc=example,dc=com**) node in the left panel, and right-click to select **New** then **Object** from the list, and then select **SambaUnixIdPool**.



When prompted, enter the following values:

OU attribute	cifsidmap
uidNumber	10000
gidNumber	10000

- 10 Click **Finish** to complete the configuration.
- 11 Once the **ou=cifsidmap,dc=idmap,dc=example,dc=com** container has been created, right-click the object, and select **properties**.
- 12 On the **Security** tab, click **Add**, and proceed to add the cifsuser user account, and grant the account Read, Write, Create All Child Objects, and Delete All Child Objects permissions.

Configuring the LDAP client for authentication using the CLI

To configure the LDAP client for authentication using the CLI

- 1 Log into the cluster CLI using the `master` account.
- 2 Configure `Network> ldap` settings.

Example settings:

```
Network> ldap set basedn dc=idmap,dc=example,dc=com
Network> ldap set binddn cn=cifsuser,dc=example,dc=com
Network> ldap set rootbinddn cn=cifsuser,cn=users,dc=example,dc=com
Network> ldap set server adserver.example.com
Network> ldap enable
```

Setting Active Directory trusted domains

To enable Active Directory (AD) trusted domains

- 1 If the server is running, enter the following:

```
CIFS> server stop
```

- 2 To enable trusted domains, enter the following:

```
CIFS> set allow_trusted_domains yes
```

- 3 To start the CIFS server, enter the following:

```
CIFS> server start
```

To disable trusted domains

- 1 If the server is running, enter the following:

```
CIFS> server stop
```

- 2 To disable trusted domains, enter the following:

```
CIFS> set allow_trusted_domains no
```

- 3 To start the CIFS server, enter the following:

```
CIFS> server start
```

About storing account information

Access Appliance maps between the domain users and groups (their identifiers) and local representation of these users and groups. Information about these mappings can be stored locally on Access Appliance or remotely using the DC directory service. Access Appliance uses the `idmap_backend` configuration option to decide where this information is stored.

This option can be set to one of the following:

<code>rid</code>	Maps SIDs for domain users and groups by deriving UID and GID from RID on the Access Appliance CIFS server.
<code>ldap</code>	Stores the user and group information in the LDAP directory service.
<code>hash</code>	Maps SIDs for domain users and groups to 31-bit UID and GID by the implemented hashing algorithm on the Access Appliance CIFS server.
<code>ad</code>	Obtains unique user IDs (UIDs) or group IDs (GIDs) from domains by reading ID mappings from an Active Directory server that uses RFC2307/SFU schema extensions.

Note: SID/RID are Microsoft Windows concepts that are described at: [http://msdn.microsoft.com/en-us/library/aa379602\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/aa379602(VS.85).aspx).

The `rid` and `hash` values can be used in any of the following modes of operation:

- Standalone
- AD domain

`rid` is the default value for `idmap_backend` in all of these operational modes.

The `ldap` value can be used if the AD domain mode is used.

When security is set as "user" `idmap_backend` is irrelevant.

Table 9-3 Store account information commands

Command	Definition
<code>set idmap_backend rid</code>	Configures Access Appliance to store information about users and groups locally. Trusted domains are allowed if <code>allow_trusted_domains</code> is set to yes. The <code>uid_range</code> is set to 10000-1000000 by default. Change the default range in cases where it is not appropriate to accommodate local Access Appliance cluster users, Active Directory, or trusted domain users. Do not attempt to modify <code>LOW_RANGE_ID</code> (10000) if user data has already been created or copied on the CIFS server. This may lead to data access denied issues since the UID changes. See “Storing user and group accounts” on page 140.
<code>set idmap_backend hash</code>	Allows you to obtain the unique SID to UID/GID mappings by the implemented hashing algorithm. Trusted domains are allowed if <code>allow_trusted_domains</code> is set to yes. See “Storing user and group accounts” on page 140.
<code>set idmap_backend ad</code>	Allows you to obtain unique user IDs (UIDs) or group IDs (GIDs) from domains by reading ID mappings from an Active Directory server that uses RFC2307/SFU schema extensions. See “Storing user and group accounts” on page 140.

Table 9-3 Store account information commands (*continued*)

Command	Definition
<code>set</code> <code>idmap_backend</code> <code>ldap</code>	Configures Access Appliance to store information about users and groups in a remote LDAP service. You can only use this command when Access Appliance is operating in the AD domain mode. The LDAP service can run on the domain controller or it can be external to the domain controller. Note: For Access Appliance to use the LDAP service, the LDAP service must include both RFC 2307 and proper schema extensions. See “Configuring the LDAP client for authentication using the CLI” on page 137. This option tells the CIFS server to obtain SID to UID/GID mappings from a common LDAP backend. This option is compatible with multiple domain environments. So <code>allow_trusted_domains</code> can be set to <code>yes</code>. If <code>idmap_backend</code> is set to <code>ldap</code>, you must first configure the Access Appliance LDAP options using the <code>Network> ldap</code> commands. See “About configuring LDAP settings” on page 64. See “Storing user and group accounts” on page 140.

Storing user and group accounts

To set `idmap_backend` to `rid`

- 1 If the server is running, enter the following:

```
CIFS> server stop
```

- 2 To store information about user and group accounts locally, enter the following:

```
CIFS> set idmap_backend rid [uid_range]
```

where `uid_range` represents the range of identifiers that are used by Access Appliance when mapping domain users and groups to local users and groups. The default range is `10000-1000000`.

- 3 To start the CIFS server, enter the following:

```
CIFS> server start
```

To set idmap_backend to LDAP

- 1 To make sure that you have first configured LDAP, enter the following:

```
Network> ldap show
```

- 2 If the CIFS server is running, enter the following:

```
CIFS> server stop
```

- 3 To use the remote LDAP store for information about the user and group accounts, enter the following:

```
CIFS> set idmap_backend ldap [idmap_ou]
```

where *idmap_ou* represents the CIFS idmap Organizational Unit Name (OU) configured on the LDAP server, which is used by Access Appliance when mapping users and groups to local users and groups. The default value is *cifsidmap*.

- 4 To start the CIFS server, enter the following:

```
CIFS> server start
```

To set idmap_backend to a hash algorithm

- 1 If the CIFS server is running, enter the following:

```
CIFS> server stop
```

- 2 To store information about user and group accounts locally, enter the following:

```
CIFS> set idmap_backend hash
```

- 3 To start the CIFS server, enter the following:

```
CIFS> server start
```

To set idmap_backend to ad

- 1 If the CIFS server is running, enter the following:

```
CIFS> server stop
```

- 2 To obtain the unique UID/GID from domains by reading ID mappings from an Active Directory (AD) server, enter the following:

```
CIFS> set idmap_backend ad [uid_range]
```

where *uid_range* represents the range of identifiers that are used by Access Appliance when mapping domain users and groups to local users and groups. The default range is 10000-1000000. Change it in cases where there are more than 1,000,000 users existing on a local Access Appliance cluster where there are joined Active Directory domains or trusted domains.

Note: The *uid_range* is adjusted automatically according to the search results of the defined UNIX IDs from the domain after a CIFS server restart.

- 3 To start the CIFS server, enter the following:

```
CIFS> server start
```

Reconfiguring the CIFS service

Sometime after you have configured the CIFS service, and used it for a while, you need to change some of the settings. For example, you may want to allow the use of trusted domains or you need to move Access Appliance from one security domain to another. To carry out these changes, set the new settings and then start the CIFS server. As a general rule, you should stop the CIFS service before making the changes.

An example where Access Appliance is moved to a new security domain (while the mode of operation stays unchanged as, AD domain) is referenced below.

This example deals with reconfiguring CIFS. So make sure that if any of the other AD services like DNS or NTP are being used by Access Appliance, that Access Appliance has already been configured to use these services from the AD server belonging to the new domain.

Make sure that the DNS service, NTP service and, if used as an ID mapping store, also the LDAP service, are configured as required for the new domain.

To reconfigure the CIFS service, do the following:

- Make sure that the server is not running.
- Set the domain user, domain, and domain controller.
- Start the CIFS server.

To set the user name, AD domain and AD server

- ◆ To set the AD client configuration details, enter the following:

```
Network> ad set domain domaincontroller workgroup domainuser
```

You can set the AD client's domain, domain controller, workgroup and domain user. A domain controller can either be an Active Directory server or a Windows NT 4.0 domain controller. A *domainuser* is the username which is used for authentication in the domain join operation.

To start the CIFS server

- 1 To start the CIFS server, enter the following:

```
CIFS> server start
```

- 2 To make sure that the service is running, enter the following:

```
CIFS> server status
```

- 3 To find the current settings, enter the following:

```
CIFS> show
```

About mapping user names for CIFS/NFS sharing

The CIFS server uses user name mapping to translate login names sent by a Windows client to local or remote UNIX user names. The CIFS server uses file lookup for mapping, and this mapping is unidirectional. You can map a CIFS user to an NFS user, but the reverse operation is not possible.

This functionality can be used for the following purposes:

- CIFS and NFS sharing by mapping CIFS users to NFS users
- File sharing among CIFS users by mapping multiple CIFS users to a single UNIX user
- Mapping between two UNIX users by using the `CIFS> mapuser add <CIFSusername> LOCAL <NFSusername>` command, where both the CIFS user and the NFS user are UNIX users

User name mapping is stored in a configuration file.

When user name mapping takes place, it is dependent on the current security configurations. If security is set to `user`, mapping is done prior to authentication, and a password must be provided for the mapped user name. For example, if there is a mapping between the users `CIFSuser1` and `NFSuser1`. If `CIFSuser1` wants to connect to the Access Appliance server, then `CIFSuser1` needs to provide a password for `NFSuser1`. In this case, `NFSuser1` must be the CIFS local user.

If security is set to `ads`, user name mapping is done after authentication with the domain controller. This means, the actual password must be supplied for the login user `CIFSuser1` in the example cited above. In this case, `NFSuser1` may not be the CIFS local user.

The domain you specify for CIFS user name mapping must be the netbios domain name (instead of the Active Directory DNS domain name) for the user. For example, a netbios domain name might be listed as `VERITASDOMAIN` instead of `VERITASDOMAIN.COM` (without the `.com` extension).

To determine the netbios domain name, login to your Active Directory Server and type the following in a command window:

```
set | findstr DOMAIN
```

The results will include:

```
USERDOMAIN netbios_domain_name  
USERDNSDOMAIN Active_Directory_DNS_domain_name
```

Use the value of `USERDOMAIN` (the netbios domain name) when you map user names.

Note: When setting quotas on home directories and using user name mapping, make sure to set the quota on the home directory using the user name to which the original name is mapped.

Note: For mapped Active Directory users to access their home directory CIFS shares, use the following convention: `\\access\realADuser` instead of `\\access\homes`.

Note: For UNIX users (LDAP/NIS/local) users, make sure to set up these users properly, so that these users are recognized by Samba. User mapping can work properly only after these users are recognized by Samba.

About the mapuser commands

The `CIFS> mapuser` commands are used to add, remove, or display the mapping between CIFS and NFS users.

Typically, a *CIFSusername* is a user coming from an AD server (with a specified domainname), or a locally created CIFS user on this system (local). An *NFSusername* is a user coming from a locally-created CIFS user on this system, or from a NIS/LDAP server configured in the network section.

Note: To make sure user mappings work correctly with a NIS/LDAP server, `Network> nsswitch` settings may need to be adjusted in the `Network> nsswitch` section. You may need to move the position of `ldap` or `nis` in the `Network> nsswitch` section, depending on which name service is being used first.

Adding, removing, or displaying the mapping between CIFS and NFS users

To add a mapping between a CIFS and an NFS user

- ◆ To add a mapping between a CIFS and an NFS user, enter the following:

```
CIFS> mapuser add CIFSusername domainname NFSusername
```

To remove a mapping between a CIFS and an NFS user

- ◆ To remove a mapping between a CIFS and an NFS user, enter the following:

```
CIFS> mapuser remove CIFSusername [domainname]
```

To display a mapping between a CIFS and an NFS user

- ◆ To display a mapping between a CIFS and an NFS user, enter the following:

```
CIFS> mapuser show [CIFSusername] [domainname]
```

Automatically mapping UNIX users from LDAP to Windows users

To automatically map UNIX users from LDAP to Windows users

- 1 Ensure that Access Appliance joins the LDAP domain using **network ldap**.
 From the LDAP server, users and groups should be visible by using the `getent passwd` or `getent group` commands .
- 2 Ensure that Access Appliance joins the Windows AD domain using **cifs**.
- 3 Use the wildcard mapping rule `CIFS> mapuser add * AD Domain Name *`.
 The effect is whenever a Windows domain user, say `DOM\foobar`, wants to access CIFS shares, the CIFS server determines if there is a local (a non-Windows) user also named `foobar`, and establishes the mapping between the Windows user and the non-Windows user.
 The user name must match between the LDAP and AD domains.

About managing home directories

You can use Access Appliance to store the home directories of CIFS users.

The home directory share name is identical to the Access Appliance user name. When Access Appliance receives a new CIFS connection request, it checks if the requested share is one of the ordinary exported shares. If it is not, Access Appliance checks if the requested share name is the name of an existing Access Appliance user (either local user or domain user, depending on the current mode of operation). If a match is found, it means that the received connection request is for a home directory share.

You can access your home directory share the same way you access the file system ordinary shares. A user can connect only to his or her own home directory.

Note: The internal directories structure of home directory file systems is maintained by Access Appliance. It is recommended not to use a file system as a homedirs that has been used by a normal share in the past or vice versa.

Note: The `CIFS> homedir` commands are supported only from the Access Appliance command-line interface.

Setting the home directory file systems

Home directory shares are stored in one or more file systems. A single home directory can exist only in one of these file systems, but a number of home directories can exist in a single home directory file system. File systems that are to be used for home directories are specified using the `CIFS> set homedirfs` command.

When a file system is exported as a homedirfs, its mode is set to a 0755 value. This takes place when you start the CIFS server after setting the homedirfs list.

Note: Snapshots cannot be shared as home directory file systems.

To specify one or more file systems as the home directories

- 1 To reserve one or more file systems for home directories, enter the following:

```
CIFS> set homedirfs [filesystemlist]
```

where *filesystemlist* is a comma-separated list of names of the file systems which are used for the home directories.

- 2 If you want to remove the file systems you previously set up, enter the command again, without any file systems:

```
CIFS> set homedirfs
```

- 3 To find which file systems (if any) are currently used for home directories, enter the following:

```
CIFS> show
```

After you select one or more of the file systems to be used in this way, you cannot export the same file systems as ordinary CIFS shares.

If you want to change the current selection, for example, to add an additional file system to the list of home directory file systems or to specify that no file system should be used for home directories, you have to use the same `CIFS> set homedirfs` command. In each case you must enter the entire new list of home directory file systems, which may be an empty list when no home directory file systems are required.

Access Appliance treats home directories differently from ordinary shares. The differences are as follows:

- An ordinary share is used to export a file system, while a number of home directories can be stored in a single file system.

- The file systems used for home directories cannot be exported as ordinary shares.
- Exporting a home directory share is done differently than exporting an ordinary share. Also, removing these two kinds of shares is done differently.
- The configuration options you specify for an ordinary share (such as read-only or use of opportunistic locks) are different from the ones you specify for a home directory share.

Setting up home directories

You can set the home directory for the specified user with the `CIFS> homedir set` command. If the home directory does not exist for the specified user, the `CIFS> homedir set` command creates that user's home directory.

Use the `Storage> quota cifshomedir set` command to set the quota value for the specified user. Otherwise, the value set from the `Storage> quota cifshomedir setdefault` command is used to configure the quota limit. If either the user or default quota is not set, 0 is used as the default value for the unlimited quota.

Once the global quota value is specified, the value applies to the automatically created homedir. For example, if you set the global quota value to `Storage> quota cifshomedir setdefault 100M`, and you then create a new homedir in Windows, then the 100M quota value is assigned to that homedir.

To set the home directory for the specified user

- 1 To set the home directory for the specified user, enter the following:

```
CIFS> homedir set username [domainname] [fsname]
```

<i>username</i>	The name of the CIFS user. If a CIFS user name includes a space, enter the user name with double quotes.
<i>domainname</i>	The domain for the new home directory.
<i>fsname</i>	The home directory file system where the user's home directory is created. If no file system is specified, the user's home directory is created on the home directory file system that has the fewest home directories.

- 2 To find the current settings for a home directory, enter the following:

```
CIFS> homedir show [username] [domainname]
```

<i>username</i>	The name of the CIFS user. If a CIFS user name includes a space, enter the user name with double quotes.
<i>domainname</i>	The Active Directory/Windows NT domain name or specify <code>local</code> for the Access Appliance local user <code>local</code> .

- 3 To find the current settings for all home directories, enter the following:

```
CIFS> homedir show
```

Because the `CIFS> homedir show` command takes a long time when there are more than 1000 CIFS home directories to display, you will be prompted if you want to continue displaying CIFS home directories or not.

When you connect to your home directory for the first time, and if the home directory has not already been created, Access Appliance selects one of the available home directory file systems and creates the home directory there. The file system is selected in a way that tries to keep the number of home directories balanced across all available home directory file systems. The automatic creation of a home directory does not require any commands, and is transparent to both the users and the Access Appliance administrators.

The quota limits the amount of disk space you can allocate for the files in a home directory.

You can set the same quota value for all home directories using the `Storage> quota cifshomedir setall` command.

Displaying home directory usage information

You can display information about home directories using the `CIFS> homedir show` command.

Note: Information about home directory quotas is up-to-date only when you enable the use of quotas for the home directory file systems.

To display information about home directories

- 1 To display information about a specific user's home directory, enter the following:

```
CIFS> homedir show [username] [domainname]
```

username The name of the CIFS user. If a CIFS user name includes a space, enter the user name with double quotes.

domainname The domain where the home directory is located.

- 2 To display information about all home directories, enter the following:

```
CIFS> homedir show
```

Deleting home directories and disabling creation of home directories

You can delete a home directory share. This also deletes the files and sub-directories in the share.

After a home directory is deleted, if you try to access the same home directory again, a new home directory will automatically be created.

If you have an open file when the home directory is deleted, and you try to save the file, a warning appears:

```
Warning: Make sure the path or filename is correct.
```

```
Save dialog?
```

Click on the **Save** button which saves the file to a new home directory.

To delete a home directory share

- ◆ To delete the home directory of a specific user, enter the following:

```
CIFS> homedir delete username [domainname]
```

username The name of the CIFS user. If a CIFS user name includes a space, enter the user name with double quotes.

Respond with *y(es)* or *n(o)* to confirm the deletion.

domainname The domain it is located in.

You can delete all of the home directory shares with the `CIFS> homedir deleteall` command. This also deletes all files and subdirectories in these shares.

After you delete the existing home directories, you can again create the home directories manually or automatically.

To delete the home directories

- ◆ To delete all home directories, enter the following:

```
CIFS> homedir deleteall
```

Respond with *y(es)* or *n(o)* to confirm the deletion.

After you delete the home directories, you can stop Access Appliance serving home directories by using the `CIFS> set homedirfs` command.

To disable creation of home directories

- ◆ To specify that there are no home directory file systems, enter the following:

```
CIFS> set homedirfs
```

After these steps, Access Appliance does not serve home directories.

About CIFS clustering modes

The following clustering modes are supported by Access Appliance:

- Normal
- Clustered Trivial Database (CTDB) - a cluster implementation of the TDB (Trivial database) based on the Berkeley database API

The following operating modes are supported by Access Appliance:

- User

- Domain
- ADS

Each clustering mode supports all of the three operating modes.

About switching the clustering mode

You can switch from normal to ctdb clustering mode or from ctdb to normal clustering mode. You must stop the CIFS server prior to switching to any cluster mode.

About migrating CIFS shares and home directories

You can migrate CIFS shares and home directories from normal to ctdb clustering mode and from ctdb to normal clustering mode.

Access Appliance automatically migrates all CIFS shares and home directories while switching from one clustering mode to another.

Automatic migration of the content of users (that is, users' home directories) from one file system to another file system while switching home directories is not supported. So, if a Access Appliance administrator changes home directories from fs1 to fs2, then users' home directories are not migrated from fs1 to fs2 automatically.

Migrating CIFS shares and home directories from normal to ctdb clustering mode

To migrate CIFS shares and home directories from normal to ctdb clustering mode

- 1 To check the CIFS server status to confirm that the current cluster mode is set to normal, enter the following:

```
CIFS> server status
```

- 2 To list the CIFS shares and home directories, enter the following:

```
CIFS> share show
```

- 3 To stop the CIFS server before changing the clustering mode to ctdb, enter the following:

```
CIFS> server stop
```

```
CIFS> set clustering_mode ctdb
```

- 4 To start the CIFS server in ctdb clustering mode and check the CIFS server status, enter the following:

```
CIFS> server start
CIFS> server status
```

- 5 To verify that all the CIFS shares and home directories are properly migrated to the ctdb clustering mode, enter the following:

```
CIFS> share show
CIFS> homedir show
```

Migrating CIFS shares and home directories from ctdb to normal clustering mode

To migrate a CIFS share and home directory from ctdb to normal clustering mode

- 1 To check the status of the CIFS server, enter the following:

```
CIFS> server status
```

- 2 To list the CIFS shares and home directories, enter the following:

```
CIFS> share show
CIFS> homedir show
```

- 3 To stop the CIFS server to switch the clustering mode to normal, enter the following:

```
CIFS> server stop
CIFS> set clustering_mode normal
```

- 4 To start the CIFS server in normal clustering mode, enter the following:

```
CIFS> server start
```

- 5 To list the CIFS shares and home directories after migrating to normal clustering mode, enter the following:

```
CIFS> share show
CIFS> homedir show
```

Setting the CIFS aio_fork option

The `CIFS> set aio_size` option allows you to set an Asynchronous I/O (AIO) read/write size with an unsigned integer.

To set the aio_fork option

- ◆ To set the `aio_fork` option, enter the following:

```
CIFS> set aio_size size
```

where *size* is the AIO read/write size.

If *size* is not set to 0, then enable the `aio_fork` option, and set it as an AIO read/write size. If *size* is set to 0, then disable the `aio_fork` option, and set 0 to an AIO read/write size.

About managing local users and groups

When Access Appliance is operating in the standalone mode, only the local users and groups of users can establish CIFS connections and access the home directories and ordinary shares. The Access Appliance local files store the information about these user and group accounts. Local procedures authenticate and authorize these users and groups based on the use of names and passwords. You can manage the local users and groups as described in the rest of this topic.

Accounts for local users can be created, deleted, and information about them can be displayed using the `CIFS> local user` commands.

Creating a local CIFS user

To create the new local CIFS user

- ◆ To create a local CIFS user, enter the following:

```
CIFS> local user add username [grouplist]
```

Where *username* is the name of the user.

The *grouplist* is a comma-separated list of group names.

To reset the local user password

- ◆ To set the local password, enter the following:

```
CIFS> local password username
```

Where *username* is the name of the user whose password you want to change.

To display the local CIFS user(s)

- 1 To display local CIFS users, enter the following:

```
CIFS> local user show [username]
```

Where *username* is the name of the user.

- 2 To display one local user, enter the following:

```
CIFS> local user show usr1
```

To delete the local CIFS user

- ◆ To delete a local CIFS user, enter the following:

```
CIFS> local user delete username
```

Where *username* is the name of the local user you want to delete.

To change a user's group membership

- ◆ To change a user's group membership, enter the following:

```
CIFS> local user members username grouplist
```

Where *username* is the local username being added to the *grouplist*.

Group names in the *grouplist* must be separated by commas.

Configuring a local group

A local user can be a member of one or more local groups. This group membership is used in the standalone mode to determine if the given user can perform some file operations on an exported share. You can create, delete, and display information about local groups using the `CIFS> local group` command.

To create a local group

- ◆ To create a local group, enter the following:

```
CIFS> local group add groupname
```

where *groupname* is the name of the local group.

To list all local groups

- ◆ To list all existing local groups, enter the following:

```
CIFS> local group show [groupname]
```

where *groupname* lists all of the users that belong to that specific group.

To delete the local CIFS groups

- ◆ To delete the local CIFS group, enter the following:

```
CIFS> local group delete groupname
```

where *groupname* is the name of the local CIFS group.

Enabling CIFS data migration

Access Appliance provides the following command for enabling CIFS data migration:

```
CIFS> set data_migration yes|no
```

To enable data migration for the CIFS server

- 1 To enable data migration for the CIFS server, enter the following:

```
CIFS> set data_migration yes
```

- 2 Restart the CIFS server by entering the following command:

```
CIFS> server start
```

- 3 Map the CIFS share on the Windows domain using the *isa_Cluster_Name\root* by the Domain Administrator.

- 4 Copy the data with ROBOCOPY by entering the following command in a Windows command prompt:

```
C:\> ROBOCOPY /E /ZB /COPY:DATSO [windows_source_dir] [CIFS_target_dir]
```

Make sure you have the Windows Resource Kit Tools installed.

- 5 Disable the CIFS data migration option after migration completes for CIFS server security by entering the following command:

```
CIFS> set data_migration no
```

- 6 Restart the CIFS server by entering the following command:

```
CIFS> server start
```

Using Access Appliance as an Object Store server

This chapter includes the following topics:

- [About the Object Store server](#)
- [Use cases for configuring the Object Store server](#)
- [Configuring the Object Store server](#)
- [About buckets and objects](#)
- [File systems used for objectstore buckets](#)
- [Enabling WORM on buckets](#)
- [S3 with NFS use case](#)
- [S3 with NSP use case](#)
- [Configuring the S3 server using GUI](#)

About the Object Store server

The Access Appliance Object Store server lets you store and retrieve the data that is stored in Access Appliance using the Amazon Simple Storage Service (S3) compatible protocol. The protocol implements the RESTful API interface over standard HTTP or HTTPS protocol.

See the ObjectAccess service (S3) APIs section in the *Access Appliance Object Access API Guide* for more information on the APIs supported by the Access Appliance Object Store server.

Features of the Object Store server include the following:

- High availability
- Customization of storage layouts as per requirement
- Concurrent access from multiple nodes
- Scalability with multiple nodes and buckets
- Sharing of file systems and customization of file system layout using groups

Use cases for configuring the Object Store server

You can configure the Object Store server depending on different use cases.

Use Case 1: Large number of objects per bucket are required.

- The admin can configure a default pool without using the `fs_sharing` option.
- The file system is not shared across buckets. A bucket can have large number of objects. The choice of file system sharing limits the number of buckets created.

Use Case 2: Admin needs large number of buckets but does not expect large number of objects per bucket.

- The admin can create a group in its authentication server and configure this group in Object Store using the `objectaccess> group set` command.
- The grouping provides options like choosing the disk pool to use, file system type, file system sharing, file system size, other file system options.
- The admin can use the `fs_sharing` option to configure the Object Store server to share a file system across all buckets that are created by a user of that particular group.
- The file system sharing allows the Object Store server to create a large number of buckets but limits the total number of objects present across the bucket.

Use Case 3: Admin wants to control the file system used for a bucket.

- The admin has to pre-create the required file system using the `storage> fs` commands.
- The admin can use the `objectaccess> map` command to map a directory of the existing file system as a bucket for a particular user.

Configuring the Object Store server

This section describes how to configure the Object Store server.

To configure the Object Store server

- 1 Log on to Access Appliance using the Access Appliance command-line interface.
- 2 You can either use an existing pool or create a default storage pool (at least one) on the cluster.

You can see the list of existing pools using the `storage pool list` command.

You can create a new pool using the `storage create pool` command.

```
storage> pool create pool1 disk1,disk2,disk3,disk4
```

- 3 Use the storage pool that was created in Step 2 as the default object access pool.

You need to set the default pool, as it is required for enabling the Object Store server.

```
objectaccess> set pools pool1
```

Note: Based on your storage requirements, you can configure different types of storage pools by using the Object Store group commands.

- 4 Verify the configured storage pool.

```
objectaccess> show
```

- 5 Enable and start the Object Store server.

```
objectaccess> server enable
```

```
objectaccess> server start
```

- 6 Configure the cluster using any authentication server (AD, LDAP, or NIS).

See the following manual pages for more information on configuring AD, LDAP, or NIS:

- `CLISH> network man ldap`

- `CLISH> network man ad`

- `CLISH> network man nis`

- 7 Create the access and secret keys for the authorized user, or any user in the authentication server.

You have two options for creating the access and the secret keys, either using the Access Appliance RESTful APIs or by using the Access Appliance helper script.

Create the access and secret keys using the Access Appliance RESTful APIs:

- Before using the Access Appliance RESTful APIs, set the host name resolution for the host as shown in the `objectaccess> show` output against `ADMIN_URL`.
- See the *Access Appliance Object Access API Guide* on the [SORT](#) site for accessing the Object Store server (S3) user management APIs.
- After creating your access and secret key, you can create a bucket using the S3 API.

Create the access and the secret keys using the Access Appliance helper script:

- Add the `ADMIN_URL` name in your `/etc/hosts` file.
Where the `ADMIN_URL` is `admin.<cluster_name>` and the port is 8144.
This url should point to the Access Appliance management console IP address.
- Location of the helper script:
`/opt/VRTSnas/scripts/utils/objectaccess/objectaccess_client.py`
- The Access Appliance helper script can be used from any client system that has Python installed.
- To run the script, your S3 client needs to have the `argparse` and `requests` Python modules.
If these modules are missing, install both these modules using `pip` or `easy_install`.
- If the Object Store server is enabled without the `SSL` option, you need to add the `--insecure` option.

```
clus_01 ~# ./objectaccess_client.py --server  
admin.clus:8144 --username <uname> --create_key --insecure
```

- Create the access and the secret key using the Access Appliance helper script by providing the username, password, and `ADMIN_URL` (check the online Help of the Access Appliance helper script for all of the provided operations like `list key` and `delete key`).

Create a secret key:

```
clus_01:~ # ./objectaccess_client.py --create_key
--server admin.clus:8144 --username localuser1 --password root123
--insecure
UserName                : localuser1
AccessKeyId              : Y2FkODU2NTU2MjVhYzV
Status                   : Active
SecretAccessKey          : ODk0YzQxMDhkMmRjM2M5OTUzNjI5OWIzMdGyNzY
```

The `<localuser1>` is the local user created on both the Access Appliance cluster nodes with same unique ID.

List a secret key for the specified user:

```
clus_01:~ # ./objectaccess_client.py --list_key --server
admin.clus:8144 --username localuser2 --password root123 --insecure
```

Delete a secret key for the specified user:

```
clus_01:~ # ./objectaccess_client.py --delete_key
ZTkYNDdjZTViM2EyMWZ --server admin.clus:8144 --username localuser2
--password root123 --insecure
```

- 8 Use the following `objectaccess` command to see all the existing access and secret keys in the Access Appliance cluster:

```
objectaccess> account user show
```

Changing the Object Store server options

It is possible to change an already set parameter or set new parameters by specifying different options. For example, you can change the other Object Store server defaults, such as `fs_type`, `fs_size`, and other options.

After setting the defaults, you can verify whether the proper value is assigned or not.

```
objectaccess> set fs_type
```

```
mirrored mirrored-stripe simple striped striped-mirror
```

```
objectaccess> set fs_type simple
```

```
ACCESS ObjectAccess INFO V-288-0 Set fs_type successful.
```

```
objectaccess> set fs_size 2G
```

```
ACCESS ObjectAccess INFO V-288-0 Set operation successful.
```

```
objectaccess> show
Name                Value
=====
Server Status      Enabled
Admin_URL           http://endpoint1:8144
S3_URL              http://dataendpoint:8143
admin_port          8144
s3_port             8143
ssl                 no
max_s3_threads      8
pools               pool_default
fs_size             2g
fs_type             simple
fs_blksize          8192
fs_pdirenable       yes
fs_encrypt          off
fs_worm             yes
retention_min       3600s
retention_max       36000s
```

Using the group option for bucket creation

If you have multiple users, and you want to set different default values for different sets of users, you can use the `group` option.

You can also use the `group` option to use the existing file systems for bucket creation instead of creating a new file system for every bucket. If you set the `group fs_sharing` option to **yes**, and if any request for bucket creation comes from a user who is part of that group, then the S3 server searches for any existing file system created by the specific group user. If an existing file system is found, it uses the existing file system. Otherwise, it creates a new file system for the bucket.

To use the group option

- 1 Create a group in the authentication server (AD/LDAP/NIS) and add the required users to that group.
- 2 Set the group specific configuration for the group created in the authentication server.
- 3 Set or unset the defaults per your requirements.

```
objectaccess> group set fs_type simple VRTS-grp
ACCESS ObjectAccess INFO V-288-0 Group set fs-type successful.
```

```
objectaccess> group set pool VRTS-grp pool1
ACCESS ObjectAccess INFO V-288-0 Success.
```

```
objectaccess> group show
```

Group Name	Fs Sharing	Fs Size	Fs Type	Pool (s)
VRTS-grp	-	-	simple	pool1

```
objectaccess> group show
```

Group Name	Fs Sharing	Fs Size	Fs Type	Pool (s)
VRTS-grp	-	-	-	pool1

About buckets and objects

The Object Store server consists of a collection of objects. The container of an object is known as a bucket. In Access Appliance Object Store, the buckets are stored on file systems as directories and objects are stored as files.

Buckets and objects are resources which can be managed using the APIs.

Once the Object Store Server is configured, you can create buckets and objects and perform the required operations.

Access Appliance supports the following methods for accessing the buckets and the objects:

- Path-style method
- Virtual-hosted-style method

When using the virtual hosted-style method, the **bucket_name.s3.cluster_name** should be DNS resolvable. The bucket name is part of the endpoint URL (bucket_name.<endpoint_name>). For example, bucket_name.s3.cluster_name. You can send a request to the S3 server only

if this URL is DNS resolvable by the client's DNS servers. A separate entry in the DNS server is required for every bucket.

See the `objectaccess_bucket(1)` manual page for more information.

Note: Access S3 does not support creation of object with and without "/" at the same time. For example, you cannot create two objects such as, `bucket1/object1` and `bucket1/object1/`, where the latter represents a directory. The user should use separate names for directories and real objects.

Note: Object name containing `...::` substring is not supported.

See the `objectaccess` manual pages for all of the Access Appliance Object Store server operations that can be performed.

Buckets are created by S3 clients by calling the standard S3 APIs to the Access Appliance S3 server. For creating a bucket, you need the endpoint of the Access Appliance server, access key, and the secret key. The endpoint of the Access Appliance Object Store server is `s3.cluster_name:8143`.

The Access Appliance Object Store server can also be accessed using the fully qualified domain name:

`s3.cluster_name.fqdn:8143`

Make sure that you associate one (or more) of the VIPs of the Access Appliance cluster to `s3.cluster_name.fqdn` in the client's DNS server.

You can delete a bucket using the standard API or by using the `objectAccess bucket delete` command.

```
objectaccess> bucket delete bucket_name
```

[Table 10-1](#) describes the restrictions enforced by the Access Appliance Object Storage Server. Configure your S3 clients within these limitations to ensure that Access Appliance works correctly.

Table 10-1 Object and bucket restrictions

Description	Limit
Maximum recommended parallel threads	8

Table 10-1 Object and bucket restrictions (*continued*)

Description	Limit
Maximum number of buckets per file system without enabling fs_sharing and without MAP command	50 The maximum number of supported file systems is 50.
Maximum number of buckets per file system with fs_sharing enabled and with MAP command	50
Maximum number of objects per file system	1 billion
Maximum length for bucket names	63
Maximum length for '/' separated object name segment	255
Maximum supported size of an object that can be uploaded using a single PUT	100 MB
Maximum supported size range of an object that can be downloaded using a single GET	8 EB
Maximum number of parts supported for multipart upload	10,000
Maximum number of objects returned per list objects request	1,000
Maximum number of objects returned per list object versions request	1,000
Maximum number of parts returned per list parts request	1,0000
Maximum number of multipart uploads returned per list multipart uploads request	1,0000
Maximum number of grantees supported for setting ACL on buckets/objects	128

Table 10-1 Object and bucket restrictions (*continued*)

Description	Limit
Maximum retention time that can be set on objects	60 years in future or less than year 2106 whichever is less Must be between the minimum and maximum retention period of the bucket
Supported clustering modes to create WORM/lock-enabled bucket	Compliance and Enterprise modes
Maximum number of supported S3 data and admin endpoints	4

Note: The maximum number of buckets supported is 50 since the maximum number of supported file systems is 50. If file systems are created for some other purposes, the number of available file systems for S3 buckets gets reduced.

File systems used for objectstore buckets

Access Appliance supports the following file systems for creating buckets:

- Mirrored
- Mirrored-stripe
- Simple
- Striped
- Striped-mirror

Enabling WORM on buckets

You can enable WORM on a bucket using the `objectaccess bucket worm` commands.

```
objectaccess> bucket worm set [minret] [maxret] <bucket_name>
```

Where

bucket_name	Specifies the name of the bucket
minret	Specifies the minimum time period of retention. The format can be: [1-9](s S h H d D m M y Y)
maxret	Specifies the maximum time period of retention. The format can be: [1-9](s S h H d D m M y Y)

Note: This command is not supported in Normal lockdown mode.

You can also list the WORM status, minimum and maximum retention period for a specified bucket. Retention values are in seconds.

```
objectaccess> bucket worm get
```

Note: In Normal lockdown mode, the WORM status appears as "No".

Setting retention period for objects

You can set the retention period on objects using the `objectaccess object retention` commands. If the version ID is provided, retention is set on the object with specified version ID. If version ID is not passed, retention is set on the object with the latest version. In Enterprise lockdown mode, if retention time is given as 0, retention is removed on that object.

```
objectaccess> object retention set <object_path> <retention_time>
[version_id]
```

Where

object_path	Specifies the path of the object starting with bucket name
-------------	--

retention_time	Specifies the retention period. The format can be: [1-9](s S h H d D m M y Y) or mm-dd-yyyy or mm-dd-yyyy:hh:mm:ss or 0
version_id	Specifies the version ID of the object

You can also display the retention period on objects. If the version ID is provided, retention period of the object with the specified version ID is displayed. If version ID is not passed, retention period of the object with the latest version is displayed.

```
objectaccess> object retention get <object_path> [version_id]
```

Note: In Normal lockdown mode, the retention status appears as "No retention set".

Note: Versioning is only supported for WORM buckets.

S3 locking privileges

Table 10-2 describes the supported S3 mode with cluster lockdown mode for a non MSDP-C use case.

Table 10-2 Non MSDP-C use case

	Normal cluster mode	Enterprise cluster mode	Compliance cluster mode
S3 Lock request with Compliance mode	Not allowed	Supported	Supported
S3 Lock request with Governance mode	Not allowed	Not supported	Not supported

Table 10-3 describes the supported S3 mode with cluster lockdown mode for an MSDP-C use case.

Table 10-3 MSDP-C use case

	Normal cluster mode	Enterprise cluster mode	Compliance cluster mode
S3 Lock request with Compliance mode	Not allowed	Not supported	Supported

Table 10-3 MSDP-C use case (*continued*)

	Normal cluster mode	Enterprise cluster mode	Compliance cluster mode
S3 Lock request with Governance mode	Not allowed	Not supported	Not supported

Note: An administrator can create an S3 bucket which is not WORM-enabled in both enterprise and compliance mode. Administrator can also change mode of existing bucket to WORM-enabled by using the `lockdown-mode set` command.

S3 with NFS use case

Access Appliance supports multi-protocol support for NFS with S3. If an NFS share is present (and objects may be present in the exported path), the storage admin can map that path as an S3 bucket (S3 over NFS). In addition, a normal file system path can also be mapped as an S3 bucket. The buckets created by S3 APIs cannot be exported as an NFS share (NFS over S3).

Obtaining the path to map as S3 bucket

The path has the following characteristics:

- Path is the absolute path inside a file system.
- The name of the bucket is the name of the directory of the path which should be S3 compliant.
- The path can be either NFS exported path or any directory in the normal file system. You cannot use the ObjectAccess file systems (file system having S3 bucket created by S3 APIs).
- No other bucket should exist with the same name.
- No other bucket should be present either inside or outside the given path. You can verify this using the following command:

```
objectaccess> bucket show
```

- NFS share should not be present before or after that directory. You can verify using the following command:

```
NFS> share show
```

Creating an S3 user

You can configure the cluster with any authentication server like AD/LDAP/NIS. Then, all the users present in the authentication server can be used as S3 users.

The S3 user should be authorized to access the S3 bucket (access key and secret key should be present for that user). You can verify using the following command:

```
objectaccess> account user show
```

See [“Configuring the Object Store server”](#) on page 158.

Mapping the path to the S3 bucket for the user

You can map the path to a bucket. The path can be either a directory inside a file system or an NFS-exported path. If the specified path is not present, then the `map` command creates the directory and sets the default S3 bucket permissions.

You can map the path to the S3 bucket for the user using the following command:

```
objectaccess> map <path> <user>
```

The storage admin can verify the bucket creation using the following command:

```
objectaccess> bucket show
```

Using the multi-protocol feature

The storage admin can use the NFS share at the same time when the S3 user uses the bucket. Existing objects inside the bucket retain the permissions set by the owner of those objects.

Unmapping the S3 bucket

In multi-protocol case, an S3 user can delete the bucket without deleting all the objects. Deleting the bucket is equivalent to unmapping or unreferencing the bucket.

Limitations

The following limitations apply for multi-protocol support:

- An S3 user cannot access a bucket if the bucket ownership or permissions from the NFS client is changed.
- Permissions that are set or modified from protocols like NFS are not honored by S3 and vice versa.
- Object ETag is inaccurate whenever object is created or modified from the NFS client. An incorrect ETag is corrected when a GET or HEAD request is performed on the object.

- Accessing the same object from different protocol in exclusive mode is not supported.

S3 with NSP use case

If an S3 server bucket is provisioned for NetBackup SaaS Protection (NSP) storage, the administrator must increase the maximum size of the object of the S3 server in the S3 configuration file. Change the value of `cf_max_obj_size` to 5497558138880 in the `/opt/VRTSnas/conf/portald.conf` file on all the nodes.

Configuring the S3 server using GUI

You can configure the S3 server using the Access Appliance GUI.

Setting Object Access server default parameters

You can set the Object Access server configurable and default file system parameters. The object access server creates a new file system for every new bucket. Use these default parameters to create a new file system of a different configuration.

To edit the default object server parameters

1 Prerequisites:

Configure at least one default storage pool.

2 In **Settings > S3 Management** for **Default parameters**, click **Edit**.

3 You can edit the following fields:

File system size

Use the file system size option to set the default file system for creating the buckets.
Units can be MB, GB, or TB.

Block size

The default block size is 8192 bytes.

Block size can affect the file size. For example, to create a file system greater than 32 TB, the block size needs to be 8192.

Enable Partition Directory	The partition directory feature distributes entries into hash directories. The hash directories are not visible in the name space view of the user or the operating system. By default, partition directory is disabled.
Encryption	Optional parameter that lets you encrypt all the data on the file system.
Number of mirrors	For layouts that use mirroring, the minimum number of mirrors is 2 (the default).
Protection Disk	For layouts that use mirroring, the protection option determines where the mirrors are created. ■ Disk (default) - creates mirrors on separate disks ■ Pool - creates mirrors in separate pools.
Set default pools	The pools option lets you choose the pool on which the bucket is created. You must specify at least one default storage pool.
Enable SSL	It is recommended to enable SSL for S3. SSL enables HTTPS communication. See the <i>Access Appliance RESTful API Guide</i> for more information.

4 View the **Recent Activity** panel for the status of the task.

Setting up the Object Access server group-specific parameters

You can set default file system parameters that can be used for creating file systems for all the users in a particular group.

Creating a user group lets you do the following:

- Specify that all the S3 buckets for a particular group get created on the specified pool.
- Specify that all the S3 buckets for a particular group get created with the same file system size.

- Specify that all the S3 buckets for a particular group get created with the same file system layout.

To set file system parameters for a group of users

- 1** In the **Settings > S3 Management** for **Group related parameters**, click **Add Group**.
- 2** Specify a name for the group.
- 3** Specify the file system size for the group.
- 4** Specify the file system layout.
- 5** Specify the pools for the group.
- 6** View the **Recent Activity** panel for the status of the task.

Viewing information about S3

From the **Settings > S3 Management** page, you can view the following information:


- Default parameters for creating S3 buckets
- Group-related parameters for associating a group with the S3 buckets
- S3 users and their access keys

Provisioning the S3 bucket using GUI


This section describes the steps to start the S3 server and provision the S3 bucket using the Access GUI.

To provision the S3 bucket

- 1 Go to **Quick Actions** and click **Provision storage**.



- 2 Select **Storage type** as **Storage for Buckets** and select **S3 without cloud** for generic S3 bucket.



Note: The generic S3 bucket also supports cloud tiering. Select **S3 without cloud option** under **Storage for NetBackup** storage type for configuring NetBackup with S3. For more details, see the *Veritas Access Appliance Solutions Guide for NetBackup*.

- 3 In the **Configure Storage** tab, *pool_default* is added as the default pool. You can create new pools or edit existing pools.

Provision storage

S3 Setup

Configure storage (active) | Configure S3 Server | Activate Storage Policy | Generate S3 Keys | Provision Storage

Create/edit Pool

This step enables creation of a new pool or editing an existing pool.

Storage pool	Total size	Available size
pool_default	250.00 GB	203.87 GB

Items per page: 10 | 1 - 1 of 1 | < >

Add pool

Create pool

- 4 Select the appropriate storage pool to configure the S3 server and click **Start S3 server** to make it online.

Note: SSL can be enabled, if required.

S3 Setup

Configure storage | Configure S3 Server (active) | Activate Storage Policy | Generate S3 Keys | Provision Storage

Start S3 server

This step configures S3 server and make it online. This configuration enables the Veritas Access server to create S3 bucket(s) on-premises.
This step will also start the server and bring it online after the default settings are configured.
S3 server is online. You can change the configuration.

Select storage pool :

	Storage pool	Total size	Available size
☐	pool_default	250.00 GB	203.87 GB

Items per page: 2 | 1 - 1 of 1 | < >

Access keys and secret keys should be regenerated if this setting is changed in future.

☐ Enable SSL

Start S3 server

- 5 In the **Activate Storage Policy** page, choose the policy for storage provisioning,

Provision storage

S3 Setup

Configure storage

Configure S3 Server

Activate Storage Policy

Generate S3 Keys

Provision Storage

Select and activate the policy for storage provisioning.

Complete prerequisites

Policy name	Description	Capabilities	Prerequisite
☒ S3 bucket	This policy provisions an S3 bucket to store data over ...	Fault Tolerance	☒ Storage pool created
☐ S3 bucket (WORM)	This policy provisions an S3 bucket to store data over ...	Fault Tolerance,WORM	☒ Storage pool created ☒ Enterprise/Compliance mode configured

Items per page: 5
1 - 2 of 2

⏪

<

>

⏩

Activate policy

- For S3 bucket (non-WORM), click **Complete prerequisites**. Choose a storage pool from the list of pools available.

Complete following prerequisite to activate S3 bucket

Select storage pool :Selected storage pools:1, total size:206.50 GB

☒	Storage pool	Total size	Available size
☒	pool_default	250.00 GB	206.50 GB

Items per page: 5
1 - 1 of 1

⏪

<

>

⏩

Activate policy

- For S3 bucket (WORM), click **Complete prerequisites**. The cluster must be in Enterprise or Compliance mode. Choose a storage pool from the list of pools available.

Complete following prerequisite to activate S3 bucket (WORM)

Select storage pool :Selected storage pools:1, total size:206.50 GB

☒	Storage pool	Total size	Available size
☒	pool_default	250.00 GB	206.50 GB

Items per page: 5
1 - 1 of 1

⏪

<

>

⏩

Set WORM

Activate policy

Set the minimum and maximum retention time.

Complete following prerequisite to activate S3 bucket (WORM) ✕

Select storage pool :Selected storage pools:1, total size:206.50 GB ▼

Set WORM ▲

Configure the retention time range that Appliance Administrators will be allowed to select from when setting WORM retention time. The minimum retention time you can set is 1 hour and the maximum is 60 years.

Minimum retention time

0

Maximum retention time

0

Hours ▼

Activate policy

Note: The retention values can be set on objects within the range of the minimum retention period and maximum retention period provided in this step.

Click **Activate policy**.


- 6 In the **Generate S3 keys** tab, you can use any local/AD/LDAP user credentials to generate the S3 keys. Enter the username for which you want to create keys and specify the user password. Click **Generate S3 keys**.



The screenshot shows the 'S3 Setup' window with the 'Generate S3 Keys' tab selected. The window has a title bar 'S3 Setup' and a navigation bar with buttons: 'Configure storage' (checked), 'Configure S3 Server', 'Activate Storage Policy', 'Generate S3 Keys' (active), and 'Provision Storage'. The main content area is titled 'Generate S3 access key and secret key.' and contains a message: 'You need an access and secret key to access the S3 buckets. Enter user credentials to generate the keys.' Below this are two input fields: 'Name *' and 'Password *'. At the bottom right is a 'Generate S3 Keys' button.


Note down the access key and secret key generated in this step for future use.

- 7 In the **Provision Storage** tab, select the storage policy as **S3 bucket** and enter the bucket size. Enter the access key and secret key generated in step 6 and click **Configure storage**.



The screenshot shows the 'S3 Setup' window with the 'Provision Storage' tab selected. The window has a title bar 'S3 Setup' and a navigation bar with buttons: 'Configure storage' (checked), 'Configure S3 Server' (checked), 'Activate Storage Policy' (checked), 'Generate S3 Keys' (checked), and 'Provision Storage' (active). The main content area is titled 'Provision storage' and contains a message: 'This step provides storage for S3 buckets. You have to enter the access and secret keys created in the previous step.' Below this are two sections: 'On-premises storage options' and 'Bucket size'. The 'On-premises storage options' section has a 'Select storage policy' dropdown menu with 'S3 bucket' selected. The 'Bucket size' section has a 'Bucket size' dropdown menu with 'Available - 146.50 GB' selected. Below these are two input fields: 'Access key' and 'Secret key'. At the bottom right is a 'Configure storage' button.

You can view the buckets by navigating to **Provisioned Storage > Buckets**.



The screenshot shows the 'Provisioned storage' window with the 'Buckets' tab selected. The window has a title bar 'Provisioned storage' and a navigation bar with buttons: 'Veritas Data Deduplication', 'Buckets' (active), 'Shares', and 'iSCSI'. The main content area is a table with the following columns: 'Name', 'File system', 'Storage pool', 'Owner', and 'WORM enabled'. The table contains one row with the following data: 'Name: e4-4d78-8841-903ee958165a3bucket', 'File system: ext4', 'Storage pool: pool_default', 'Owner: root', and 'WORM enabled: No'. Below the table is a pagination bar with the text 'Items per page: 5' and '1 - 1 of 1'.

Name	File system	Storage pool	Owner	WORM enabled
e4-4d78-8841-903ee958165a3bucket	ext4	pool_default	root	No

Managing Access Appliance security

- [Chapter 11. Managing security](#)
- [Chapter 12. Setting up FIPS mode](#)
- [Chapter 13. Configuring STIG](#)
- [Chapter 14. Setting the banner](#)
- [Chapter 15. Setting the password policy](#)
- [Chapter 16. Immutability in Access Appliance](#)

Managing security

This chapter includes the following topics:

- [Security overview](#)

Security overview

Access Appliance includes multiple features to ensure the security of your data. Each element of the appliance is tested for vulnerabilities using both industry standards and advanced security products. These measures ensure that exposure to unauthorized access and resulting data loss or theft is minimized.

Note: IP forwarding is enabled in Access Appliance by design; it is used to facilitate network communication between application instances and external networks.

Setting up FIPS mode

This chapter includes the following topics:

- [FIPS 140-2 conformance for Access Appliance](#)
- [Viewing FIPS status for Access Appliance](#)
- [Enabling FIPS for Access Appliance](#)
- [To enable FIPS using the command-line interface](#)

FIPS 140-2 conformance for Access Appliance

The Federal Information Processing Standards (FIPS) define U.S. and Canadian Government security and interoperability requirements for computer systems. The National Institute of Standards and Technology (NIST) issued the FIPS 140 Publication Series to coordinate the requirements and standards for validating cryptography modules. The FIPS 140-2 standard specifies the security requirements for cryptographic modules and applies to both the hardware and the software components. It also describes the approved security functions for symmetric and asymmetric key encryption, message authentication, and hashing.

For more information about the FIPS 140-2 standard and its validation program, see the following links:

<https://csrc.nist.gov/csrc/media/publications/fips/140/2/final/documents/fips1402.pdf>

<https://csrc.nist.gov/projects/cryptographic-module-validation-program>

Starting with Access Appliance 7.4.3, FIPS 140-2 standard is enabled by default for the Veritas Operating System (VxOS). The FIPS mode for VxOS is enabled with the default factory settings. After FIPS for VxOS is enabled, the `sshd` uses the following FIPS approved ciphers:

- `aes128-ctr`

- aes192-ctr
- aes256-ctr

Older SSH clients are likely to prevent access to the appliance after FIPS for VxOS is enabled. Ensure that your SSH client supports the listed ciphers, and upgrade to the latest version if necessary. Default cipher settings are not typically FIPS-compliant, which means you might need to select them manually in your SSH client configuration.

You can enable the FIPS 140-2 standard for NetBackup MSDP to increase appliance security. See [“Enabling FIPS for Access Appliance”](#) on page 183.

Viewing FIPS status for Access Appliance

The Federal Information Processing Standards (FIPS) 140-2 standard is enabled for Veritas Operating system (VxOS) with the default factory settings. You can view the FIPS status of VxOS and NetBackup MSDP for a cluster node.

The Federal Information Processing Standards (FIPS) 140-2 standard is enabled for Veritas Operating system (VxOS) with the default factory settings. You can view the FIPS status of VxOS for a cluster node and the FIPS status of NetBackup MSDP for each Veritas Data Deduplication server configured for the cluster.

To view the FIPS status:

- 1 Log in to the Access web interface of the configured cluster by opening a supported browser and typing:

`http://console-ip:14161`

where *console-ip* is the management console IP address where the web interface is hosted.
- 2 In the navigation pane, click **Settings**, and then click **Security management**.

On the **FIPS** tab, the FIPS status for VxOS and MSDP is displayed for each of the cluster nodes:
 - **Enabled:** The FIPS mode is enabled.
 - **Disabled:** The FIPS mode is disabled.
 - **Unknown:** The FIPS status cannot be retrieved because the node is stopped, shut down, or unreachable.

Enabling FIPS for Access Appliance

You can enable the FIPS mode for NetBackup MSDP for increased appliance security.

When you enable FIPS for MSDP, all backup and restore jobs that are in progress are terminated. After you enable the FIPS mode, restart the NetBackup services to restart the jobs.

To enable the FIPS mode for NetBackup MSDP, complete the following steps:

- 1** Log in to the Access web interface of the configured cluster by opening a supported browser and typing:

`http://console-ip:14161`

where *console-ip* is the management console IP address where the web interface is hosted.
- 2** In the navigation pane, click **Settings**, and then click **Security management**.
- 3** On the **FIPS** tab, in the MSDP node status area click **Enable FIPS**.

Settings / Security management

FIPS

STIG

VxOS node status

Nodes	Status
vmclu-01	Enabled
vmclu-02	Enabled

Showing 1-2 of 2

MSDP node status

Enable FIPS

Nodes	Status
vmclu-01	Disabled
vmclu-02	Disabled

Showing 1-2 of 2

- 4 Review the compatibility list to ensure that all components that communicate with the FIPS service are FIPS compliant. Confirm that you have reviewed the compatibility list and click **Continue**.
- 5 To monitor the progress, click **View details** on the **Security** page. The ongoing and completed tasks for the operation are also displayed in **Recent activity**.

After the operation is complete, you can view the FIPS status for both the cluster nodes. If the FIPS mode is enabled for a node, the status is displayed as **Enabled**. If the FIPS mode cannot be enabled for a node, the status is displayed as **Disabled**, and if the FIPS status cannot be retrieved because the node is stopped, shut down, or unreachable, the status is displayed as **Unknown**. For nodes that display **Unknown** status, you can enable the FIPS mode again. If the FIPS mode is enabled on the cluster and a node is down, it automatically synchronize its status with the cluster after it is up.

Settings / Security management

FIPS

STIG

VxOS node status

Nodes	Status
vmncl-01	Enabled
vmncl-02	Enabled

Showing 1-2 of 2

MSDP node status

Enable FIPS

Nodes	Status
vmncl-01	Enabled
vmncl-02	Enabled

Showing 1-2 of 2

To enable FIPS using the command-line interface

Use the following commands to enable the FIPS mode from the command-line interface.

- 1** 1. Open an SSH session and log on to the appliance as an administrator using the Access Appliance shell menu.

- 2** 2. To open an interactive shell menu, in the Access Appliance shell menu, enter the following command:

```
support elevate
```

Enter the password for the Maintenance user account.

- 3** 3. In the command shell, run the following command to find the node where the Access GUI is running:

```
hastatus -sum | grep -I ManagementConsole
```

- 4** 4. Log on to the node where the ManagementConsole is ONLINE and go to the Access command-line interface:

```
LOGNAME=admin /opt/VRTSnas/clish/bin/clish -u admin
```

- 5** 5. In the Access command-line interface, run the following command:

```
system
```

- 6** 6. Type the following command:

```
System> security fips
```

- 7** 7. To enable FIPS, enter the following command:

```
System> security fips enable MSDP
```

Successfully enabled FIPS mode.

Restart NetBackup services (media/master - servers and clients for FIPS mode to come into effect.

Configuring STIG

This chapter includes the following topics:

- [STIG overview for Access Appliance](#)
- [Enabling OS STIG hardening for Access Appliance](#)
- [Viewing STIG status of an Access Appliance](#)
- [Enabling STIG using the command-line interface](#)

STIG overview for Access Appliance

The Security Technical Implementation Guides (STIGs) provide technical guidance for increasing the security of information systems and software to help prevent malicious computer attacks. This type of security is also referred to as hardening.

In version 8.0, you can enable OS STIG hardening rules for increased security. These rules are based on the following profile from the Defense Information Systems Agency (DISA):

`STIG for Red Hat Enterprise Linux 7 Server - Version V3R3`

The STIG option is enabled at cluster level. If the STIG option is enabled, the STIG rules are enforced on both the nodes in a cluster.

See [“Enabling OS STIG hardening for Access Appliance”](#) on page 186.

Enabling OS STIG hardening for Access Appliance

You can enable OS STIG hardening rules for increased security. These rules are based on the following profile from the Defense Information Systems Agency (DISA):

`STIG for Red Hat Enterprise Linux 7 Server - Version V3R3`

After the STIG option is enabled, a STIG-compliant password policy is automatically enforced. All the current local user passwords that were created under the default password policy remain valid. Once a password expires, you must follow the STIG-compliant policy rules when you change the password.

Review the following guidelines before enabling STIG:

- When you enable STIG, the STIG option is configured for all the nodes in a cluster. The cluster must be configured before you enable the STIG option.
- The STIG option does not allow individual rule control.
- After the option is enabled, a factory reset is required to disable the associated rules.
- If Lightweight Directory Access Protocol (LDAP) is configured, it is recommended that you set it up to use the Transport Layer Security (TLS) protocol before you enable the option.
- Before you enable STIG, it is recommended that you complete the following prerequisites. However, not completing the prerequisites does not prevent you from enabling STIG. You can complete these requirements after you enable the STIG option.
 - Configure two NTP servers to synchronize the time across cluster nodes.
 - Configure two DNS servers for the cluster.
 - Configure an SMTP server to enable notifications.

To enable the STIG hardening rules, complete the following steps:

- 1 Log in to the web interface of a configured Access Appliance cluster by opening a supported browser and typing:


```
http://console-ip:14161
```

where *console-ip* is the management console IP address where the web interface is hosted.

- 2 In the navigation pane, click **Settings**.
- 3 Click **Security management**.

4 On the **STIG** tab, click **Enable STIG**.

If the prerequisites are not met, you are prompted to resolve the errors. However you can choose to ignore these errors and proceed by clicking **Continue**. You can complete the prerequisites later after you enable the STIG option. If the requirements are met, review the displayed guidelines and click **Enable STIG**.



- 5 To monitor the progress, click **View details** on the **Security** page. The ongoing and completed tasks for the operation are also displayed in **Recent activity**.

After the operation is complete, you can view the STIG status for both the cluster nodes. If the STIG option is enabled for a node, the status is displayed as **Enabled**. If the STIG option cannot be enabled for a node, the status is displayed as **Disabled** and if the STIG status cannot be retrieved because the node is stopped, shut down, or unreachable, the status is displayed as **Unknown**.

For nodes that display **Unknown** status, you can enable the STIG option again. If STIG is enabled on the cluster and the node is down, it automatically synchronizes its status with the cluster after it is up.

The screenshot shows the 'Settings / Security management' page with the 'STIG' tab selected. At the top, there are three summary boxes: '2 Nodes' (black text), '2 STIG enabled' (green text), and '0 STIG not enabled' (red text). To the right of these boxes is a blue 'Enable STIG' button with an information icon. Below this is a section titled 'STIG Node status' containing a table with two columns: 'Nodes' and 'Status'. The table lists two nodes, 'vmclu-01' and 'vmclu-02', both with a green checkmark and the status 'Enabled'. At the bottom, it says 'Showing 1-2 of 2'.

Nodes	Status
vmclu-01	✓ Enabled
vmclu-02	✓ Enabled

Viewing STIG status of an Access Appliance

You can use the Access Appliance web interface to view the STIG status.

- 1 Log in to the Access web interface of the configured Access Appliance cluster by opening a supported browser and typing:

```
http://console-ip:14161
```

where *console-ip* is the management console IP address where the web interface is hosted.

- 2 In the navigation pane, click **Settings**, and then click **Security management**.

On the STIG tab, status of both the cluster nodes is displayed:

- **Enabled:** The STIG option was successfully enabled for the node.
- **Disabled:** The STIG option is not enabled for the node.
- **Unknown:** The STIG status cannot not be retrieved because the node is stopped, shut down, or not reachable.

Enabling STIG using the command-line interface

Use the following commands to enable STIG from the command-line interface.

Before you enable STIG, it is recommended that you complete the following prerequisites. However, not completing the prerequisites does not prevent you from enabling STIG. You can complete these requirements after you enable the STIG option.

- Configure two NTP servers to synchronize the time across cluster nodes.
- Configure two DNS servers for the cluster.
- Configure an SMTP server to enable notifications.

To enable STIG:

- 1 Log on to the Access command-line interface by opening an SSH session to the management console IP as an administrator.
- 2 In the Access command-line interface, run the following command:

```
system
```

- 3 Type the following command:

```
System> security stig
```

The enable, show, and precheck options are displayed.

- 4** To run the precheck, enter the following command:

```
System> security stig precheck
```

```
NTP Check : PASS
```

```
Email Check : PASS
```

```
DNS Check : PASS
```

- 5** To enable STIG on both the nodes:

```
System> security stig enable
```

Setting the banner

This chapter includes the following topics:

- [Managing the login banner using the UI](#)
- [Managing the banner from the command-line interface](#)

Managing the login banner using the UI

You can create a customized text banner that appears when you sign in to Access Appliance UI or the system console. You can use the login banner to communicate various kinds of messages to users. Typical uses for the login banner include legal notices, warning messages, and company policy information.

If the Security Technical Implementation Guide (STIG) mode is enabled for the cluster, the login banner cannot be modified.

To set the login banner:

- 1 Log in to the web interface of a configured Access Appliance cluster by opening a supported browser and typing:

```
http://console-ip:14161
```

where *console-ip* is the management console IP address where the web interface is hosted.

- 2 In the left navigation pane, click **Settings** and then click **User management**.
- 3 Click **Manage sign-in banner**.

If a banner is already set, the Manage sign-in banner page displays the current banner.

- 4 On the Manage sign-in banner page, click **Edit**.

Note: The Edit button is disabled and the banner cannot be edited if the STIG mode is enabled.

- 5 (Optional) Under **Sign-in banner heading**, enter the banner heading. The heading can be a maximum of 250 characters.
- 6 Under **Sign in banner text**, enter the text for the banner message. The message can be a maximum of 4000 characters.
- 7 To review the changes, click **Preview**.
- 8 To confirm the changes, click **Save**.

Managing the banner from the command-line interface

You can create a customized text banner that appears when you sign in to Access Appliance UI or the system console. You can use the login banner to communicate various kinds of messages to users. Typical uses for the login banner include legal notices, warning messages, and company policy information.

If the Security Technical Implementation Guide (STIG) mode is enabled for the cluster, the login banner cannot be modified.

To display the banner:

- 1 Use SSH to log in to one of the nodes with admin credentials.
The Veritas Appliance Shell is displayed.
- 2 Enter `support elevate` and enter the maintenance account password.
- 3 To go to the Access command-line interface, enter `LOGNAME=admin`
`/opt/VRTSnas/clish/bin/clish -u admin`
- 4 To display the banner, enter the following command

```
access-clus> system banner get
No login banner is currently set for the appliance.
```

Note: No banner is set by default.

To set the banner:

- 1** Use SSH to log in to one of the nodes with admin credentials.
The Veritas Appliance Shell is displayed.
- 2** Enter `support elevate` and enter the maintenance account password.
- 3** To go to the Access command-line interface, enter `LOGNAME=admin`
`/opt/VRTSnas/clish/bin/clish -u admin`
- 4** To set the banner, run the following command:

```
# access-clus> system banner set
Enter Header (Type <enter> and <Ctrl + d> to complete the body. Max 250
characters)
This is header
Enter Body (Type <enter> and <Ctrl + d> to complete the body. Max 4000
characters)
This is body text line 1.
This is body text line 2.
```

The banner is made up of two parts, the header and the body. The header is the heading for the banner and the body is the banner content. The header can be a maximum of 250 characters and body can be a maximum of 4000 characters.

After the banner is set, you can view the banner using the `system banner get` command:

```
access-clus> system banner get
*****
This is header
*****
This is body text line 1.
This is body text line 2.
```

Note: If STIG is enabled on the system, you cannot change the default STIG banner.

```
accessclus> system banner set
ACCESS Banner ERROR V-493-10-0 The banner cannot be set as STIG is enable
on the cluster.
```

Setting the password policy

This chapter includes the following topics:

- [Managing the password policy using the UI](#)
- [Managing the password policy from the command-line interface](#)

Managing the password policy using the UI

You can customize the password policies by setting rules for the passwords that are used by the Access Appliance local users. You can set rules for password complexity, password age, and password lockout. Password complexity specifies the number and type of characters a password must include. Password age defines the duration for which the password is valid. Password lockout specifies the number of failed attempts because of incorrect usage of passwords after which a user is prevented from logging in to the account.

If the Security Technical Implementation Guide (STIG) mode is enabled for the cluster, the password policy cannot be modified.

The default password policy for a local user is as follows:

Password complexity:

- Minimum characters: 8
- Minimum numbers: 1
- Minimum lowercase characters: 1
- Minimum uppercase characters: 1
- Minimum special characters: 1

To change the password policy:

- 1 Log in to the web interface of a configured Access Appliance cluster by opening a supported browser and typing:

`http://console-ip:14161`

where *console-ip* is the management console IP address where the web interface is hosted.

- 2 In the left navigation pane, click **Settings** and then click **User management**.
- 3 Click **Manage password policy**.
- 4 On the Manage password policy page, click **Edit**.

Note: The Edit button is disabled and the password policy cannot be edited if the STIG mode is enabled.

- 5 If you want your password policy to comply with STIG, select **Reset to STIG default values** to fill in the default values for all the parameters.

Selecting this option enforces a higher security password policy.

- 6 Edit the parameters as required. To ignore a rule, leave the corresponding parameter blank. After making the changes click **Save**.

Table 15-1

Parameter	Description
Minimum characters	Minimum number of characters to include in a password
Minimum uppercase characters	Minimum number of uppercase characters to include in a password
Maximum repetitive characters of the same class	Maximum number of consecutive uppercase, lowercase, numeric, and special characters
Minimum numbers	Minimum number of numeric characters
Minimum special characters	Minimum number of special characters in a password
Minimum character classes	Minimum character classes to include in a password. Character classes include uppercase, lowercase, numeric, and special characters.
Minimum lowercase characters	Minimum number of lowercase characters

Table 15-1 *(continued)*

Parameter	Description
Maximum repetitive characters	Maximum number of characters that can be repeated in a password.
Character difference with old password	Number of characters the new password must differ by from the previous password
Days after which password can be changed	Number of days after which a password can be changed
Days after which password must be changed	Number of days after which a password must be changed
Days before warning message	Number of days before the password expires to display a warning
Minimum different passwords before allowing reuse	Number of unique passwords before a previous password can be reused
Number of incorrect login attempts before lockout	Number of failed login attempts after which the account gets locked
Time before locked account is reenabled	Duration in seconds the account remains locked
Time between login failures before account lockout	Number of seconds between consecutive failed login attempts

Managing the password policy from the command-line interface

You can customize the password policies by setting rules for the passwords that are used by the Access Appliance local users. You can set rules for password complexity, password age, and password lockout. Password complexity specifies the number and type of characters a password must include. Password age defines the duration for which the password is valid. Password lockout specifies the number of failed attempts because of incorrect usage of passwords after which a user is prevented from logging in to the account.

Commands to view and set the password policy

To view the password policy, use the following command:

```
system password-policy get
```

To set the password policy, use the following command:

```
system password-policy set minlen ucredit maxclassrepeat dcredit
ocredit minclass lcredit maxrepeatdifok pass_min_days pass_max_days
pass_warn_age remember deny unlock_time fail_interval
```

where

Table 15-2

Parameter	Description
minlen	Minimum characters. Range is 6 - 100.
ucredit	Minimum upper case characters. Range is 1 - 100 .
maxclassrepeat	Maximum repetitive characters of same class. Range is 1 - 100.
dcredit	Minimum numbers. Range is 1 - 100.
ocredit	Minimum special characters. Range is 1 - 100.
minclass	Minimum character classes. Range is 1 - 4.
lcredit	Minimum lower case characters. Range is 1 - 100.
maxrepeat	Maximum repetitive characters. Range is 1 - 100.
difok	Character difference with old password. Range is 1 - 100.
pass_min_days	Days after which password can be changed. Range is 1 - 100.
pass_max_days	Days after which password must be changed. Range is 1 - 100.
pass_warn_age	Days before warning message Range is 1 - 100.
remember	Minimum different password before allowing reuse. Range is 1 - 100.
deny	Number of incorrect login attempts before lockout. Range is 1 - 100.
unlock_time	Time before locked account is reenabled(seconds). Range is 1 - 604800.
fail_interval	Time before login failures before account locked out (seconds). Range is 1 - 3600.

To display the current password policy:

- 1** Use SSH to log in to one of the nodes with admin credentials.
The Veritas Appliance Shell is displayed.
- 2** Enter `support elevate` and enter the maintenance account password.
- 3** To go to the Access command-line interface, enter `LOGNAME=admin`
`/opt/VRTSnas/clish/bin/clish -u admin`
- 4** To view the password policy, use the following command:

```
access-clus> system password-policy get
Password policy setup on the system...
Password complexity:
=====
Minimum characters: 8
Minimum upper case characters: 1
Maximum repetitive characters of the same class: -
Minimum numbers: 1
Minimum special characters: 1
Minimum character classes: -
Minimum lower case characters: 1
Maximum repetitive characters: -
Character difference with old password: -
Password age:
=====
Days after which password can be changed: -
Days after which password must be changed: -
Days before warning message: -
Minimum different password before allowing reuse: -
Password lockout:
=====
Number of incorrect login attempts before lockout: -
Time before locked account is reenabled(seconds): -
Time before login failures before account locked out(seconds): -
```

Note: Initially, the default set rules are displayed.

To set the password policy:

- 1** Use SSH to log in to one of the nodes with admin credentials.
The Veritas Appliance Shell is displayed.
- 2** Enter `support elevate` and enter the maintenance account password.

- 3** To go to the Access command-line interface, enter `LOGNAME=admin`
`/opt/VRTSnas/clish/bin/clish -u admin`

- 4 To set the password policy rules, enter the `system password-policy set` command:

For example:

```
access-clus> system password-policy set minlen=8 ucredit=1
maxclassrepeat=4 dcredit=1 ocredit=1 minclass=4 lcredit=1
maxrepeat=2 difok=8 pass_min_days=1 pass_max_days=60
pass_warn_age=7 remember=7 deny=3 unlock_time=300
fail_interval=900
Access Appliance password-policy SUCCESS V-493-10-0 Password
policy updated successfully.
```

The newly set policy can be displayed using the `system password-policy get` command:

```
access-clus> system password-policy get
Password policy setup on the system...
Password complexity:
=====
Minimum characters:                        8
Minimum upper case characters:             1
Maximum repetitive characters of the same class: 4
Minimum numbers:                          1
Minimum special characters:                1
Minimum character classes:                 4
Minimum lower case characters:              1
Maximum repetitive characters:              2
Character difference with old password:     8
Password age:
=====
Days after which password can be changed:  1
Days after which password must be changed: 60
Days before warning message:                7
Minimum different password before allowing reuse: 7
Password lockout:
=====
Number of incorrect login attempts before lockout: 3
Time before locked account is reenabled(seconds): 300
Time before login failures before account locked out(seconds): 900
```

Note: If STIG is enabled on the system, you cannot change the custom password-policy rules.

```
accessclus> system password-rules set maxrepeat=3
maxclassrepeat=vxdefault dcredit=vxdefault minlen=15
ucredit=vxdefault ocredit=vxdefault lcredit=vxdefault
difok=vxdefault minclass=5 pass_min_days=vxdefault
pass_max_days=vxdefault pass_warn_age=vxdefault deny=3
unlock_time=vxdefault fail_interval=vxdefault remember=vxdefault
ACCESS PasswordRules ERROR V-493-10-0 The password rules cannot
be set as the cluster is STIG enabled.
```

Note: Setting the parameter to **vxdefault** is equivalent to setting the value to **no** or **None**.

Immutability in Access Appliance

This chapter includes the following topics:

- [Support for immutability in Access Appliance](#)
- [About lockdown modes](#)
- [Selecting or changing the lockdown mode](#)
- [Accessing the root shell in lockdown mode](#)
- [Configuring immutability using GUI](#)

Support for immutability in Access Appliance

Immutability support for backup images requires locking down the appliance and not permitting any operations that can lead to data destruction. When the appliance is placed in lockdown mode, administrators are prevented from making any changes to the operating system and the internal components.

Important features:

- Immutable data support with retention locking
- Retention lock deletion for backup images
- Access to Remote Management Platform (HP ILO)
- Transition between different modes
- Retention lock extension

About lockdown modes

Lockdown mode protects your cluster data from internal and external threats by securing all the external endpoints from unauthorized access. Access to all the services is protected and authenticated.

Access Appliance lockdown mode offers additional security levels to protect your appliance and data, in addition to the hardened, secure operating environment that comes out of the box.

Lockdown mode provides the following benefits:

- It prevents unauthorized access or modification to the underlying operating system (OS). Once the lockdown mode is enabled, administrators cannot make changes to the OS or the internal components. If you need access to the OS for emergency operations, you must contact Veritas Technical Support to obtain a One-Time Password and temporarily unlock the appliance. This functionality prevents unauthorized changes even if a malicious actor gained access to stolen credentials.
- It gives the appliance users options for managing WORM (Write Once Read Many) data. Your data is protected from being encrypted, modified, and deleted using WORM properties.

Different lockdown modes provide different level of granularity for WORM and retention. The Access Appliance support three lockdown modes.

- **Normal mode:**
This is the default mode of the cluster if the lockdown mode is not specified during installation. In this mode, WORM and retention capabilities are disabled. User cannot create WORM STUs and WORM-enabled files/objects in this mode.
- **Enterprise mode:**
In this mode, WORM and data retention features are enabled. User can choose to create WORM enabled STUs, files and objects. Also, in this mode user has the option to remove the retention locks and expire image data. The user can extend the retention period but cannot reduce the retention period.
The retention time period can be extended from the WORM enabled STUs and files/objects within them only if the user has the Appliance administrator role.
- **Compliance mode:**
In this mode, WORM and data retention features are enabled. The user can extend the retention period. The user does not have the option to remove retention locks and expire image data from underlying files/objects of WORM STUs and backup images before the predefined time. Once the appliance lockdown mode is set to compliance, the user does not have the option to delete data until it is expired.

Veritas strongly recommends that you enable enterprise lockdown mode to prevent unauthorized access to the OS, even if you do not plan to create WORM storage instances.

Selecting or changing the lockdown mode

The user can select the lockdown mode during initial configuration. After cluster configuration, user has the option to see/change the lockdown mode using the GUI as well as CLISH. The user can switch between the following modes without any restriction:

- From Normal to Enterprise mode
- From Normal to Compliance mode
- From Enterprise to Compliance mode

You can change the mode from Enterprise to Normal, from Compliance to Normal or from Compliance to Enterprise only if:

- Locked data is not present in deduplication storage or deduplication is not configured in WORM mode.
- WORM enabled file system for any other use cases, such as NFS and CIFS are not present.
- WORM policies are not activated in the GUI.
- All the file systems are not in offline state.

The user can set minimum and maximum retention time for backup images in Enterprise and Compliance mode only. The retention period range is between 1 hour and 60 years. The retention period can be in second(s) or hour(s) if you use CLISH. The retention period can be in hour(s), day(s), month(s), or year(s) if you use the GUI. Creation of images with retention time less than the minimum retention time or greater than the maximum retention time is not allowed. This minimum and maximum retention time should be set by the appliance administrator as per the retention requirement of their use case.

If Enterprise or Compliance mode has been configured, retention values can be set on files and objects within the range of the minimum retention period and maximum retention period of the WORM-enabled shares or S3 buckets in which they are present.

- Once the lockdown mode is set, only Appliance administrators can change the lockdown mode.
- The lockdown modes are maintained during upgrade.

- Only the Appliance administrator can remove the retention locks if the lockdown mode is enterprise.
- The user cannot change the mode if any existing operation is in progress.

To change the lockdown mode using the GUI

- 1 Go to **Settings > Security management > Immutability** and click **Lockdown mode**.
- 2 On the Lockdown mode page, click **Edit**.
- 3 Select the mode that you want to enable and click **Save**.

You can also modify the lockdown mode using the `cluster lockdown-mode` commands from CLISH.

```
cluster> lockdown-mode set <mode> [minret] [maxret]
```

Where

mode	Specifies the lockdown mode [normal compliance enterprise]
minret	Specifies the minimum retention value range
maxret	Specifies the maximum retention value range

You can also list the lockdown configuration of a cluster using the `cluster lockdown-mode get` command.

Restrictions in different modes

- If lockdown mode is set to compliance or enterprise for any node, it is not available for factory reset.
- During add and replace node operations, the new node is automatically placed in the existing lockdown mode of the cluster. The lockdown mode of the replaced node is set to normal and the node is available for factory reset.
- Cluster maintenance shell is enabled with two-factor authentication (2FA).
- The lockdown mode settings are done at a cluster level and are applicable for all the services, such as NFS and CIFS that are configured on that cluster.

Accessing the root shell in lockdown mode

If your appliance is in lockdown mode and you need assistance from Veritas Support, you need to generate a One-Time Password (OTP) to allow your representative to access the root shell. The OTP has a two-hour expiration period, after which you

need to generate the OTP again; so make sure that you have opened a support case with Veritas Support and your support representative is ready before you generate the OTP.

To access the root shell when the appliance is in lockdown mode:

- 1 From any one of the nodes in the cluster, log in to Veritas Appliance Shell menu.
- 2 Run the `support generate-otp` command to generate the OTP. A 10-digit number, which is the OTP, is displayed as shown in the following example:

```
[access 8.0] appnode-01 > support generate-otp
One-time password: 3279459335
Operation completed successfully
```


- 3 Mention the OTP in the support case. If you forget the OTP, you can use the `support show-otp` command to view it again. Your support representative uses this OTP to generate a security key, which you later need to specify when prompted.
- 4 When your representative asks you to, enter the `support unlock` command. You are prompted for the security key, which your representative must generate using the OTP. Enter the security key and press Enter.
- 5 To get root access for the current node, run the `support elevate` command. The other node remains locked. Enter the passphrase, which your representative specified while generating the security key. Press Enter.
- 6 Enter the maintenance password to access the root shell.
- 7 When your support representative is done troubleshooting the issue, enter the `support lock` command to close root access to the node. Root access to the node is closed automatically after 12 hours and all root sessions are terminated.

Configuring immutability using GUI

You can configure immutability using the Access Appliance GUI.


To configure immutability using GUI

- 1 Go to **Settings > Security management > Immutability**. Click **Edit**.



- 2 Choose any of the lockdown modes under **Mode Selection**. You can choose from Normal, Enterprise and Compliance. The supported features for each mode are listed.

- 3
- If you choose Enterprise or Compliance mode, you also have to configure the retention time range from which the appliance administrators will be allowed to select when creating backup images. These values will also be used as default values while provisioning storage for different use-cases (NFS, CIFS) from GUI. Specify the maximum and minimum retention time. These modes also lock down the system and disable immutable data destruction operations. Click **Save**.



Note: Ensure that the image retention period in backup policies is within the lockdown mode retention time to avoid any error during backup.

Monitoring and troubleshooting

- [Chapter 17. Configuring event notifications and audit logs](#)
- [Chapter 18. Appliance log files](#)

Configuring event notifications and audit logs

This chapter includes the following topics:

- [About troubleshooting](#)
- [Monitoring command activity](#)
- [Monitoring alerts](#)
- [About alert management](#)
- [Monitoring events](#)
- [Viewing reports](#)
- [Viewing cluster storage usage](#)
- [Viewing file system usage](#)
- [About event notifications](#)
- [About severity levels and filters](#)
- [About SNMP notifications](#)
- [Configuring a syslog server](#)
- [Displaying events on the console](#)

About troubleshooting

Troubleshooting procedures for Access Appliance include the following types of procedures:

- Alert and log message review
- Routine maintenance tasks
- Commonly used recovery procedures
- Feature-specific problems and resolutions

Each of these procedures are described in the remaining chapters of this guide.

Some of the troubleshooting procedures in this guide require that you log on as the `maintenance user`.

Monitoring command activity

The **Recent activity** icon shows the status of current tasks and tasks that recently completed. For more details, click **View all activities**.

The **View all activities** page displays the following information for each task:

- Name
- Status
- Starting time
- Ending time

Select a task to view the details of the command that was executed and the command output.

Monitoring alerts

Access Appliance monitors the status and health of various network and storage components, and generates alerts to notify the administrator.

The **Alerts** icon shows the status of current events and events that recently completed. For more details, select **View all alerts**.

The following details are displayed:

- Alert Name
- Severity
- Object Type
- Received Time

About alert management

Access Appliance displays alerts or sends notifications for the current problems or critical conditions that take place in the system.

You can use **Alert management** option to view, enable, or disable the alerts as required.

The alerts are displayed or shared in the following ways:

- Email
- SNMP

See [“Managing alerts”](#) on page 213.

See [“Viewing information about alert management”](#) on page 213.

Viewing information about alert management

From the **Settings > Alert management** page, you can view all the alerts and manage them. You can view the following details for the alerts:

- Name
- Severity type
- Object type
- Date on which the alert is generated

See [“Managing alerts”](#) on page 213.

See [“About alert management”](#) on page 213.

Managing alerts

You can disable or enable the alerts.

To manage the alerts

- 1 Go to **Settings > Alert management**, and then click **Manage alerts**.
The **Manage alerts** dialog box is displayed.
- 2 Select the check box for the alert that you want to disable, and then click **Next**.
The page displays that the task is initiated. Click **Finish**.
- 3 View the **Recent Activity** panel for the status of the task.

Note: To enable the alerts, you can clear the check boxes for the alerts.

See [“About alert management”](#) on page 213.

See [“Viewing information about alert management”](#) on page 213.

Monitoring events

The **Events** icon shows the status of current events and events that recently completed.

For more details, click **View all events**.

- Date and Time
- Severity
- Object Type
- Source
- Message

Select an event icon (**All**, **Error**, **Critical**, **Warning**, **Information**) to view the details of the event that was executed and the accompanying message.

Viewing reports

Access Appliance provides the ability to report on usage trends over a designated time period. You can export a report as a file in .CSV format.

The following reports are available:

- The Cluster storage usage report shows the storage usage trends for cloud storage and on-premises storage.
See [“Viewing cluster storage usage”](#) on page 214.
- The File System Usage report compares the storage usage patterns for up to 5 file systems.
See [“Viewing file system usage”](#) on page 215.

Viewing cluster storage usage

These reports show usage trends for on-premises and cloud storage. Depending on how the file system is configured, it may use only storage from on-premises devices or on-premises devices and cloud storage.

You can export all of these graphs in a comma-separated values (CSV) file or in a spreadsheet. The Access Appliance GUI gathers daily, weekly, and monthly data

for all the graphs. It also samples data every five minutes and the sampling rate is not configurable.

Viewing file system usage

These reports show usage trends for all the file systems configured in the cluster. It shows only the usage pertaining to on-premises devices.

You can simultaneously view up to five graphs.

There is a subtle difference between storage pool usage and file system usage. If there is a storage pool of size 1 TB and a file system of 100 GB is created, storage pool usage is 10%. If only 5% of a 100-GB file system is in use, then file system usage is 5%.

You can export all of these graphs in a comma-separated values (CSV) file or in a spreadsheet.

Access Appliance gathers daily, weekly, and monthly data for all the graphs. Data is sampled every five minutes and you cannot configure the sample intervals.

About event notifications

Access Appliance monitors the status and health of various network and storage components, and generates events to notify the administrator. Access Appliance provides a mechanism to send these events to external event monitoring applications such as, Veritas NetInsights console, syslog server, SNMP trap logger, and mail servers. This section explains how to configure Access Appliance so that external event monitoring applications are notified of events on the Access Appliance cluster.

About severity levels and filters

Access Appliance monitors events of different severity levels. Set the severity to a particular level to specify the severity level to include in notifications. Notifications are sent for events having the same or higher severity.

[Table 17-1](#) describes the valid Access Appliance severity levels in descending order of severity.

Table 17-1 Severity levels

Valid value	Description
emerg	Indicates that the system is unusable

Table 17-1 Severity levels (*continued*)

Valid value	Description
alert	Indicates that immediate action is required
crit	Indicates a critical condition
err	Indicates an error condition
warning	Indicates a warning condition
notice	Indicates a normal but a significant condition
info	Indicates an informational message
debug	Indicates a debugging message

Access Appliance also classifies event notifications by type. Set the event filter to specify which type of events to include in notifications. Notifications are sent only for events matching the given filter.

The filter is set to one of the following options:

- Admin
- Backup
- CIFS
- Cluster
- Network
- NFS
- Replication
- Report
- Storage
- Support
- System
- Upgrade
- All - resets the filter to show all events.

For example, if the filter is set to `network`, a network event triggers a notification. A storage-related event would not trigger a notification.

About SNMP notifications

Simple Network Management Protocol (SNMP) is a network protocol to simplify the management of remote network-attached devices such as servers and routers. SNMP is an open standard system management interface. Information from the Management Information Base (MIB) can also be exported.

SNMP traps enable the reporting of a serious condition to a management station. The management station is then responsible for initiating further interactions with the managed node to determine the nature and extent of the problem.

See [“About severity levels and filters”](#) on page 215.

Configuring a syslog server

Access Appliance can be configured to send syslog messages to syslog servers based on set severities and filters.

In Access Appliance, options include specifying the external system log (syslog) server for event reporting, and setting the interval of messages. Event notifications matching configured severity levels and filters are logged to those external syslog servers.

See [“About severity levels and filters”](#) on page 215.

To display the list of syslog servers

- ◆ To display the list of syslog servers, enter the following:

```
Report> syslog show
```

To add a syslog server to receive event notifications

- ◆ To add a syslog server to receive event notifications, enter the following:

```
Report> syslog add syslog-server-ipaddr
```

where *syslog-server-ipaddr* specifies the host name or the IP address of the external syslog server.

To delete a syslog server from receiving message notifications

- ◆ To delete a syslog server from receiving message notifications, enter the following:

```
Report> syslog delete
```

Displaying events on the console

To display events on the console

- ◆ To display events on the console, enter the following:

```
Report> showevents [number_of_events] [severities]
```

where

number_of_events specifies the number of events that you want to display. If you leave *number_of_events* blank, or if you enter **0**, Access Appliance displays all of the events in the system.

severities specifies the type of severity to show [alert/err/warn].

Appliance log files

This chapter includes the following topics:

- [About appliance log files](#)
- [Gathering device logs with the DataCollect command](#)
- [Downloading logs using the Log Transfer Console](#)
- [Forwarding logs to an external server](#)

About appliance log files

Log files help you to identify and resolve any issues that you may encounter with the appliance.

The Veritas Access Appliance captures software-, system-, and performance-related data in log files. This data offers information about things like appliance operation, issues, and processes.

Appliance logs also help with the diagnosis of hardware-related issues, such as unconfigured volumes or arrays, temperature irregularities, and battery malfunction.

[Table 18-1](#) describes the methods you can use to access the appliance log files.

Table 18-1 Methods of accessing appliance logs

Method	Log details
Download logs to a remote computer. Log onto the Access Appliance shell menu and download the appliance unified logs or appliance device logs.	■ Appliance unified logs ■ Appliance storage device logs

Table 18-1 Methods of accessing appliance logs *(continued)*

Method	Log details
You can use the <code>support data-collect</code> commands to collect storage device logs. See “Gathering device logs with the DataCollect command” on page 220.	Appliance storage device logs

Gathering device logs with the DataCollect command

You can use the `support legacy-data-collect` command to gather device logs. You can share these device logs with the Veritas Support team to resolve device-related issues.

Along with the operating system, IPMI, and storage logs, the DataCollect command now collects the following logs as well:

- Access product logs
- Command output logs
- Patch logs
- File System logs
- Test hardware logs
- CPU information
- Disk performance logs
- Memory information
- Hardware information

To gather device logs with the DataCollect command

- 1** Log on to the Access Appliance shell menu of the desired node.

2 Type the `support legacy-data-collect` command to gather storage device logs.

```
Gathering release information
Gathering disk performance logs
Gathering command output logs
Gathering dmidecode logs
Gathering ipmitool sel list logs
Gathering ipmitool sel writeraw logs
Gathering fwtermlog logs
Gathering AdpEventLog logs
Gathering smartctl logs
Gathering ipmiutil command output
Gathering BMC Debug logs
Gathering Seagate storage array logs
Gathering cpu information
Gathering memory information
Gathering os logs
Gathering dfinfo logs
Gathering vxprint logs
Gathering patch logs
Gathering autosupport logs
Gathering sysinfo logs
Gathering sdr logs
Gathering adpallinfo logs
Gathering encinfo logs
Gathering cfgdsply logs
Gathering cfgdsply logs for LTR
Gathering ldpdinfo logs
Gathering pdlist logs
Gathering fru logs
Gathering adpbucmd logs
Gathering sas3ircu logs
Gathering sas3ircu display logs
Gathering adpalilog logs
Gathering Test Hardware logs
Gathering Access product logs
```

All logs have been collected in `/log/DataCollect.zip`

Log file can be collected from the appliance shared folder - `\\logs\APPL`

Share can be opened using 'support collect-logs share open'

The Log files `DataCollect-XXXX9000918-20180201012500.tar.gz` collected from

appliance are sent to AutoSupport servers.

The data transmission may complete in several minutes or longer due to file size, network status, appliance performance and other reasons.

=====End of DataCollect=====

The appliance generates the device log in the `/log/DataCollect.zip` file.

- 3 Copy the `/log/DataCollect.zip` file to your local folder by using the following command:

```
[access 8.0]access-appliance> support collect-logs share open
- [INFO] Created an NFS share for sharing the logs.
- [INFO] You can access the NFS share at access-appliance:/log.
To ensure appliance security close the share when you are done
downloading the necessary logs.
- [INFO] Created a CIFS share for sharing the logs.
- [INFO] You can access the CIFS share at \\access-appliance\logs.
To ensure appliance security close the share when you are done
downloading the necessary logs.
Operation completed successfully
```

- 4 You can send the `DataCollect.zip` file to the Veritas Support team to resolve your issues.

See [“About appliance log files”](#) on page 219.

Collecting default and additional diagnostic logs

You can collect the default or additional diagnostic logs and sanitize sensitive data from the generated log packages using the `support data-collect` command.

The command provides the following options:

- Collect default and additional diagnostic logs
- Delete the log package that were generated using the `data-collect` command.
- List all the generated log packages from the `/log/data-collect` directory.
- Display the status of the ongoing `data-collect` command.
- Upload the log package that were generated using the `data-collect` command.
- Enable, disable, and view the current setting for sanitizing sensitive data.

To collect additional diagnostic logs:

```
support data-collect advanced
```

To view generated log packages in the `/log/data-collect` directory :

```
support data-collect list
```

To upload the generated log package:

```
support data-collect upload file=
```

To mask sensitive data such as personally identifiable information (PII) in the generated data-collect log package:

```
support data-collect sanitize-data enable
```

Downloading logs using the Log Transfer Console

The Log Transfer Console provides a web-based interface to download logs from a node that is a part of a cluster or from a node that is not yet configured in a cluster. The logs are downloaded from the node to your local system using a web service. You can access the Log Transfer Console using the IP address that is assigned to the eth1 network interface of the node. The Log Transfer Console is compatible with Google Chrome browser version 57.0.2987 or later and Mozilla Firefox browser version 52.0.0 or later.

You can download a single log file or multiple log files at a time. To download multiple log files, you need to create a log package, which is a compressed tar file, and then download the log package to your system.

To collect logs from a node:

- 1 Use the admin user credentials to log in to the node using the Access Appliance Shell Menu.
- 2 To start the log transfer service, enter the following command:

```
support logbrowser start
```

If the service starts successfully, the URL to access the Log Transfer Console web interface is displayed. The URL is similar to <https://eth1-IPaddress:8514/appliance/logs/> and uses the IP address that is assigned to the eth1 interface.

- 3 Open a supported browser and use the displayed URL to access the Log Transfer Console. If the node is not a part of a cluster, sign in using the admin user account, which is the local administrator account. If the node is a part of a cluster, sign in using the admin user account or an account that is authenticated with LDAP or Active Directory.
- 4 Click the **Select files to download** tab.
All the log files and directories in the `/log` directory of the node are displayed.

- 5 To download a single log file, select the file, and from the menu options click **Download**.

The file is downloaded to your system.

- 6 To download multiple log files and directories, complete the following steps:
 - Select the files that you want download and click **Include in the log package**. To add a single file at a time, select the file and click **Include**.
 - Click the **View selected files** tab.

The files that you selected earlier are displayed. Optionally, to exclude a selected file from the log package, select the file and from the menu options click **Exclude**. To exclude multiple files, select the files and click **Exclude selected files**.
 - Review the list of files and click **Generate and download log package**.

A compressed tar file is created and downloaded to your system.

Forwarding logs to an external server

Starting with version 8.0, you can forward system logs to an external log management server. System logs (syslog) contain event and notification messages in a specific format. Forwarding the appliance syslogs to an external log management server provides system administrators a centralized location for viewing logs and for further analysis and troubleshooting. The following log servers are supported:

- HP ArcSight
Splunk

Access appliances use the Rsyslog client to forward logs. In addition to HP ArcSight and Splunk, other log management servers that support the Rsyslog client can also be used to receive syslogs from the appliance.

To secure the log transmission from the appliance to the log management server, you can use the TLS (Transport Layer Security) option. The Access Appliance currently supports only TLS Anonymous Authentication for log forwarding.

Setting up log forwarding

You must use the Veritas Access command-line interface to set up log forwarding. Use the `Report> syslog` commands on the Access command-line interface to set up the log forwarding. For more details, see the *Veritas Access Command Reference Guide*.

Viewing the configured settings

You can view the configured settings for an appliance node by using the `show log-forwarding` command. The following details are displayed:

- IP address of the log management server.
- Port number of the log management server.
- Protocol used for forwarding the logs to the log management server.
- Time interval in minutes for forwarding logs. The options are 0, 15, 30, 45, or 60. The default is 15. If the interval is set to 0, appliance continuously forwards syslogs to the log management server.
- Whether TLS is enabled for secure log transmission.

Provisioning and managing Access Appliance file systems

- [Chapter 19. Creating and maintaining file systems](#)

Creating and maintaining file systems

This chapter includes the following topics:

- [About creating and maintaining file systems](#)
- [About encryption at rest](#)
- [Considerations for creating a file system](#)
- [Creating a file system](#)
- [Bringing the file system online or offline](#)
- [Listing all file systems and associated information](#)
- [Modifying a file system](#)
- [Managing a file system](#)
- [Destroying a file system](#)
- [Upgrading disk layout versions](#)

About creating and maintaining file systems

An Access Appliance environment consists of multiple nodes that can access and update files in the same Veritas file system at the same time. Many file systems can be supported at the same time. You create file systems on groups of disks called storage pools.

File systems consist of both metadata and file system data. Metadata contains information such as the last modification date, creation time, permissions, and so on. The total amount of the space that is required for the metadata depends on the

number of files in the file system. A file system with many small files requires more space to store metadata. A file system with fewer larger files requires less space for handling the metadata.

When you create a file system, you need to set aside some space for handling the metadata. The space that is required is generally proportional to the size of the file system. For this reason, after you create the file system, a small portion of the space appears to be used. The space that is set aside to handle metadata may increase or decrease as needed. For example, a file system on a 1-GB volume takes approximately 35 MB (about 3%) initially to store metadata. In contrast, a file system of 10 MB requires approximately 3.3 MB (30%) initially for storing the metadata.

Veritas recommends that you create a maximum of 50 file systems in a cluster.

File systems can be increased or decreased in size. SmartTier functionality is also provided at the file system level.

For a newly created file, the file system name can have a maximum of 25 characters. If you create a space-optimized or full-sized rollback for a specified file system, then the file system name can have a maximum of 19 characters because additional strings are added to its volume name and a volume name can have a maximum of 31 characters.

About encryption at rest

Access Appliance provides advanced security for data at rest by the encryption of data volumes. Encryption is a technology that converts data or information into code that can be decrypted only by authorized users.

You can encrypt Access Appliance data volumes to:

- Protect sensitive data from unauthorized access.
- Retire disks from use or ship them for replacement without the overhead of secure wiping of content.

Encryption is implemented using the Advanced Encryption Standard (AES) cryptographic algorithm with 256-bit key size validated by the Federal Information Processing Standard (FIPS) Publication 140-2, (FIPS PUB 140-2) security standard.

When you create file systems in Access Appliance on encrypted volumes using this feature, Access Appliance generates a volume encryption key at the time of file system creation. This encryption key is encrypted (wrapped) using a different key that is retrieved from a Key Management Server (KMS). The wrapped key is stored with the volume record. The volume encryption key is not stored on disk.

Access Appliance supports the use of a KMS that conforms to the OASIS Key Management Interoperability Protocol (KMIP) version 1.1 specification.

During creation of encrypted volumes:

1. Access Appliance sends a key generation request to the configured KMS using the Key Management Interoperability Protocol (KMIP) protocol.
2. The KMS responds with a unique identifier. Access Appliance sends the identifier to the KMS to obtain the key that is generated by the KMS.
3. The KMS responds with the key. Access Appliance generates the random volume encryption key, and encrypts it using the key that is provided by the KMS.
4. Access Appliance stores the encrypted key and the KMS identifier in the volume record.

During startup of encrypted volumes:

1. Access Appliance retrieves the encrypted key and the KMS identifier from the volume record.
2. Access Appliance sends the identifier to the KMS to obtain the key.
3. The KMS responds with the key. Access Appliance decrypts the encrypted key (stored in the volume record) with the key provided by the KMS.

Note: Veritas recommends that you use CPUs designed to support Advanced Encryption Standard Instruction Set (or the Intel Advanced Encryption Standard New Instructions (AES-NI) to improve performance.

Veritas recommends that you use IBM Secure Key Lifecycle Manager (SKLM), which supports KMIP protocol version 1.1, as a KMS server for this feature.

To register a Access Appliance cluster with the IBM SKLM KMS server

1. Install the IBM SKLM server on any system in your environment. You can visit this [URL](#) to find the supported IBM SKLM servers with Access Appliance. Obtain the KMS server's public certificate in base64 format using its admin GUI console or the CLI.
2. In the Access Appliance GUI management console, go to **Settings > Services Management** to register the Access Appliance cluster with the KMS server.
3. Ensure that the time on the Access Appliance server and IBM SKLM server are in sync.
4. Select **Provide Key & Certificates** to generate self-sign certificates for the Access Appliance cluster. Provide the KMS server's public SSL certificate in the same window.

- 5 **Configure KMS Server** gets activated now. Select this tab to enter the KMS server-related details.
- 6 Use the IBM SKLM server's GUI-based management to accept the client request from the Access Appliance cluster and to accept its SSL keys.

You can use the `Storage> fs create` command to create the file system with the `encrypt=on` option.

```
storage> fs create mirrored fs2 lg 2 pool1 protection=disk blksize=8192
pdir_enable=no encrypt=on
```

You can use the storage encryption feature in the GUI by activating the secure data storage policy. You can add new NFS and CIFS shares using the activated policy.

Note: Use the `encrypt=on` option for all the file systems.

Considerations for creating a file system

The following sections describe the considerations and best practices for creating file systems.

Best practices for creating file systems

The following are the best practices for creating file systems:

- Ensure all the disks (LUNs) in each storage pool have an identical hardware configuration.
 Best performance results from a striped file system that spans similar disks. The more closely you match the disks by speed, capacity, and interface type, the better the performance you can expect. When striping across several disks of varying speeds, performance is no faster than that of the slowest disk.
- Create striped file systems rather than simple file systems when creating your file systems.
 See [“About striping file systems”](#) on page 235.
- In a given storage pool, create all the file systems with the same number of columns.
- Ensure that the number of disks in each storage pool is an exact multiple of the number of columns used by the file systems created in that storage pool.
- Consider how many disks you need to add to your storage pool to grow your striped file systems.

A 5-TB file system using five columns cannot be grown in a storage pool containing 8*1-TB disks, despite having 3 TB of disk space available. Instead create the file system with either four or eight columns, or else add 2*1-TB disks to the pool. See further examples in the table.

Use case	Action	Result
storage pool with eight disks of the same size (1 TB each)	Create a 5 TB striped file system with five columns.	You cannot grow the file system greater than 5 TB, even though there are three unused disks.
storage pool with eight disks of the same size (1 TB each)	Create a 5 TB striped file system with eight columns.	You can grow the file system to 8 TB.
storage pool with eight disks of the same size (1 TB each)	Create a 4 TB striped file system with four columns.	You can grow the file system to 8 TB.
storage pool with eight disks of the same size (1 TB each)	Create a 3 TB striped file system with three columns.	You cannot grow the file system to 8 TB.
storage pool with eight disks of the different sizes (3 are 500 GB each, and 5 are 2 TB each)	Create an 8 TB striped file system with eight columns.	You cannot create this 8-TB file system.

- Consider the I/O bandwidth requirement when determining how many columns you require in your striped file system.
Based on the disks you have chosen, I/O throughput is limited and potentially restricted. [Figure 19-1](#) describes the LUN throughput restrictions.
- Consider populating each storage pool with the same number of disks from each HBA. Alternatively, consider how much of the total I/O bandwidth that the disks in the storage pool can use.
If you have more than one card or bus to which you can connect disks, distribute the disks as evenly as possible among them. That is, each card or bus must have the same number of disks attached to it. You can achieve the best I/O

performance when you use more than one card or bus and interleave the stripes across them.

- Use a stripe unit size larger than 64 KB. Performance tests show 512 KB as the optimal size for sequential I/O, which is the default value for the stripe unit. A greater stripe unit is unlikely to provide any additional benefit.
- Do not change the operating system default maximum I/O size of 512 KB.
- Veritas recommends that you do not create a file system whose name format is such as *<file system name_integer>*. This is because such file names are reserved for internal objects and may lead to file system creation errors.
- It is recommended that you keep the size of the file system such that there are a maximum of 100 million files on the file system.

Figure 19-1 LUN throughput - details on the LUN throughput restrictions



Reading from a single LUN, the performance team achieved **170 MB/sec.**



Reading from two LUNs, the performance team achieved **219 MB/sec** because of the contention at the modular array level.



Reading from two LUNs, where each LUN is located in a different modular storage array, the performance is **342 MB/sec.**

Each modular array controller cache is shared when I/O is generated to multiple LUNs within the same array.

Choosing a file system layout type

Access Appliance allows you to create file systems with several layout types. The following table describes the layout types and their advantages.

Table 19-1 Types of volume layout

Layout type	Description
Simple	Arranges the disks sequentially and contiguously. A simple layout allows a file system to be created from multiple regions of one or more disks if there is not enough space on a single region of a disk.
Striped	Spreads the data evenly across multiple disks. Stripes are equal-sized fragments that are allocated alternately and evenly to the disks. Throughput increases with the number of disks across which a file system is striped. Striping helps to balance I/O load in cases where high traffic areas exist on certain disks.
Mirrored	Mirrors the information contained in the file system to provide redundancy of data. For the redundancy to be useful, each mirror should contain disk space from different disks.
Mirrored-stripe	Configures a striped file system and then mirrors it. This requires at least two disks for striping and one or more other disks for mirroring (depending on whether the mirror is simple or striped). The advantages of this layout are increased performance by spreading data across multiple disks and redundancy of data.
Striped-mirror	Configures several mirrors as the columns of a striped file system. This layout offers the same benefits as a mirrored-stripe file system. In addition, it provides faster recovery as the failure of single disk does not force an entire striped mirror offline.

Determining the initial extent size for a file system

Veritas File System (VxFS) determines the size of the first extent that is allocated based on the first write to a new file. Normally, the first extent is the smallest power of 2 that is larger than the size of the first write. If that power of 2 is less than 8 KB (the default file system block size), the first extent that is allocated is 8 KB. After the initial extent is allocated, the file system increases the size of subsequent extents with each allocation as the file size is increased using extending writes.

The initial extent size is tunable, and can be changed using the `System> option modify tuneftab` command.

Increasing the initial extent size to a larger value helps to reduce file system fragmentation and improves I/O performance.

The best value for the initial extent size depends on the expected file sizes that are created by the application. The maximum value is 32768, which equates to a 256 MB extent allocation using the default 8 KB file system block size. Any over allocation of space is returned to the free space pool after the file is closed.

If the application creates a lot of small files with an exact size of 1 MB, then the initial extent size can be set to 128 (1 MB). If 1 MB is an approximate file size, then the initial extent size can be set to 64 (512 KB) instead. If most files are approximately 1 GB or greater in size, then the maximum value of 32768 can be used.

About striping file systems

You can obtain huge performance benefits by striping (RAID-0) using software-defined storage (SDS). You achieve performance benefits regardless of the choice of LUN configuration in your storage hardware. Striping is useful if you need large amounts of data that is written to or read from physical disks, and consistent performance is important. SDS striping is a good practice for all Access Appliance use cases and workloads.

Veritas strongly recommends that you create striped file systems when creating your file system for the following reasons:


- Maximize the I/O performance.
- Proportion the I/O bandwidth available from the storage layer.
- Balance the I/O load evenly across multi-user applications running on multiple nodes in the cluster.

However there are also pitfalls to avoid.

The following information is essential before selecting the disks to include in your striped file system:

- Understanding of your hardware environment
- Storage capabilities and limitations (bottlenecks)
- Choice of LUNs (each LUN, or disk, equates to a column in a SDS-striped volume)


Figure 19-2 An example hardware configuration



An extreme example might be if one column (equal to one LUN) is composed of only hard disk drives (HDDs) in the storage array. All of the other columns in the same striped volume are composed of only SSDs in the storage array. The overall I/O performance bottlenecks on the single slower HDD LUN.


Understanding the LUN configuration and ensuring that all of the LUNs have an identical configuration is therefore essential for maximizing performance and achieving balanced I/O across all the LUNs.

Figure 19-3 LUN configuration



All 24 LUNs have an identical hardware configuration.

Figure 19-4 Volume configuration



The performance team created a volume with 24 columns to use the entire storage bandwidth in one file system. Veritas does not advise changing the operating system default maximum I/O size of 512 KB. The optimum stripe-unit size is 512 KB.

About FastResync

The FastResync feature performs quick and efficient resynchronization of stale mirrors (a mirror that is not synchronized). FastResync optimizes mirror resynchronization by keeping track of updates to stored data that have been missed by a mirror.

When FastResync has been enabled, it does not alter how you administer mirrors. The only visible effect is that repair operations conclude more quickly.

About fsck operation

The `fsck` operation does the following:

- Checks the consistency of the metadata container, the data container, and the database, and repairs any inconsistencies.
- Checks if the metadata container and the data container are marked for full fsck. If yes, the `fsck` operation performs a full file system check of the corresponding file systems. Based on the actions taken by fsck on the individual file systems, the `fsck` operation repairs the inconsistencies in other parts of the file system.
- Goes through all the file handles present in the database, and checks if the corresponding metadata container and the data container file handles are consistent with each other.

In some cases, full fsck might delete files from the data container. To maintain consistency, the corresponding files from the metadata container and the data container are removed, and the corresponding key is removed from the database.

```
Storage> fs fsck fs1
```

Enabling WORM on a file system

You can enable WORM on a file system using the `storage fs worm` commands.

```
storage> fs worm set <fs_name> [minret] [maxret]
```

Where

<code>fs_name</code>	Specifies the name of the file system
<code>minret</code>	Specifies the minimum time period of retention. It can contain: [1-9](s S h H d D m M y Y)
<code>maxret</code>	Specifies the maximum time period of retention. It can contain: [1-9](s S h H d D m M y Y)

The minimum and maximum retention period of a file system can be increased or decreased only in the Enterprise and Compliance lockdown mode.

You can also check if a particular file system is WORM-enabled.

```
storage> fs worm get
```

You cannot rename a folder on a WORM-enabled file system.

Note: If episodic replication is configured on a file system, the `storage fs worm set <fs_name>` and `storage fs worm set <fs_name> <min_retention> <max_retention>` commands are successful on the source cluster only if the destination cluster is reachable and vice-versa. After these commands are successfully executed, WORM is enabled and/or retention is set on the remote cluster as well.

Note: The `storage fs worm set <fs_name>` and `storage fs worm set <fs_name> <min_retention> <max_retention>` commands fail to set WORM on the file system if continuous replication is configured on a file system. If you want to set WORM on the file system, you have to disable continuous replication.

You can enable WORM on a file system using the Access Appliance GUI.

Setting retention in files

The retention feature provides a way to ensure that the files are not deleted or modified until the retention period is applied on the files. You can set and show the retention period on files from the Access command-line interface.

The file system should be created with `worm=yes` option to use the retention feature. See the `Storage> fs create` man page for more information.

To set retention:

```
Storage> fs retention set [path] [retention_period]
```

Where *path* is the specified file or directory on which retention is set. If the specified path is a directory, then retention is set on all the files that are currently present in the directory.

retention_period is the duration for which retention is set. It can be in [1-9](s|S|h|H|d|D|m|M|y|Y) or mm-dd-yyyy or mm-dd-yyyy:hh:mm:ss format.

The retention period cannot be reduced once it is set.

The retention period range is between 1 hour and 60 years. The retention period can be in hour(s), day(s), month(s), or year(s) if you use CLISH. The retention period can be in second(s), hour(s), day(s), month(s), or year(s) if you use the GUI. The retention period of a file should be within the minimum and maximum retention period of the WORM-enabled file system, share or bucket within which the file is present.

Note: The retention period in Access Appliance is based on the file system compliance clock and not based on the system clock.

To show retention:

```
Storage> fs retention show [path]
```

Where *path* is the specified file on which retention is set.

In Enterprise lockdown mode, if retention period is given as 0, retention is removed on the specified file. You can also extend the retention period on a file by specifying a value which is greater than the currently set retention period on a file.

See the `Storage> fs` man page for detailed examples.

You can set retention on a file system using the Access Appliance GUI.

Setting WORM over NFS

When a file is committed as Write-Once-Read-Many (WORM), the data in the file can be read but cannot be altered. The retention time for a WORM file specifies the time period for which the file must be retained after it is committed to WORM storage. The file cannot be deleted till the retention period expires. Once the retention time period has expired, the storage system allows the deletion of the file.

The retention period cannot be reduced once it is set.

The file system on the server should be created with `worm=yes` option as the per-file WORM feature is supported only on file systems created with this option.

See the `Storage> fs create` man page for more details.

Export the file system from the server and mount it on the client.

To enable WORM on a file over NFS

- 1 Change the access time of the file so that it has the same value as the period of retention.

```
# touch -at YYYYMMDDhhmm.ss <filename>
```

For example, if a file named `foo` has to be retained till 14th July, 2035 10:37:42pm, run the following command:

```
# touch -at 203507141037.42 foo
```

- 2 Mark the file as read-only by changing the permissions of the file.

For example, to make the file `foo` read-only, run the following command:

```
# chmod -w foo
```

On successful execution of the above two steps, WORM is enabled on the file, `foo` with *14th July, 2035 10:37:42pm* as the retention time.

Manually setting WORM-retention on a file over CIFS

You can set WORM-retention on a file manually over CIFS.

Perform the following steps:

- Set the file's access time (*atime*) to the required retention time
- Set the `read-only` attribute to file

The file is marked WORM with assigned retention after the successful completion of the above steps.

Note: Veritas does not support extending the existing retention period of a WORM file over CIFS by changing the access time of the file as it is not possible to change the access time of a read-only file.

Note: The retention period cannot be reduced once it is set.

To manually set WORM-retention on a file over CIFS from a Windows client

1 Create a WORM-enabled file system.

```
Storage> fs create simple fs_worm_ct 15g pool1 blksize=1024
pdir_enable=no encrypt=off worm=yes
100% [#] Creating simple filesystem
ACCESS fs SUCCESS V-288-0 Created simple file system fs_worm_ct
```

2 Create a CIFS share of the file system in full_acl mode.


```
CIFS> share add fs_worm_ct sh_fs_worm_ct allow=evlab\
vaultadmin,rw,full_acl
Exporting CIFS filesystem : sh_fs_worm_ct ...Success.
```

3 You can see the CIFS shares using the following command:


```
CIFS> share show
```

ShareName	FileSystem	ShareOptions
access_partition	ev_vs_partition	owner=root,group=root,fs_mode=1777,allow=EVLAB\ vaultadmin,rw
access_worm_share	ev_worm_fs	owner=root,group=root,fs_mode=1777,allow=evlab\ vaultadmin,rw,full_acl
sh_fs_worm_ct	fs_worm_ct	owner=root,group=root,fs_mode=1777,allow=evlab\ vaultadmin,rw,full_acl


- 4 Verify the CIFS share from Windows explorer using the virtual IP associated with the share.




5 Create some files on the CIFS share.



6 Map the share to the network drive.



7 Click **Finish**.



The image shows a Windows 'Map Network Drive' dialog box. The title bar is blue with a close button (X) in the top right corner. Below the title bar, there is a question: 'What network folder would you like to map?'. Underneath, it says 'Specify the drive letter for the connection and the folder that you want to connect to:'. There are two input fields: 'Drive:' with a dropdown menu showing 'Z:' and 'Folder:' with a text box containing '\\[redacted]\\sh_fs_worm_ct' and a 'Browse...' button to its right. Below these fields, there is an example: 'Example: \\server\\share'. There are two checkboxes: 'Reconnect at sign-in' and 'Connect using different credentials', both of which are currently unchecked. At the bottom, there is a link: 'Connect to a Web site that you can use to store your documents and pictures.' and two buttons: 'Finish' and 'Cancel'.

Map Network Drive

What network folder would you like to map?

Specify the drive letter for the connection and the folder that you want to connect to:

Drive: Z: ▼

Folder: \\[redacted]\\sh_fs_worm_ct ▼ Browse...

Example: \\server\\share


☐ Reconnect at sign-in

☐ Connect using different credentials

[Connect to a Web site that you can use to store your documents and pictures.](#)

Finish Cancel

- 8 The drive which is associated with the share is Z:.



- 9 Use Microsoft PowerShell to set WORM-retention on a file manually over CIFS.

- Go to Z: drive. Check the current access time of the file.

```
(Get-Item .\testf1.txt).LastAccessTime
```

Thursday, February 1, 2018 12:03:54 AM"

Where *testf1.txt* is the name of the file.

- Set the file's access time (*atime*) to the required retention time using the following command:

```
(Get-Item .\testf1.txt).LastAccessTime="February 10, 2018 12:03:54 AM"
```

- Set the `read-only` attribute to file.

```
Set-ItemProperty -Path "Z:\testf1.txt" -Name IsReadOnly -Value $true
```

```
PS C:\Windows\system32> cd Z:
PS Z:\> ls

Directory: Z:\

Mode                LastWriteTime         Length Name
----                -
d-----          2/1/2018 12:01 AM             lost+found
d-----          2/1/2018 12:04 AM              dir1
-a----          2/1/2018 12:05 AM           174 testf1.txt
-a----          2/1/2018 12:05 AM           148 testf2.txt

PS Z:\> (Get-item testf1.txt).LastAccessTime
Thursday, February 1, 2018 12:03:54 AM

PS Z:\> (Get-item testf1.txt).LastAccessTime="February 10, 2018 12:03:54 AM"
PS Z:\> (Get-item testf1.txt).LastAccessTime
Saturday, February 10, 2018 12:03:54 AM

PS Z:\> Set-ItemProperty -Path .\testf1.txt -Name IsReadOnly -Value $true
```

The file is marked as WORM with retention time of *Feb 10, 2018 12:03:54 AM*.

Note: If you execute a delete operation on a WORM-enabled file from Windows, then it appears that the WORM file is deleted even though the file is not deleted from the storage. This is Windows behavior, and is seen because the file delete command is issued from the Windows client over CIFS which asynchronously marks the file for deletion and returns success to the user. But the asynchronously issued delete command actually fails as the file is marked as WORM and you can see the file by refreshing directory contents.

About managing application I/O workloads using maximum IOPS settings

When multiple applications use a common storage subsystem, it is important to ensure that a particular application does not monopolize the storage bandwidth thereby impacting all the other applications using the same storage. It is also important to balance the application I/O requests in a way that allows all the applications to co-exist in a shared environment. You can address this need by setting a maximum threshold on the I/O operations per second (MAXIOPS) for the file system.

The MAXIOPS limit determines the maximum number of I/Os processed per second collectively by the storage underlying the file system.

When an I/O request comes in from an application, it is serviced by the storage underlying the file system until the application I/O reaches the MAXIOPS limit. When the limit is exceeded for a specified time interval, further I/O requests on the

application are queued. The queued I/Os are taken up on priority in the next time interval along with new I/O requests from the application.

You should consider the following factors when you set the MAXIOPS threshold:

- Storage capacity of the shared subsystem
- Number of active applications
- I/O requirements of the individual applications

Only application-based I/Os can be managed with MAXIOPS.

MAXIOPS addresses the use case environment of multiple applications using a common storage subsystem where an application is throttled because of insufficient storage bandwidth while another less critical application uses more storage bandwidth.

See the `maxiops` man pages for detailed examples.

Setting the maximum IOPS

The maximum IOPS limit determines the maximum number of I/Os processed per second by a file system.

To set the maximum IOPS

1 Prerequisites:

Configure a storage pool.

2 File Systems > Current Performance: At the ellipses, select **Set Max. IOPs**.

You can also edit the maximum IOPS or reset the maximum IOPS.

3 Follow the wizard prompts.

4 View the **Recent Activity** panel for the status of the task.

See [“About storage provisioning and management”](#) on page 80.

Creating a file system

Access Appliance supports the following types of file systems:

- Cluster File System (CFS) - creates a standard CFS file system. The CFS file system is the default file system when creating a new file system.

Access Appliance supports the following file system layouts:

- Simple - creates a file system of a specified size, and you can specify a block size for the file system.

- Mirrored - creates a mirrored file system with a specified number of mirrors along with a list of pools and online status.
- Mirrored-stripe - creates a mirrored-stripe file system with a specified number of columns, mirrors, pools, and protection options.
- Striped - creates a striped file system. A striped file system is a file system that stores its data across multiple disks rather than storing the data on just one disk.
- Striped-mirror - creates a striped-mirror file system with a specified number of columns, mirrors, pools, and protection options.

The default block size is determined based on the size of the file system when the file system is created. For example, 1 KB is the default block size for up to a 2-TB file system size. Other default block sizes, 2 KB, 4 KB, and 8 KB are available for different ranges of file system sizes. If you create a 1-TB file system, and then increase it to 3 TB, the file system block size remains at 1 KB.

See the `storage_fs(1)` man page for detailed examples.

For details on the limitations on the length of the file system name, See [“About creating and maintaining file systems”](#) on page 228.

You can also create a file system for customized workloads.

Note: If the configuration file system creation fails, the originating file system request also fails. Access Appliance requires at least two disks for the mirrored configuration file system, and in case of Flexible Shared Storage (FSS), at least two nodes are required to be part of the storage pool.

Note: Veritas recommends that you do not create a file system whose name format is such as `<file system name_integer>`. This is because such file names are reserved for internal objects and may lead to file system creation errors.

To create a simple file system of a specified size

- ◆ Create a simple file system with a specified size.

```
Storage> fs create simple fs_name size pool1[,disk1,...]  
[blksize] [pdir_enable] [encrypt] [worm] [minret] [maxret]
```

To create a mirrored file system

- ◆ Create a mirrored file system.

```
Storage> fs create mirrored fs_name size nmirrors pool1[,disk1,...]
[protection=disk|pool] [blksize=bytes] [pdir_enable] [encrypt]
[worm] [mediatype] [multi_volume] [minret] [maxret]
```

To create a mirrored-stripe file system

- ◆ Create a mirrored-stripe file system.

```
Storage> fs create mirrored-stripe fs_name size nmirrors
ncolumns pool1[,disk1,...] [protection=disk|pool]
[stripeunit=kilobytes] [blksize=bytes] [pdir_enable] [encrypt]
[worm] [mediatype] [minret] [maxret]
```

To create a striped-mirror file system

- ◆ Create a striped-mirror file system.

```
Storage> fs create striped-mirror fs_name size nmirrors
ncolumns pool1[,disk1,...] [protection=disk|pool]
[stripeunit=kilobytes] [blksize=bytes] [pdir_enable] [encrypt]
[worm] [mediatype] [multi_volume] [minret] [maxret]
```

To create a striped file system

- ◆ Create a striped file system.

```
Storage> fs create striped fs_name size ncolumns pool1[,disk1,...]
[stripeunit=kilobytes] [blksize=bytes] [pdir_enable] [encrypt] [worm]
[minret] [maxret]
```

fs_name

Specifies the name of the file system being created. The file system name should be a string. If you enter a file that already exists, you receive an error message and the file system is not created.

size	Specifies the size of a file system. To create a file system, you need at least 10 MB of space. You can create a file system in the following units: ■ MB ■ GB ■ TB You can enter the units with either uppercase (10 M) or lowercase (10 m) letters. To see how much space is available on a pool, use the <code>Storage> pool free</code> command. See “About configuring storage pools” on page 81.
nmirrors	Specifies the number of mirrors for the file system. You must enter a positive integer.
ncolumns	Specifies the number of columns for the striped file system. The number of columns represents the number of disks to stripe the information across. If the number of columns exceeds the number of disks for the entered pools, an error message is displayed. This message indicates that there is not enough space to create the striped file system.
pool1[,disk1,...]	Specifies the pool(s) or disk(s) for the file system. If you specify a pool or disk that does not exist, you receive an error message. You can specify more than one pool or disk by separating the name with a comma. Do not include a space between the comma and the name. To find a list of pools and disks, use the <code>Storage> pool list</code> command. To find a list of disks, use the <code>Storage> disk list</code> command. The disk must be part of the pool or an error message is displayed.

protection	If you do not specify a protection option, the default is "disk." The available options for this field are: ■ disk - Creates mirrors on separate disks. ■ pool - Creates mirrors in separate pools. If there is not enough space to create the mirrors, an error message is displayed, and the file system is not created. The protection=pool option is not supported for an isolated pool. If the file system is created with the protection=pool option, then both the pools in question cannot be merged to create a single pool, which defeats the purpose of the protection=pool option.
stripeunit=kilobytes	Specifies a stripe unit (in kilobytes). The following are the possible values: ■ 64 ■ 128 ■ 256 ■ 512 (default) ■ 1024 ■ 2048
blksize=bytes	Specifies the block size for the file system. The following are the possible values in bytes: ■ 1024 ■ 2048 ■ 4096 ■ 8192 (default) Block sizes can affect the file size. For example, to create a file system greater than 32 TB, the block size needs to be 8192.
pdir_enable	Specifies if you want to enable a partition directory for the file system. By default, this feature is not enabled.
encrypt	Enable encryption. You can set encrypt=on/encrypt=off.
worm	Enables WORM. The possible values are worm=yes/worm=no.
mediatype	Specifies the disk type for the file system.
multi_volume	Create multiple volumes.

minret	Specifies the minimum file system retention period. It can contain: [1-9](s S h H d D m M y Y).
maxret	Specifies the maximum file system retention period. It can contain: [1-9](s S h H d D m M y Y).

After a file system is created, the file system reserves some space for internal logging. Internal logging provides additional data integrity. Due to the space that is reserved for internal logging, the file system may appear to be used immediately after file system creation. The space that is reserved for internal logging increases with the number of nodes in the Access Appliance cluster.

Log file sizes for the file systems are as follows:

10 GB to 100 GB	Log size = 60 MB per node
100 GB to 1 TB	Log size = 100 MB per node
1 TB and above	Log size = 256 MB per node

Bringing the file system online or offline

The `Storage> fs online` or `Storage> fs offline` command lets you mount (online) or unmount (offline) a file system. You cannot access an offline file system from a client.

To change the status of a file system

- ◆ To change the status of a file system, enter one of the following, depending on which status you use:

```
Storage> fs online fs_name
Storage> fs offline fs_name
```

where *fs_name* specifies the name of the file system that you want to mount (online) or unmount (offline). If you specify a file system that does not exist, an error message is displayed.

Listing all file systems and associated information

To list all file systems and associated information

- ◆ To list all file systems and associated information, enter the following:

```
Storage> fs list [fs_name]
```

where *fs_name* is optional.

If you do not enter a specified file system, a list of file systems is displayed.

Modifying a file system

You can modify a file system in the following ways:

See [“Adding or removing a mirror from a file system”](#) on page 256.

See [“Adding or removing a column from a file system”](#) on page 258.

See [“Increasing the size of a file system”](#) on page 259.

See [“Decreasing the size of a file system”](#) on page 261.

Adding or removing a mirror from a file system

A mirrored file system is one that has copies of itself on other disks or pools.

To add a mirror to a file system

- ◆ To add a mirror to a file system, enter the following:

```
Storage> fs addmirror tier_name fs_name pool1[,disk1,...]  
[protection=disk|pool] [iosize]
```

<i>tier_name</i>	Specifies which tier name. If the specified file system does not exist, an error message is displayed.
------------------	--

<i>fs_name</i>	Specifies which file system to add the mirror. If the specified file system does not exist, an error message is displayed.
----------------	--

pool1[,disk1,...]	Specifies the pool(s) or disk(s) to use for the file system. If the specified pool or disk does not exist, an error message is displayed, and the file system is not created. You can specify more than one pool or disk by separating the name with a comma, but do not include a space between the comma and the name. To find a list of existing pools and disks, use the <code>Storage> pool list</code> command. See “About configuring storage pools” on page 81. To find a list of the existing disks, use the <code>Storage> disk list</code> command. The disk needs to be part of the pool or an error message is displayed.
protection	The default value for the protection field is <code>disk</code>. Available options are: ■ <code>disk</code> - if the protection is set to <code>disk</code>, then mirrors are created on separate disks. This flag only works for file systems of type mirrored, mirrored-striped, and striped-mirror. The disks may or may not be in the same pool. ■ <code>pool</code> - if the protection is set to <code>pool</code>, then mirrors are created in separate pools. This flag only works for file systems of type mirrored, mirrored-striped, and striped-mirror. If not enough space is available, then the file system creation operation fails.
iosize	Size of the IO request.

To remove a mirror from a file system

- ◆ To remove a mirror from a file system, enter the following:

```
Storage> fs rmmirror tier_name fs_name [pool_or_disk_name]
```

fs_name	Specifies the file system from which to remove the mirror. If you specify a file system that does not exist, an error message is displayed.
pool_or_disk_name	Specifies the pool or the disk name to remove from the mirrored file system that spans the specified pools or disks. If a pool name is the same as the disk name, then the mirror present on the pool is deleted.
tier_name	Specifies the tier name.

For a striped-mirror file system, if any of the disks are bad, the `Storage> fs rmmirror` command disables the mirrors on the disks that have failed. If no disks have failed, Access Appliance chooses a mirror to remove.

Adding or removing a column from a file system

You may want to add or remove a column from a file system in specific situations. Adding columns can help to perform more I/Os in parallel, so you may want to increase the number of columns in the file system.

Note: For a striped file system when you add a column, the layout that is displayed when you issue the `Storage> fs list` and `Storage> fs list fsname` commands may be different than the original layout of the file system while the relayout (addition of new columns) operation is in progress. The original file system layout is displayed when the relayout operation is completed.

Note: Adding and removing a column to and from a file system involves a volume-level relayout. This is an I/O intensive operation. For a large file system, adding or removing columns takes a long time and can hurt application performance during this relayout period.

To add a specified number of columns to a file system

- ◆ To add a specified number of columns to a file system, enter the following:

```
Storage> fs addcolumn tier_name fs_name ncolumns pool_or_disk_name
```

fs_name Specifies the file system for which you want to add additional columns.

tier_name Specifies the tier name.

ncolumns Specifies the number of columns that you want to add to the file system.

Note: In the case of a striped file system, the number of the disks that are specified should be equal to the number of columns (*ncolumns*).

Note: In the case of a mirrored-stripe and a striped-mirrored file system, the disks should be equal to (*ncolumns* * *number_of_mirrors_in_fs*).

pool_or_disk_name Specifies the pool or the disk name for the file system.

To remove a column from a file system

- ◆ To remove a column from a file system, enter the following:

```
Storage> fs rmcolumn tier_name fs_name
```

where *fs_name* is the name of the file system for which you want to remove the column and *tier_name* is the name of the tier.

Increasing the size of a file system

To increase (grow) the size of a file system, it must be online. If the file system is not online, an error message is displayed, and no action is taken.

To increase the size of a file system to a specified size

- ◆ To increase the size of a file system to a specified size, enter the following:

```
Storage> fs growto tier_name fs_name new_length
[pool1[,disk1,...]] [protection=disk|pool] [balanced] [num_vols]
```

If no pool is specified with the command, the disks for growing the file system can be taken from any available pool. The protection flag takes the default value of `disk` in this case. The value of the `protection` field cannot be set to `pool` when no pool is specified with the command. This operation may convert the layout of the file system if the command determines that the new file system is too large for the original layout.

If `balanced=yes`, all the volumes of the file system are grown by the same percentage which is calculated based on the specified `new_length`. If `balanced=no`, the file system is grown linearly starting from the last volume.

To increase the size of a file system by a specified size

- ◆ To increase the size of a file system by a specified size, enter the following:

```
Storage> fs growby tier_name fs_name length_change
[pool1[,disk1,...]] [protection=disk|pool] [balanced] [num_vols]
```

If no pool is specified with the command, the disks for growing the file system can be taken from any available pool. The protection flag takes the default value of `disk` in this case. The value of the `protection` field cannot be set to `pool` when no pool is specified with the command. This operation may convert the layout of the file system if the command determines that the new file system is too large for the original layout.

If `balanced=yes`, all the volumes of the file system are grown by the same percentage which is calculated based on the specified `length_change`. If `balanced=no`, the file system is grown linearly starting from the last volume.

fs_name Specifies the file system whose size is increased. If you specify a file system that does not exist, an error message is displayed.

tier_name Specifies the tier name.

new_length Expands the file system to a specified size. The size that you specify must be a positive number, and it must be bigger than the size of the existing file system. If the new file system is not larger than the size of the existing file system, an error message is displayed, and no action is taken.

This variable is used with the `Storage> fs growto` command.

length_change	Expands the file system by a specified size. The size that you specify must be a positive number, and it must be lesser than the available space. If it exceeds the available space, an error message is displayed, and no action is taken. This variable is used with the <code>Storage> fs growby</code> command.
pool1[,disk1,...]	Specifies the pool(s) or disk(s) to use for the file system. If you specify a pool or disk that does not exist, an error message is displayed, and the file system is not resized. You can specify more than one pool or disk by separating the name with a comma; however, do not include a space between the comma and the name. To find a list of existing pools and disks, use the <code>Storage> pool list</code> command. See “About configuring storage pools” on page 81. To find a list of the existing disks, use the <code>Storage> disk list</code> command. The disk needs to be part of the pool or an error message is displayed.
protection	The default value for the protection field is <code>disk</code>. Available options are: ■ <code>disk</code> - if the protection is set to <code>disk</code>, then mirrors are created on separate disks. This flag only works for file systems of type mirrored, mirrored-striped, and striped-mirror. The disks may or may not be in the same pool. ■ <code>pool</code> - if the protection is set to <code>pool</code>, then mirrors are created in separate pools. This flag only works for file systems of type mirrored, mirrored-striped, and striped-mirror. If not enough space is available, then the file system creation operation fails.
num_vols	Number of volumes to be created for <code>growto</code> and <code>growby</code> operations.

Decreasing the size of a file system

You can decrease (shrink) the size of the file system.

To decrease the size of the file system, it must be online. If the file system is not online, an error message is displayed, and no action is taken.

You cannot decrease the size of a file system if a rollback exists. Delete the rollback first before using the `Storage> fs shrinkto` or `Storage> fs shrinkby` commands.

To decrease the size of a file system to a specified size

- ◆ To decrease the size of a file system, enter the following:

```
Storage> fs shrinkto tier_name fs_name new_length [balanced]
```

Note: If you have configured continuous replication, the *storage fs shrinkto* command does not reduce the SRL size even if the file system size is reduced.

To decrease the size of a file system by a specified size

- ◆ To decrease the size of a file system, enter the following:

```
Storage> fs shrinkby tier_name fs_name length_change [balanced]
```

Note: If you have configured continuous replication, the *storage fs shrinkby* command does not reduce the SRL size even if the file system size is reduced.

fs_name	Specifies the file system whose size decreases. If you specify a file system that does not exist, an error message is displayed.
tier_name	Specifies the tier name.
new_length	Specifies the size to decrease the file system to. The size that you specify must be a positive number, and it must be smaller than the size of the existing file system. If the new file system size is not smaller than the size of the existing file system, an error message is displayed, and no action is taken.
length_change	Decreases the file system by a specified size. The size that you specify must be a positive number, and it must be smaller than the size of the existing file system. If the new file system size is not smaller than the size of the existing file system, an error message is displayed, and no action is taken.

Note: Decreasing the size of a file system can take a long time if there are many extents allocated in the shrink area, as these extents have to be relocated to other areas in the file system.

Managing a file system

The following sections detail about the managing a file system.

Defragmenting a file system

You can either defragment a file system now or you can schedule a defragment job for a file system.

To defragment a file system

- ◆ To defragment a file system, enter the following:

```
Storage> fs defrag now fs_name time [defrag_level]
```

<i>fs_name</i>	Specifies the name of the file system that you want to defragment. Note: The specified file system must be online before attempting to defragment the file system.
<i>time</i>	Specifies the maximum time to run. The defragmentation options are processed until defragmentation is complete, or until the time limit expires. The time value should be larger than one minute. Potential time value output and what the values mean: ■ 10M - indicates 10 minutes ■ 1H20M - indicates 1 hour and 20 minutes ■ Infinite - indicates the defragmentation process continues to run until the defragmentation process is done completely. There is no limit time.
<i>defrag_level</i>	Specifies the defragmentation level such as <code>dir</code>, <code>extent</code>, or <code>all</code>.

To schedule a defragment job for a file system.

- 1 Create a defrag schedule job for a file system that reoccurs once a week:

```
Storage> fs defrag schedule create sched_name sched_duration \
minute [hour] [day_of_the_month] \
[month] [day_of_the_week]
```

<i>sched_name</i>	Specifies the the name of the schedule.
<i>sched_duration</i>	Specifies the duration of the defragmenatation job.
<i>minute</i>	Specifies the minute (0-59).
<i>hour</i>	Specifies the hour (0-23).
<i>day_of_the_month</i>	Specifies the day of the month (1-31).
<i>month</i>	Specifies the month of the year (1-12).
<i>day_of_the_week</i>	Specifies the day of the week (0-6 with 0=Sunday).

For example:

Create a defrag schedule called `schedule1` that runs at 11:00 pm every Saturday for a duration of 2 hours.

```
Storage> fs defrag schedule create schedule1 2 0 23 * * 6
```

The number 2 after `schedule1` is the duration of how long the defrag schedule will run. 0 indicates minutes and 23 is the hour for which the defrag schedule will run.

- 2 Show the defrag schedule details:

```
Storage> fs defrag schedule show sched_name
```

For example:

- 3 Start the defrag schedule job for a file system:

```
Storage> fs defrag schedule start fs_name sched_name
```

- 4 List the scheduled defrag job status for a file system:

```
Storage> fs defrag schedule list fs_name
```

Checking and repairing a file system

The `Storage> fs fsck` command lets you check and repair a file system while the file system is offline.

The `Storage> fs fsck` command tries to perform a normal fsck (check and repair) of the file system first, but if the `fullfsck` option is set, the command proceeds depending on the input that is provided by the user.

In most cases, a normal fsck (only log replay) is sufficient to repair a file system. In cases where there is structural damage to the file system's metadata, a full fsck of the file system may be necessary to repair the file system.

Warning: Using the `Storage> fs fsck` command on an online file system can damage the data on the file system. Only use the `Storage> fs fsck` command on a file system that is offline.

Note: When running the `Storage> fs fsck` command, you may encounter a process of `Unknown`. The `Unknown` process is normal, since there is no process printed as output when running a normal fsck using `Storage> fsck fs_name`. Full fsck is run only if the normal fsck fails. In the support mode, if you are running a full fsck, Access Appliance records that status in an internal database/file.

To check and repair a file system

- ◆ To check and repair a file system, enter the following:

```
Storage> fs fsck fs_name
```

where `fs_name` specifies the file system for which you want to check and repair.

Configuring FastResync for a file system

If the power fails or a switch fails, mirrors in a file system may not be in a consistent state.

The `Storage> fs setfastresync(FastResync)` command performs quick and efficient resynchronization of stale mirrors.

Note: You must have at least two mirrors on the file system to enable FastResync.

To enable the FastResync option

- ◆ To enable FastResync for a file system, enter the following:

```
Storage> fs setfastresync tier_name fs_name [pool_or_disk_name]
```

fs_name	Specifies the name of the file system for which to enable FastResync.
pool_or_disk_name	Specifies the pool or the disk name to resynchronize from the mirrored file system that spans the specified pool or disk.
tier_name	Specifies the name of the tier.

You can also enable FastResync for a tier of a file system.

```
Storage> tier setfastresync fs_name [pool_or_disk_name]
```

fs_name	Specifies the name of the file system for which to enable FastResync.
pool_or_disk_name	Specifies the pool or the disk name to resynchronize from the mirrored file system that spans the specified pool or disk.

Disabling the FastResync option for a file system

You can disable the FastResync option for a file system.

Note: When instant rollbacks exist for a volume, you cannot disable the FastResync option for a file system.

To disable the FastResync option

- ◆ To disable the FastResync option for a file system, enter the following:

```
Storage> fs unsetfastresync tier_name fs_name
```

where:

fs_name Specifies the name of the file system for which to disable FastResync. If you specify a file system that does not exist, an error message is displayed.

tier_name Specifies the tier name.

You can also disable FastResync for a tier of a file system.

```
Storage> tier unsetfastresync fs_name
```

where:

fs_name Specifies the name of the file system for which to disable FastResync. If you specify a file system that does not exist, an error message is displayed.

Checking and resynchronizing stale mirrors

You can check if there is a stale mirror on any of your file systems. If there is a stale mirror, the stale mirror needs to be resynchronized, and the resynchronization process needs to be verified.

To check if there are stale mirrors on your file systems

- ◆ To check if your file systems contain a stale mirror, enter the following:

```
Storage> fs checkmirror
```

To resynchronize all stale mirrors or a stale mirror for a specified file system

- ◆ To resynchronize all stale mirrors or a stale mirror for a specified file system, enter the following:

```
Storage> fs resync [fs_name]
```

where *fs_name* is the name of the specified file system where you want to resynchronize for stale mirrors.

If you do not include *fs_name*, you resynchronize all the stale mirrors for all your file systems.

To verify the resynchronization process for your stale mirrors

- ◆ To verify the resynchronization process for your stale mirrors, enter the following:

```
Storage> fs checkresync
```

Note: If a column addition to a file system is in progress, the output of the `Storage> fs checkresync` command will include RELAYOUT status.

Setting file system alerts

For a file system to run efficiently, you should always reserve some space for the file system rather than using 100% of the space. You can set file system alerts based on file system or snapshot usage. You can set the alert based on the number of inodes used, file system space used, or snapshot usage.

File system alerts can be displayed by using the `Report> showevents` command.

To set file system alerts

- ◆ To set file system alerts, enter the following:

```
Storage> fs alert set numinodes | numspace | fullspace | fullinodes
value [fs_name,...] [snapshot_name]
```

numinodes	When setting the alert for <code>numinodes</code> , <i>value</i> is the number of inodes used. The default alert value for <code>numinodes</code> is set at 0. An alert will not be sent until you set it to a different value.
numspace	When setting the alert for <code>numspace</code> , <i>value</i> is the percentage you want to set to trigger the alert. By default, the alert is sent at 80%. If you do not specify a file system name, the default value is modified.
fullspace	<code>fullspace</code> is the tunable for setting an alert if the file system becomes full. When file system usage is above the limit set by the <code>fullspace</code> tunable, all the NFS/CIFS shares on the file system are automatically changed to read-only to prevent the file system from becoming full again. When you grow the file system or delete some files to free up space, the NFS/CIFS shares are automatically changed back to read-write (there might be a delay of up to five minutes) for the change to occur.

Note: The file system size is checked every five minutes. During this five-minute interval, if the usage of the file system grows to more than 80%, the NFS/CIFS shares are changed to read-only. If the file system is small and write I/O is fast, then the file system can be filled up to 100% before being changed to read-only. This is by design.

By default, the `fullspace` tunable is set to 0, which means that the `fullspace` tunable is disabled.

fullinodes	<code>fullinodes</code> is the tunable for setting an alert if the file system becomes full. When inodes on a file system reach the limit of <code>fullinodes</code> , the NFS/CIFS shares on the file system are automatically changed to read-only. After the file system is changed to read-only, you need to delete some files from the file system, and a remount of the file system may be required for the NFS/CIFS shares to be changed to read-write. By default, the <code>fullinodes</code> tunable is set to 0, which indicates that the <code>fullinodes</code> tunable is disabled.
fs_name	Name of the file system for which you want to set the file system alerts. <i>fs_name</i> is optional. To specify multiple file systems, use commas to separate the file system names.
snapshot_name	Name of the snapshot for which you want to set the file system alert. Note: The following are reserved words for <code>snapshot_name</code> : <code>flags</code> , <code>ctime</code> , and <code>mtime</code> .

Displaying file system alert values

You can display the current disk space usage and the set alert value. A `D` beside the value indicates that the value is the default value used throughout the system.

To display file system alert values

- ◆ To display file system alert values, enter the following:

```
Storage> fs alert show
```

Removing file system alerts

You can remove the alerts set on a file system. If you remove an alert on any file system, you receive alerts for the file systems based on the default values.

To remove file system alerts

- ◆ To remove file system alerts, enter the following:

```
Storage> fs alert unset numnodes | numspace | fullspace | fullnodes
[fs_name,...] [snapshot_name]
```

fs_name Name of the file system for which you want to remove the file system alert. *fs_name* is optional.

When the fullspace/fullnodes tunables are unset (set to 0), the shares that were changed to read-only due to file system high usage are changed back to read-write mode immediately.

snapshot_name Name of the snapshot for which you want to remove the file system alert.

Note: The following are reserved words for *snapshot_name*: *flags*, *ctime*, and *mtime*.

See “[Setting file system alerts](#)” on page 268.

Destroying a file system

The `Storage> fs destroy` command unmounts a file system and releases its storage back to the storage pool. You cannot destroy the file systems that CIFS or NFS share.

To destroy a file system

- ◆ To destroy a file system, enter the following:

```
Storage> fs destroy fs_name
```

where *fs_name* specifies the name of the file system that you want to destroy.

Upgrading disk layout versions

In this release, you can create and mount only file systems with disk layout Version 8 and above.

To upgrade the disk layout versions

- 1** Ensure the file system is online before proceeding.

You can find information about the file system version, enter the following:

```
Storage> fs list file_system_name
```

- 2** To upgrade a file system to the current layout, enter the following:

```
Storage> fs upgrade file_system_name
```

Provisioning and managing Access Appliance shares

- [Chapter 20. Creating shares for applications](#)
- [Chapter 21. Creating and maintaining NFS shares](#)
- [Chapter 22. Creating and maintaining CIFS shares](#)
- [Chapter 23. Integrating Access Appliance with Data Insight](#)

Creating shares for applications

This chapter includes the following topics:

- [About file sharing protocols](#)
- [About concurrent access](#)
- [Sharing directories using CIFS and NFS protocols](#)
- [Sharing a file system as a CIFS home directory](#)
- [About concurrent access with NFS and S3](#)

About file sharing protocols

Access Appliance provides support for multiple file sharing protocols.

Access Appliance offers unified access, which provides the option to share a file system or a directory in a file system with more than one protocol. For unified access, only certain protocols combinations are supported.

See [“About concurrent access”](#) on page 274.

Table 20-1 Protocols

Protocol	Definition
Amazon S3	The object server lets you store and retrieve the data that is stored in Access Appliance using the Amazon Simple Storage Service (Amazon S3) protocol. See “About the Object Store server” on page 157.

Table 20-1 Protocols (*continued*)

Protocol	Definition
CIFS	CIFS is active on all nodes within the Access Appliance cluster. The specific shares are read/write on the node they reside on, but can failover to any other node in the cluster. Access Appliance supports CIFS home directory shares. See “About configuring Access Appliance for CIFS” on page 120.
NFS	All the nodes in the cluster can serve the same NFS share at the same time in read-write mode. This creates very high aggregated throughput rates, because you can use the sum of the bandwidth of all the nodes. Cache-coherency is maintained throughout the cluster. Access Appliance supports the NFS kernel-based server. See “About using the NFS server with Access Appliance” on page 112.

About concurrent access

Access Appliance provides support for multi-protocol file sharing where the same file system can be exported to both Windows and UNIX users using the Common Internet File System (CIFS), Network File System (NFS), and Simple Storage Service (S3) protocols. The result is an efficient use of storage by sharing a single data set across multiple application platforms.

Note: When a share is exported over both NFS and CIFS protocols, the applications running on the NFS and CIFS clients may attempt to concurrently read or write the same file. This may lead to unexpected results, such as reading stale data, since the locking models used by these protocols are different. For this reason, Access Appliance warns you when the share export is requested over NFS or CIFS and the same share has already been exported for write access over CIFS or NFS.

The following sections describe concurrent access with multiple protocols.

See [“Sharing directories using CIFS and NFS protocols”](#) on page 275.

See [“About concurrent access with NFS and S3”](#) on page 277.


See [“About file sharing protocols”](#) on page 273.

Sharing directories using CIFS and NFS protocols

Access Appliance provides support for multi-protocol file sharing where the same directory or file system can be exported to both Windows and UNIX users using the CIFS and NFS protocols. The result is an efficient use of storage by sharing a single data set across multi-application platforms.

Figure 20-1 shows how the directory sharing for the two protocols works.

Figure 20-1 Exporting and/or sharing CIFS and NFS directories



It is recommended that you disable the `oplocks` option when the following occurs:

- A file system is exported over both the CIFS and NFS protocols.

- Either the CIFS and NFS protocol is set with read and write permission.

To export a directory to Windows and UNIX users

- 1 To export a directory to Windows and UNIX users with read-only and read-write permission respectively, enter the `CIFS` mode and enter the following commands:

```
CIFS> show
```

Name	Value
----	-----
netbios name	Pei60
ntlm auth	yes
allow trusted domains	no
homedirfs	
aio size	0
idmap backend	rid:10000-1000000
workgroup	PEI-DOMAIN
security	ads
Domain	PEI-DOMAIN.COM
Domain user	Administrator
Domain Controller	10.200.107.251
Clustering Mode	normal

```
CIFS> share add fs1 share1 ro
Exporting CIFS filesystem : share1...
CIFS> share show
```

ShareName	FileSystem	ShareOptions
share1	fs1	owner=root,group=root,ro

Exit CIFS mode:

```
CIFS> exit
```

- 2 Enter the `NFS` mode and enter the following commands:

```
NFS> share add rw fs1
ACCESS nfs WARNING V-288-0 Filesystem (fs1)
is already shared over CIFS with 'ro' permission.
Do you want to proceed (y/n): y
Exporting */vx/fs1 with options rw
..Success.
NFS> share show
```

/vx/fs1	*	(rw)
---------	---	------

Sharing a file system as a CIFS home directory

When the file system in CIFS is set to `homedirfs`, the Access Appliance software assumes that the file system is exported to CIFS users in read and write mode. Access Appliance does not allow you to export the same file system as an CIFS share and a home directory file system (`homedirfs`). For example, if the file system `fs1` is already exported as a CIFS share, then you cannot set it as `homedirfs`.

To export a file system set as `homedirfs`

- ◆ To request that a file system be used for home directories, you need to export the file system. Go to the `CIFS` mode and enter the following:

```
CIFS> share show
ShareName      FileSystem    ShareOptions
share1         fs1          owner=root,group=root,rw
CIFS> set homedirfs fs1
ACCESS cifs ERROR V-288-615 Filesystem (fs1) is already exported
by another CIFS share.
```

About concurrent access with NFS and S3

Access Appliance supports concurrent access to a shared file system or a directory from both NFS and S3. The supported configurations are:

- Applications or users write data to NFS shares, while other applications or users read the data over S3.
- Applications or users write data to S3 shares, while other applications or users read the data over NFS.

Creating and maintaining NFS shares

This chapter includes the following topics:

- [About NFS file sharing](#)
- [About the NFS shares](#)
- [Displaying file systems and snapshots that can be exported](#)
- [Exporting an NFS share](#)
- [Displaying exported directories](#)
- [About managing NFS shares using netgroups](#)
- [Unexporting a directory or deleting NFS options](#)
- [Exporting an NFS share for Kerberos authentication](#)
- [Mounting an NFS share with Kerberos security from the NFS client](#)
- [Exporting an NFS snapshot](#)

About NFS file sharing

The Network File System (NFS) protocol enables exported directories (including all files under the directory that reside on the exported directory's file system) hosted by an NFS server to be accessed by multiple UNIX and Linux client systems.

Using NFS, a local system can mount and use a disk partition or a file system from a remote system (an NFS server), as if it were local. The Access Appliance NFS server exports a directory, with selected permissions and options, and makes it available to NFS clients.

The selected permissions and options can also be updated, by adding or removing permissions, to restrict or expand the permitted use.

The Access Appliance NFS service is clustered. The NFS clients continuously retry during a failover transition. Even if the TCP connection is broken for a short time, the failover is transparent to NFS clients, and NFS clients regain access transparently as soon as the failover is complete.

See [“About using the NFS server with Access Appliance”](#) on page 112.

About the NFS shares

You can create an NFS share to provision the storage.

See [“Creating an NFS share without replication and encryption”](#) on page 283.

See [“Creating an NFS share with continuous replication”](#) on page 279.

See [“Creating an NFS share with episodic replication”](#) on page 280.

See [“Creating an NFS share with encryption”](#) on page 281.

See [“Creating an NFS share with episodic replication and encryption”](#) on page 282.

Creating an NFS share with continuous replication

You can create an NFS share with continuous replication.

To create an NFS share with continuous replication

1 Prerequisites:

Configure storage.

Start the NFS server.

Configure continuous replication.

Activate a storage policy.

2 Quick Actions: Go to **Provision Storage > Storage for NFS > With Continuous Replication**, and then click **Next**.

3 Click the **Configure Storage** tab and specify the storage options.

4 Click the **Start NFS Server** tab, and then click **Start NFS Server** to start the NFS server.

5 Click **Configure Replication > Go to Replication** link to set up the replication virtual IP on the source cluster.

Go to **Settings > Replication management** on destination cluster's GUI and set up the continuous replication virtual IP on the destination cluster.

- 6 Click **Set Up Replication Link**.
- 7 Click the **Activate Storage Policy** tab, select a policy for the NFS share, and then click **Activate Policy**.
- 8 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the NFS share.

- 9 View the **Recent Activity** panel for the status of the task.
- See [“Creating an NFS share without replication and encryption”](#) on page 283.
- See [“Creating an NFS share with episodic replication”](#) on page 280.
- See [“Creating an NFS share with encryption”](#) on page 281.
- See [“Creating an NFS share with episodic replication and encryption”](#) on page 282.

Creating an NFS share with episodic replication

You can create an NFS share with episodic replication.

To create an NFS share with episodic replication

- 1 **Prerequisites:**
 - Configure storage.
 - Start the NFS server.
 - Configure episodic replication.
 - Activate a storage policy.
- 2 **Quick Actions:** Go to **Provision Storage > Storage for NFS > With Episodic Replication**, and then click **Next**.
- 3 Click the **Configure Storage** tab and specify the storage options.
- 4 Click the **Start NFS Server** tab, and then click **Start NFS Server** to start the NFS server.
- 5 Click **Configure Replication > Go to Replication** link to set up the replication virtual IP on the source cluster.

Go to **Settings > Replication management** on destination cluster's GUI and set up the episodic replication virtual IP on the destination cluster.
- 6 Click **Set Up Replication Link**.

- 7 Click the **Activate Storage Policy** tab, select a policy for the NFS share, and then click **Activate Policy**.
- 8 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the NFS share.

- 9 View the **Recent Activity** panel for the status of the task.
- See [“Creating an NFS share with continuous replication”](#) on page 279.
- See [“Creating an NFS share without replication and encryption”](#) on page 283.
- See [“Creating an NFS share with encryption”](#) on page 281.
- See [“Creating an NFS share with episodic replication and encryption”](#) on page 282.

Creating an NFS share with encryption

You can create an NFS share with encryption.

To create an NFS share with encryption

- 1 **Prerequisites:**
 - Configure storage.
 - Start the NFS server.
 - Configure KMS server.
 - Activate an encryption policy.
- 2 **Quick Actions:** Go to **Provision Storage > Storage for NFS > With Encryption**, and then click **Next**.
- 3 Click the **Configure Storage** tab and specify the storage options.
- 4 Click the **Start NFS Server** tab, and then click **Start NFS Server** to start the NFS server.
- 5 Click **Configure KMS Server**: Set up the KMS server if not done already by clicking on **Go to Configure KMS**.
- 6 Click the **Activate Encryption Policy** tab, select a policy for the NFS share, and then click **Activate Policy**.

- 7 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the NFS share.

- 8 View the **Recent Activity** panel for the status of the task.
- See [“Creating an NFS share with episodic replication ”](#) on page 280.
- See [“Creating an NFS share with continuous replication ”](#) on page 279.
- See [“Creating an NFS share without replication and encryption ”](#) on page 283.
- See [“Creating an NFS share with episodic replication and encryption ”](#) on page 282.

Creating an NFS share with episodic replication and encryption

You can create an NFS share with episodic replication and encryption.

To create an NFS share with episodic replication and encryption

- 1 **Prerequisites:**
 - Configure storage.
 - Start the NFS server.
 - Configure episodic replication.
 - Configure KMS server.
 - Activate an encryption policy.
- 2 **Quick Actions:** Go to **Provision Storage > Storage for NFS > With Episodic Replication and Encryption**, and then click **Next**.
- 3 Click the **Configure Storage** tab and specify the storage options.
- 4 Click the **Start NFS Server** tab, and then click **Start NFS Server** to start the NFS server.
- 5 Click **Configure Replication > Go to Replication** link to set up the replication virtual IP on the source cluster.

Go to **Settings > Replication management** on destination cluster's GUI and set up the episodic replication virtual IP on the destination cluster.
- 6 Click **Set Up Replication Link**.
- 7 Click **Configure KMS Server**: Set up the KMS server if not done already by clicking on **Go to Configure KMS**.

- 8 Click the **Activate Encryption Policy** tab, select a policy for the NFS share, and then click **Activate Policy**.
- 9 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the NFS share.

- 10 View the **Recent Activity** panel for the status of the task.

See [“Creating an NFS share with encryption ”](#) on page 281.

See [“Creating an NFS share with episodic replication ”](#) on page 280.

See [“Creating an NFS share with continuous replication ”](#) on page 279.

See [“Creating an NFS share without replication and encryption ”](#) on page 283.

Creating an NFS share without replication and encryption

You can create an NFS share without replication and encryption.

To create an NFS share without replication and encryption

1 Prerequisites:

Configure storage.

Start the NFS server.

Activate a storage policy.

- 2 **Quick Actions:** Go to **Provision Storage > Storage for NFS > Without Replication and Encryption**, and then click **Next**.
- 3 Click the **Configure Storage** tab and specify the storage options.
- 4 Click the **Start NFS Server** tab, and then click **Start NFS Server** to start the NFS server.
- 5 Click the **Activate Storage Policy** tab, select a policy for the NFS share, and then click **Activate Policy**.

- 6 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the NFS share.

- 7 View the **Recent Activity** panel for the status of the task.
 See [“Creating an NFS share with continuous replication ”](#) on page 279.
 See [“Creating an NFS share with episodic replication ”](#) on page 280.
 See [“Creating an NFS share with encryption ”](#) on page 281.
 See [“Creating an NFS share with episodic replication and encryption ”](#) on page 282.

Displaying file systems and snapshots that can be exported

To display a file system and snapshots that can be exported

- ◆ To display online file systems and the snapshots that can be exported, enter the following:

```
NFS> show fs
```

For example:

```
NFS> show fs
FS/Snapshot
=====
fs1
```

Exporting an NFS share

You can export an NFS share with the specified NFS options that can then be accessed by one or more client systems.

If you add a directory that has already been exported with a different NFS option (rw, ro, async, or secure, for example), Access Appliance provides a warning message saying that the directory has already been exported. Access Appliance updates (overwrite) the old NFS options with the new NFS options.

Directory options appear in parentheses.

If a client was not specified when the `NFS> share add` command was used, then `*` is displayed as the system to be exported to, indicating that all clients can access the directory.

Directories that have been exported to multiple clients appear as separate entries. Directories that are exported to `<world>` and other specific clients also appear as separate entries.

For example:

Consider the following set of exported directories where only the client (1.1.1.1) has **read-write** access to directory (fs2), while all other clients have **read** access only.

```
share1 /vx/fs2          * (ro)

share2 /vx/fs2  1.1.1.1 (rw)
```

When sharing a directory, Access Appliance does not check whether the client exists or not. If you add a share for an unknown client, then an entry appears in the `NFS> show` command output.

The `NFS> show fs` command displays the list of exportable file systems. If a directory does not exist, the directory is automatically created and exported when you try to export it.

Valid NFS options include the following:

<code>rw</code>	Grants read and write permission to the directory (including all files under the directory that reside on the exported directory's file system). Hosts mounting this directory will be able to make changes to the directory.
<code>ro (Default)</code>	Grants read-only permission to the directory. Hosts mounting this directory will not be able to change it.
<code>sync (Default)</code>	Grants synchronous write access to the directory. Forces the server to perform a disk write before the request is considered complete.
<code>async</code>	Grants asynchronous write access to the directory. Allows the server to write data to the disk when appropriate.
<code>secure (Default)</code>	Grants secure access to the directory. Requires that clients originate from a secure port. A secure port is between 1-1024.
<code>insecure</code>	Grants insecure access to the directory. Permits client requests to originate from unprivileged ports (those above 1024).

<code>secure_locks</code> (Default)	Requires authorization of all locking requests.
<code>insecure_locks</code>	Some NFS clients do not send credentials with lock requests, and therefore work incorrectly with <code>secure_locks</code> , in which case you can only lock world-readable files. If you have such clients, either replace them with better ones, or use the <code>insecure_locks</code> option.
<code>root_squash</code> (Default)	Prevents the root user on an NFS client from having root privileges on an NFS mount. This effectively "squashes" the power of the remote root user to the lowest local user, preventing remote root users from acting as though they were the root user on the local system.
<code>no_root_squash</code>	Disables the <code>root_squash</code> option. Allows root users on the NFS client to have root privileges on the NFS server.
<code>wdelay</code> (Default)	Causes the NFS server to delay writing to the disk if another write request is imminent. This can improve performance by reducing the number of times the disk must be accessed by separate write commands, reducing write overhead. Note: The <code>wdelay</code> option is deprecated, and is supported for backward-compatibility only.
<code>no_wdelay</code>	Disables the <code>wdelay</code> option. The <code>no_wdelay</code> option has no effect if the <code>async</code> option is also set. Note: The <code>no_wdelay</code> option is deprecated, and is supported for backward-compatibility only. Using the <code>no_wdelay</code> option is always effective.
<code>subtree_check</code>	Verifies that the requested file is in an exported subdirectory. If this option is turned off, the only verification is that the file is in an exported file system.
<code>no_subtree_check</code> (Default)	Sometimes subtree checking can produce problems when a requested file is renamed while the client has the file open. If many such situations are anticipated, it might be better to set <code>no_subtree_check</code> . One such situation might be the export of the home directory. Most other situations are best handled with <code>subtree_check</code> .
<code>fsid</code> (Default)	Allows the Access Appliance administrator to associate a specific number as <code>fsid</code> with the share.

<code>nordirplus</code>	Allows you to disable a readdirplus remote procedure call (RPC).
<code>sec</code>	Specifies the Kerberos security options for exporting an NFS share. The value can be <code>krb5</code> , <code>krb5i</code> , <code>krb5p</code> , or <code>sys</code> . The <code>sys</code> option does not provide Kerberos authentication. The other options use Kerberos V5 to authenticate users to the NFS server.

Note: With `root_squash`, the root user can access the share, but with 'nobody' permissions.

To export a directory/file system

- 1 To see your exportable online file systems and snapshots, enter the following:

```
NFS> show fs
```

- 2 To see your NFS shares and their options, enter the following:

```
NFS> share show
```

- 3 To export a directory, enter the following command:

```
NFS> share add nfsoptions export_dir [share_name] [client]
```

<code>nfsoptions</code>	Comma-separated list of export options from the set.
<code>export_dir</code>	Specifies the name of the directory you want to export. The directory name should start with <code>/vx</code> , and only <code>a-zA-Z0-9_/@+.=:-</code> characters are allowed for <code>export_dir</code> .
<code>share_name</code>	Specifies the share name. If the share name is not specified, then the default value is <code>vxdefault</code> .

client

Clients may be specified in the following ways:

- Single host - specify a host either by an abbreviated name that is recognized by the resolver (DNS is the resolver), the fully qualified domain name, or an IP address.
- Netgroups - specify netgroups as *@group*. Only the host part of each netgroup member is considered for checking membership.
- IP networks - specify an IP address and netmask pair (address/netmask) to simultaneously export directories to all hosts on an IP sub-network. Specify the netmask as a contiguous mask length. You can specify either an IPv4 address or an IPv6 address.

If the client is not given, then the specified directory can be mounted or accessed by any client. To re-export new options to an existing share, the new options will be updated after the command is run.

Displaying exported directories

You can display the exported directories and the NFS options that are specified when the directory was exported.

To display exported directories

To display exported directories, enter the following:

```
NFS> share show
```

The command output displays the following columns:

First column	Displays the share name.
Second column	Displays the share path.
Third column	Displays the client and share options.

About managing NFS shares using netgroups

A netgroup defines a network-wide group of hosts and users. You use netgroups for restricting access to shared NFS file systems and to restrict remote login and shell access.

Each line in the netgroup file consists of a netgroup name followed by a list of members, where a member is either another netgroup name, or a comma-separated list of host, user, or a domain. Host, user, and domain are character strings for the corresponding components. Any of these three fields can be empty, which indicates a wildcard, or may consist of the string "-" to indicate that there is no valid value for the field. The domain field must either be the local domain name or empty for the netgroup entry to be used. This field does not limit the netgroup or provide any security. The domain field refers to the domain in which the host is valid, not the domain containing the trusted host.

When exporting a directory by NFS with the specified options, clients may be specified using netgroups. Netgroups are identified using @group. Only the host part of each netgroup member is considered when checking for membership.

```
NFS> share add rw,async /vx/fs1/share share1 @client_group
```

Unexporting a directory or deleting NFS options

You can unexport the share of the exported directory.

Note: You will receive an error message if you try to remove a directory that does not exist.

To unexport a directory or delete NFS options

- 1 To see your existing exported resources, enter the following command:

```
NFS> share show
```

Only the directories that are displayed can be unexported.

For example:

```
NFS> share show
```

```
share1 /vx/fs2          * (sync)
share2 /vx/fs3          * (secure,ro,no_root_squash)
```

- 2 To delete a share, enter the following command:

```
NFS> share delete share_name
```

share_name Specifies the name of the share you want to delete.

Exporting an NFS share for Kerberos authentication

Kerberos provides three types of security options for exporting an NFS share:

- `krb5`
- `krb5i`
- `krb5p`

Access Appliance also provides a `sys (sec=sys)` export option, which does not provide Kerberos authentication. Access Appliance supports all of the three types of Kerberos security options. All of the security options use Kerberos V5 to authenticate users to NFS servers.

`krb5i` computes a hash on every remote procedure (RPC) call request to the server and every response to the client. The hash is computed on an entire message: RPC header, plus NFS arguments or results. Since the hash information travels with the NFS packet, any attacker modifying the data in the packet can be detected. Thus `krb5i` provides integrity protection.

`krb5p` uses encryption to provide privacy. With `krb5p`, NFS arguments and results are encrypted, so a malicious attacker cannot spoof on the NFS packets and see file data or metadata.

Note: Since `krb5i` and `krb5p` perform an additional set of computations on each NFS packet, NFS performance decreases as compared with `krb5`.

Performance decreases in the following order: `krb5` > `krb5i` > `krb5p`.

`krb5` provides better performance and `krb5p` gives the least performance.

Additional export options are available.

See [“Exporting an NFS share ”](#) on page 284.

To export a directory using only the `krb5` mount option

- ◆ Export a directory using only the `krb5` mount option:

```
NFS> share add sec=krb5 /vx/fs1 share1 client1
Exporting /vx/fs1 with options sec=krb5
Success.
```

If the client name is not mentioned while creating the share, the default value is `*`.

To export a directory using `krb5`, `krb5i`, `krb5p`, and `sys` options

- ◆ Export a directory using `krb5`, `krb5i`, `krb5p`, and `sys` options.

```
NFS> share add sec=krb5:krb5i:krb5p:sys /vx/fs1 share1 client1
Exporting /vx/fs1 with options sec=krb5:krb5i:krb5p:sys
Success.
```

If the client name is not mentioned while creating the share, the default value is `*`.

Different clients can use different levels of security in this case. Client A can mount with `krb5`, and client B can mount with `krb5p`. If no mount option is given at the client side, security to be chosen is negotiated, and the highest level of security is chosen. In this case, it is `krb5p`.

Mounting an NFS share with Kerberos security from the NFS client

This section explains how the NFS client does an NFS mount with the Kerberos mount options. This procedure assumes that the NFS service principal of the NFS client is added to the KDC server, and the keytab is copied at the appropriate location on the client.

The steps may differ depending on the operating system and version of the client. On a Red Hat Enterprise Linux (RHEL) client, Kerberos can be configured as follows.

To mount the NFS client with the Kerberos mount options

- 1 Create the NFS service principal for the client on the KDC server and copy it to the client system at `/etc/krb5.keytab`.
- 2 Configure the `/etc/krb5.conf` file with the KDC details.
- 3 Enable `SECURE_NFS=yes` in the `/etc/sysconfig/nfs` file.
- 4 Start the `rpcgssd` service.

```
# service rpcgssd start
```

- 5 Keep the clocks of the KDC server, the Access Appliance server, and the NFS client in sync.

A maximum of a five-minute variation is accepted, or otherwise the Kerberos NFS mount fails.

```
[root@krb-client]# mount -o vers=4,sec=krb5 10.209.107.24:/vx/fs2/share1 /mnt/share1
```

Make sure that the virtual IP that is used for mounting can use reverse name lookup to the Access Appliance cluster name. For example, if `access_ga` is the cluster name, then in the above example, `access_ga` should look up to 10.209.107.24 and vice versa. If the IP 10.209.107.24 can be looked up by multiple host names, make sure that the entry `access_ga` is first in the reverse lookup.

- 6 Ensure that the user accessing the NFS share:

- Is already added on the KDC server.

Use `kinit` to get the ticket granting ticket from the KDC server on the NFS client.

```
[root@krb-client]# su - sfuuser2

[sfuuser2@krb-client ~]$ kinit
Password for sfuuser2@TESTKDC.COM:
[sfuuser2@krb-client ~]$ cd /mnt/share1
[sfuuser2@krb-client share1]$ touch test.txt
[sfuuser2@krb-client share1]$
[sfuuser2@krb-client share1]$ ls -al total 4
drwxrwxrwx 2 root root 96 May 14 16:03 .
drwxr-xr-x. 17 root root 4096 May 7 19:41 ..
-rw-r--r-- 1 sfuuser2 sfugroup1 0 May 14 16:03 test.txt
```

Exporting an NFS snapshot

To export an NFS snapshot

- 1 For example, to create an NFS snapshot, enter the following:

```
Storage> snapshot create fs5sp1 FS5
```

See [“About snapshots”](#) on page 384.

- 2 For example, to export the NFS snapshot, enter the following:

```
NFS> share add rw /vx/FS5:fs5sp1 share1 client1
```

If the client name is not mentioned while creating the share, the default value is *.

See [“Exporting an NFS share ”](#) on page 284.

Creating and maintaining CIFS shares

This chapter includes the following topics:

- [About managing CIFS shares](#)
- [About the CIFS shares](#)
- [Exporting a directory as a CIFS share](#)
- [Configuring a CIFS share as secondary storage for an Enterprise Vault store](#)
- [Exporting the same file system/directory as a different CIFS share](#)
- [About the CIFS export options](#)
- [Setting share properties](#)
- [Displaying CIFS share properties](#)
- [Hiding system files when adding a CIFS normal share](#)
- [Allowing specified users and groups access to the CIFS share](#)
- [Denying specified users and groups access to the CIFS share](#)
- [Exporting a CIFS snapshot](#)
- [Deleting a CIFS share](#)
- [Modifying a CIFS share](#)
- [Making a CIFS share shadow copy aware](#)
- [About managing CIFS shares for Enterprise Vault](#)

About managing CIFS shares

You can export the Access Appliance file systems to clients as CIFS shares. When a share is created, it is given a name. The share name is different from the file system name. Clients use the share name when they import the share.

You create and export a share with one command. The same command binds the share to a file system, and you can also use it to specify share properties.

In addition to exporting file systems as CIFS shares, you can use Access Appliance to store user home directories. Each of these home directories is called a home directory share. Shares that are used to export ordinary file systems (that is, file systems that are not used for home directories), are called ordinary shares to distinguish them from home directory shares.

About the CIFS shares

You can create a CIFS share to provision the storage.

See [“Creating a CIFS share with encryption”](#) on page 297.

See [“Creating a CIFS share with episodic replication”](#) on page 296.

See [“Creating a CIFS share with episodic replication and encryption”](#) on page 298.

See [“Creating a CIFS share without replication and encryption”](#) on page 299.

See [“Creating a CIFS share with continuous replication”](#) on page 295.

Creating a CIFS share with continuous replication

You can create a CIFS share with continuous replication.

To create a CIFS share with continuous replication

1 Prerequisites:

Configure storage.

Start the CIFS server.

Activate a storage policy.

Set up the continuous replication link.

2 Quick Actions: Go to **Provision Storage > Storage for CIFS > With Continuous Replication**, and then click **Next**.

3 Click the **Configure Storage** tab and specify the storage options.

4 Click the **Start CIFS Server** tab, and then click **Start CIFS Server** to start the CIFS server.

- 5 Click **Configure Replication > Go to Replication** link to set up the replication virtual IP on the source cluster.
 Go to **Settings > Replication management** on destination cluster's GUI and set up the continuous replication virtual IP on the destination cluster.
- 6 Click **Set Up Replication Link**.
- 7 Click the **Activate Storage Policy** tab, select a policy for the CIFS share, and then click **Activate Policy**.
- 8 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the CIFS share.

- 9 View the **Recent Activity** panel for the status of the task.
 See [“Creating a CIFS share with encryption ”](#) on page 297.
 See [“Creating a CIFS share with episodic replication ”](#) on page 296.
 See [“Creating a CIFS share with episodic replication and encryption ”](#) on page 298.
 See [“Creating a CIFS share without replication and encryption ”](#) on page 299.

Creating a CIFS share with episodic replication

You can create a CIFS share with episodic replication.

To create a CIFS share with episodic replication

- 1 **Prerequisites:**
 Configure storage.
 Start the CIFS server.
 Configure replication.
 Activate a storage policy.
 Set up the episodic replication link.
- 2 **Quick Actions:** Go to **Provision Storage > Storage for CIFS > With Episodic Replication**, and then click **Next**.
- 3 Click the **Configure Storage** tab and specify the storage options.
- 4 Click the **Start CIFS Server** tab, and then click **Start CIFS Server** to start the CIFS server.

- 5 Click **Configure Replication** > **Go to Replication** link to set up the replication virtual IP on the source cluster.
 Go to **Settings** > **Replication management** on destination cluster's GUI and set up the episodic replication virtual IP on the destination cluster.
- 6 Click **Set Up Replication Link**.
- 7 Click the **Activate Storage Policy** tab, select a policy for the CIFS share, and then click **Activate Policy**.
- 8 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the CIFS share.

- 9 View the **Recent Activity** panel for the status of the task.
 See [“Creating a CIFS share with continuous replication”](#) on page 295.
 See [“Creating a CIFS share with encryption ”](#) on page 297.
 See [“Creating a CIFS share with episodic replication and encryption ”](#) on page 298.
 See [“Creating a CIFS share without replication and encryption ”](#) on page 299.

Creating a CIFS share with encryption

You can create a CIFS share with encryption.

To create a CIFS share with encryption

- 1 **Prerequisites:**
 Configure storage.
 Start the CIFS server.
 Configure KMS server.
 Activate an encryption policy.
- 2 **Quick Actions:** Go to **Provision Storage** > **Storage for CIFS** > **With Encryption**, and then click **Next**.
- 3 Click the **Configure Storage** tab and specify the storage options.
- 4 Click the **Start CIFS Server** tab, and then click **Start CIFS Server** to start the CIFS server.

- 5 Click **Configure KMS Server**: Set up the KMS server if not done already by clicking on **Go to Configure KMS**.
- 6 Click the **Activate Encryption Policy** tab, select a policy for the CIFS share, and then click **Activate Policy**.
- 7 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the CIFS share.

- 8 View the **Recent Activity** panel for the status of the task.
- See [“Creating a CIFS share with continuous replication”](#) on page 295.
- See [“Creating a CIFS share with episodic replication ”](#) on page 296.
- See [“Creating a CIFS share with episodic replication and encryption ”](#) on page 298.
- See [“Creating a CIFS share without replication and encryption ”](#) on page 299.

Creating a CIFS share with episodic replication and encryption

You can create a CIFS share with episodic replication and encryption.

To create a CIFS share with episodic replication and encryption

- 1 **Prerequisites:**
 - Configure storage.
 - Start the CIFS server.
 - Configure episodic replication.
 - Configure the KMS server.
 - Activate an encryption policy.
- 2 **Quick Actions:** Go to **Provision Storage > Storage for CIFS > With Episodic Replication and Encryption**, and then click **Next**.
- 3 Click the **Configure Storage** tab and specify the storage options.
- 4 Click the **Start CIFS Server** tab, and then click **Start CIFS Server** to start the CIFS server.

- 5 Click **Configure Replication > Go to Replication** link to set up the replication virtual IP on the source cluster.
 Go to **Settings > Replication management** on destination cluster's GUI and set up the episodic replication virtual IP on the destination cluster.
 - 6 Click **Set Up Replication Link**.
 - 7 Click the **Activate Encryption Policy** tab, select a policy for the CIFS share, and then click **Activate Policy**.
 - 8 Click **Configure KMS Server**: Set up the KMS server if not done already by clicking on **Go to Configure KMS**.
-
- 9 **Note:** If a storage policy is not activated, you cannot provision storage for the CIFS share.
-

- 10 View the **Recent Activity** panel for the status of the task.
 See [“Creating a CIFS share with continuous replication”](#) on page 295.
 See [“Creating a CIFS share with encryption ”](#) on page 297.
 See [“Creating a CIFS share with episodic replication ”](#) on page 296.
 See [“Creating a CIFS share without replication and encryption ”](#) on page 299.

Creating a CIFS share without replication and encryption

You can create a CIFS share without replication and encryption.

To create a CIFS share without replication and encryption

- 1 **Prerequisites:**
 Configure storage.
 Start the CIFS server.
 Activate a storage policy.
- 2 **Quick Actions:** Go to **Provision Storage > Storage for CIFS > Without Replication and Encryption**, and then click **Next**.
- 3 Click the **Configure Storage** tab and specify the storage options.
- 4 Click the **Start CIFS Server** tab, and then click **Start CIFS Server** to start the CIFS server.
- 5 Click the **Activate Storage Policy** tab, select a policy for the CIFS share, and then click **Activate Policy**.

- 6 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the CIFS share.

- 7 View the **Recent Activity** panel for the status of the task.
- See [“Creating a CIFS share with continuous replication”](#) on page 295.
- See [“Creating a CIFS share with encryption ”](#) on page 297.
- See [“Creating a CIFS share with episodic replication ”](#) on page 296.
- See [“Creating a CIFS share with episodic replication and encryption ”](#) on page 298.

Exporting a directory as a CIFS share

To export a directory as a CIFS share

- 1 To export a directory as a CIFS share, enter the following:

```
CIFS> share add fs1/access share1 rw,full_acl
```

If the directory name contains a space, enter the directory name with double quotes (" ").

- 2 To list the CIFS shares, enter the following:

```
CIFS> share show
```

Configuring a CIFS share as secondary storage for an Enterprise Vault store

You can use Access Appliance as secondary storage with Enterprise Vault 12.0 by exporting the file system over the CIFS protocol.

Note: Before configuring the CIFS share path as secondary storage, you need to verify that the CIFS share path is accessible. Confirm that I/O operations can occur on the CIFS share.

Configuring a CIFS share as secondary storage for an Enterprise Vault store

- 1 On the Access Appliance cluster, you export the file system over the CIFS protocol using the following CIFS export options: `fs_mode=1777,rw,full_acl`.
See [“About the CIFS export options”](#) on page 302.
- 2 On the Enterprise Vault server, open the Enterprise Vault console.
- 3 Right-click on the partition that is created on **Vault Store > Properties**.
Enterprise Vault brings up the **Vault Store Partition Properties** window.
- 4 In the **Vault Store Partition Properties** window, select the **Migration** tab.
- 5 Specify the path of the CIFS share in the **Secondary storage location** text box.

Example:

`\\IP address of the CIFS share\name of file system`

- 6 Press **Apply**.

Exporting the same file system/directory as a different CIFS share

In ctdb clustering mode, you can export the same file system or directory as a different CIFS share with different available CIFS options. This features allows you more granular control over CIFS shares for different sets of users.

If the same file system is exported as different shares in ctdb clustering mode, then after switching to normal clustering mode only one share out of these is available.

Note: If the same file system or directory is exported as different shares, then the `fs_mode` value is the same for all of these shares; that is, the last modified `fs_mode` value is applicable for all of those shares. This applies to `fs_mode`, owner, and group.

Note: This feature is only supported in the ctdb clustering mode.

To export a directory with read access to everyone, but write access to the limited set of users who need to be authenticated

- ◆ To export a directory with read access to everyone, but write access to the limited set of users who need to be authenticated, enter the following:

CIFS> share add "fsl/Veritas isa" share1 rw,noguest
CIFS> share add "fsl/Veritas isa" share2 ro,guest
CIFS> share show

The above example illustrates that the same directory is exported as a different CIFS share for `guest` and `noguest` users with different sets of permissions.

About the CIFS export options

The following are the CIFS export options.

Table 22-1 CIFS export options

CIFS export option	Definition
rw	There is a share option which specifies if the files in the share will be read-only or if both read and write access will be possible, subject to the authentication and authorization checks when a specific access is attempted. This share option can be given one of these values, either <code>rw</code> or <code>ro</code>. Grants read and write permission to the exported share.
ro (Default)	Grants read-only permission to the exported share. Files cannot be created or modified.
guest	This configuration option specifies if a user trying to establish a CIFS connection with the share must always provide the user name and password, or if they can connect without it. In this case, only restricted access to the share will be allowed. The same kind of access is allowed to <code>anonymous</code> or <code>guest</code> user accounts. This share option can have one of the following values, either <code>guest</code> or <code>noguest</code>. Access Appliance allows restricted access to the share when no user name or password is provided.
noguest (Default)	Access Appliance always requires the user name and password for all of the connections to this share.

Table 22-1 CIFS export options (*continued*)

CIFS export option	Definition
full_acl	All Windows Access Control Lists (ACLs) are supported except in the case when you attempt using the Windows Explorer folder Properties > Security GUI to inherit down to a non-empty directory hierarchy while denying all access to yourself.
no_full_acl (Default)	Some advanced Windows Access Control Lists (ACLs) functionality does not work. For example, if you try to create ACL rules on files saved in a CIFS share using Windows explorer while allowing some set of file access for <code>user1</code> and denying file access for <code>user2</code> , this is not possible when CIFS shares are exported using <code>no_full_acl</code> .
hide_unreadable	Prevents clients from seeing the existence of files and directories that are not readable to them. The default is: <code>hide_unreadable</code> is set to off.
veto_sys_files	To hide some system files (lost+found, quotas, quotas.grp) from displaying when using a CIFS normal share, you can use the <code>veto_sys_files</code> CIFS export option. For example, when adding a CIFS normal share, the default is to display the system files. To hide the system files, you must use the <code>veto_sys_files</code> CIFS export option.
fs_mode	When a file system or directory is exported by CIFS, its mode is set to an <code>fs_mode</code> value. It is the UNIX access control set on a file system, and CIFS options like <code>rw/ro</code> do not take precedence over it. This value is reset to <code>0755</code> when the CIFS share is deleted. The default is: <code>fs_mode = 1777</code> .
dir_mask	When a directory is created under a file system or directory exported by CIFS, the necessary permissions are calculated by mapping DOS modes to UNIX permissions. The resulting UNIX mode is then bit-wise 'AND'ed with this parameter. Any bit not set here is removed from the modes set on a directory when it is created. The default is: <code>dir_mask = 0775</code> .

Table 22-1 CIFS export options (*continued*)

CIFS export option	Definition
create_mask	When a file is created under a file system or directory exported by CIFS, the necessary permissions are calculated by mapping DOS modes to UNIX permissions. The resulting UNIX mode is then bit-wise 'AND'ed with this parameter. Any bit not set here is removed from the modes set on a file when it is created. The default is: <code>create_mask = 0775</code>.
oplocks (Default)	Access Appliance supports the CIFS opportunistic locks. You can enable or disable them for a specific share. The opportunistic locks improve performance for some workloads, and there is a share configuration option which can be given one of the following values, either oplocks or nooplocks. Access Appliance supports opportunistic locks on the files in this share.
nooplocks	No opportunistic locks will be used for this share. Disable the oplocks when: ■ 1) A file system is exported over both CIFS and NFS protocols. ■ 2) Either CIFS or NFS protocol has read and write access.
owner	There are more share configuration options that can be used to specify the user and group who own the share. If you do not specify these options for a share, Access Appliance uses the current values as default values for these options. You may want to change the default values to allow a specific user or group to be the share owner. Irrespective of who are owner and group of the exported share, any CIFS clients can create folders and files in the share. However, there are some operations that require owner privileges; for example, changing the owner itself, and changing permissions of the top-level folder (that is, the root directory in UNIX terms). To enable these operations, you can set the owner option to a specific user name, and this user can perform the privileged operations.

Table 22-1 CIFS export options (*continued*)

CIFS export option	Definition
group	By default, the current group is the primary group owner of the root directory of the exported share. This lets CIFS clients create folders and files in the share. However, there are some operations that require group privileges; for example, changing the group itself, and changing permissions of the top-level folder (that is, the root directory in UNIX terms). To enable these operations, you can set the <code>group</code> option to a specific group name, and this group can perform the privileged operations.
ip	Access Appliance lets you specify a virtual IP address. If you set <code>ip=virtualip</code>, the share is located on the specified virtual IP address. This address must be part of the Access Appliance cluster, and is used by the system to serve the share internally. Note: <code>ip</code> is not a valid CIFS option when using the <code>ctdb</code> clustering mode. See “About CIFS clustering modes” on page 151.
max_connections	Specify the maximum limit for concurrent CIFS connections for a CIFS share. The default value is 0, indicating that there are no limited connections.
shadow_copy	Indicates that this is a <code>shadow_copy</code> capable CIFS share. See “Making a CIFS share shadow copy aware” on page 313.
enable_encryption	If <code>enable_encryption</code> is set, then all the traffic to a share must be encrypted once the connection has been made to the share. The server will return an <code>access denied</code> message to all unencrypted requests on such a share. As SMB3 is the max protocol, only SMB3 clients supporting encryption will be able to connect to the share.
disable_encryption	If <code>disable_encryption</code> is set, then encryption cannot be negotiated by the client. SMB1, SMB2, and SMB3 clients can connect to the share.
enable_durable_handles	Enables support for durable handles for CIFS shares. Enabling this option disables use of POSIX/fcntl locks. Exporting the same CIFS share using NFS may result in data corruption. For support for durable handles on CIFS shares, you must specify this option.

Table 22-1 CIFS export options (continued)

CIFS export option	Definition
allow	Gives allow access to mentioned local or AD users/groups only. If this option is not specified, all the existing users within the domain group are allowed implicitly.
deny	Restricts access to mentioned local or AD users/groups.

Setting share properties

After a file system is exported as a CIFS share, you can change one or more share options. This is done using the same `share add` command, giving the name of an existing share and the name of the file system exported with this share. Access Appliance will realize the given share has already been exported and that it is only required to change the values of the share options.

For example, to export the file system `fs1` with the name `share1`, enter the following:

```
CIFS> share add fs1 share1 "owner=administrator,group=domain users,rw"
CIFS> share show
```

To export a file system

- ◆ Export a file system, enter the following:

```
CIFS> share add filesystem sharename \
[@virtual_ip] [cifsoptions]
```

filesystem	A Access Appliance file system that you want to export as a CIFS share. The given file system must not be currently used for storing the home directory shares. The file system or directory path should always start with the file system name, not with the file system mount point <code>/vx</code> .
sharename	The name for the newly-exported share. Names of the Access Appliance shares can consist of the following characters: lower and uppercase letters "a" - "z" and "A" - "Z," numbers "0" - "9" and special characters: "_" and "-". ("-" cannot be used as the first character in a share name). Note: A share name cannot exceed 256 characters.

@virtual_ip Specifies an optional full identifier allowing a virtual IP to access the specified CIFS share.

Access Appliance provides unified access to all shares through virtual IPs.

Shares of this kind are called segregated shares. Their share name is of the type `share@vip`

If a CIFS share is added with the `@virtual_ip` full identifier, the CIFS share is created by allowing only this virtual IP to access this CIFS share.

```
CIFS> share show
```

cifsoptions A comma-separated list of CIFS export options. This part of the command is optional.

If a CIFS export option is not provided, Access Appliance uses the default value.

See [“About the CIFS export options”](#) on page 302.

For example, an existing file system called `FSA` being exported as a share called `ABC`:

```
CIFS> share add FSA ABC rw,guest,owner=john,group=abcdev
```

Displaying CIFS share properties

To display share properties

- 1 To display the information about all of the exported shares, enter the following:

```
CIFS> share show
```

- 2 To display the information about one specific share, enter the following:

```
CIFS> share show sharename
```

Hiding system files when adding a CIFS normal share

When adding a CIFS normal share, the default is to display the system files (lost+found, quotas, quotas.grp). To hide the system files, you must use the `veto_sys_files` CIFS export option.

See [“About the CIFS export options”](#) on page 302.

To hide system files when adding a CIFS normal share

- ◆ To hide system files when adding a CIFS normal share, enter the following:

```
CIFS> share add filesystem  

sharename [cifsoption]
```

Use the `veto_sys_files` CIFS export option to hide system files.

Allowing specified users and groups access to the CIFS share

To allow specified users and groups access to the CIFS share

- ◆ To allow specified users and groups access to the CIFS share, enter the following:

```
CIFS> share allow sharename
      @group1 \
      [, @group2, user1, user2, ...]
```

sharename	Name of the CIFS share for which you want to allow specified users and groups access. Names of the Access Appliance shares are non case sensitive and can consist of the following characters: lower and uppercase letters "a" - "z" and "A" - "Z," numbers "0" - "9" and special characters: "_" and "-". ("-", cannot be used as the first character in a share name).
group	If the CIFS server joined a domain, and there is a space in the user or group name, the user or group name needs to be entered with double quotes (for example, "@domain users"). By default, all groups are allowed to access the shares. In the case where a CIFS share has joined a domain, and the domain contains trusted domains, and <code>allow_trusted_domains</code> is set to <code>yes</code> on the CIFS server, if you want to allow/deny users or groups from the trusted domains, the user or group needs to be prefixed with the trusted domain name. Separate the domain and user/group with a double backslash. For example: <pre>CIFS> share allow sharename "@domain name\\group name"</pre>
user	Name of the CIFS user allowed access to the CIFS share. By default, all users are allowed to access the shares.

If `all` is specified, then default access restrictions are restored on the CIFS share.

```
CIFS> share allow share1 user1, @group1
```

Denying specified users and groups access to the CIFS share

To deny specified users and groups access to the CIFS share

- ◆ To deny specified users and groups access to the CIFS share, enter the following:

```
CIFS> share deny sharename \  
@group1[,@group2,user1,user2,...]
```

<i>sharename</i>	Name of the CIFS share for which you want to deny specified users and groups access. Names of the Access Appliance shares are non case sensitive and can consist of the following characters: lower and uppercase letters "a" - "z" and "A" - "Z," numbers "0" - "9" and special characters: "_" and "-". ("-", cannot be used as the first character in a share name).
<i>group</i>	If the CIFS server joined a domain, and there is a space in the user or group name, the user or group name needs to be entered with double quotes (for example, "@domain users"). By default, all groups are allowed to access the shares. In the case where a CIFS share has joined a domain, and the domain contains trusted domains, and CIFS is set to trusted domains as true, if you want to allow/deny users or groups from the trusted domains, the user or group needs to be prefixed with the trusted domain name. Separate the domain and user/group with a double backslash. For example: <pre>CIFS> share deny sharename "@domain name\\user name"</pre>
<i>user</i>	Name of the CIFS user denied access to the CIFS share. By default, all users are allowed to access the shares.

If *all* is specified, then all the users and groups are not able to access the share.

```
CIFS> share deny share1 user1,@group1
```

Exporting a CIFS snapshot

To export a CIFS snapshot

- 1 To create a CIFS snapshot, enter the following for example:

```
Storage> snapshot create cf11sp1 CF11
```

See [“About snapshots”](#) on page 384.

- 2 To export the CIFS snapshot, enter the following for example:

```
CIFS> share add CF11:cf11sp1 cf11sp1 rw,guest
```

A client can access the CIFS snapshot by the CIFS share name, `cf11sp1`.

Deleting a CIFS share

To delete a CIFS share

- 1 To delete a share, enter the following:

```
CIFS> share delete sharename [@virtual_ip]
```

sharename Specifies the name of the share that you want to delete.

@virtual_ip Specifies an optional full identifier allowing a virtual IP to access the specified CIFS share.

For example:

```
CIFS> share delete share1
```

- 2 To confirm the share is no longer exported, enter the following:

```
CIFS> share show
```

In the case of any remanent sessions (sessions that are not closed while deleting a CIFS share), Access Appliance displays the following output:

```
CIFS> share delete share2
```

The following remanent sessions are present:

```
pid nodename
19293 clust_01
```

Clients may still access `share2` unless the relevant processes are killed.

This is a rare situation, and it occurs if the following conditions are met:

- CIFS server is online
- CIFS share that is being deleted is ONLINE
- There are some existing client connections with that CIFS share
- While deleting the share, some remanent sessions are left

If any of the conditions fail, then the `CIFS> share delete` command output displays as usual.

Modifying a CIFS share

You can re-export the file system with the given share name. The new options are updated after the command is run.

To modify a CIFS share

- ◆ To modify a CIFS share, enter the following:

```
CIFS> share modify sharename[@virtual_ip] [cifsoptions]
```

sharename	The name of the CIFS share that you want to modify. Names of the Access Appliance shares can consist of the following characters: lower and uppercase letters "a" - "z" and "A" - "Z," numbers "0" - "9" and special characters: "_" and "." ("-" cannot be used as the first character in a share name).
@virtual_ip	Specifies an optional full identifier allowing a virtual IP to access the specified CIFS share. Access Appliance provides unified access to all shares through virtual IPs.
cifsoptions	A comma-separated list of CIFS export options. If a CIFS export option is not provided, Access Appliance uses the default value. See “About the CIFS export options” on page 302.

For example:

```
CIFS> share modify share2 ro,full_acl
```

```
CIFS> share show
```

Making a CIFS share shadow copy aware

Shadow Copy (Volume Snapshot Service or Volume Shadow Copy Service or VSS) is a technology included in Microsoft Windows that allows taking manual or automatic backup copies or snapshots of data on a specific volume at a specific point in time over regular intervals.

To make a CIFS share shadow copy aware

- ◆ Add the CIFS export option `shadow_copy` to the CIFS share.

For example:

```
CIFS> share add fs1 share1 rw,shadow_copy
```

```
CIFS> share show share1
```

See [“About the CIFS export options”](#) on page 302.

About managing CIFS shares for Enterprise Vault

You can create a CIFS share to provision the storage for Enterprise Vault.

See [“Creating a CIFS share for Enterprise Vault with replication”](#) on page 313.

See [“Creating a CIFS share for Enterprise Vault without replication”](#) on page 314.

Creating a CIFS share for Enterprise Vault with replication

You can create a CIFS share for Enterprise Vault with replication.

To create a CIFS share for Enterprise Vault with replication

1 Prerequisites:

Configure storage.

Start the CIFS server.

Configure episodic replication.

Activate an archival policy.

2 Quick Actions: Go to **Provision Storage > Storage for Enterprise Vault > With Replication**, and then click **Next**.

3 Click the **Configure Storage** tab and specify the storage options.

4 Click the **Start CIFS Server** tab, and then click **Start CIFS Server** to start the CIFS server.

- 5 Click **Configure Replication** > **Go to Replication** link to set up the replication virtual IP on the source cluster.
 Go to **Settings** > **Replication management** on destination cluster's GUI and set up the episodic replication virtual IP on the destination cluster.
- 6 Click **Set Up Replication Link**.
- 7 Click the **Activate Archival Policy** tab, select a policy for the CIFS share, and then click **Activate Policy**.
- 8 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the CIFS share.

- 9 View the **Recent Activity** panel for the status of the task.
- See [“Creating a CIFS share for Enterprise Vault without replication”](#) on page 314.

Creating a CIFS share for Enterprise Vault without replication

You can create a CIFS share for Enterprise Vault without replication.

To create a CIFS share for Enterprise Vault without replication

- 1 **Prerequisites:**
 Configure storage.
 Start the CIFS server.
 Activate an archival policy.
- 2 **Quick Actions:** Go to **Provision Storage** > **Storage for Enterprise Vault** > **Without Replication**, and then click **Next**.
- 3 Click the **Configure Storage** tab and specify the storage options.
- 4 Click the **Start CIFS Server** tab, and then click **Start CIFS Server** to start the CIFS server.
- 5 Click the **Activate Archival Policy** tab, select a policy for the CIFS share, and then click **Activate Policy**.

- 6 Click the **Provision Storage** tab, enter the required details for shares and its options and then click on **Provision Storage**.

Note: If a storage policy is not activated, you cannot provision storage for the CIFS share.

- 7 View the **Recent Activity** panel for the status of the task.

See [“Creating a CIFS share for Enterprise Vault with replication”](#) on page 313.

Integrating Access Appliance with Data Insight

This chapter includes the following topics:

- [Access Appliance integration with Data Insight](#)

Access Appliance integration with Data Insight

The integration of Access Appliance with Data Insight enables organizations to improve data governance through insights into the ownership and usage of unstructured data, including files such as documents, spreadsheets, and emails.


Access Appliance allows DI to display NFS share-level activities on the DI GUI and thus helps the users to make better business decisions and risk-management strategy.

The integration of Access Appliance with Data Insight has the following functionalities:

- You can monitor NFS share activities and get details such as who uses the data, who owns the data and who has access to the data.
- You can gather audit information on operations such as create, delete, read, write, make directories, remove directories, and rename type.
- You have full visibility into data access, which helps drive security remediation, and auditing and compliance efforts using the web-based graphical user interface (GUI).

- You can automate the movement of stale, less accessed, and other unimportant data to cloud or cheaper storage alternative with the help of the DI reporting and Access Appliance policy engine.
- You can have a cluster-level view of audit logs. Access Appliance has multiple nodes, so you can access the same share or even the same file from multiple nodes in the cluster. You can merge the audit logs from different nodes to give a unified view to DI.
- You can monitor data for any potential security breaches or internal misuse of information. You can ensure that proper document protocols are followed consistently, and you can also prevent and track down fraud.
- You can enhance your risk-management capabilities by running reports on a given end user and examine their activity. You can see which records have been accessed, edited, added, or deleted. You can also get details on deleted records like when it was deleted, who deleted it, and so on. You can track what goes on in the database.

Figure 23-1 High-level design



Access Appliance helps the DI administrator to:

- Add Access Appliance in the DI GUI.
- Enable and disable auditing on a given NFS share.
- Mount and scan the share as and when required.
- Fetch the audit logs periodically from Access Appliance.
- Allow showing the audit statistics on the DI console.

Prerequisites to configuring Access Appliance in the Data Insight GUI

- An NFS server should be running on all the nodes.
- `xprtld` (Access Appliance UI REST server) should be running on all the nodes.
- `sfsdg` disk group should be present

Configuring Access Appliance in the Data Insight GUI

- You can add Access Appliance as one of the file system servers.
- You can administer Access Appliance from the DI GUI and perform the following functions:
 - Scan the Access Appliance share from DI. For NFS scanning, the mount and scan approach is used.
 - Enable and disable audit logs for the share.
 - List file systems on which auditing is enabled.

You can get the information (path, permissions, and other details) about the shares. The Veritas File System Information Management Infrastructure (IMI) is used for collecting audit logs. The audit logs are created for each node and merged to give a cluster view.

- DI can collect the audit logs periodically. `xprtld` (Access Appliance UI REST server) is used for transferring the logs. The audit logs of Access Appliance are shared with DI using the pull model.

Figure 23-2 Access Appliance integration with Data Insight

Veritas™ Data Insight Workspace Policies Reports Workflows **Settings**

Filter... System Overview x Filers x Add New Ver x

Health & Monitoring

- System Overview
- Scan Status
- Events
- Performance
- Installation Status

Inventory

- Data Insight Servers
- Filers
- SharePoint Web Applications
- Cloud Sources
- Directory Services
- Containers
- Data Insight Users
- Saved Credentials

Remediation

- Permissions
- Data Management
- Action Status

Classification

- Configuration
- Requests
- Policy Manager

Global Settings

- SMTP Settings
- Advanced Analytics
- Exclude Rules
- Scanning and Event Monitoring
- Event Notifications
- Watchlist Settings
- Data Retention
- Data Loss Prevention
- File Groups
- Workspace Data Owner Policy
- Custodian Manager
- Upload Manager
- Console Settings

Add New Veritas File System Server

Connection Details

☒ This is a VCS Clustered File Server

VCS Cluster Name:

Cluster Node IP Addresses:

Collector: [Select Collector](#)

Indexer: [Select Indexer](#)

Domain: ☒ Filers do not belong to a domain

Login credentials: [Test Credentials](#)

Share Configuration

☒ Discover shares automatically (Discovery for new shares will happen twice a day)

Exclude following shares from discovery (e.g. "\$\$,snap",tmpshare)

Event Monitoring

☒ Enable File System Event Monitoring

Time to live: [?](#)

Records per file: [?](#)

File System Scanning

☒ Enable Filers Scanning

☒ Scan newly added shares immediately [?](#)

Scanner credentials: [Test Credentials](#)

Scanning Schedule (Full Scan)


☒ Use collector's default scanning schedule ☐ Use custom schedule

Default scan schedule 7pm last friday of each month

[Save](#) [Cancel](#)

Logged in User: Administrator@SAMGWIN

Figure 23-3 Audit information



Limitation

- Only NFS shares scanning is supported for this release.
- In case of a node restart, monitoring of shares does not work for that node.

Managing Access Appliance storage services

- [Chapter 24. Configuring episodic replication](#)
- [Chapter 25. Configuring continuous replication](#)
- [Chapter 26. Using snapshots](#)
- [Chapter 27. Using instant rollbacks](#)

Configuring episodic replication

This chapter includes the following topics:

- [About Access Appliance episodic replication](#)
- [How Access Appliance Replication works](#)
- [Starting Access Appliance episodic replication](#)
- [Setting up communication between the source and the destination clusters](#)
- [Setting up the file systems to replicate](#)
- [Setting up files to exclude from an episodic replication unit](#)
- [Scheduling the episodic replication](#)
- [Defining what to replicate](#)
- [About the maximum number of parallel episodic replication jobs](#)
- [Managing an episodic replication job](#)
- [Replicating compressed data](#)
- [Displaying episodic replication job information and status](#)
- [Synchronizing an episodic replication job](#)
- [Behavior of the file systems on the episodic replication destination target](#)
- [Accessing file systems configured as episodic replication destinations](#)
- [Episodic replication job failover and fallback](#)

About Access Appliance episodic replication

The Access Appliance episodic replication solution provides high performance, scalable data replication and is ideal for use as a content distribution solution, and for use to create hot standby copies of important data sets.

Access Appliance episodic replication lets you asynchronously replicate a file system from one node in a source cluster to another node in a destination cluster at regularly timed intervals. This allows for content sharing, replication, and distribution.

The Access Appliance episodic replication functionality allows episodic replication with a minimum timed interval update of 15 minutes and no set maximum. Unlike many replication solutions, Access Appliance episodic replication also allows the destination file system to be online for reads while replication is active.

Major features of Access Appliance episodic replication include:


- Online access (read-only) to replicated data.
- Immediate read/write access to destination-replicated data in the unlikely event that the source file system goes offline for a sustained period of time.
- Load balancing across replication links.
- Transport failover of episodic replication service from one node to another.
- No limit on the number of episodic replication jobs that are configured, though the number of simultaneous/parallel jobs that can run at any time depends on the amount of memory available.

The Access Appliance episodic replication feature is designed to copy file systems only between Access Appliance clusters.

Note: The Access Appliance episodic replication feature does not support user modifications to the target file system if episodic replication is configured.

[Figure 24-1](#) describes the workflow for configuring episodic replication between two Access Appliance clusters.

Figure 24-1 Episodic replication workflow



How Access Appliance Replication works

Access Appliance Replication is an incremental file-level replication service that runs on top of the Cluster File System that is used by Access Appliance which is, in turn, based on the Veritas File System (VxFS). Access Appliance Replication uses two file system specific features: File Change Log (FCL) and Storage Checkpoint services, to retrieve file changes between replication periods.

For a given period, the FCL records every change that is made to the file system. By scanning the FCL, Access Appliance Replication quickly identifies the file(s) that have changed and generates the modified file list. This list prevents the expensive file system scanning that is normally associated with file-based replication, and which typically results in sub-optimal performance.

Next, Access Appliance Replication uses VxFS Storage Checkpoint's metadata comparison feature to retrieve the modified extent list of each changed file. It does not need to access the file data. Access Appliance Replication replicates files data, metadata, access control lists (ACLs), extended attributes, and modification timestamp changes. It does not replicate any changes made to access time and change time.

The Access Appliance Replication transport layer works with and interfaces to the well-known rsync remote file synchronization tool. Using this existing network transportation program makes the network configuration much easier in the enterprise domain: the Secure Socket Shell (SSH) port (22) required by rsync is opened by default on almost all enterprise firewalls. rsync is also a reliable solution for a low bandwidth or unreliable link environment.

Note: Access Appliance uses the rsync protocol to provide transportation of Access Appliance Replication encapsulated files. The use of rsync is not exposed in Access Appliance, and cannot be administered outside of the Access Appliance Replication feature set.

Starting Access Appliance episodic replication

This section lists the specific commands that are needed to run Access Appliance episodic replication on your clusters.

Ensure the following before starting episodic replication:

- Before you set up your clusters for episodic replication, you must first identify which is the source cluster and which is the destination cluster. All of the commands are performed on the source cluster first.
- Make sure both the source cluster and the destination cluster have the same version of Access Appliance.
- To use Access Appliance episodic replication, you must first create an online file system on the Access Appliance source cluster and an online file system on the Access Appliance destination cluster.
- Assign a virtual IP (VIP) address to both the source and the destination clusters. The Access Appliance episodic replication service requires VIP addresses not already in use for the two clusters to communicate.

To start Access Appliance episodic replication on the source cluster

- 1 To bind a virtual IP address for the episodic replication service on the source cluster, enter the following:

```
Replication> episodic config bind ip_addr [device] [netmask]
```

ip_addr Virtual IP address for the episodic replication service on the source cluster.

device The public network interface name that you want the episodic replication IP address to use.

netmask Netmask for the episodic replication IP address.

- 2 To start the episodic replication service, enter the following on the source node:

```
Replication> episodic service start [nodename]
```

nodename The name of the node in the local cluster where you want to start the episodic replication service.

- 3 To check the status of the episodic replication service, enter the following:

```
Replication> episodic service status
```

- 4 To confirm that the IP address is up and running, enter the following:

```
Replication> episodic config show ip
```

The definitions of the headings are as follows:

Note: Alternately, you can use the `Network> ip addr show` command to confirm that the IP address is up and running.

To start Access Appliance episodic replication on the destination cluster

- 1 To bind a virtual IP address for the episodic replication service on the destination cluster, enter the following:

```
Replication> episodic config bind ip_addr [device] [netmask]
```

ip_addr Virtual IP address for the episodic replication service on the destination cluster.

device The public network interface name that you want the episodic replication IP address to use.

netmask Netmask for the episodic replication IP address.

- 2 To start the episodic replication service, enter the following on the destination node:

```
Replication> episodic service start [nodename]
```

nodename The name of the node in the local cluster where you want to start the episodic replication service.

- 3 To check the status of the episodic replication service, enter the following:

```
Replication> episodic service status
```

- 4 To confirm that the IP address is up and running, enter the following:

```
Replication> episodic config show ip
```

You next need to set up communication between the source and the destination clusters.

Setting up communication between the source and the destination clusters

You need to set up communication between your source and your destination clusters.

Make sure that you already created an online file system on the Access Appliance source cluster and an online file system on the Access Appliance destination cluster.

Access Appliance episodic replication authentication strategy is based on RSA-key authentication, and both the source and the destination clusters have to export their

episodic replication public keys. The source cluster imports the destination cluster's public key and the destination cluster imports the source cluster's public key.

After you have determined which two Access Appliance clusters to use, you need to authenticate them.

The `replication episodic config` commands must be run in a specific order.

- Use the `replication episodic config del_keys` after the `replication episodic config deauth` command, or it fails.
- You can only run the `replication episodic config unbind` command (to unbind the virtual IP) after you have run the `replication episodic service stop` command.
- You need to run the `replication episodic config bind` command (to bind the virtual IP) before you can run the `replication episodic service start` command.
- You need to run the `replication episodic config export_keys` and `Replication> episodic config import_keys` to export and import the keys of both the source and the destination clusters.
- You can only run the `replication episodic config auth` command after both the source and destination have imported each other's keys.
- You need to run the `replication episodic config auth` command to create a link from every cluster to any remaining cluster that is used for episodic replication irrespective of their role as a source or a destination cluster.

After the source and the destination clusters have successfully imported each other's public keys, you need to run the `Replication> episodic config auth` command on the source cluster to complete the authentication between the two clusters. This command checks the two-way communication between the source and the destination cluster, and authenticates the clusters allowing the Access Appliance episodic replication service to begin.

Note: The `replication episodic config auth` command must be run from the source cluster.

This section provides a walk-through for the creation and export/import of these encrypted keys for both the source and the destination cluster.

Note: Without the correct authentication of the source and the destination encryption keys, Access Appliance episodic replication does not function correctly.

To export the source cluster's key to the destination cluster

- 1 To export the source cluster's key to the destination cluster, enter the following:

```
Replication> episodic config export_keys [URL]
```

URL The location you want to copy the public keys to.

If you do not want to enter a URL, you can copy the output from the `Replication> episodic config export_keys` command into the `replication episodic config import_keys` command at the destination cluster.

By default, the output is displayed to your computer screen.

The SCP protocol is supported.

- 2 To import the source cluster's key to the destination cluster, enter the following:

```
Replication> episodic config import_keys [URL/keyfile]
```

URL The location you want to copy the public keys from.

keyfile The file name of the key that is generated by the export.

If you did not enter a URL during the `replication episodic config export_keys` command, you can cut and paste the output and enter it into the `replication episodic config import_keys` command.

- 3 To verify that the key has been imported correctly, enter the following:

```
Replication> episodic config show
```

To export the destination cluster's key to the source cluster

- 1 To export the destination cluster's key to the source cluster, enter the following:

```
Replication> episodic config export_keys [URL]
```

URL The location you want to copy the public keys to.

The SCP protocol is supported.

If you do not want to enter a URL, you can cut and paste the output from the `Replication> episodic config export_keys` command to the `Replication> episodic config import_keys` command. By default, the output is displayed to your computer screen.

- 2 To import the destination cluster's key to the source cluster, enter the following:

```
Replication> episodic config import_keys [URL/keyfile]
```

URL Enter the URL of the location you want to copy the public keys from.

keyfile Enter the file name of the key that is generated by the export.

If you did not enter a URL during the `replication episodic config export_keys` command, you can cut and paste the output and enter it into the `replication episodic config import_keys` command.

- 3 To verify that the key has been imported correctly, enter the following:

```
Replication> episodic config show
```

To authenticate source cluster and destination clusters for episodic replication

- 1 This command should be run on the source cluster as well as on the destination cluster. To authenticate the public keys on the source cluster and the destination clusters, enter the following:

```
Replication> episodic config auth conIP link_name
```

conIP Enter the destination cluster console IP address.

link_name Both the source cluster and the destination cluster need to be assigned a unique identifier (name). This identifier is used to identify the link that is established between the source and the destination clusters. You can use the link name instead of the virtual IP addresses of the source and the destination clusters when using the other episodic replication commands. For example:
Pune_Shanghai.

- 2 To confirm the authentication, enter the following:

```
Replication> episodic config show
```

Note: These steps must be run on the destination side cluster to authenticate the public keys on the source and the destination cluster.

Once you have configured the clusters and links, you need to set up the file systems you want to replicate.

Setting up the file systems to replicate

You need to set up the file systems you want to replicate using the `Replication> episodic repunit` commands. The `Replication> episodic repunit` commands let you define the type of data that you replicate from the source cluster to the destination cluster. All files and folders belonging to an episodic replication unit are replicated together from the source cluster to the destination cluster.

Note: The maximum number of episodic replication units supported in Access Appliance episodic replication is 128.

Make sure that you already set up communication between your source and the destination clusters.

See [“Setting up communication between the source and the destination clusters”](#) on page 328.

An episodic replication unit is defined as an ordered set of entries, where each entry is one of the following:

- A single file system
- A single subdirectory
- A single file

Note: The episodic replication source has to be one of the entry types shown. It cannot be a snapshot or a Storage Checkpoint (ckpt).

Access Appliance episodic replication requires that the source and the destination episodic replication units of a job definition have the same type of ordered entries, that is, every entry pair (one entry from the source and one entry from the destination episodic replication unit) must be of a similar type.

Both can be files, or both can be directories, as shown in the following example:

Replication unit Name	Replication unit Entries
=====	=====
ru1	fs1/dir1, fs1/dir1
ru2	fs2/f1, fs2/f2

The entry is identified by the file system name, optionally followed by a slash '/', followed by the path of the directory or the file inside the file system. Member entries are ordered inside an episodic replication unit and such ordering information is used to determine the episodic replication entity pair mapping from the source episodic replication unit to the destination episodic replication unit.

Note: Make sure that the paths in the destination episodic replication unit exist in the destination cluster.

Note: The commands in this section apply only to the source episodic replication unit.

To create an episodic replication unit

- 1 From the source cluster, to create an episodic replication unit, enter the following:

```
Replication> episodic repunit create repunit_name
            repunit_entry[,repunit_entry,...]
```

repunit_name The name of the episodic replication unit you want to create.

repunit_entry The file system file, file, folder, or directory.

Note: Destination episodic replication units should be created only at the source cluster using the `replication episodic repunit create` command.

- 2 To confirm the creation of the episodic replication unit, enter the following:

```
Replication> episodic repunit show verbose
```

You can use the `replication episodic repunit add_entry`, `replication episodic repunit modify_entry`, `replication episodic repunit remove_entry`, and `replication episodic repunit destroy` commands to manage your episodic replication units.

Note: The above-mentioned commands are restricted for a job that is enabled with a replication unit containing a WORM-enabled file system with user-defined data.

Note: The `replication episodic repunit destroy` operation is not allowed for the episodic replication units that are included in any job definitions.

Setting up files to exclude from an episodic replication unit

Once you have set up the file systems you want to replicate, you can define a set of directories or files to exclude from an episodic replication unit. This step is optional. The `exclunit` entry has higher priority over the `repunit` entry. If any file name matches the `exclunit` entry, the file is not replicated to the target.

To work with exclusion units:

- Use the `Replication> episodic exclunit create` command to name the excluding unit and configure the directories and files you want to exclude from an episodic replication. The excluding unit you create can be used in multiple episodic replication jobs. A single excluding unit can span across multiple directories.
- Use the `Replication> episodic job exclude` command to add the excluding unit to an episodic replication job. You cannot add an excluding unit to a job that is active. You must disable the job first.
- You can use the following commands: `Replication> episodic exclunit add_entry`, `Replication> episodic exclunit modify_entry`, and `Replication> episodic exclunit remove_entry` to make changes to an excluding unit, provided the excluding unit you want to modify is not included in any job definitions.
- Use the `Replication> episodic job show` command to show which excluding units are configured for a job. Use the `Replication> episodic exclunit show` command to show the names and contents of all excluding units that are defined for the cluster.
- Use the `Replication> episodic exclunit destroy` command to permanently delete the excluding unit. You can only destroy an excluding unit if the excluding unit you want to destroy is not included in any job definitions.

If an episodic replication is defined for a directory, an excluding unit should be a subset of that directory. The excluding unit cannot be the same directory as the episodic replication and it cannot be a parent directory of the episodic replication. For example, if a episodic replication is configured for `fs1/dir1/dir2`, a valid exclusion could be `dir1/dir2/file` or `dir1/dir2/dir3`, but not `/dir1` (the parent directory for the episodic replication).

By default, Access Appliance excludes some common directories and files from all episodic replication units. These directories and files include:

- `lost+found`
- `.placement_policy.xml`
- `quotas`
- `quotas.grp`
- `quotas.64`
- `quotas.grp.64`

In addition, you can use the `Replication> episodic exclunit` commands to specify additional directories and files to exclude.

The directories and files you specify for an excluding unit are applied based on the overall definition of the episodic replication. For example, an episodic replication job that contains an `fs1` episodic replication unit and an `dir3` excluding unit, replicates all the files in `fs1`, except for the files in `fs1/dir3`.

To create an excluding unit:

- 1 To create an excluding unit, enter the following:

```
Replication> episodic exclunit create exclunit_name
exclunit_entry [exclunit_entry,...]
```

exclunit_name Enter the name of the excluding unit.

exclunit_entry Enter the comma-separated list of directories and files you want to exclude from an episodic replication.

- 2 To confirm the creation of the excluding unit enter the following:

```
Replication> episodic exclunit show verbose
```

You can use the `Replication> episodic exclunit add_entry`, `Replication> episodic exclunit modify_entry`, `Replication> episodic exclunit remove_entry`, and `Replication> episodic exclunit destroy` commands to manage your excluding units.

Note: The `Replication> episodic exclunit add_entry`, `Replication> episodic exclunit modify_entry`, `Replication> episodic exclunit remove_entry`, and `Replication> episodic exclunit destroy` operations are not allowed for excluding units that are included in any job definitions.

Scheduling the episodic replication

You use the `Replication> episodic schedule` commands to create a schedule for replicating files from the source to the destination cluster.

Access Appliance episodic replication supports periodic replications, where the data gets replicated from the source to the destination cluster at regular intervals as defined by the schedule. Access Appliance episodic replication uses the following parameters to schedule the episodic replication jobs: minute, hour, day-of-the-month, month, and day-of-the-week.

Make sure that you already set up the file systems you want to replicate.

See [“Setting up the file systems to replicate”](#) on page 332.

To create an episodic replication schedule

- ◆ To create an episodic replication schedule, enter the following:

```
Replication> episodic schedule create schedule_name minute
[hour] [day_of_the_month] [month] [day_of_the_week]
```

<i>schedule_name</i>	Specify the name of the schedule to be created.
<i>minute</i>	Enter a numeric value between 0-59, or an asterisk (*), which represents every minute. This variable is not optional.
<i>hour</i>	Enter a numeric value between 0-23, or an asterisk (*), which represents every hour.
<i>day_of_the_month</i>	Schedule the day of the month you want to run the replication. Enter a numeric value between 1-31, or an asterisk (*), which represents every day of the month.
<i>month</i>	Schedule the month you want to run the replication. Enter a numeric value between 1-12, or an asterisk (*), which represents every month. You can also use the names of the month. Enter the first three letters of the month (not case sensitive).
<i>day_of_the_week</i>	Schedule the day of the week you want to run the replication. Enter a numeric value between 0-6, or an asterisk (*), which represents every day of the week. Sunday is interpreted as 0. You can also enter the first three letters of the week (you must use lower case letters).

You can enter an interval (two numbers separated by a hyphen) for the *minute*, *hour*, *day-of-month*, *month*, and *day-of-week*. If you want to run the schedule between 1:00 a.m. and 4:00 a.m., you can enter a value of 1-4 for the hour variable. The range is inclusive.

The parameters also accept a set of numbers separated by a comma. For example, 1,3,5,7 or 1-4,5-10.

To display the list of schedules

- ◆ To display the schedule you have set up for episodic replication, enter the following:

```
Replication> episodic schedule show
```

You can also use the `Replication> episodic schedule modify` and `Replication> episodic schedule delete` to manage your episodic replication schedules.

Note: The `Replication> episodic schedule modify` and `Replication> episodic schedule delete` operations are not allowed for the schedules that are included in any job definition.

You next need to define what is replicated.

See [“Defining what to replicate”](#) on page 338.

Defining what to replicate

You use the `Replication> episodic job` commands to set up a job definition. This defined job determines what to replicate and when, using the settings from the previous commands.

Make sure that you created a schedule for replicating files from the source to the destination cluster.

See [“Scheduling the episodic replication”](#) on page 336.

To set up the episodic replication job

- 1 To create an episodic replication job, enter the following:

```
Replication> episodic job create job_name src_repunit tgt_repunit  
link_name schedule_name [evpsn] [metadata_only]
```

<i>job_name</i>	Specify a name for the episodic replication job you want to create.
<i>src_repunit</i>	Specify the source episodic replication unit. The episodic replication unit determines the exact item (such as a file system) that you want to replicate.
<i>tgt_repunit</i>	Specify target episodic replication units.
<i>link_name</i>	Specify the link name used when you ran the <code>Replication> episodic config auth</code> command between the local cluster and the remote cluster. Both the source cluster and the destination cluster need to be assigned a unique identifier (name). This identifier is used to identify the link that is established between the source and the destination clusters. You can use the link name instead of the virtual IP addresses of the source and the destination clusters when using the other episodic replication commands.
<i>schedule_name</i>	Specify the name of the episodic replication schedule you want to apply to the episodic replication job.
<i>evpsn</i>	Enable or disable Enterprise Vault partition secure notifications.
<i>metadata_only</i>	Enable or disable metadata-only episodic replication. This feature is not supported with consistency groups and tunables.

Note: If you configure episodic replication on a WORM-enabled file system, the job configuration fails if either the source or destination clusters are configured with Normal lockdown mode. It is recommended that the source and destination clusters should be configured with the same lockdown modes.

- 2 To add an excluding unit to the job, enter the following command. This step is optional.

```
Replication> episodic job exclude job_name exclunit_name
```

- 3 By default, the job is disabled. To enable the job, enter the following:

```
Replication> episodic job enable job_name
```

- 4 To check if the job was enabled, enter the following:

```
Replication> episodic job show [job_name]
```

About the maximum number of parallel episodic replication jobs

The maximum number of episodic replication jobs is 64, but there are stricter limits on the number of episodic replication jobs that can be running in parallel at the same time. Episodic replication uses a RAM-based file system for storing the transit messages. Each GB of this RAM-based file system can accommodate up to eight parallel running jobs. The default size of this file system depends upon the amount of physical memory of the node on which episodic replication is running. If the physical memory is less than 5 GB, episodic replication limits its maximum usage for storing messages to 1 GB of memory, which means the user can run up to eight episodic replication jobs in parallel at the same time. If the physical memory is between 5 GB to 10 GB, episodic replication limits its maximum usage for storing messages to 2 GB of memory, which means you can run up to 16 episodic replication jobs in parallel. If the physical memory is greater than 10 GB, episodic replication limits its maximum usage for storing messages to 4 GB of memory, which means you can run up to 32 episodic replication jobs in parallel at the same time.

Managing an episodic replication job

You can manage an episodic replication job using the `replication episodic job` commands. The commands are required only on the source system.

The `replication episodic job enable`, `replication episodic job sync`, `replication episodic job disable`, `replication episodic job abort`, `replication episodic job pause`, and `replication episodic job resume` commands change the status of an existing episodic replication job.

You can use the `replication episodic job modify`, and `replication episodic job destroy` commands to modify or destroy an episodic replication job definition.

The `replication episodic job enable` command starts episodic replication immediately and initiates episodic replication after every subsequent set frequency

interval. When an episodic replication job is created it is disabled by default, and you must enable the job to start episodic replication.

To enable an episodic replication job

- ◆ To enable an episodic replication job, type the following command:

```
Replication> episodic job enable job_name
```

job_name Specify the name of the episodic replication job you want to enable.

At each frequency interval, a fresh file system Storage Checkpoint is taken and episodic replication is started against the new Storage Checkpoint. If a previous episodic replication run has not completed, a new Storage Checkpoint is not taken and the current run is skipped.

Note: Running the `replication episodic job enable` command on a previously aborted episodic replication job automatically restarts the job.

The `Replication episodic job sync` command lets you start an episodic replication job, but then stops the episodic replication job after one iteration (full or incremental) is complete. You can use this command to recover from the secondary site in the event that the primary file system is completely destroyed. This command can also be used if you want to run an episodic replication job at a predefined time using a script or a cron job.

See [“Synchronizing an episodic replication job”](#) on page 346.

The `Replication> episodic job disable` command drops the episodic replication job from the schedule and waits for any already running iterations to complete. The `Replication> episodic job disable` command disables a job definition which is in one of these states: ENABLED, PAUSED, or FAILED. This process can take some time if the network is slow or if a large amount of data has changed since the last episodic replication run.

To disable an episodic replication job

- ◆ To disable an episodic replication job, type the following command:

```
Replication> episodic job disable job_name
```

job_name Specify the name of the episodic replication job you want to stop.

Note: If episodic replication has been configured on a WORM-enabled file system, the `replication episodic job disable` command is restricted unless all the user data on the file system is deleted.

The `replication episodic job abort` command forcefully cancels an episodic replication job even if it is in progress. Aborting an episodic replication job may leave Storage Checkpoints mounted on the source system and the target file system may be left in an intermediate state.

To abort an episodic replication job

- ◆ To abort an episodic replication job, type the following command:

```
Replication> episodic job abort job_name
```

job_name Specify the name of the episodic replication job you want to abort.

Note: If episodic replication has been configured on a WORM-enabled file system, the `replication episodic job abort` command is restricted unless all the user data on the file system is deleted.

The `Replication> episodic job pause` command immediately stops the episodic replication job. You must use the `Replication> episodic job resume` command to resume the episodic replication job from where it was paused. When episodic replication is resumed, the episodic replication job replicates the set of selected files before pausing the job, and attempts to replicate as much of the latest data as possible. This action allows the customer to have two recovery point objectives (RPO). When the episodic replication job is paused, the episodic replication frequency option is disabled. Once the episodic replication job is resumed, the frequency option resumes for subsequent iterations. The pause and the resume functions let you manage the episodic replication job based on workload requirements.

To pause and resume an episodic replication job

- 1 To pause an episodic replication job, type the following command:

```
Replication> episodic job pause job_name
```

Where *job_name* is the name of the episodic replication job you want to pause.

- 2 To resume an episodic replication job, type the following command:

```
Replication> episodic job resume job_name
```

Where *job_name* is the name of the episodic replication job you want to resume.

Note: You cannot start or sync a paused job. You can abort a paused job. However, if synchronization is performed on a paused job that has been aborted, the last RPO for the paused job is not available.

The `replication episodic job modify` command lets you modify debugging or setting tunables on an episodic replication job definition.

The addition or removal of a file system from the source episodic replication unit or the destination episodic replication unit is not supported. To remove a specific file system from the episodic replication unit you must destroy the episodic replication job and recreate the episodic replication job with the new set of file systems in the episodic replication unit. To add a specific file system from an existing episodic replication unit, you can either create a new episodic replication job with a new source episodic replication unit and target episodic replication unit, or destroy the episodic replication job and recreate it with the new set of file systems in the episodic replication unit to use the same job name.

The `replication episodic job modify debug` command lets you enable or disable debugging on a given job.

To modify debugging on an episodic replication job

- ◆ To modify debugging on an episodic replication job definition, enter the following command:

```
Replication> episodic job modify debug job_name on|off
```

job_name Specify the episodic replication job name you want to modify.

The `replication episodic job modify tunables` command lets you modify the job configuration to use multiple network connections (sockets) for replicating data from source to target. In configurations where WAN latency is high, it is

recommended to use multiple connections for significantly increased throughput. After the tunables are set for a job, only one job is supported.

To modify tunables on an episodic replication job

- ◆ To modify tunables on an episodic replication job definition, enter the following command:

```
Replication> episodic job modify tunables job_name netconn rw_count
```

job_name Specify the episodic replication job name you want to modify.

netconn Specify the number of connections.

rw_count Specify the number of threads.

The increased number of connections is effective in case of a relatively small number of large files. For large number of small files, full sync performance may be slower with increased number of connections.

The `replication episodic job destroy` command destroys a job definition. This command completely removes the specified job from the configuration, cleans up any saved job-related statistics, and removes any Storage Checkpoints. The episodic replication job must be disabled before the job definition can be destroyed.

To destroy an episodic replication job definition

- ◆ To destroy a job definition, enter the following command:

```
Replication> episodic job destroy job_name
```

Where *job_name* is the name of the job definition you want to delete. Make sure that the job is not enabled.

Using the `replication episodic job destroy` command with the `force` option removes the local job irrespective of the job state, and all episodic replication units are disassociated from the job. Cluster configurations, which are part of the job, are not modified.

Note: If episodic replication has been configured on a WORM-enabled file system, the `replication episodic job destroy` and `replication episodic job destroy force` commands are restricted unless all the user data on the file system is deleted.

Note: When setting up episodic replication, Veritas does not advise you to make any modifications or deletions on the target side of the file system. In the event that some or all of the target data is modified or deleted, you must re-create the episodic replication job from the source cluster to resume episodic replication services.

To re-create an episodic replication job

- 1 To re-create an episodic replication job, you must first delete the job definition. Enter the following command on the source cluster:

```
Replication> episodic job destroy job_name
```

Where *job_name* is the name of the job definition you want to delete. Make sure that the job is not enabled.

- 2 Re-create the job definition:

```
Replication> episodic job create job_name src_repunit tgt_repunit  
link_name schedule_name [evpsn] [metadata_only]
```

You can reuse the source episodic replication unit, target episodic replication unit, link, and schedule names.

Replicating compressed data

Using the `vxcompress` utility, episodic replication is able to replicate any compressed file that is created at the source to the target, while maintaining the same compression characteristics. The compression characteristics include the algorithm, the strength, and the block size. The data is read in the compressed format from the source, sent over the network, and written to the target system in the same format. This form of compression reduces the amount of storage that is required on the target system.

Note: Compressed files that are created using archive utilities such as `.tar` or `.zip`, are treated as normal data files and not compressed during episodic replication.

Displaying episodic replication job information and status

The `Replication> episodic job show` and `Replication> episodic job status` commands display job definition information, which allows you to confirm any changes that are made to your episodic replication job and view current job status.

The `Replication> episodic job show` command displays single job definition, or all of the job definitions for a destination cluster.

To display the job definitions

- ◆ To display the job definitions, enter the following command:

```
Replication> episodic job show [job_name]
```

job_name Enter the name of the job you want to display. If you want to list all of the job definitions, enter the command without a job name.

The `Replication> episodic job status` command displays the status of one or all of the jobs that are copied during episodic replication and the time the episodic replication occurred.

To display the status of an episodic replication job

- ◆ To display the status of an episodic replication job or all the jobs, enter the following command:

```
Replication> episodic job status job_name
```

job_name Enter the name of the job you want to display status for.

If a job is not specified, all status of all the jobs is displayed.

If the Job State displays `Trying_to_enable`, then the `job enable` command is in progress. Check the job status again after a few minutes.

Synchronizing an episodic replication job

To synchronize an enabled episodic replication job

- ◆ To synchronize an enabled episodic replication job, enter the following:

```
Replication> episodic job sync job_name
```

job_name Specify the name of the episodic replication job you want to synchronize.

Behavior of the file systems on the episodic replication destination target

Destination file systems are mounted as read-write. Read-only access is allowed, but you are not expected to modify the destination file system content. While episodic replication occurs, destination file systems may not be in a consistent state. To provide consistent images of the destination file systems at different stages of the episodic replication, the episodic replication service creates and manages Storage Checkpoints of each destination file system.

The episodic replication service creates a new destination Storage Checkpoint:

- Before the first session (before a full-sync)
- After every successful episodic replication session (after every incremental sync)

Storage Checkpoints are automatically mounted under the `.checkpoint` directory inside the target file system, for example:

```
/vx/target_mount/.checkpoint/ckpt_name
```

Where *target_mount* is the name of the target file system and *ckpt_name* is the name of the Storage Checkpoint.

You can use the `Storage> snapshot list` command to view these Storage Checkpoints and you can use Access Appliance commands to export any of these Storage Checkpoints for read-only purposes. The episodic replication Storage Checkpoint names are prefixed with `vxfsrepl_` and also contain the Storage Checkpoint creation time.

Accessing file systems configured as episodic replication destinations

Destination Storage Checkpoints are automatically mounted and therefore cannot be brought online or taken offline using the `Storage> snapshot` commands. The destination Storage Checkpoints can only be accessed through the `.checkpoint` directory. This accessibility also applies to any user created Storage Checkpoints on the episodic replication destination file system.

Episodic replication job failover and failback

Typically, the source cluster drives an episodic replication session. However, in some situations, it may be useful for the destination cluster to drive the episodic replication session. Access Appliance supports a failover and a failback feature for

episodic replication jobs. This feature enables control of episodic replication jobs to be temporarily relocated from the source cluster to the destination (target) cluster.

Job failover and failback is useful for:

- **Planned failover**
In cases where the source cluster is taken down for routine maintenance or for moving applications to another cluster, a planned failover procedure is available for moving episodic replication jobs from the source cluster to the destination cluster.
- **Unplanned failover or Disaster recovery**
In cases where the source cluster fails unexpectedly, an unplanned failover procedure is available for moving episodic replication jobs to the destination cluster.

Note: In the event of a planned or unplanned failover from the source cluster to the destination cluster, there should be at least one successful sync attempt. The successful sync ensures that a consistent point in time image is present on the destination cluster that can be used for the failover.

With job failover and failback, you use the `Replication> episodic job failover` command to move control from the source cluster to the destination cluster. You use the `Replication> episodic job failback` to restore control to the source cluster. The *link_name* is the link of one of the destination clusters. The *link_name* argument can be empty when the source cluster is not available, in which case the job failover can be run from one of the destination clusters.

Essentially, job failover takes job and episodic replication unit definitions from the episodic replication database on the source cluster and copies them to the episodic replication database on the destination cluster.

Warning: Job failover assumes that all episodic replication job names and episodic replication unit names are unique across all Access Appliance clusters on your network. Before you use the episodic replication failover feature, make sure that these names are unique.

After a job failover or failback, you must manually start or enable the episodic replication job to start preconfigured schedules. Link throttle information should be reconfigured after the job failover or failback.

Job failover does not automatically move the NFS or the CIFS share information that is associated with job failover episodic replication units from the source cluster to the destination cluster. Share information has to be done manually.

Table 24-1 Job failover and failback commands

Command	Definition
<code>episodic job failover</code>	Transfer control of an episodic replication job from the source cluster to the destination cluster.
<code>episodic job failback</code>	Return control of an episodic replication job from the destination cluster to the source cluster.

Process summary

The steps you take for job failover and failback vary depending on the type of failover or failback you perform. Failover and failback types include:

- **Planned failover:** This is performed to switch the job roles of source and destination clusters.
- **Failback after a planned failover:** This is performed after a planned failover to switch the job roles of the clusters back to the original state.
- **Unplanned failover:** This is performed when the primary cluster is down and disaster recovery has to be performed.
- **Failback after an unplanned failover:** This is performed after an unplanned failover or disaster recovery.

Each process is summarized in the following sections. Typically, you would use the planned failover and planned failback processes in most situations.

Overview of the planned failover process

Consider two Access Appliance clusters, cluster A and cluster B. Here, cluster A is the original source cluster that replicates data to cluster B which is the original destination cluster. If the source cluster A is taken down for routine maintenance or for moving applications to cluster B, you can use a planned failover operation so that the replication direction switches from cluster B (new source) to cluster A (new destination).

For planned failovers, most of the failover steps are run from the source cluster.

- **From the source cluster (cluster A):**
 - Stop all applications that access the replicated files. This step is recommended, but not required.
 - Use the `replication episodic job sync job_name` command to run the job and make sure files on the source cluster and destination cluster are synchronized.

- Use the `replication episodic job failover force=yes/no job_name current_cluster_link` command to move control of the job from the source cluster to the destination cluster.
- From the destination cluster (cluster B):
 - Use the `replication episodic job enable job_name` command to enable the job or run a sync on the destination cluster.
 - Use the `replication episodic job sync job_name` command to ensure that the episodic replication job is in a well-defined state and incremental episodic replication can be resumed.

Once the job is failed over, job control remains on the destination cluster until a planned fallback is activated.

Note: While doing a planned replication episodic job failover from source to target, ensure that no shares are configured with the source file system as the source is marked as a target post failover.

Overview of the planned fallback process

After a job failover has been accomplished and the source cluster is ready to take back control of the episodic replication task, you can use the job fallback feature to release control from the original destination cluster (cluster B) and return it to the original source cluster (cluster A).

- From the original destination cluster (cluster B):
 - Stop all applications that access the replicated files. This step is recommended, but not required.
 - Use the `replication episodic job sync job_name` command to run the job and make sure files on the source cluster and destination cluster are synchronized.
 - Use the `Replication> episodic job disable job_name` command to disable the job.
- From the original source cluster (cluster A):
 - Use the `replication episodic job fallback force=yes/no job_name current_cluster_link` command to move control of the job from the destination cluster back to the original source cluster.
 - Use the `replication episodic job enable job_name` command to enable the job or run a sync on the source cluster.

- Use the `replication episodic job sync job_name` command to ensure that the episodic replication job is in a well-defined state and incremental episodic replication can be resumed.

Once the planned fallback process is completed, cluster A becomes the source cluster again and cluster B becomes the destination cluster again.

Overview of the unplanned failover process

In some cases (for example, unexpected equipment failure), you may need to run an unplanned failover for episodic replication jobs. The unplanned failover process differs from the planned failover process.

Consider two Access Appliance clusters, cluster A and cluster B. Here, cluster A is the original source cluster that replicates data to cluster B which is the original destination cluster. If the source cluster A fails unexpectedly, you have to perform an unplanned failover operation or a disaster recovery operation. After disaster recovery is performed, cluster B is marked as the new source cluster.

This section shows an overview of the steps you take to perform an unplanned failover.

For unplanned failovers, all the commands are run from the destination cluster (cluster B).

- Make sure that you are logged into the destination cluster.
- Use the `replication episodic job failover force=yes/no job_name` command to failover the job.

Overview of the unplanned fallback process

After an unplanned failover, when the source cluster (cluster A) comes up, you can use the following unplanned fallback process to return control to the original source cluster (cluster A):

- Make sure that you are logged into the source cluster.

Note: Before starting the fallback process, verify that the episodic replication service is running on the primary node. If the service is not running on the primary node, stop the service using the `replication episodic service stop` command and start it again using the `Replication> episodic service start` command.

- Use the `replication episodic job failover force=yes/no job_name current_cluster_link` command to configure the current source cluster as a

valid target to the new source cluster. This command should be run from the original source cluster (cluster A).

- Use the `replication episodic job sync job_name` command from the new source cluster to synchronize file system data with the newly added destination cluster.
- Use the `replication episodic job failback force=yes/no job_name current_cluster_link` command to move control of the episodic replication job from the destination cluster back to the source cluster. This command should be run from the original source cluster (cluster A).
- Use the `replication episodic job sync job_name` command to ensure that the episodic replication job is in a well-defined state and incremental episodic replication can be resumed.

Once the unplanned failback is completed, cluster A becomes the source cluster again and cluster B becomes the destination cluster again.

Note: An administrator can use the `Replication> episodic job destroy force` command to clean up local job configuration. Configuration of the other clusters, which are part of the job, are not modified and any episodic replication units are disassociated from job. The `Replication> episodic job destroy force` and `Replication> episodic repunit destroy force` commands should be used in the event of an unrecoverable configuration or episodic replication direction mismatch.

Configuring continuous replication

This chapter includes the following topics:

- [About Access Appliance continuous replication](#)
- [How Access Appliance continuous replication works](#)
- [Starting Access Appliance continuous replication](#)
- [Setting up communication between the source and the destination clusters](#)
- [Setting up the file system to replicate](#)
- [Managing continuous replication](#)
- [Displaying continuous replication information and status](#)
- [Unconfiguring continuous replication](#)
- [Preserving the file system on the destination cluster](#)
- [Continuous replication failover and fallback](#)
- [Addition of multiple file systems to a Replicated Volume Group](#)

About Access Appliance continuous replication

The Access Appliance continuous replication solution provides high performance, robustness, ease of use, and synchronous replication capability which is designed to contribute to an effective disaster recovery plan.

Access Appliance continuous replication lets you replicate volumes from one node in the source cluster to another node in the destination cluster. The continuous

replication enables you to maintain a consistent copy of application data at one remote location. It replicates the application writes on the volumes at the source location to a remote location across any distance.

If a disaster occurs at the source location, you can use the copy of the application data at the remote location and restart the application at the remote location. The host at the source location on which the application is running is known as the primary host and the host at the target location is known as the secondary host. The volumes on the primary host must be synchronized initially with the volumes on the secondary host.

Major features of Access Appliance continuous replication include:

- Performs replication of volumes in synchronous as well as asynchronous mode, ensuring data integrity and consistency.
- Maintains write-order fidelity, which applies writes on the secondary host in the same order that they were issued on the primary host.
- Enables easy recovery of the application at the remote site.
- Provides a command-line interface (CLI) and a graphical user interface (GUI) for online management of the synchronous replication.

See the `continuous(1)` manual page for more information.

How Access Appliance continuous replication works

The Access Appliance continuous replication implements Cluster Volume Replication (CVR) method internally. It is based on volume level replication.

It includes the following components:

- **Replicated Volume Group (RVG)**

A Replicated Volume Group (RVG) is a group of volumes or file sets within a given Volume Manager (VxVM) disk group which is configured for replication. The volumes in an RVG are consistent while replicating data. An RVG is always a subset of a VxVM disk group. One or more related volumes in a disk group can be configured as an RVG.

Volumes that are associated with an RVG and contain application data are called data volumes. The data volumes in the RVG are under the control of an application.

During replication, the write-order is strictly maintained within an RVG to ensure that each remote volume is always consistent, both internally and with all other volumes of the group. CVR replicates data from a primary RVG on the host (where the application is running) to the secondary RVG. An RVG also contains

the Storage Replicator Log (SRL) and Replication Link (RLINK), which are used internally by CVR.

- **Storage Replicator Log (SRL)**

The Storage Replicator Log (SRL) is a circular buffer of writes for an RVG. Each RVG contains one SRL. Writes to the data volumes in the RVG are first queued in the SRL on the primary host before they are sent to the secondary. CVR uses the SRL to track the order of writes to data volumes in the RVG. The SRL enables continuous replication to maintain write-order fidelity at the secondary RVG.

- **Replication Link (RLINK)**

An RLINK is associated with an RVG and establishes the link between the primary and the secondary RVG. Each RLINK associated with a primary RVG represents one secondary. Each RLINK associated with a secondary RVG represents the primary.

- **Data Change Map (DCM)**

The Data Change Map (DCM) is another important component of CVR that is used to track writes when the SRL overflows. It enables you to avoid complete resynchronization of the data on the secondary. The DCM contains a bitmap. The DCM is active only on the primary side. The DCM becomes active only when the SRL is no longer large enough to hold accumulated updates. While the DCM is active, each bit that has been set in the DCM represents a region whose contents are different between the primary and the secondary.

- **Replicated Data Set (RDS)**

A Replicated Volume Group (RVG) on the primary host and its counterparts on the secondary hosts make up a Replicated Data Set (RDS). An RDS enables grouping of the RVG on the primary and its counterparts on the secondary. The concept of primary host and secondary host is used only in the context of a particular Replicated Data Set (RDS).


If a synchronous operation for a file system is paused, replication stops when the SRL (Storage Replicator Log) overflows and the `replication> continuous status` command displays the replication status as "logging to DCM (needs dcm resynchronization)". If the replicator log (SRL) overflows, the RDS (Replicated Data Set) begins to track the writes using the DCM (Data Change Map). If the DCM is in use, replication stops and all the new writes are tracked in the DCM on the primary. To avoid this:

- Execute the `replication continuous status <fs-name>` command on CLISH.
- Check the output for the value of **Replicated Data Set**. This refers to the RVG name.
- Resynchronize the DCM by running the `vradmin resync <rvg name>` command from bash and wait for resync to complete.

How data flows in continuous replication synchronous mode

In synchronous mode of continuous replication, CVR processes an incoming write by performing the following steps.

Figure 25-1 Data flow in continuous replication synchronous mode



1. CVR receives a write on the primary host.
2. CVR writes it to the primary SRL.
3. CVR sends the write to the secondary host and waits for the synchronous network acknowledgments from the secondary hosts. At the same time, CVR writes to the data volumes on the primary host.

4. On the secondary host, CVR receives the write, processes it, and sends a network acknowledgment to the primary host.
5. CVR sends writes to the data volumes on the secondary host. When the primary host receives a network acknowledgment from the secondary host, CVR acknowledges to the application that the write is complete.

Note: The secondary RVG sends the network acknowledgment as soon as the write is received in the CVR kernel memory. This removes the time required to write to the secondary data volumes from the application latency. On the primary host, CVR does not wait for data to be written to the secondary data volumes. This improves application performance. However, CVR tracks all such acknowledged writes that have not been written to the data volumes. CVR can replay these tracked writes if the secondary host crashes before writing to the data volumes on the secondary host or if the primary host crashes before it receives the data acknowledgment.


When the write is written to the data volumes on the secondary host, CVR on the secondary host sends a data acknowledgment to the primary host. CVR marks the write as complete in the SRL when the primary receives the data acknowledgment from all the secondary hosts.

The bandwidth available for replication between primary and secondary sites should be able to handle the rate at which data is being written on the primary site. With synchronous continuous replication, the writes on file systems on primary site are completed after it is replicated to the secondary site. Hence, Veritas recommends that you have a high bandwidth for replication with synchronous continuous replication.

How data flows in continuous replication asynchronous mode

In asynchronous mode of continuous replication, CVR processes an incoming write by performing the following steps.

Figure 25-2 Data flow in continuous replication asynchronous mode



1. CVR receives a write on the primary host.
2. CVR writes it to the primary SRL.
3. On the primary host, CVR acknowledges to the application that the write is complete.
4. CVR sends the writes to the asynchronous secondary host in the order in which they were received on the primary host. At the same time, CVR writes to the primary data volumes.
5. When the primary host receives the network acknowledgment, it knows that the write has been received in the secondary CVR memory buffer.

6. CVR sends the writes to the data volumes on the secondary host and then sends a data acknowledgment to the primary host.
7. When the primary host receives the data acknowledgment, CVR marks the write as complete in the SRL.

Starting Access Appliance continuous replication

This section lists the specific commands that are needed to run Access Appliance continuous replication on your clusters.

Ensure the following before starting replication:

- Before you set up your clusters for continuous replication, you must first identify which is the source cluster and which is the destination cluster. All of the commands are performed on the source cluster first.
- Make sure both the source cluster and the destination cluster have the same version of Access Appliance.
- To use Access Appliance continuous replication, you have to first create an online file system on the Access Appliance source cluster.
- Assign a virtual IP (VIP) address to both the source and the destination clusters. The Access Appliance continuous replication service requires VIP addresses not already in use for the two clusters to communicate.
- The name resolution for replication virtual IPs for both source and target site should get resolved correctly. If there is no DNS entry for replication virtual IPs, add appropriate entry in the `/etc/hosts/` file on both nodes of the source and destination clusters.

To start Access Appliance continuous replication on the source cluster

- 1 To bind a virtual IP address for the continuous replication service on the source cluster, enter the following:

```
Replication> continuous config bind ip_addr [device] [netmask]
```

ip_addr Virtual IP address for the continuous replication service on the source cluster. It should not be part of the network IP pool.

device The public network interface name that you want the replication IP address to use.

netmask Netmask for the replication IP address.

- 2 To start the continuous replication service, enter the following on the source node:

```
Replication> continuous service start [nodename]
```

nodename The name of the node in the local cluster where you want to start the replication service.

- 3 To check the status of the continuous replication service, enter the following:

```
Replication> continuous service status
```

- 4 To confirm that the IP address is up and running, enter the following:

```
Replication> continuous config show ip
```

Note: Alternately, you can use the `Network> ip addr show` command to confirm that the IP address is up and running.

To start Access Appliance continuous replication on the destination cluster

- 1 To bind a virtual IP address for the replication service on the destination cluster, enter the following:

```
Replication> continuous config bind ip_addr [device] [netmask]
```

ip_addr Virtual IP address for the continuous replication service on the destination cluster. It should not be part of the network IP pool.

device The public network interface name that you want the replication IP address to use.

netmask Netmask for the replication IP address.

- 2 To start the continuous replication service, enter the following on the destination node:

```
Replication> continuous service start [nodename]
```

nodename The name of the node in the local cluster where you want to start the replication service.

- 3 To check the status of the continuous replication service, enter the following:

```
Replication> continuous service status
```

- 4 To confirm that the IP address is up and running, enter the following:

```
Replication> continuous config show ip
```

You next need to set up communication between the source and the destination clusters.

Setting up communication between the source and the destination clusters

You need to set up communication between your source and your destination clusters.

Access Appliance Replication authentication strategy is based on RSA-key authentication, and both the source and the destination clusters have to export their replication public keys. The source cluster imports the destination cluster's public key and the destination cluster imports the source cluster's public key.

After you have determined which two Access Appliance clusters to use, you need to authenticate them.

The `replication continuous config` commands must be executed in a specific order.

- You need to run the `replication continuous config bind` command (to bind the virtual IP) before you can run the `replication continuous service start` command.
- You need to run the `replication continuous config export_keys` and `replication continuous config import_keys` to export and import the keys of both the source and the destination clusters.

Copy the keys displayed in the output of the `replication continuous config export_keys` command from the source cluster. Run the `replication continuous config import_keys` command on the destination cluster using the output that you copied.

Copy the keys displayed in the output of the `replication continuous config export_keys` command from the destination cluster. Run the `replication continuous config import_keys` command on the source cluster using the output that you copied.

- You can only run the `replication continuous config auth` command after both the source and destination have imported each other's keys.
- You need to run the `replication continuous config auth` command to create a link from every cluster to any remaining cluster that is used for replication irrespective of their role as a source or a destination cluster.

The command checks the two-way communication between the source and the destination cluster, and authenticates the clusters allowing the Access Appliance continuous replication service to begin.

This section provides a walk-through for the creation and export/import of these encrypted keys for both the source and the destination cluster.

Note: Without the correct authentication of the source and the destination encryption keys, Access Appliance continuous replication does not function correctly.

To export the source cluster's key to the destination cluster

- 1 To export the source cluster's key to the destination cluster, enter the following:

```
Replication> continuous config export_keys [URL]
```

URL The location you want to copy the public keys to.

If you do not want to enter a URL, you can copy the output from the `replication continuous config export_keys` command into the `Replication> continuous config import_keys` command at the destination cluster.

By default, the output is displayed on your computer screen.

The SCP protocol is supported.

- 2 To import the source cluster's key to the destination cluster, enter the following:

```
Replication> continuous config import_keys [URL/keyfile]
```

URL The location you want to copy the public keys from.

keyfile The file name of the key that is generated by the export.

If you did not enter a URL during the `Replication> continuous config export_keys` command, you can cut and paste the output and enter it into the `replication continuous config import_keys` command.

- 3 To verify that the key has been imported correctly, enter the following:

```
Replication> continuous config show
```

To export the destination cluster's key to the source cluster

- 1 To export the destination cluster's key to the source cluster, enter the following:

```
Replication> continuous config export_keys [URL]
```

URL The location you want to copy the public keys to.

The SCP protocol is supported.

If you do not want to enter a URL, you can cut and paste the output from the `Replication> continuous config export_keys` command to the `Replication> continuous config import_keys` command. By default, the output is displayed to your computer screen.

- 2 To import the destination cluster's key to the source cluster, enter the following:

```
Replication> continuous config import_keys [URL/keyfile]
```

URL Enter the URL of the location you want to copy the public keys from.

keyfile Enter the file name of the key that is generated by the export.

If you did not enter a URL during the `replication continuous config export_keys` command, you can cut and paste the output and enter it into the `replication continuous config import_keys` command.

- 3 To verify that the key has been imported correctly, enter the following:

```
Replication> continuous config show
```

To authenticate source cluster and destination clusters for replication

- 1 This command should be run on the source cluster as well as on the destination cluster. To authenticate the public keys on the source cluster and the destination clusters, enter the following:

```
Replication> continuous config auth conIP link_name
```

conIP Enter the destination cluster console IP address.

link_name Both the source cluster and the destination cluster need to be assigned a unique identifier (name). This identifier is used to identify the link that is established between the source and the destination clusters. You can use the link name instead of the virtual IP addresses of the source and the destination clusters when using the other replication commands. For example:
Pune_Shanghai.

- 2 To confirm the authentication, enter the following:

```
Replication> continuous config show
```

Once you have configured the clusters and links, you need to set up the file system you want to replicate.

Setting up the file system to replicate

You need to set up the file systems you want to replicate. The file system which you are going to replicate should be in online state.

To set up a continuous replication, enter the following:

```
Replication> continuous enable fs_name pool_name link_name  
delayed=yes/delayed=no rvg_name create_target_fs srl_size
```

fs_name The name of the file system that you want to replicate from source to destination cluster.

pool_name The name of the pool.

link_name The link name which was created during authentication time.

<i>delayed=yes/delayed=no</i>	The <i>delayed</i> parameter value is yes if you want to set up continuous asynchronous replication. By default, its value is no .
<i>rvg_name</i>	RVG name at primary site.
<i>create_target_fs</i>	The value is yes if you want to create file system on secondary cluster. Else, the value is set to no .
<i>srl_size</i>	Specify the SRL size for the file system.

The command configures the continuous replication between the source and the destination cluster. A file system is created with the same name and same size on the destination cluster. The file system is in offline state at the secondary site to maintain a consistent copy of data. You can enable continuous replication using the same RVG on multiple file systems.

It requires the pool on the destination cluster to have sufficient storage to create a file system, the Storage Replicator Log (SRL) volume and the DCM logs. The pool name must be same at the source and the destination cluster. If *delayed* mode is enabled, there can be a non-zero RPO. For asynchronous continuous replication, the SRL volume size is 20% of the file system size. For synchronous continuous replication, the SRL volume size is 5% of the file system size. By default, the number of DCM logs is 2 but for a mirrored file system, the number of DCM logs is equal to the number of mirrors.

For example, if the file system size on the source cluster is 8 GB, then at least 3 GB storage should be present in the pool at the source cluster to create the SRL volume and the DCM logs in case of asynchronous replication. For the destination cluster, 11 GB storage should be present in the pool to create the file system, the SRL volume, and the DCM logs.

Note: If you configure continuous replication on a WORM-enabled file system, the `replication continuous enable` command fails if either the source or destination clusters are configured with Normal lockdown mode. It is recommended that the source and destination clusters should be configured with the same lockdown modes.

You have an option to preserve the file system and also re-use the file system that was created.

For more details, See [“Preserving the file system on the destination cluster”](#) on page 376.

Note: Continuous replication is not supported for a file system with encrypted volume. When setting up replication, Veritas recommends that you do not make any modifications or deletions on the target side of the file system. File system grow and shrink operations are supported on file systems which are configured under continuous replication. The mode of replication (synchronous or asynchronous) cannot be changed after configuring the replication.

Managing continuous replication

You can manage continuous replication by `Replication> continuous start`, `Replication> continuous stop`, `Replication> continuous pause` and `Replication> continuous resume` commands. It is recommended that you should run all these commands from the source cluster.

To start data replication

- ◆ To start data replication, enter the following command:

```
Replication> continuous start fs_name
```

fs_name Specify the file system name that you have configured for continuous replication.

The data replication between source cluster and destination cluster starts.

To pause data replication

- ◆ To pause data replication, enter the following command:

```
Replication> continuous pause fs_name
```

fs_name Specify the file system name that you have configured for continuous replication.

The data replication between source cluster and destination cluster is paused.

To resume data replication

- ◆ To resume data replication, enter the following command:

```
Replication> continuous resume fs_name
```

fs_name Specify the file system name that you have configured for continuous replication.

The data replication between the source cluster and the destination cluster which was paused is resumed.

To stop data replication

- ◆ To stop data replication, enter the following command:

```
Replication> continuous stop fs_name
```

fs_name Specify the file system name that you have configured for continuous replication.

The data replication between source cluster and destination cluster is stopped.

Note: Veritas recommends that you should run `replication continuous resume` command when replication is in paused state. You should run `Replication> continuous stop` command when replication is in progress.

Displaying continuous replication information and status

The `Replication continuous show` and `Replication continuous status` commands display information on continuous replication which lets you confirm any changes that are made to your replication file system and view the current file system status.

To display the list of file systems which are configured under continuous replication

- ◆ To display the list of file systems which are configured under continuous replication, enter the following command:

```
Replication> continuous show
```

To display the status of a replication file system

- ◆ To display the status of a replication file system, enter the following command:

```
Replication> continuous status fs_name
```

fs_name Specify the file system name that you have configured for continuous replication.

Table 25-1 describes the important attributes displayed by the `Replication> continuous status` command.

Table 25-1

Attribute	Description
Replicated Data Set	Specifies the name of the replicated data set
Replication role	Specifies the role of the cluster in continuous replication. It is either primary or secondary.
Replication link	Specifies the link name which is created during the authentication of the source and the destination cluster.
Primary site Info	Provides the details of continuous replication related to the source cluster like host name and RVG state.
Secondary site Info	Provides the details of continuous replication related to the destination cluster like host name, configured mode, data status, replication status, current mode, and timestamp information.
Host name	For the primary site, it provides the continuous replication IP which binds on the source cluster using the <code>Replication> continuous config bind</code> command. For the secondary site, it provides the continuous replication IP which binds on the destination cluster using the <code>Replication> continuous config bind</code> command.
RVG state	Specifies the state of the primary RVG. See Table 25-2 for details on its various states.

Table 25-1 (continued)

Attribute	Description
Configured mode	Specifies the continuous replication configured mode. It may be <i>synchronous-override</i> or <i>asynchronous</i> .
Current mode	Specifies the mode of replication - <i>asynchronous</i> or <i>synchronous</i> that is used to replicate data to the secondary.
Data status	Shows the data status for the secondary. See Table 25-3 for details on its various states.
Replication status	Specifies the status of the replication to the secondary. See Table 25-4 for details on its various states.
Logging to	Indicates whether updates for the secondary are tracked on the primary using the SRL or DCM. See Table 25-5 for details on its various states.
Timestamp information	Shows the time by which secondary is lagging behind the primary.

[Table 25-2](#) describes the values for the RVG state.

Table 25-2 RVG status

Value	Description
acting_secondary	The primary RVG is currently the acting secondary as part of the fast failback process. Writes to the data volumes in this RVG are disabled independent of whether the RVG is started or stopped.
Disabled for I/O	Primary RVG is disabled for I/O. The RVG is stopped.
Enabled for I/O	Primary RVG is enabled for I/O. The RVG has been started.

Table 25-2 RVG status (*continued*)

Value	Description
Needs recovery	State of the RVG after an import or restart.
Passthru	The primary RVG is in passthru mode because the primary SRL is detached or missing.

[Table 25-3](#) describes the values for the Data status.

Table 25-3 Data status

Value	Description
consistent, behind	Secondary data is consistent but not up to date with the primary data.
consistent, stale	The data on the secondary is consistent. Replication to the secondary has been stopped. The primary RLINK is detached.
consistent, up-to-date	The secondary data is consistent and is current or up-to-date with the primary data. The primary role can be migrated to the secondary.
inconsistent	The data on the secondary volumes is not consistent and the secondary cannot take over.
N/A	Current state of the secondary data cannot be determined. This may occur because of a configuration error on the secondary.

[Table 25-4](#) describes the values for the Replication status.

Table 25-4 Replication status

Value	Description
logging to DCM	DCM is active for the secondary. New updates on primary are tracked using DCM for the secondary. The following information may be displayed: ■ <i>needs dcm resynchronization</i>—Resynchronize the secondary using DCM resynchronization to continue replication.
Needs failback synchronization	The primary RVG is acting as secondary as part of the fast failback process. Start failback resynchronization on the new primary to continue replication.
Not replicating	Data is not being replicated to secondary because primary RLINK is in <i>needs_recovery</i> state. ■ <i>primary needs_recovery</i>—Primary RLINK needs to be recovered before replication can resume.
Paused by user	Replication to secondary is paused due to some administrative action. This results in the following states: ■ <i>primary paused</i>—primary RLINK is paused. ■ <i>secondary paused</i>—secondary RLINK is paused.
Paused due to error	Replication to secondary is paused due to the following errors: ■ <i>secondary config error</i>—secondary has some configuration error. ■ <i>secondary log error</i>—secondary SRL has an I/O error.
Paused due to network disconnection	Replication to secondary is paused due to some network problem.
Replicating	Replication can take place if there are updates on the primary data volumes.

Table 25-4 Replication status (*continued*)

Value	Description
Resync in progress	Resynchronization to the secondary is in progress. ■ <i>autosync</i>—Resynchronization type is autosync. ■ <i>dcm resynchronization</i>—resynchronization after an SRL overflow. ■ <i>failback resynchronization</i>—resynchronization using failback logging. ■ <i>smartsync</i>—resynchronization type is autosync using SmartMove.
Resync paused by user.	Resynchronization to secondary is paused due to some administrative action. This results in the following states: ■ <i>primary paused</i>—primary RLINK is paused. ■ <i>secondary paused</i>—secondary RLINK is paused.
Resync paused due to error.	Resynchronization to Secondary is paused because of the following errors: ■ <i>secondary config error</i>—secondary has some configuration error. ■ <i>secondary log error</i>—Secondary SRL has an I/O error.
Resync paused due to network disconnection.	Resynchronization to secondary is paused due to some network problem.
Stopped	Replication to secondary is stopped due to the following: ■ <i>Primary detached</i>—Primary RLINK is detached. ■ <i>Secondary detached</i>—Secondary RLINK is detached.
N/A	The replication status cannot be determined.

Table 25-5 describes the values for the Logging to field.

Table 25-5 Logging to status

Value	Description
DCM (contains xxx KB) (log_type)	DCM is active (in use) for the replication to the secondary. The <i>log_type</i> can be <i>autosync</i> , <i>failback logging</i> , or <i>SRL protection logging</i> . The yyy% value can sometimes reach beyond 100%. If synchronization is restarted and the DCM map is full, new incoming writes cause the total yyy% to exceed 100%.
SRL (xxx KB behind, yyy % full)	Updates to be transferred to secondary are logged into the SRL and are currently occupying xxx KB or yyy% of the SRL.
SRL	SRL is used for logging. Check the Data status field for the status of the secondary data.

Unconfiguring continuous replication

You can unconfigure continuous replication.

To unconfigure continuous replication

- 1 Stop the replication. Before you disable continuous replication, you have to stop replication using the following command.

```
Replication> continuous stop <fs_name>
```

fs_name Specify the file system name.

Note: This command should be run from the source cluster.

- 2 Check the replication status.

```
Replication> continuous status <fs_name>
```

fs_name Specify the file system name.

- 3 Disable continuous replication. All the configuration which was done for replication configuration when you enabled continuous replication from the source and the destination cluster is destroyed.

```
Replication> continuous disable <fs_name> <link_name>
```

fs_name Specify the file system name.

link_name Specify the link name.

Note: This command should be run from the source cluster.

The RVG continues to exist until the last file system in the RVG is disabled. Once the last file system in the RVG is disabled, the RVG is deleted.

If a file system is WORM-enabled and contains data, the `replication continuous disable <fs_name> <link_name>` command fails because the default value of the `destroy_target_fs` parameter is `yes`. To disable replication, set `destroy_target_fs=no`.

You have an option to preserve the file system and also re-use the file system that was created.

For more details, See [“Preserving the file system on the destination cluster”](#) on page 376.

- 4 Delete the authentication links.

```
Replication> continuous config deauth <link_name>
```

link_name Specify the link name.

Note: The `Replication> continuous config deauth` command should be run from the source cluster. The command deletes the link from the destination to the source cluster.

- 5 Delete the keys from the source and the destination clusters.

```
Replication> continuous config del_keys <remote_console_ip>
```

remote_console_ip Specify the remote console IP address.

- 6 Stop the continuous service from the source and the destination clusters.

```
Replication> continuous service stop
```

- 7 Unbind the replication IP from the source and the destination clusters.

```
Replication> continuous config unbind <replication_ip>
```

replication_ip Specify the replication IP address.

Preserving the file system on the destination cluster

When continuous replication is disabled, the target file system is deleted. When the same file system is enabled for replication again, a full sync happens. You have an option to preserve the file system and also re-use the file system that was created during the previous `replication continuous enable` operation using the `destroy_target_fs=yes/no` parameter.

If you disable the file system and set `destroy_target_fs=no` in the `replication continuous disable`, then the file system on the destination cluster will not be deleted. The default value of the `destroy_target_fs` parameter is `yes` and if you do not pass this parameter, the file system on the destination cluster is deleted.

```
replication> continuous disable fs_name link_name destroy_target_fs=yes/no
```

Later, if you want to re-use the file system created on the destination cluster, you can set `create_target_fs=no` in the `replication continuous enable` command.

```
replication> continuous enable fs_name link_name
source delayed=yes rvg_test create_target_fs=yes
srl_size
```

If there is a mismatch in any of the file system parameters, a failure message appears to inform the user that the file system on the destination cluster does not match with the source file system.

Continuous replication failover and failback

Typically, the source cluster drives a replication session. However, in some situations, it may be useful for the destination cluster to drive the replication session. Access Appliance supports a failover and a failback feature for continuous

replication. This feature enables control of replication to be temporarily relocated from the source cluster to the destination (target) cluster.

Continuous replication failover and failback is useful for:

- **Planned failover**
In cases where the source cluster is taken down for routine maintenance or for moving applications to another cluster, a planned failover procedure is available for moving replication from the source cluster to the destination cluster.
- **Unplanned failover or Disaster recovery**
In cases where the source cluster fails unexpectedly, an unplanned failover procedure is available for moving replication to the destination cluster.

With failover and failback, you can use the `replication continuous failover` command to move control from the source cluster to the destination cluster. You use the `replication continuous failback` to restore control to the source cluster.

Continuous replication failover does not automatically move the NFS or the CIFS share information that is associated with file system from the source cluster to the destination cluster. Share information has to be done manually.

CIFS shares should be removed before failover and failback operations. After failover and failback operation are complete, add the CIFS shares again.

Process summary

The steps you take for failover and failback vary depending on the type of failover or failback you perform. Failover and failback types include:

- **Planned failover:** This is performed to switch the job roles of source and destination clusters.
- **Failback after a planned failover:** This is performed after a planned failover to switch the roles of the clusters back to the original state.
- **Unplanned failover:** This is performed when the primary cluster is down and disaster recovery has to be performed.
- **Failback after an unplanned failover:** : This is performed after an unplanned failover or disaster recovery.

Each process is summarized in the following sections. Typically, you would use the planned failover and planned failback processes in most situations.

Overview of the planned failover process

Consider two Access Appliance clusters, cluster A and cluster B. Here, cluster A is the original source cluster that replicates data to cluster B which is the original

destination cluster. If the source cluster A is taken down for routine maintenance or for moving applications to cluster B, you can use a planned failover operation so that the replication direction switches from cluster B (new source) to cluster A (new destination). You should perform the planned failover before cluster A is taken down.

For planned failover, run the following command on the source cluster (cluster A):

```
Replication> continuous failover fs_name
```

Where *fs_name* is the name of the file system which is configured under continuous replication.

Once a planned failover happens, the roles of primary and secondary are switched. It will online the file system at new primary site and offline the file system at new secondary site. Now, cluster B is the primary site and cluster A is the secondary site. The file systems come online at the new primary site (cluster B) and the file systems at new secondary site (cluster A) are offline.

Note: Planned failover command should be run when both the source and the destination clusters are reachable from each other. It should be run from the source cluster and replication should be in progress. If you have NFS/CIFS shares on the source cluster, it is recommended that you should stop the NFS/CIFS server before planned failover.

Overview of the planned failback process

After a planned failover has been accomplished and the original source cluster (cluster A) is ready to take back control of the replication task, you can use the failback feature to release control from the original destination cluster (cluster B) and return it to the original source cluster (cluster A).

For planned failback, run the following command:

```
Replication> continuous failback fs_name
```

Where *fs_name* is the name of the file system which is configured under continuous replication.

Once a planned failback happens, the roles of primary and secondary are switched back to the original state (the pre-failover state). Now, cluster A is the primary site again and cluster B is the secondary site again.

Note: Planned failback command should be run when both the source and the destination clusters are reachable from each other. It should be run from the destination cluster (which was the original source cluster) and replication should be in progress. If you have NFS/CIFS shares on the source cluster, it is recommended that you should stop the NFS/CIFS server before planned failback.

Overview of the unplanned failover process

In some cases (for example, unexpected equipment failure), you may need to run an unplanned failover for replication. The unplanned failover process differs from the planned failover process.

If cluster A is the original source cluster that replicates data to cluster B which is the original destination cluster and if cluster A fails unexpectedly, you can perform an unplanned failover or a disaster recovery. This marks the cluster B as the new source cluster and you can access the data/applications replicated to it.

For unplanned failover, run the following command from the original destination cluster (cluster B):

```
Replication> continuous failover fs_name
```

Where *fs_name* is the name of the file system which is configured under continuous replication.

Once an unplanned failover happens, the destination cluster (cluster B) becomes the new source cluster. It online the file system at the new source cluster (cluster B).

Note: Though the commands used for planned and unplanned failover are the same, the intention and pre-requisites are different. For unplanned failover, the source cluster should be unreachable from the destination cluster.

Overview of the unplanned failback process

After an unplanned failover, when the original source cluster (cluster A) comes up, you can use the following unplanned failback process to retain the replication between source and destination cluster.

When the original source cluster (cluster A) comes up, it still acts as the primary cluster. At this time, both the source and the destination clusters show the same status.

For unplanned failback, run the following command from the original source cluster (cluster A):

```
Replication> continuous failback fs_name
```

Where *fs_name* is the name of the file system which is configured under continuous replication.

Once, an unplanned failback happens, the original source cluster (cluster A) becomes the new destination cluster. The file system is still online at new source cluster (cluster B).

Note: Unplanned failback command should be run from the original source cluster. If you have NFS/CIFS shares on the original source cluster, it is recommended that you stop the NFS/CIFS server before unplanned failback.

Addition of multiple file systems to a Replicated Volume Group

Access Appliance supports the addition of multiple file systems to a Replicated Volume Group (RVG). You can enable replication for many file systems in a single RVG.

Considerations while configuring Access Appliance for continuous replication:

- The number of RVGs configured under continuous replication on Access appliance should be restricted to 4.
- The number of file systems under one RVG should be limited to 12.
- Veritas recommends that you to add one file system at a time to an RVG. Once the file system is added and the status of continuous replication becomes consistent (up-to-date), add another file system to the RVGs.
See [“Setting up the file system to replicate”](#) on page 365.
- Ensure that continuous replication IPs configured on Access Appliance are properly resolved by the configured DNS server or by `/etc/hosts` file.
- Addition of multiple file systems to an existing Replicated Volume Group is supported only through CLISH.

To configure an RVG with multiple file systems

- 1 Create an RVG. You can specify the name in the parameter.

```
replication> continuous enable fs_name pool_name link_name
[delayed=yes|no] rvg_name create_target_fs
srl_size
```

fs_name	Specifies the name of the file system. It should be present on the source cluster.
pool_name	Specifies the pool name. It should be present on the source as well as destination cluster.
link_name	Specifies the link name. This link is created during authentication of the source and the destination cluster,
[<i>delayed=yes delayed=no</i>]	Specifies if the continuous replication should happen in delayed mode. This is an optional parameter, and the default value is 'no'.
rvg_name	Specifies the name of the RVG.
create_target_fs	The value is yes if you want to create file system on secondary cluster. Else, the value is set to no .
srl_size	Specify the SRL size for the file system.

For example: To create an RVG, `rvg_test`:

```
replication> continuous enable fs1 pool1 src_link delayed=yes rvg_test
```

You can specify the RVG name in the command. If the RVG name is not specified, the RVG is created with the default name. The default name of the RVG is `rvg_fs_name`.

- 2 Before adding another file system, ensure that the data is synced between the primary and secondary sites and the second file system is offline.

```
replication> continuous status fs_name
```

3 Add another file system to the same RVG.

```
replication> continuous enable fs_name pool_name link_name
[delayed=yes|no] rvg_name
```

Note: If you have configured any CIFS or NFS shares on the file system, you must remove them before you add the file system to an existing RVG.

For example: To add a file system. *fs2* to the existing RVG, *rvg_test*:

```
replication> continuous enable fs2 pool1 src_link delayed=yes rvg_test
```

After the execution of this command, *rvg_test* is created with two file systems in it, *fs1* and *fs2*.

You can add more file systems in the RVG by repeating steps 2 to 3.

You can view the mapping between the RVG and the file systems using the following command:

```
replication> continuous show
RVG                      FS List
-----
rvg_test                 fs1,fs2
```

The `replication continuous` commands such as `start`, `stop`, `pause`, `resume`, `failover`, and `failback` accept the file system name as the parameter, but the changes are applicable for the entire RVG (including all the file systems which are part of the RVG).

For example, if you have multiple file systems in an RVG and if you failover using the `replication continuous failover fs1` command, then the entire RVG will fail over.

The `replication continuous disable` command does not delete the RVG until the file system is the last file system in the RVG.

For each file system that is added to the RVG, if the file system is configured in async mode, the Storage Replicator Log (SRL) grows by 20% of the file system size, and if the file system is configured in sync mode, it grows by 5% of the file system size. If the SRL has been grown more than 12 times(either as a result of the `fs grow` command or addition of file systems to the RVG), then the SRL is not grown further.

You can also set the maximum SRL size for an RVG.

```
replication> continuous config set_max_srl_size rvg_name size
```

For example, the following command sets the SRL size to maximum of 200GB for *rvg_test*.

```
replication> continuous config set_max_srl_size rvg_test 200G
```

You can also use the `replication continuous enable` command for setting the RVG size.

```
replication> continuous enable fs1 pool1 source
delayed=yes rvg_test create_target_fs=yes 100G
```

Here, 100G is the SRL size for *rvg_test*, which is an optional parameter.

This option is useful when setting the SRL size for the first file system when an RVG does not exist.

You can also view the current SRL and maximum SRL size. Appropriate messages are shown after the completion of each command to inform the user about successful completion or reason for failure.

You can verify whether the file systems are added correctly and the appropriate SRL size is set.

```
replication> continuous show
```

RVG	FS List	SRL Size
rvg_test	fs1, fs2, fs3	204m

You can also verify whether the maximum SRL size is set correctly.

```
replication> continuous config show
```

Link name	Remote Console IP	Remote Replication VIP	Time of Key Import
clus_repl	10.221.35.102	10.221.35.104	Tue Oct 12 23:22:19 PDT 2021

Time of Authorization	RVG	Max SRL Size
Tue Oct 12 23:23:32 PDT 2021	rvg_test	-
	rvg_test2	19g
	rvg_test3	10g
	rvg_test4	11g

Using snapshots

This chapter includes the following topics:

- [About snapshots](#)
- [Enabling WORM on storage snapshots](#)
- [Creating snapshots](#)
- [Displaying snapshots](#)
- [Managing disk space used by snapshots](#)
- [Bringing snapshots online or taking snapshots offline](#)
- [Restoring a snapshot](#)
- [About snapshot schedules](#)
- [Configuring snapshot schedules](#)
- [Managing automated snapshots](#)

About snapshots

A snapshot is a virtual image of the entire file system. You can create snapshots of a parent file system on demand. Physically, it contains only data that corresponds to the changes that are made in the parent, and so consumes significantly less space than a detachable full mirror.

Snapshots are used to recover from data corruption. If files, or an entire file system, are deleted or become corrupted, you can replace them from the latest uncorrupted snapshot. You can mount a snapshot and export it as if it were a complete file system. Users can then recover their own deleted or corrupted files. You can limit the space snapshots consume by setting a quota on them. If the total space that

snapshots consume exceeds the quota, Access Appliance rejects attempts to create additional ones.

You can create a snapshot by either using the `snapshot create` command or by creating a schedule to create the snapshot at a specified time.

Enabling WORM on storage snapshots

You can enable WORM on storage snapshots which are created on WORM enabled filesystem to make the snapshot data unmodifiable. WORM storage snapshot allows you to retain the snapshot data in read-only form for a specified period for security or governance purposes. The retention period for a WORM storage snapshot specifies the time period for which the snapshot must be retained. The snapshot can be deleted after the retention period has expired. To protect the snapshot data from ransomware attack, you must keep the snapshots in read-only state. You cannot perform write operations to the storage snapshot.

Considerations for enabling WORM on snapshots:

- This feature is supported only in Compliance and Enterprise lockdown mode.
- The retention period is based on the compliance clock. The minimum retention period is 1 day and the maximum period is 2038-01-19.
- You can extend the retention period in Compliance and Enterprise mode, but retention period can be reduced only in Enterprise mode.
- You cannot set the retention period on a snapshot if snapshot is not WORM-enabled.
- Snapshots which are WORM-enabled cannot be online/mounted in read-write mode.
- Snapshots with active retention period cannot be destroyed. You can set WORM and retention period to existing snapshots. Snapshots which are created on WORM enabled file system cannot be restored.
- You can create and modify snapshot schedules if WORM is enabled and retention period is set.
- The `schedule destroyall` command does not destroy snapshots which have active retention period.

To enable WORM on a snapshot

- ◆ To enable WORM on an existing non-WORM snapshot, enter the following:

```
Storage> snapshot worm set snapshot_name fs_name
```

snapshot_name Specifies the name for the snapshot.

fs_name Specifies the name for the file system.

To set the retention period on a snapshot

- ◆ To set retention period on WORM-enabled snapshot, enter the following:

```
Storage> snapshot retention set snapshot_name fs_name  
retention_period
```

snapshot_name Specifies the name for the snapshot.

fs_name Specifies the name for the file system.

retention_period Specifies the retention period. The format can be
[1-9](d|D|m|M|y|Y) or yyyy-mm-dd.

You can also extend the retention period or set a retention date in the past using this command.

Creating snapshots

The `snapshot create` command quickly creates a persistent image of a file system at an exact point in time. Snapshots minimize the use of disk space by using a Storage Checkpoint within the same free space available to the file system. After you create a snapshot of a mounted file system, you can also continue to create, remove, and update files on the file system without affecting the logical image of the snapshot. A snapshot preserves not only the name space (directory hierarchy) of the file system, but also the user data as it existed at the moment the file system image was captured.

You can use a snapshot in many ways. For example, you can use them to:

- Create a stable image of the file system that can be backed up to tape.
- Provide a mounted, on-disk backup of the file system so that end users can restore their own files in the event of accidental deletion. This is especially useful in a home directory, engineering, or email environment.

- Create an on-disk backup of the file system that can be used in addition to a traditional tape-based backup to provide faster backup and restore capabilities.

To create a snapshot

- ◆ To create a snapshot, enter the following:

```
Storage> snapshot create snapshot_name fs_name [removable]
<worm> <retention_period>
```

If the file system is not WORM-enabled, the snapshot cannot be WORM-enabled. Retention period can be set only if the snapshot is WORM-enabled.

snapshot_name	Specifies the name for the snapshot. Note: The following are reserved words for snapshot name: flags, ctime, and mtime.
fs_name	Specifies the name for the file system.
removable	Valid values are: ■ yes ■ no If the removable attribute is yes, the snapshot is removed automatically if the file system runs out of space. The default value is removable=no.
worm	Enables WORM. The possible values are worm=yes/no. Default value of this parameter is no.
retention_period	Specifies the retention period. The format can be [1-9](d D m M y Y) or yyyy-mm-dd.

Displaying snapshots

You can display all snapshots, or the snapshots taken of a specific file system or specific schedule of a file system. The output displays the snapshot name and the properties of the snapshots such as creation time and size.

If the snapshot is WORM-enabled, the retention period is also displayed.

To display snapshots

- ◆ To display snapshots, enter the following:

```
Storage> snapshot list [fs_name] [schedule_name]
```

fs_name Displays all of the snapshots of the specified file system. If you do not specify a file system, snapshots of all of the file systems are displayed.

schedule_name Displays the schedule name. If you do not specify a schedule name, then snapshots created under *fs_name* are displayed.

Managing disk space used by snapshots

To manage the disk space used by snapshots, you can set a snapshot quota or capacity limit for the file system. When all of the snapshots for the file system exceed the capacity limit, snapshot creation is disabled for the file system.

You can also remove unnecessary snapshots to conserve disk space.

To enable snapshot quotas

- 1 To display snapshot quotas, enter the following:

```
Storage> snapshot quota list
```

FS	Quota	Capacity Limit
==	=====	=====
fs1	on	1G
fs2	off	0
fs3	off	0

- 2 To enable a snapshot quota, enter the following:

```
Storage> snapshot quota on fs_name [capacity_limit]
```

fs_name	Specifies the name of the file system.
capacity_limit	Specifies the number of blocks used by all the snapshots for the file system. Enter a number followed by K, M, G, or T (for kilo, mega, giga, or terabyte). The default value is 0.

- 3 If necessary, you can disable snapshot quotas. You can retain the value of the capacity limit. To disable a snapshot quota, enter the following:

```
Storage> snapshot quota off [fs_name] [remove_limit]
```

fs_name	Specifies the name of the file system.
remove_limit	Specifies whether to remove the capacity limit when you disable the quota. The default value is true, which means that the quota capacity limit is removed. The value of false indicates that the quota is disabled but the value of the capacity limit remains unchanged for the file system.

To destroy a snapshot

- ◆ To destroy a snapshot, enter the following:

```
Storage> snapshot destroy snapshot_name fs_name
```

snapshot_name	Specifies the name of the snapshot to be destroyed.
fs_name	Specifies the name of the file system from which the snapshot was taken. Snapshots with the same name could exist for more than one file system. In this case, you must specify the file system name.

Note: If you want to destroy a snapshot with active retention period, you have to set a retention date in the past on that snapshot using the `storage snapshot retention set` command, but you can set a retention period in the past only if you have configured Enterprise lockdown mode.

Bringing snapshots online or taking snapshots offline

If you want to mount a snapshot through NFS or export a CIFS snapshot, you must bring the snapshot online. You can then create a CIFS or an NFS share using the snapshot name as the path. For example: `/vx/fs1:snap1`. The snapshot can only be mounted through NFS or exported through CIFS if it is online.

To bring a snapshot online

- ◆ To bring a snapshot online:

```
Storage> snapshot online snapshot_name fs_name mode
```

snapshot_name	Specifies the name of the snapshot.
fs_name	Specifies the name of the file system from which the snapshot was taken. Snapshots with the same name could exist for more than one file system. In this case, you must specify the file system name.
mode	Specifies the mode. The possible values are <code>read-only</code> and <code>read-write</code> . Default value for mode parameter is <code>mode=read-only</code> . WORM-enabled snapshots can be online in only <code>read-only</code> mode.

To take a snapshot offline

- ◆ To take a snapshot offline:

```
Storage> snapshot offline snapshot_name fs_name
```

snapshot_name Specifies the name of the snapshot.

fs_name Specifies the name of the file system from which the snapshot was taken. Snapshots with the same name could exist for more than one file system. In this case, you must specify the file system name.

Restoring a snapshot

This operation restores the file system to the state that is stored in the specified snapshot. When you restore the file system to a particular snapshot, snapshots taken after that point in time are no longer relevant. The restore operation also deletes these snapshots.

The restore snapshot operation prompts you for confirmation. Be sure that you want to restore the snapshot before responding yes.

Snapshots cannot be restored if the file system is WORM-enabled or if Veritas Data Deduplication is configured on the file system.

To restore a snapshot

- ◆ To restore a snapshot, enter the following:

```
Storage> snapshot restore snapshot_name fs_name
```

snapshot_name Specifies the name of the snapshot to be restored.

fs_name Specifies the name of the file system to be restored.

About snapshot schedules

The `Storage> snapshot schedule` commands let you automatically create or remove snapshots for a file system at a specified time. The schedule indicates the time for the snapshot operation as values for minutes, hour, day-of-the-month, month, and day-of-the-week. The schedule stores these values in the crontab along with the name of the file system.

For example, `snapshot schedule create schedule1 fs1 30 2 * * *` automatically creates a snapshot every day at 2:30 AM, and does not create snapshots every two and a half hours. If you wanted to create a snapshot every two and a half hours with at most 50 snapshots per schedule name, then run `snapshot schedule create schedule1 fs1 50 */30 */2 * * *`, where the value `*/2` implies that the schedule runs every two hours. You can also specify a step value for the other parameters, such as day-of-month or month and day-of-week as well, and you can use a range along with a step value. Specifying a range in addition to the `numeric_value` implies the number of times the crontab skips for a given parameter.

Automated snapshots are named with the schedule name and a time stamp corresponding to their time of creation. For example, if a snapshot is created using the name `schedule1` on February 27, 2016 at 11:00 AM, the name is:

`schedule1_Feb_27_2016_11_00_01_IST.`

Note: If the primary node is being rebooted, snapshot schedules will be missed if scheduled during the reboot of the primary node.

Configuring snapshot schedules

You can use snapshot schedules to automate creation of snapshots at regular intervals. The snapshot limit defines how many snapshots to keep for each schedule.

In some instances, snapshots may skip scheduled runs.

This may happen because of the following:

- When a scheduled snapshot is set to trigger, the snapshot needs to gain a lock to begin the operation. If any command is issued from the CLI or is running through schedules, and if the command holds a lock, the triggered snapshot schedule is not able to obtain the lock, and the scheduled snapshot fails.
- When a scheduled snapshot is set to trigger, the snapshot checks if there is any instance of a snapshot creation process running. If there is a snapshot creation process running, the scheduled snapshot aborts, and a snapshot is not created.

To create a snapshot schedule

- ◆ To create a snapshot schedule, enter the following:

```
Storage> snapshot schedule create schedule_name fs_name
max_snapshot_limit minute [hour] [day_of_the_month]
[month] [day_of_the_week] [max_num_of_parallel_snapshot_removal]
[worm] [retention_period]
```

For example, to create a schedule for an automated snapshot creation of a given file system at 3:00 am every day, enter the following:

```
Storage> snapshot schedule create schedule1 fs1 100 0 3 * * *
```

When an automated snapshot is created, the entire date value is appended, including the time zone.

schedule_name	Specifies the name of the schedule corresponding to the automatically created snapshot. The <i>schedule_name</i> cannot contain an underscore ('_') as part of its value. For example, <i>sch_1</i> is not allowed.
fs_name	Specifies the name of the file system. The file system name should be a string.
max_snapshot_limit	Specifies the number of snapshots that can be created for a given file system and schedule name. The value is a numeric value between 1-366. When the number of snapshots reaches the limit, then the oldest snapshot is destroyed. If you decrease the limit for an existing schedule, then multiple snapshots may be destroyed (oldest first) until the number of snapshots is less than the maximum snapshot limit value. Note: If you need to save daily snapshots for up to one year, the <i>max_snapshot_limit</i> is 366.
minute	This parameter may contain either an asterisk like <i>*/15</i>", which implies every 15 minutes, or a numeric value between 0-59. Note: If you are using the <i>*/xx</i>' format, the smallest value for <i>'xx'</i> is 15. You can enter <i>*/(15-59)</i> or a range such as 23-43. An asterisk (*) is not allowed.

hour	This parameter may contain either an asterisk, (*), which implies "run every hour," or a number value between 0-23. You can enter */(0-23), a range such as 12-21, or just the *.
day_of_the_month	This parameter may contain either an asterisk, (*), which implies "run every day of the month," or a number value between 1-31. You can enter */(1-31), a range such as 3-22, or just the *.
month	This parameter may contain either an asterisk, (*), which implies "run every month," or a number value between 1-12. You can enter */(1-12), a range such as 1-5, or just the *. You can also enter the first three letters of any month (must use lowercase letters).
day_of_the_week	This parameter may contain either an asterisk (*), which implies "run every day of the week," or a numeric value between 0-6. Crontab interprets 0 as Sunday. You can also enter the first three letters of the week (must use lowercase letters).
max_num_of_parallel_snapshot_removal	Specifies the number of snapshots that can be removed in parallel for a given file system and schedule name.
worm	Enables WORM. The possible values are <code>worm=yes/no</code>. Default value of this parameter is <code>no</code>. To create a schedule with <code>worm=yes</code>, the file system should be WORM-enabled.
retention_period	Specifies the retention period. The format can be [1-9](d D m M y Y) or yyyy-mm-dd.

For example, the following command creates a schedule `schedule1` for automated snapshot creation of the `fs1` file system every 3 hours each day, and maintains only 30 snapshots:

```
Storage> snapshot schedule create schedule1 fs1 30 0 */3 * * *
```

To modify a snapshot schedule

- ◆ To modify a snapshot schedule, enter the following:

```
Storage> snapshot schedule modify schedule_name fs_name  
max_snapshot_limit minute [hour] [day_of_the_month]  
[month] [day_of_the_week] [max_num_of_parallel_snapshot_removal]  
[worm] [retention_period]
```

You can modify the worm and retention period values for a schedule.

For example, to modify the existing schedule so that a snapshot is created at 2:00 am on the first day of the week, enter the following:

```
Storage> snapshot schedule modify schedule1 fs1 *2**1
```

To display a snapshot schedule

- ◆ To display all of the schedules for automated snapshots, enter the following:

```
Storage> snapshot schedule show [fs_name] [schedule_name]
```

<code>fs_name</code>	Displays all of the schedules of the specified file system. If no file system is specified, schedules of all of the file systems are displayed.
<code>schedule_name</code>	Displays the schedule name. If no schedule name is specified, then all of the schedules created under <i>fs_name</i> are displayed.

Managing automated snapshots

You can remove all of the automated snapshots created by a schedule, specify that certain snapshots be preserved, or delete a schedule for a file system.

To remove all snapshots

- ◆ To automatically remove all of the snapshots created under a given schedule and file system name (excluding the preserved and online snapshots), enter the following:

```
Storage> snapshot schedule destroyall schedule_name
fs_name
```

The `destroyall` command only destroys snapshots that are offline. If some of the snapshots in the schedule are online, the command exits at the first online snapshot.

This command does not destroy snapshots which are preserved or online or have active retention period.

Note: The `Storage> snapshot schedule destroyall` command may take a long time to complete depending on how many snapshots are present that were created using schedules.

Preserved snapshots are never destroyed automatically or as part of the `destroyall` command.

Example 1: If you try to destroy all automated snapshots when two of the automated snapshots are still mounted, Access Appliance returns an error. No snapshots under the given schedule and file system are destroyed.

```
Storage> snapshot schedule destroyall schedule1 fs1
ACCESS snapshot ERROR V-288-1074 Cannot destroy snapshot(s)
schedule1_7_Dec_2009_17_58_02_UTC schedule1_7_Dec_2009_16_58_02_UTC
in online state.
```

Example 2: If you try to destroy all automated snapshots (which are in an offline state), the operation completes successfully.

```
Storage> snapshot schedule destroyall schedule2 fs1
100% [#] Destroy automated snapshots
```

To preserve snapshots

- ◆ To preserve the specified snapshots corresponding to an existing schedule and specific file system name, enter the following:

```
Storage> snapshot schedule preserve schedule_name
fs_name snapshot_name
```

snapshot_name is a comma-separated list of snapshots..

To delete a snapshot schedule

- ◆ To delete a snapshot schedule, enter the following:

```
Storage> snapshot schedule delete fs_name [schedule_name]
```

Using instant rollbacks

This chapter includes the following topics:

- [About instant rollbacks](#)
- [Creating a space-optimized rollback](#)
- [Creating a full-sized rollback](#)
- [Listing Access Appliance instant rollbacks](#)
- [Restoring a file system from an instant rollback](#)
- [Refreshing an instant rollback from a file system](#)
- [Bringing an instant rollback online](#)
- [Taking an instant rollback offline](#)
- [Destroying an instant rollback](#)
- [Creating a shared cache object for Access Appliance instant rollbacks](#)
- [Listing cache objects](#)
- [Destroying a cache object of a Access Appliance instant rollback](#)

About instant rollbacks

Instant rollbacks are volume-level snapshots. All rollback commands take a file system name as an argument and perform operations on the underlying volume of that file system.

Note: If you plan to add a tier to the file system, add the tier first and then create the rollback. If you add the tier after a rollback exists, the rollback hierarchy would have inconsistencies because the rollback is not aware of the tier.

Both space-optimized and full-sized rollbacks are supported by Access Appliance. Space-optimized rollbacks use a storage cache, and do not need a complete copy of the original volume's storage space. However, space-optimized rollbacks are not suitable for write-intensive volumes, because the copy-on-write mechanism may degrade the performance of the volume. Full-sized rollbacks use more storage, but that has little impact on write performance after synchronization is completed.

Both space-optimized rollbacks and full-sized rollbacks can be used instantly after operations such as create, restore, or refresh.

Note: When instant rollbacks exist for a volume, you cannot disable the FastResync option for a file system.

When creating instant rollbacks for volumes bigger than 1T, there may be error messages such as the following:

```
ACCESS instant_snapshot ERROR V-288-1487 Volume prepare for full-fs1-1
failed.
```

An error message may occur because the default amount of memory allocated for a Data Change Object (DCO) may not be large enough for such big volumes. You can use the `vxtune` command to change the value. The default value is 6M, which is the memory required for a 1T volume.

To change it to 15M, use the following command:

```
vxtune volpagemod_max_memsz `expr 15 \* 1024 \* 1024`
```

Creating a space-optimized rollback

To create a space-optimized rollback

- ◆ To create a space-optimized rollback for a specified file system, enter the following:

```
Storage> rollback create space-optimized rollback_name  
fs_name [cacheobj]
```

rollback_name Indicates the name of the rollback.

fs_name Indicates the name of the file system for which to create the space-optimized rollback.

cacheobj Indicates the cache object name. If the cache object is specified, then the shared cache object is used. Otherwise, Access Appliance automatically creates a cache object for the rollback.

Creating a full-sized rollback

To create a full-sized rollback for a specified file system

- ◆ To create a Access Appliance full-sized rollback for a specified file system, enter the following:

```
Storage> rollback create full-sized rollback_name  
fs_name pool
```

rollback_name Indicates the name of the rollback.

fs_name Indicates the name of the file system for which to create the full-sized rollback.

pool Indicates the name of the pool on which to create the full-sized rollback.

The disks used for the rollback are allocated from the specified pool.

Listing Access Appliance instant rollbacks

To list Access Appliance instant rollbacks

- ◆ To list Access Appliance instant rollbacks, enter the following:

```
Storage> rollback list [fs_name]
```

where *fs_name* is the name of the file system where you want to list the instant rollbacks.

If no file system is specified, instant rollbacks are displayed for all the file systems.

Restoring a file system from an instant rollback

Prior to restoring a file system by a specified rollback, the file system should be offline.

See [“Taking an instant rollback offline”](#) on page 402.

To restore a file system from an instant rollback

- 1 To restore a file system from an instant rollback, enter the following:

```
Storage> rollback restore fs_name rollback_name
```

fs_name Indicates the name of the file system that you want to restore.

rollback_name Indicates the name of the rollback that you want to restore.

- 2 Bring the file system back online.

See [“Bringing an instant rollback online”](#) on page 402.

Bringing the file system online may take some time depending on the size of the file system.

Refreshing an instant rollback from a file system

To refresh an instant rollback from a file system

- ◆ To refresh an instant rollback from a file system, enter the following:

```
Storage> rollback refresh rollback_name fs_name
```

rollback_name Indicates the name of the rollback that you want to refresh.

fs_name Indicates the name of the file system that you want to refresh.

Bringing an instant rollback online

You can choose to bring an instant rollback online and use it as a live file system. If the original file system is offline for some reason, the instant rollback can be used as a backup.

When an instant rollback is mounted and written to with new data, the instant rollback may no longer be suitable for use in restoring the contents of the original volume. If you chose to write to an instant rollback, create another instant rollback as a backup of the original file system.

Bringing an instant rollback online

- ◆ To bring an instant rollback online, enter the following:

```
Storage> rollback online rollback_name
```

rollback_name Indicates the name of the rollback that you want to bring online.

The instant rollback is available for read/write access just as the file system.

Taking an instant rollback offline

Taking an instant rollback offline

- ◆ To take an instant rollback offline, enter the following:

```
Storage> rollback offline rollback_name
```

rollback_name Indicates the name of the rollback that you want to take offline.

For example:

Destroying an instant rollback

The instant rollback must be in the offline state before it can be destroyed.

See [“Taking an instant rollback offline”](#) on page 402.

To destroy an instant rollback

- ◆ To destroy an instant rollback, enter the following:

```
Storage> rollback destroy rollback_name fs_name
```

rollback_name Indicates the name of the rollback that you want to destroy.

fs_name Indicates the name of the file system that you want to destroy.

For example:

Creating a shared cache object for Access Appliance instant rollbacks

You can create a shared cache object for Access Appliance instant rollbacks.

Space-optimized rollbacks use a storage cache to save the data. Using a shared cache object, cache storage can be shared by all the space-optimized rollbacks.

To create a shared cache object for Access Appliance instant rollbacks

- ◆ To create a shared cache object for Access Appliance instant rollbacks, enter the following:

```
Storage> rollback cache create cache_name [cache_size] [pool]
```

cache_name	Indicates the name of the cache object you want to create.
cache_size	Indicates the cache size for the cache object. Cache size can be specified in any units, such as M, G, or T. The size of the shared cache object should be sufficient to record changes to the file system during intervals between instant rollback refreshes. By default, the size of the cache object for an instant rollback is 20% of the total size of the parent file system. The size of the cache object is dependent on your environment.
pool	Indicates the pool for storing the cache object. For better performance, the pool used for the space-optimized rollback should be different from the pool used by the file system.

To convert an existing file system into a cache object

- 1 Select or create a file system with the layout that you want to use for the cache object. In this way, you can create cache objects with any kind of file system type. If you use an existing file system, the data on the file system is lost when you convert it to a cache object.

The following example shows how to create a file system with a file system type of striped:

```
Storage> fs create striped cobj1 100m 2 pool0
100% [#] Creating striped filesystem
```

- 2 Run the `Storage> rollback cache create` command without the `cache_size` and `pool` parameters:

```
Storage> rollback cache create cache_name
```

`cache_name` is the name of the file system from step 1 that you want to convert.

A confirmation message in the Access Appliance CLI asks if you want to convert the specified file system to a cache object.

For example, to convert the striped file system `cobj1` to a cache object:

```
Storage> rollback cache create cobj1
ACCESS rollback WARNING V-288-0 Filesystem cobj1 will be converted to
cache object.
All data on Filesystem cobj1 will be lost
ACCESS rollback WARNING V-288-0 Are you sure you want to convert cobj1
to a cache object? (yes/no)
yes
100% [#]
```

- 3 Verify that the new cache object exists:

```
Storage> rollback cache list
```

CACHE NAME	TOTAL (Mb)	USED (Mb)	(%)	AVAIL (Mb)	(%)	SDCNT
cache1	15	15 (100)		0 (0)		2
cobj1	100	4 (4)		96 (96)		0

Listing cache objects

The `Storage> rollback cache list` command allows you to list the Access Appliance instant rollbacks that are using a cache object.

To list cache objects for Access Appliance instant rollbacks

- ◆ To list cache objects for Access Appliance instant rollbacks, enter the following:

```
Storage> rollback cache list [cache_name]
```

where *cache_name* is the name of the cache object you want to display. When *cache_name* is specified, the instant rollbacks that are using the cache object are listed.

A disabled cache object is listed with '-' as the attribute. *cache2* and *mycache* are in the DISABLED state.

For example:

```
Storage> rollback cache list
```

CACHE NAME	TOTAL (Mb)	USED (Mb)	(%)	AVAIL (Mb)	(%)	SDCNT
cache1	15	15	(100)	0	(0)	2
cobj1	100	4	(4)	96	(96)	0
cache2	-	-	-	-	-	-
mycache	-	-	-	-	-	-

SDCNT is the number of subdisks that have been created on the cache object.

If the cache object is disabled for some reason, it will automatically be restarted when the `Storage> rollback cache list cache_name` command is run.

For example:

```
Storage> rollback cache list cache2
rollbacks located on cache cache2:
roll3
ACCESS rollback WARNING V-288-0 Cache object cache2 was DISABLED,
trying to restart it.
ACCESS rollback INFO V-288-0 Cache object cache2 started successfully.
```

You can choose to start the cache object, or destroy it after destroying all the instant rollbacks located on it.

See [“Destroying a cache object of a Access Appliance instant rollback”](#) on page 407.

If you did not assign a cache object, a cache object is internally created for the instant rollback.

Destroying a cache object of a Access Appliance instant rollback

To destroy a cache object of a Access Appliance instant rollback

- ◆ To destroy a cache object of a Access Appliance instant rollback, enter the following:

```
Storage> rollback cache destroy cache_name
```

where *cache_name* is the name of the cache object that you want to destroy.

You can only destroy the cache object if there is no instant rollback that is using this cache object.

Reference

- [Appendix A. Access Appliance documentation](#)

Access Appliance documentation

This appendix includes the following topics:

- [Using the Access Appliance product documentation](#)
- [About accessing the online man pages](#)

Using the Access Appliance product documentation

The latest version of the Access Appliance product documentation is available on the Veritas Services and Operations Readiness Tools (SORT) website.

<https://sort.veritas.com/documents>

You need to specify the product and the platform and apply other filters for finding the appropriate document.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The documents are updated periodically for errors or corrections.

The following documents are available for Access Appliance on the SORT site:

- *Access Appliance Administrator's Guide*
- *Access Appliance Cloud Storage Tiering Solutions Guide*
- *Veritas Access Command Reference Guide*
- *Access Appliance Release Notes*
- *Veritas Access RESTful API Guide*

- *Access Appliance Solutions Guide for Enterprise Vault*
- *Access Appliance Solutions Guide for NetBackup*
- *Access Appliance Troubleshooting Guide*
- *Access Appliance Command Reference Guide*
- *Access Appliance Hardware Installation Guide*
- *Access Appliance Initial Configuration Guide*
- *Access Appliance Product Description*
- *Access Appliance Safety and Maintenance Guide*
- *Access Appliance Third-party Legal Notices Guide*
- *Access Appliance Upgrade Guide*

About accessing the online man pages

You access the online man pages by typing `man name_of_command` at the command line.

The example shows the result of entering the `Network> man ldap command`.

```
Network> man ldap
```

```
NAME
```

```
ldap - configure LDAP client for authentication
```

```
SYNOPSIS
```

```
ldap enable
```

```
ldap disable
```

```
ldap show [users|groups|netgroups]
```

```
ldap set {server|port|basedn|binddn|ssl|rootbinddn|users-basedn|  
groups-basedn|netgroups-basedn|password-hash} value
```

```
ldap get {server|port|basedn|binddn|ssl|rootbinddn|  
users-basedn|groups-basedn|netgroups-basedn|password-hash}
```

You can also type a question mark (?) at the prompt for a list of all the commands that are available for the command mode that you are in. For example, if you are within the `admin` mode, if you type a question mark (?), you will see a list of the available commands for the `admin` mode.

```
ACCESS> admin ?
```

```
Entering admin mode...
```

```
ACCESS.Admin>
```

<code>exit</code>	--return to the previous menus
<code>logout</code>	--logout of the current CLI session
<code>man</code>	--display on-line reference manuals
<code>passwd</code>	--change the administrator password
<code>show</code>	--show the administrator details
<code>supportuser</code>	--enable or disable the support user
<code>user</code>	--add or delete an administrator

To exit the command mode, enter the following: `exit`.

For example:

```
ACCESS.Admin> exit  
ACCESS>
```

To exit the system console, enter the following: `logout`.

For example:

```
ACCESS> logout
```

Index

A

About

- concurrent access 274
- FastResync 238
- SNMP notifications 217

about

- Access Appliance continuous replication 353
- Access Appliance episodic replication 324
- Active Directory (AD) 69
- bonding Ethernet interfaces 43
- buckets and objects 163
- changing share properties 306
- configuring Access Appliance for CIFS 120
- configuring CIFS for AD domain mode 123
- configuring disks 81
- configuring in IPv4 and IPv6 mixed mode 59
- configuring routing tables 53
- configuring storage pools 81
- Continuous replication failover and failback 376
- creating and maintaining file systems 228
- episodic replication job failover and failback 347
- Ethernet interfaces 48
- leaving AD domain 125
- managing CIFS shares 295
- Multi-protocol support for NFS with S3 169
- NFS file sharing 278
- setting trusted domains 127
- shares 273
- snapshot schedules 391
- snapshots 384
- storage provisioning and management 80
- storing account information 138
- striping file systems 235
- the IP addresses for the Ethernet interfaces 48

about Access Appliance

- as an iSCSI target 100

About Configuring

- network 42

about maximum IOPS 249

about S3

- viewing information 173

Access

- accessing the CLISH 21

Access Appliance

- about 16
- key features 16
- product documentation 409

Access Appliance continuous replication

- about 353
- description of feature 354
- starting 359

Access Appliance episodic replication

- about 324
- compression 345
- scheduling 336
- starting 326

Access Appliance Replication

- description of feature 325

Access Appliance to the Kerberos realm

- adding and configuring 117

accessing

- Access Appliance product documentation 409
- CLISH 21
- disk details 98
- episodic replication destinations 347
- man pages 410

Active Directory

- setting the trusted domains for 137

Active Directory (AD)

- about 69
- joining Access Appliance to 123
- verifying Access Appliance has joined successfully 124

activity

- monitoring 212

AD domain mode

- changing domain settings 125
- configuring CIFS 123
- security settings 125

AD interface

- using 126

AD trusted domains

- disabling 137

- adding
 - a column to a file system 258
 - a syslog server 217
 - CIFS share 306
 - IP address to a cluster 49
 - mirror to a file system 256
 - VLAN interfaces 52
- adding a mapping
 - between CIFS and NFS users 145
- adding and configuring
 - Access Appliance to the Kerberos realm 117
- aio_fork option
 - setting 154
- alerts
 - file system removing 269
- allowing
 - specified users and groups access to the CIFS share 309
- as an iSCSI target
 - about Access Appliance 100
- Authenticate
 - NFS clients 115
- authenticating
 - NFS clients using Kerberos 116
- authentication
 - configuring the LDAP client using the CLI 137

B

- best practices
 - creating file systems 231
- bind distinguished name
 - setting for LDAP server 66
- bonding
 - Ethernet interfaces 44
- bonding Ethernet interfaces
 - about 43
- bringing
 - file system online or offline 255
- buckets and objects
 - about 163

C

- cache object
 - destroying for an instant rollback 407
- cache objects
 - listing 405

- changing
 - an IP address to online
 - on any running node 49
 - domain settings for AD domain mode 125
 - local CIFS user password 154
 - security settings 123
 - share properties about 306
- checking
 - and repairing a file system 265
 - for stale mirrors on file systems 267
 - on the status of the NFS server 113
- CIFS
 - allowing specified users and groups access to the CIFS share 309
 - configuring schema extensions 133
 - denying specified users and groups access to the CIFS share 310
 - export options 302
 - mapuser commands 145
 - standalone mode 121
 - using multi-domain controller support 124
- CIFS aio_fork option
 - setting 154
- CIFS and NFS protocols
 - share directories 275
- CIFS clustering modes
 - about 120
- CIFS data migration
 - enabling 156
- CIFS home directories
 - quotas 89
- CIFS operating modes
 - about 120
- CIFS server
 - starting 142
 - trusted domains that are allowed access 128
- CIFS server status
 - standalone mode 122
- CIFS service
 - standalone mode 122
- CIFS share
 - adding 306
 - deleting 311
 - exporting as a directory 300
 - exporting the same file system/directory as a different CIFS share 301
 - making shadow copy aware 313
 - modifying 312

- CIFS share and home directory
 - migrating from ctdb to normal clustering mode 153
- CIFS share for Enterprise Vault with replication
 - creating 313
- CIFS share for Enterprise Vault without replication
 - creating 314
- CIFS share with continuous replication
 - creating 295
- CIFS share with encryption
 - creating 297
- CIFS share with episodic replication
 - creating 296
- CIFS share with episodic replication and encryption
 - creating 298
- CIFS share without replication and encryption
 - creating 299
- CIFS shares and home directories
 - migrating from ctdb clustering modes 152
 - migrating from normal to ctdb clustering mode 152
- CIFS snapshot
 - exporting 311
- CIFS/NFS sharing
 - mapping user names 143
- clearing
 - DNS domain names 46
 - DNS name servers 46
 - LDAP configured settings 66
- CLI
 - configure and manage storage 91
- client configurations
 - displaying 68
 - LDAP server 68
- cluster
 - adding an IP address to 49
 - changing an IP address to online for any running node 49
 - displaying all the IP addresses for 49
- clustering modes
 - ctdb 151
- collect logs
 - datacollect 220
- columns
 - adding or removing 258
- command history
 - displaying 24
- Command-Line Interface (CLI)
 - getting help on how to use 22
- communicating
 - source and destination clusters 328, 361
- Concurrent access
 - about 274
- concurrent access
 - NFS and S3 277
- Configure and manage storage
 - CLI 91
- Configuring
 - Object Store server 158
- configuring
 - Access Appliance for CIFS 120
 - AD schema with CIFS-schema extensions 133
 - CIFS for standalone mode 121
 - event notifications 212, 215
 - IP routing 53
 - IPv4 and IPv6 mixed mode 59
 - job resynchronization 346
 - NFS client for ID mapping 115
 - NIC devices 53
 - NSS lookup order 75
 - VLAN interfaces 52
 - Windows Active Directory as an IDMAP backend 132
- Configuring Access Appliance
 - ID mapping for NFS version 4 114
- configuring CIFS share
 - secondary storage for an Enterprise Vault store 300
- configuring disks
 - about 81
- configuring Ethernet interfaces
 - about 48
- configuring routing tables
 - about 53
- configuring storage pools
 - about 81
- Configuring the Object Store server
 - use case 158
- Considerations
 - creating a file system 231
- continuous replication
 - display status 368
 - setting up file system 365
 - unconfiguring 374
- Continuous replication failover and failback
 - about 376
- converting
 - existing file system into a cache object 403

- creating
 - CIFS share for Enterprise Vault with replication 313
 - CIFS share for Enterprise Vault without replication 314
 - CIFS share with continuous replication 295
 - CIFS share with encryption 297
 - CIFS share with episodic replication 296
 - CIFS share with episodic replication and encryption 298
 - CIFS share without replication and encryption 299
 - full-sized rollback 400
 - iSCSI target 109
 - local CIFS group 155
 - local CIFS user 154
 - mirrored file systems 250
 - mirrored-stripe file systems 250
 - NFS share with continuous replication 279
 - NFS share with encryption 281
 - NFS share with episodic replication 280
 - NFS share with episodic replication and encryption 282
 - NFS share without replication and encryption 283
 - shared cache object 403
 - simple file systems 250
 - snapshot schedules 392
 - snapshots 386
 - space-optimized instant rollbacks 400
 - storage pools 82, 96
 - striped file systems 250
 - striped-mirror file systems 250
- Creating a file system
 - considerations 231
- creating and maintaining file systems
 - about 228
- creating file systems
 - best practices 231
- ctdb clustering mode
 - about 151
 - directory-level share support 300
 - switching the clustering mode 152
- current Ethernet interfaces and states
 - displaying 48
- D**
- data flow
 - continuous replication synchronous mode 356
- Data Insight
 - integrating Access Appliance 316
- datacollect
 - device logs 220
- decreasing
 - size of a file system 261
- defining
 - what to replicate 338
- defrag 263
- defragmenting
 - file systems 263
- deleting
 - CIFS share 311
 - home directories 150
 - home directory of given user 150
 - local CIFS group 155
 - local CIFS user 154
 - NFS options 289
 - route entries from routing tables of nodes in cluster 53
 - snapshot schedules 395
 - syslog server 217
 - VLAN interfaces 52
- denying
 - specified users and groups access to the CIFS share 310
- description of Access Appliance continuous replication 354
- description of Access Appliance Replication 325
- destroying
 - a file system 270
 - cache object of an instant rollback 407
 - instant rollbacks 403
 - snapshots 388
 - storage pools 82
- directories
 - displaying exported 288
 - unexporting the share 289
- directory-level share support
 - ctdb clustering mode 300
- disabling
 - AD trusted domains 137
 - creation of home directories 150
 - DNS settings 46
 - FastResync option 266
 - LDAP clients
 - configurations 68
 - NIS clients 73
 - NTLM 126

- disabling *(continued)*
 - quota limits used by snapshots 388
- disk layout versions
 - upgrading 270
- Disk quotas
 - CIFS 85
 - file systems 85
 - usage 85
- disks
 - accessing details about 98
 - viewing information about 97
- displaying
 - all the IP addresses for cluster 49
 - command history 24
 - current Ethernet interfaces and states 48
 - DNS settings 46
 - events on the console 218
 - exported directories 288
 - file system alert values 269
 - file systems that can be exported 284
 - home directory usage information 150
 - information for all disk devices for the nodes in a cluster 92
 - LDAP client configurations 68
 - LDAP configured settings 66
 - list of syslog servers 217
 - local CIFS group 155
 - local CIFS user 154
 - NFS statistics 114
 - NIS-related settings 73
 - NSS configuration 75
 - routing tables of the nodes in the cluster 53
 - share properties 307
 - snapshot quotas 388
 - snapshot schedules 395
 - snapshots 387
 - snapshots that can be exported 284
 - values of the configured syslog server 217
 - VLAN interfaces 52
- displaying a mapping
 - between CIFS and NFS users 145
- displaying WWN information 93
- DNS
 - domain names
 - clearing 46
 - name servers
 - clearing 46
 - specifying 46
- DNS *(continued)*
 - settings
 - disabling 46
 - displaying 46
 - enabling 46
- domain
 - setting 142
 - setting user name 142
- domain controller
 - setting 142
- domain name
 - for the DNS server
 - setting 46
- E**
 - enabling
 - CIFS data migration 156
 - DNS settings 46
 - FastResync for a file system 265
 - LDAP client configurations 68
 - NIS settings 73
 - NTLM 126
 - quota limits used by snapshots 388
 - enabling, disabling, and displaying
 - file system quotas 86
 - encryption
 - KMS 229
 - Enterprise Vault store
 - configuring CIFS share as secondary storage 300
 - episodic replication destination file system behavior
 - about 347
 - episodic replication destinations
 - accessing 347
 - episodic replication job
 - displaying status 345
 - enabling compression 345
 - managing 340
 - show job 345
 - episodic replication jobs
 - maximum number of parallel 340
 - episodic replication unit
 - setting up files to exclude 334
 - Ethernet interfaces
 - bonding 44
 - event notifications
 - configuring 212, 215
 - events
 - displaying on the console 218
 - monitoring 214

- excluding directories and files
 - setting up 334
- export options
 - CIFS 302
- exporting
 - an NFS share 284
 - CIFS snapshot 311
 - directory as a CIFS share 300
 - NFS snapshot 293
 - same file system/directory as a different CIFS share 301
- exporting for Kerberos authentication
 - NFS share 290

F

- failover and failback
 - about 376
 - about episodic replication 347
- FastResync
 - about 238
- file system
 - converting into a cache object 403
- file system alert values
 - displaying 269
- file system alerts
 - removing 269
 - setting 268
- file system quotas
 - enabling, disabling, and displaying 86
 - setting and displaying 87
- file systems
 - adding a mirror to 256
 - bringing online or offline 255
 - checking and repairing 265
 - checking for stale mirrors 267
 - creating 250
 - decreasing the size of 261
 - defragmenting 263
 - destroying 270
 - disabling FastResync option 266
 - enabling FastResync 265
 - increasing the size of 259
 - listing with associated information 256
 - removing a mirror from 256
 - restoring from an instant rollback 401
 - that can be exported
 - displayed 284
 - types of layout 234

- filter
 - about 212, 215
- firewall setting 56
- forcefully
 - importing new LUNs for new or existing pools 94
- fsck
 - about 239

G

- group membership
 - managing 154

H

- hiding
 - system files when adding a CIFS normal share 308
- history command
 - using 24
- home directories
 - setting up 148
- home directory file systems
 - setting 147
- home directory of given user
 - deleting 150
- home directory usage information
 - displaying 150
- hostname or IP address
 - setting for LDAP server 66
- how to use
 - Command-Line Interface (CLI) 22

I

- ID mapping for NFS version 4
 - configuring Access Appliance 114
- importing
 - new LUNs forcefully for new or existing pools 94
- increasing
 - size of a file system 259
- initiating host discovery of LUNs 94
- instant rollbacks
 - about 398
 - bringing online 402
 - creating a shared cache object 403
 - creating full-sized 400
 - creating space-optimized 400
 - destroying 403
 - listing 401
 - refreshing from a file system 402

- instant rollbacks *(continued)*
 - restoring a file system from 401
 - taking offline 402
- integrating Access Appliance with Data Insight 316
- IP addresses
 - adding to a cluster 49
 - displaying for the cluster 49
 - modifying 49
 - removing from the cluster 49
- IP addresses for the Ethernet interfaces
 - about 48
- IP routing
 - configuring 53
- iSCSI target
 - creating 109
- iSCSI target service
 - managing 101
- iSCSI targets
 - managing 102

J

- job resynchronization
 - configuring 346
- joining
 - Access Appliance to Active Directory (AD) 123

K

- Kerberos authentication
 - authenticating NFS clients 116
- Kerberos realm
 - adding and configuring Access Appliance for 117
- Kerberos share
 - mounting from the NFS client 291
- Kernel-based
 - NFS server 113

L

- layouts
 - types of file system 234
- LDAP
 - before configuring 64
- LDAP client
 - configuring for authentication using the CLI 137
- LDAP password hash algorithm
 - setting password for 66
- LDAP server
 - clearing configured settings 66
 - disabling client configurations 68

LDAP server *(continued)*

- displaying client configurations 68
- displaying configured settings 66
- enabling client configurations 68
- setting over SSL 66
- setting port number 66
- setting the base distinguished name 66
- setting the bind distinguished name 66
- setting the hostname or IP address 66
- setting the root bind DN 66
- setting the users, groups, and netgroups base DN 66
- leaving
 - AD domain 125
- listing
 - all file systems and associated information 256
 - cache objects 405
 - free space for storage pools 82
 - instant rollbacks 401
 - storage pools 82
- local CIFS groups
 - creating 155
 - deleting 155
 - displaying 155
- local CIFS user
 - creating 154
 - deleting 154
 - displaying 154
- local CIFS user password
 - changing 154
- local user and groups
 - managing 154
- log files
 - introduction 219
- LUNs
 - initiating host discovery 94
 - managing 104
 - provisioning 109

M

- man pages
 - how to access 410
- managing
 - CIFS shares 295
 - continuous replication 367
 - group membership 154
 - home directories 146
 - iSCSI target service 101
 - iSCSI targets 102

- managing *(continued)*
 - local users and groups 154
 - LUNs 104
 - mappings with iSCSI initiators 108
 - users 109
- managing NFS shares
 - using netgroups 288
- mapping
 - of UNIX users from LDAP to Windows users 146
- mappings with iSCSI initiators
 - managing 108
- mapuser commands
 - about 145
- maximum IOPS
 - setting 250
- maximum number
 - parallel episodic replication jobs 340
- migrating
 - CIFS share and home directory from ctdb to normal clustering mode 153
 - CIFS shares and home directories 152
 - CIFS shares and home directories from normal to ctdb clustering mode 152
- mirrored file systems
 - creating 250
- mirrored-stripe file systems
 - creating 250
- modifying
 - an IP address 49
 - CIFS share 312
 - snapshot schedules 395
- monitoring
 - activity 212
 - events 214
- more command
 - using 24
- mounting
 - NFS share from the NFS client 291
- mounting snapshots 390
- Multi-protocol support for NFS with S3
 - limitations 169

N

- naming requirement for new users 35
- navigating CLISH
 - Access 22
- NFS client
 - configuring for ID mapping 115

- NFS clients
 - authenticating 115
- NFS file sharing
 - about 278
- NFS options
 - deleting 289
- NFS server
 - about 112
 - checking on the status 113
 - kernel-based 113
 - starting 113
 - stopping 113
- NFS share
 - exporting 284
 - exporting for Kerberos authentication 290
- NFS share with continuous replication
 - creating 279
- NFS share with encryption
 - creating 281
- NFS share with episodic replication
 - creating 280
- NFS share with episodic replication and encryption
 - creating 282
- NFS share without replication and encryption
 - creating 283
- NFS shares
 - managing using netgroups 288
- NFS snapshot
 - exporting 293
- NFS statistics
 - displaying 114
 - resetting 114
- NIC devices
 - configuring 53
- NIS
 - clients
 - disabling 73
 - enabling 73
 - domain name
 - setting on all the nodes of cluster 73
 - related settings
 - displaying 73
 - server name
 - setting on all the nodes of cluster 73
- node
 - in a cluster
 - displaying information for all disk devices 92
- NSS
 - displaying configuration 75

NSS *(continued)*
 lookup order
 configuring 75

NTLM
 disabling 126
 enabling 126

O

object server 157
 object server group-specific parameters
 setting up 172
 Object Store server
 configuring 158
 objectstore buckets 166
 offline
 taking an instant rollback offline 402
 online
 bringing an instant rollback online 402

P

planned failback
 process overview 350, 378
 planned failover
 process overview 349, 377
 preserving
 snapshot schedules 395
 Private network
 configure 42
 provisioning
 LUNs 109
 Public network
 configure 42

Q

quota commands
 setting and displaying file system quotas 87
 quota limits
 enabling or disabling snapshot 388
 quotas
 CIFS home directories 89
 setting user quotas for users of specified
 groups 89

R

refreshing
 instant rollbacks from a file system 402
 removing
 a column from a file system 258

removing *(continued)*
 IP address from the cluster 49
 mirror from a file system 256
 snapshot schedules 395
 removing a mapping
 between CIFS and NFS users 145
 renaming
 storage pools 82
 replicating file systems
 setting up 332
 resetting
 NFS statistics 114
 restoring
 a file system from an instant rollback 401
 snapshots 391
 resynchronizing
 stale mirrors on file systems 267
 route entries
 deleting from routing tables 53
 routing tables
 of the nodes in the cluster
 displaying 53

S

scheduling
 episodic replication 336
 security
 standalone mode 122
 security settings
 AD domain mode 125
 changing 123
 setting
 aio_fork option 154
 base distinguished name for the LDAP server 66
 bind distinguished name for LDAP server 66
 domain 142
 domain controller 142
 domain name for the DNS server 46
 domain user name 142
 file system alerts 268
 filter of the syslog server 217
 home directory file systems 147
 IDMAP backend to ad for access to CIFS 131
 IDMAP backend to hash for accessing CIFS 131
 IDMAP backend to ldap for trusted domain access
 to CIFS 130
 IDMAP backend to rid for access to CIFS 129
 LDAP server hostname or IP address 66
 LDAP server over SSL 66

setting *(continued)*

- LDAP server port number 66
- LDAP users, groups, and netgroups base DN 66
- maximum IOPS 250
- NIS domain name on all the nodes of cluster 73
- prior to configuring LDAP 64
- retention 240
- root bind DN for the LDAP server 66
- severity of the syslog server 217
- the NIS server name on all the nodes of cluster 73
- trusted domains 127
- trusted domains for the Active Directory 137
- user quotas for users of specified groups 89
- WORM over NFS 241
- WORM-retention over CIFS 242

setting up

- home directories 148
- object server group-specific parameters 172
- replicating file systems 332

setting up an episodic replication unit

- to exclude directories and files 334

severity levels

- about 212, 215

shadow copy

- making a CIFS share aware 313

share directories

- CIFS and NFS protocols 275

share properties

- displaying 307

shared cache object

- creating 403

shares

- about 273

showing

- snapshot schedules 395

snapshot schedules

- about 391
- creating 392
- deleting 395
- displaying 395
- modifying 395
- preserving 395
- removing 395
- showing 395

snapshots

- about 384
- creating 386
- destroying 388

snapshots *(continued)*

- displaying 387
- displaying quotas 388
- enabling or disabling quota limits 388
- mounting 390
- restoring 391
- that can be exported
 - displayed 284
- unmounting 390

SNMP notifications

- about 217

source and destination clusters

- communicating 328, 361

specifying

- DNS name servers 46

SSL

- setting the LDAP server for 66

standalone mode

- CIFS server status 122
- CIFS service 122
- security 122

starting

- Access Appliance continuous replication 359
- Access Appliance episodic replication 326
- CIFS server 142
- NFS server 113

stopping

- NFS server 113

storage pools

- creating 82, 96
- destroying 82
- listing 82
- listing free space 82
- renaming 82

storage provisioning and management

- about 80

storing

- account information 138
- user and group accounts in LDAP 140
- user and group accounts locally 140

striped file systems

- creating 250

striped-mirror file systems

- creating 250

striping file systems

- about 235

switching

- ctdb clustering mode 152

- syslog server
 - adding 217
 - deleting 217
 - displaying the list of 217
 - displaying the values of 217
 - setting the filter of 217
 - setting the severity of 217
- system files
 - hiding when adding a CIFS normal share 308

T

- troubleshooting
 - about 211
- trusted domains
 - allowing access to CIFS when setting an IDMAP backend to ad 131
 - allowing access to CIFS when setting an IDMAP backend to hash 131
 - allowing access to CIFS when setting an IDMAP backend to ldap 130
 - allowing access to CIFS when setting an IDMAP backend to rid 129
 - specifying which are allowed access to the CIFS server 128

U

- unexporting
 - share of exported directory 289
- UNIX users from LDAP to Windows users
 - automatic mapping of 146
- unmounting snapshots 390
- unplanned failback
 - process overview 351, 379
- unplanned failover
 - process overview 351, 379
- upgrading
 - disk layout versions 270
- Use case
 - configuring the Object Store server 158
- user and group accounts in LDAP
 - storing 140
- user and group accounts locally
 - storing 140
- user names
 - mapping for CIFS/NFS sharing 143
- users
 - managing 109

- using
 - AD interface 126
 - history command 24
 - more command 24
 - multi-domain controller support in CIFS 124
- Using NFS Server 112

V

- verifying
 - Access Appliance has joined Active Directory (AD) 124
- viewing
 - information about disks 97
- viewing information
 - about S3 173
- VLAN
 - adding interfaces 52
 - configuring interfaces 52
 - deleting interfaces 52
 - displaying interfaces 52

W

- what to replicate
 - defining 338
- Windows Active Directory
 - configuring as an IDMAP backend 132
- workflow
 - object server 157
- WORM over NFS
 - setting 241
- WORM-retention over CIFS
 - setting 242
- WWN information
 - displaying 93