NetBackup™ Access Appliance 8.6 Administrator's Guide

April 2026

VERITAS
now part of Cohesity

Access Appliance Administrator's Guide

Last updated: 2026-04-14

Legal Notice

Copyright © 2026 Cohesity All rights reserved.

© 2026 Cohesity All Rights Reserved. Veritas, the Veritas Logo and other Veritas Marks are trademarks of Cohesity in the US and/or internationally. The information supplied herein is the confidential and proprietary information of Veritas and may only be used (a) by the intended recipients and (b) in conjunction with validly licensed Veritas software and services. Find the terms of Veritas licenses at www.cohesity.com/agreements.

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. COHESITY 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.

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to Cohesity Support, to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the Cohesity Support Portal to create a new case.
- Click the (**?**) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the Cohesity Knowledge Base.
- Log in to the Cohesity Support Portal to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

   Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Contents

Section 1 Introducing Access Appliance 14

Chapter 1 Introducing Access Appliance 15

About Access Appliance 15
About the Access Appliance Dashboard 18
Getting started with the Access CLISH 20
Accessing the Access CLISH 20
Navigating the Access CLISH 20
Getting help using the Access CLISH 20
Displaying the command history 22
Using the more command 23
Getting started with Access Appliance APIs 23

Section 2 Configuring Access Appliance 30

Chapter 2 Managing users 31

About user management 31
About the naming requirements for adding new users 33
About the naming requirements when assigning roles to domain users 34
Adding and removing user roles using GUI 35
Performing user management using CLISH 36

Chapter 3 Managing licenses 39

About Access Appliance licenses 39
Managing licenses from the GUI 40
Managing licenses from the command-line interface 41

Chapter 4 Configuring the network 43

About configuring the Access Appliance network 44
About bonding Ethernet interfaces 44
Bonding Ethernet interfaces 45
Considerations for configuration a LACP bond 47

Configuring DNS settings 50
About Ethernet interfaces 52
Displaying current Ethernet interfaces and states 52
Configuring IP addresses 52
Configuring IP addresses and FQDNs in a non-DNS environment 56
Configuring VLAN interfaces 56
Configuring NIC devices 57
About configuring routing tables 58
Configuring routing tables 58
Changing the firewall settings 60
Configuring Access Appliance in IPv4 and IPv6 mixed mode 64
Support for multiple data subnets 67
Adding console FQDN to the network and accessing the GUI using the console FQDN 68

Chapter 5 Configuring authentication services 70

About configuring LDAP settings 70
Configuring LDAP server settings 71
Administering the Access Appliance cluster's LDAP client 74
About Active Directory (AD) 75
Configuring AD server settings 75
Configuring entries for Access Appliance DNS for authenticating to Active Directory (AD) 77
Configuring AD/LDAP using the GUI 78
Configuring NSS lookup order 78
Sign-in options for the Access Appliance UI 80
Configuring user authentication using digital certificates or smart cards 80
Adding CA certificates for smart card authentication 82
Deleting CA certificates 82

Section 3 Managing Access Appliance storage 83

Chapter 6 Configuring storage 84

About storage provisioning and management 84
About configuring disks 85
About configuring storage pools 85
Configuring storage pools 86
About quotas for usage 89

Enabling, disabling, and displaying the status of file system quotas 90
Setting and displaying file system quotas 91
Setting user quotas for users of specified groups 93
Displaying information for all disk devices associated with the nodes in a cluster 94
Displaying WWN information 95
Importing new LUNs forcefully for new or existing pools 96
Initiating host discovery of LUNs 96
Adding a node to the cluster using the GUI 97

Chapter 7 Managing disks 100

Adding disks to a storage pool 100
Removing disks from a storage pool 101
Viewing information about disks 101
Accessing disk details 102
Discovering disks 103
Formatting a disk 103

Section 4 Managing Access Appliance file access services 104

Chapter 8 Configuring the NFS server 105

About using the NFS server with Access Appliance 105
Using the kernel-based NFS server 106
Accessing the NFS server 106
Displaying and resetting NFS statistics 107
Configuring Access Appliance for ID mapping for NFS version 4 108
Configuring the NFS client for ID mapping for NFS version 4 108
About authenticating NFS clients 109
Setting up Kerberos authentication for NFS clients 109
Adding and configuring Access Appliance to the Kerberos realm 110

Chapter 9 Using Access Appliance as an Object Store server 112

About the Object Store server 112
Use cases for configuring the Object Store server 113
Configuring the Object Store server 113
About buckets and objects 117

File systems used for objectstore buckets 120
Enabling WORM on buckets 121
Object Access SSL certificate 123
Object Access endpoints 123
S3 with NFS use case 124
S3 with NSP use case 126
Configuring the S3 server using GUI 126
Setting Object Access server default parameters 126
Setting up the Object Access server group-specific parameters 127
Viewing information about S3 128
Provisioning the S3 bucket using GUI 128

Section 5 Managing Access Appliance security 133

Chapter 10 Managing security 134

Security overview 134
About the security meter 134
Setting the BIOS password 136

Chapter 11 Setting up FIPS mode 137

FIPS 140-2 conformance for Access Appliance 137
Viewing FIPS status for Access Appliance 138
Enabling FIPS for Access Appliance 139
Disabling FIPS mode for NetBackup MSDP 140
Managing the FIPS mode using the command-line interface 141

Chapter 12 Configuring STIG 143

STIG overview for Access Appliance 143
Enabling OS STIG hardening for Access Appliance 144
Viewing STIG status of an Access Appliance 147
Enabling STIG using the command-line interface 148

Chapter 13 Setting the banner 150

Managing the login banner using the UI 150
Managing the banner from the command-line interface 151

Chapter 14 Setting the password policy 153

Managing the password policy using the UI 153
Managing the password policy from the command-line interface 156

Chapter 15 Immutability in Access Appliance 162

Support for immutability in Access Appliance 162
About lockdown modes 163
Selecting or changing the lockdown mode 164
Accessing the root shell in lockdown mode 165
Configuring immutability using GUI 166

Chapter 16 Deploying certificates on Access Appliance 169

About system certificates on Access Appliance 169
About external certificates on Access Appliance 170
Deploying ECA using the GUI 172
Log locations 173

Chapter 17 Single Sign-On (SSO) 175

About single sign-on (SSO) configuration 175
Configuring SSO on Access Appliance 175
Limitations and log locations 177

Chapter 18 Configuring multifactor authentication 178

About multifactor authentication 178
Considerations when configuring multifactor authentication 179
Configuring multifactor authentication for your user account 179
Disabling multifactor authentication for your user account 180
Enforcing multifactor authentication for all users 181
Configuring multifactor authentication for your user account when it is enforced in the cluster 182
Resetting multifactor authentication for a user 183

Section 6 Monitoring and troubleshooting 184

Chapter 19 Monitoring the appliance 185

Testing the appliance hardware 185
About Veritas AutoSupport on the Access Appliance 186
Setting up AutoSupport on the appliance 186
Using a proxy server with the appliance 188

Setting up email notifications on the appliance 188
Setting up SNMP notifications on the appliance 190

Chapter 20 Configuring event notifications and audit logs 191

About troubleshooting 191
Monitoring command activity 192
Monitoring alerts 192
About alert management 193
Viewing information about alert management 193
Managing alerts 193
Monitoring events 194
Viewing reports 194
Viewing cluster storage usage 194
Viewing file system usage 195
About event notifications 195
About severity levels and filters 195
About SNMP notifications 196
Configuring a syslog server 197
Displaying events on the console 198

Chapter 21 Appliance log files 199

About appliance log files 199
Gathering device logs with the DataCollect command 200
Collecting default and additional diagnostic logs 203
Downloading logs using the Log Transfer Console 204
Forwarding logs to an external server 205
Configuring log forwarding using the UI 206
Modifying log forwarding settings using the UI 209
Removing log forwarding using the UI 210
Setting up log forwarding using the command-line interface 211
Viewing the configured settings from the appliance shell menu 211
Collecting logs using the UI 212
Uploading log packages to Veritas Support 214
Downloading log packages 215
Deleting log packages 215

Section 7 Provisioning and managing Access Appliance file systems 217

Chapter 22 Creating and maintaining file systems 218

About creating and maintaining file systems 218
Considerations for creating a file system 219
Best practices for creating file systems 219
Choosing a file system layout type 222
Determining the initial extent size for a file system 222
About striping file systems 223
About FastResync 226
About fsck operation 227
Enabling WORM on a file system 227
Setting retention in files 228
Setting WORM over NFS 229
About managing application I/O workloads using maximum IOPS settings 230
Creating a file system 231
Bringing the file system online or offline 235
Listing all file systems and associated information 236
Modifying a file system 236
Adding or removing a mirror from a file system 236
Adding or removing a column from a file system 238
Increasing the size of a file system 239
Decreasing the size of a file system 241
Managing a file system 243
Defragmenting a file system 243
Checking and repairing a file system 245
Configuring FastResync for a file system 245
Disabling the FastResync option for a file system 246
Checking and resynchronizing stale mirrors 246
Setting file system alerts 247
Displaying file system alert values 248
Removing file system alerts 249
Destroying a file system 249
Upgrading disk layout versions 249

Section 8 Provisioning and managing Access Appliance shares 251

Chapter 23 Creating shares for applications 252

About file sharing protocols 252
About concurrent access 253
About concurrent access with NFS and S3 253

Chapter 24 Creating and maintaining NFS shares 254

About NFS file sharing 254
About the NFS shares 255
Creating an NFS share with continuous replication 255
Creating an NFS share without replication and encryption 256
Displaying file systems and snapshots that can be exported 257
Exporting an NFS share 257
Displaying exported directories 260
About managing NFS shares using netgroups 261
Unexporting a directory or deleting NFS options 261
Exporting an NFS share for Kerberos authentication 262
Mounting an NFS share with Kerberos security from the NFS client 263
Exporting an NFS snapshot 265

Chapter 25 Integrating Access Appliance with Data Insight 266

Access Appliance integration with Data Insight 266

Section 9 Managing Access Appliance storage services 272

Chapter 26 Configuring continuous replication 273

About Access Appliance continuous replication 273
How Access Appliance continuous replication works 274
How data flows in continuous replication synchronous mode 276
How data flows in continuous replication asynchronous mode 277
Starting Access Appliance continuous replication 279

Setting up communication between the source and the destination clusters 281
Setting up the file system to replicate 285
Managing continuous replication 287
Displaying continuous replication information and status 290
Unconfiguring continuous replication 295
Preserving the file system on the destination cluster 297
Configuring a continuous replication job using the GUI 298
Stopping or starting a continuous replication job 302
Pausing and resuming a continuous replication job 302
Failing over or failing back a continuous replication job 303
Unconfiguring a continuous replication job 303
Continuous replication failover and failback 303
Process summary 304
Overview of the planned failover process 304
Overview of the planned failback process 305
Overview of the unplanned failover process 306
Overview of the unplanned failback process 306
Addition of multiple file systems to a Replicated Volume Group 307

Chapter 27 Using snapshots 311

About snapshots 311
Enabling WORM on storage snapshots 312
Creating snapshots 313
Displaying snapshots 314
Managing disk space used by snapshots 315
Bringing snapshots online or taking snapshots offline 317
Restoring a snapshot 318
About snapshot schedules 318
Configuring snapshot schedules 319
Managing automated snapshots 322

Chapter 28 Using instant rollbacks 325

About instant rollbacks 325
Creating a space-optimized rollback 327
Creating a full-sized rollback 327
Listing Access Appliance instant rollbacks 328
Restoring a file system from an instant rollback 328
Refreshing an instant rollback from a file system 329
Bringing an instant rollback online 329
Taking an instant rollback offline 329
Destroying an instant rollback 330

Creating a shared cache object for Access Appliance instant rollbacks 330
Listing cache objects 332
Destroying a cache object of a Access Appliance instant rollback 334

Section 10 Reference 335

Appendix A Access Appliance documentation 336

Using the Access Appliance product documentation 336
About accessing the online man pages 337

Index 339

Section 1

Introducing Access Appliance

- Chapter 1. Introducing Access Appliance

Chapter 1

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 *Access Appliance RESTful API Guide* 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 112. ■ NFS See "About using the NFS server with Access Appliance" on page 105.
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 230.
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 311.
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 *Access Appliance Solutions Guide for NetBackup* for more information.
Quotas	Support for setting file system quotas, user quotas, and hard quotas. See "About quotas for usage" on page 89.
Replication	Synchronous replication of data over IP networks See "About Access Appliance continuous replication" on page 273. See the `continuous(1)` man page for more information.

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

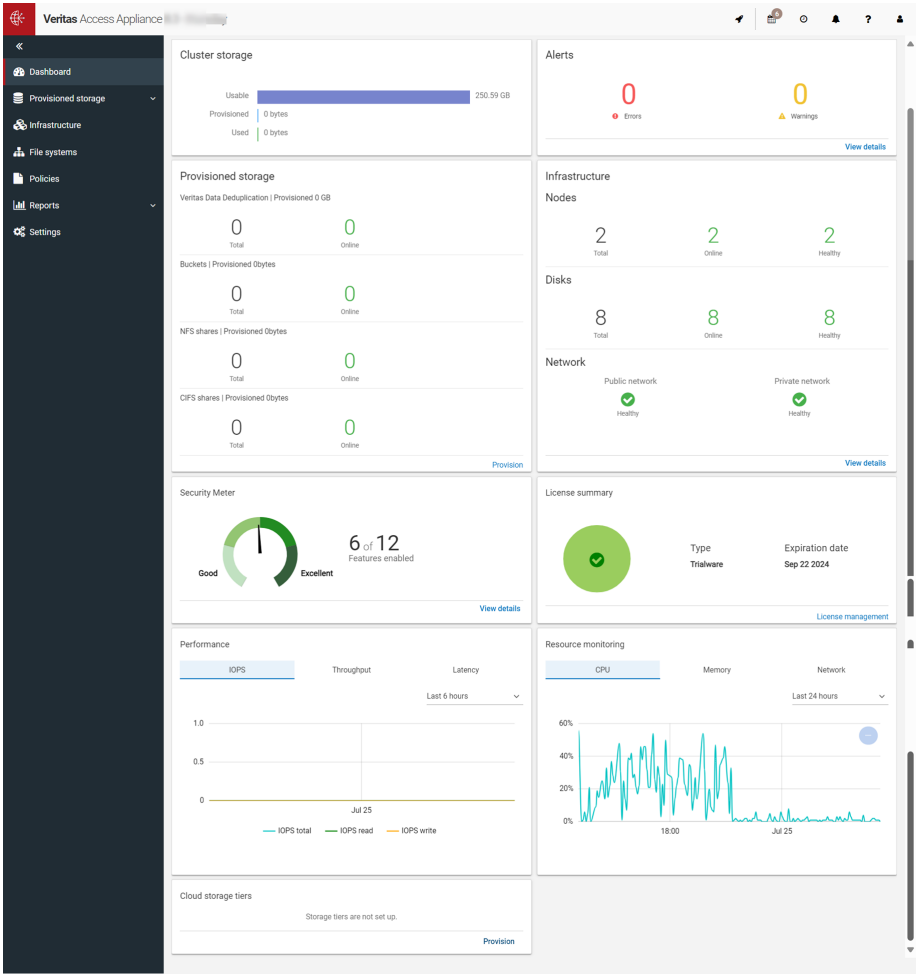
Feature	Description
Support for LDAP and AD	You can configure LDAP and AD authentication services with Access Appliance. See "About configuring LDAP settings" on page 70. See "About Active Directory (AD)" on page 75.
Partition Directory	With support for partitioned directories, directory entries are 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 `storage_fs`(1) 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 *Access Appliance Solutions Guide for NetBackup*.
FIPS	FIPS 140-2 standard is enabled by default for the Veritas Operating System (VxOS). See "Enabling FIPS for Access Appliance" on page 139.
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 144. For more information, see the *Appliance security* chapter in the *Veritas Access Appliance Initial Configuration Guide*.
Separation of management and data network	Ability to configure a separate management and data network during cluster configuration. For more information, see the *Veritas Access Appliance Initial Configuration Guide*.

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 67.
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 163.
Single node configuration	Support for 3360 single node configuration, which provides the same functionality as a two-node configuration. For more information, see the *Veritas Access Appliance Initial Configuration Guide*.

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, S3 buckets, and NFS 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 **Security meter** shows the security status of the appliance and offers recommendations to improve the appliance security.
- The **License summary** area displays details on the type of license installed - whether it a valid trialware or the type of license (perpetual/ subscription) installed.

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 Acess 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.

Table 1-4 Access command-line keyboard shortcut keys for deletions *(continued)*

Shortcut key	Description
[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 `history` 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 NFS 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
```

`enable`	Enables the more filter on all of the nodes in the cluster.
`disable`	Disables the more filter on all of the nodes in the cluster.
`status`	Displays the status of the `more` 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": "eyJ0eXAiOiJKV1QiLCJhbGciOiJSUzI1NiIsInppcCI6IkRFRiJ9.
   eNp0VF1v2jAU_S9-nHiArXQrb15yaT2cOLIdEJ2mKGWZmnXARMKEVvW_13EciO3weu65H
   -d-vaIyr9FsMp18Gt_dTsefb6bjESqrCs3Q7umYbfOqLg5ohH7XpYIWqwVfzNM1k6tHvCK
   ```

```
JMhSnv63_5MvHO-V_cztC1fFJkQ_b3bGaHar6UGxeitm_k2Lnx_o5q_cvxW44Wlll-c9t2V
jrw7FokW2-eS53hcJ-5X-qwoT5n232u7o4qfJfEQchMZfZN_YVzb6jD-jHSGOMQ0YifA9nd
Elg5dJAZg9MyDMUcMASsmUkgC-96jGYrFDoKCwgFEKgSQsPlsoETILgEsyJ4HyE2dLxEIyX
9txDJYwSoK1Xa_F01EFZcLnZJgS3MOBNxEhDhNGYi9Tr7KLdJZ4JXBYsgUM0nUtnFFfmirQ
7ZBUUZzm6EwEHClDmbTBKczEtacZK55kkXJuwrjcflm6gq45Tgne7I1_o9XF8CNPqdMTnIZE
UnYvXLIjQZPnEIoQS2yjAwPVuJLrIFioFfYy4VQ-6JY70hLgERGiv6gad2SYVR9Ya822h6m
OKI3sSXRnwSGEWBJMnTLsszTZhpbfarnh9cfI9_saN08DbzZF87m6RcSxyqAKoK63o9Rsr
JWnPd2zdq-5A8MRqUjUNllN6I4BgpQTuU44S7zzHOqQyeJdboDjAKiVor2YK8_nMutL7DQ
Jux-m_KIIYukvnrthumPCbaS9BCbwkB5j8vW0cYZGfPVrXZHqrHrvU7tfsWmY9wf0Onakt
7d3AAAA__8.VVE25rQqbrC-isGOqbRTqMPoK4ts5-9_6zSgz0fUg11m9GCClq10PS9u1D
laXye-S2MYYyHVEHSVs6uKcPVvN2WGBHkv7t-c4Hixc9O8zrJYhJaP979wF_gn08YnRlX
7_o4Qj6muc1IWHjK0hPMIgq0X-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
I1NiIsInppcCI6IkRFRiJ9.eNp0VF1v2jAU_S9-nHiArXQrb15yaT2cOLIdEJ2mKGWZmnX
ARMKEVvW_13EciO3weu65H-d-vaIyr9FsMp18Gt_dTsefb6bjESqrCs3Q7umYbfOqLg5oh
H7XpYIWqwVfzNM1k6tHvCKJMhSnv63_5MvHO-V_cztC1fFJkQ_b3bGaHar6UGxeitm_k2
Lnx_o5q_cvxW44Wlll-c9t2Vjrw7FokW2-eS53hcJ-5X-qwoT5n232u7o4qfJfEQchMZfZ
N_YVzb6jD-jHSGOMQ0YifA9ndElg5dJAZg9MyDMUcMASsmUkgC-96jGYrFDoKCwgFEKgS
QsPlsoETILgEsyJ4HyE2dLxEIyX9txDJYwSoK1Xa_F01EFZcLnZJgS3MOBNxEhDhNGYi
9Tr7KLdJZ4JXBYsgUM0nUtnFFfmirQ7ZBUUZzm6EwEHClDmbTBKczEtacZK55kkXJuw
rjcflm6gq45Tgne7I1_o9XF8CNPqdMTnIZEUnYvXLIjQZPnEIoQS2yjAwPVuJLrIFi
oFfYy4VQ-6JY70hLgERGiv6gad2SYVR9Ya822h6mOKI3sSXRnwSGEWBJMnTLsszTZh
pbfarnh9cfI9_saN08DbzZF87m6RcSxyqAKoK63o9RsrJWnPd2zdq-5A8MRqUjUNl
lN6I4BgpQTuU44S7zzHOqQyeJdboDjAKiVor2YK8_nMutL7DQJux-m_KIIYukvnr
thumPCbaS9BCbwkB5j8vW0cYZGfPVrXZHqrHrvU7tfsWmY9wf0Onakt7d3AAAA__
8.VVE25rQqbrC-isGOqbRTqMPoK4ts5-9_6zSgz0fUg11m9GCClq10PS9u1Dla
```

```
Xye-S2MYYyHVEHSVs6uKcPVvN2WGBHkv7t-c4Hixc9O8zrJYhJaP979wF_gn
08YnRlX7_o4Qj6muc1IWHjK0hPMIgq0X-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.eyJzdWIiOi
   I5OWU5ZGRlMzky NWIiLCJpYXQiOjE2NDk5MzQyMDZ9.Eyv2lmRz7pM__PL
   UdgkyevuV8dM78-ie65rr5xzVgcshWS25HuN7bmELscNIIC13n1EKILb953ne
   BfW_o9R9RA8IalmL_m6NsBvNWb6AfXbAFxVXHRAPAZnDpkkt6hls6koONkC0Ld
   I1707w7wrceZMgWMe-F8cGmffN3dSiCvG89JkSLOwL0K6gQNcwqkkzJnI5S7by
   R_YMrB5hbgWV2zJMbWIIwAJoGChWE2uhhGobKqztxo0Y1ZN-XrCk4E_AFtLrbZ
   4GjO_AddvbpYcV1Q7DdCZsKgSMpHG5N8FFwDtk3n_8zriVV0XPLnMGMadbSM
   9qm5YmO-MkuydSI07K1Q
   ```

   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

Section 2

Configuring Access Appliance

- Chapter 2. Managing users
- Chapter 3. Managing licenses
- Chapter 4. Configuring the network
- Chapter 5. Configuring authentication services

Chapter 2

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`.
- Nested LDAP group for role assignment is not supported.
- AD/LDAP users/groups must have UIDs and GIDs.
- Veritas recommends that LDAP and AD 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.
- Veritas recommends that LDAP and AD GIDs start from 10000. Otherwise, when you assign a role to the AD/LDAP user, the GID of any local group may conflict with the GID of a group associated with the user from the directory server.
- The LDAP admin should modify the home directory of the LDAP user on the LDAP server to `/home/<directory_name>` to create a home directory for that LDAP user on Access Appliance.

- Assigning the appliance administrator role to an AD user from trusted domain 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 See "Adding and removing user roles using GUI" on page 35.

To perform user management using CLISH, see See "Performing user management using CLISH" on page 36.

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 (_).
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`.

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 78.

Chapter 3

Managing licenses

This chapter includes the following topics:

- About Access Appliance licenses
- Managing licenses from the GUI
- Managing licenses from the command-line interface

About Access Appliance licenses

Access Appliance supports the following types of licenses:

- Trialware: A built-in evaluation license, which is automatically activated after the appliance initial configuration is completed and is valid for 60 days.
- Subscription: A time-bound license that must be renewed at the end of the expiration date. Support and maintenance is included in the time period.
- Perpetual: An unlimited validity license with a one-time cost for using the appliance.

Starting with version 8.2, you can add multiple subscription licenses to increase the total cumulative licensed capacity and the subscription period. You can also license a subset of the total usable storage capacity based on your storage requirements. With partial licensing, you can purchase a license based on your used capacity rather than the usable capacity of the appliance.

The following terminology is used to define the storage capacity:

- **Licensed**: Storage capacity for which a license is purchased.
- **Used**: Storage capacity that is used for data. Alerts are generated if the used storage is more than the licensed storage.
- **Available**: Total storage capacity of the appliance.
- **Overutilized**: Storage utilization that exceeds the licensed capacity.

See “Managing licenses from the GUI” on page 40.

See “Managing licenses from the command-line interface” on page 41.

Managing licenses from the GUI

You can add licenses, stack licenses, and view license details from the UI. You cannot delete licenses from the UI. You must use the Access command-line interface to delete licenses.

To add a license:

1. Sign in to the Access Appliance UI by using the `https://console-ip:14161` URL where *console-ip* is the management console IP address.
2. Click **Settings > Licensing management**.
3. Do one of the following:
   - If a trialware license exists on your system, on the Licensing management page, click **Upload license**.
   - If a subscription or a perpetual license is already installed and you want to stack additional licenses, click **Add license key**.
4. Browse to the location of your `.slf` license file and click **Add**.

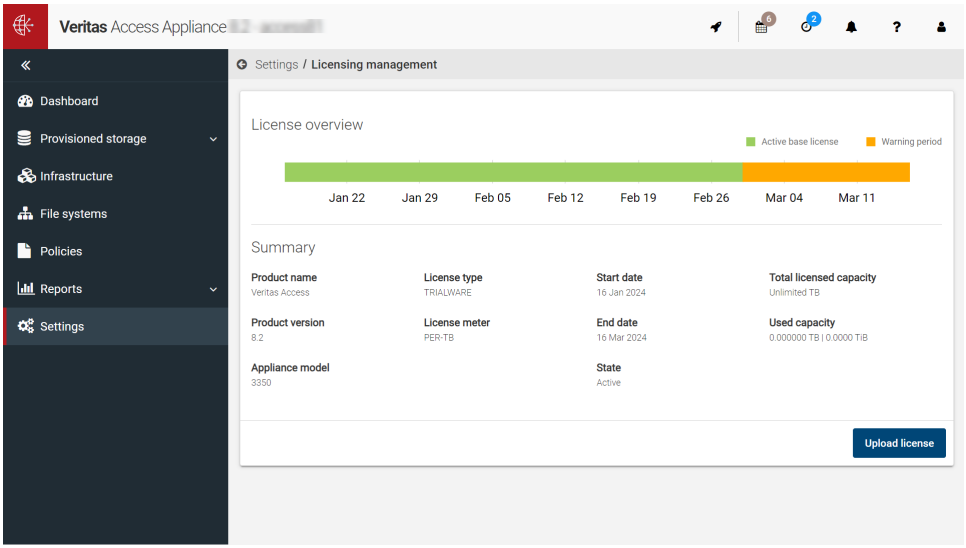
   A notification is displayed on the top of the page. To monitor the task, click **View details**. The ongoing and completed tasks for the add license operation are also displayed in **Recent activity**.

The following details are displayed on the **Settings > Licensing management** page:

- Product name: Name of the product.
- License type: Type of installed license: subscription, perpetual, or trialware.
- Start date: Start date of the license. Not shown for perpetual licenses.
- Total licensed capacity: Storage capacity for which the license is purchased.
- Product version: Version of the product
- License meter: Per-TB
- End date: Expiry date of the license. Not shown for perpetual licenses.
- Used capacity: Storage capacity used for data. Alerts are generated if the used storage is more than the licensed storage.
- Appliance model: Model number
- State: Current state of the license, active or expired.

- Available capacity: Total storage capacity of the appliance. Not shown for trialware licenses.

The following figure shows the details displayed for a trialware license:



Managing licenses from the command-line interface

You can add, view, and delete licenses from the command-line interface.

Adding licenses

To add a license use the following command:

```
system license add absolutepath_slf_license_file
```

Viewing license details

To view the details for each of the registered license keys:

```
system license list
```

To view the cumulative licensed storage for all the registered license keys:

```
system license list details
```

Deleting licenses

Use the Access command-line interface to delete a license. You cannot delete licenses from the GUI.

To delete a license:

```
system license delete slf_license_file
```

You can also specify comma separated license keys to delete multiple license keys.

Chapter 4

Configuring the network

This chapter includes the following topics:

- About configuring the Access Appliance network
- About bonding Ethernet interfaces
- Bonding Ethernet interfaces
- Considerations for configuration a LACP bond
- Configuring DNS settings
- About Ethernet interfaces
- Displaying current Ethernet interfaces and states
- Configuring IP addresses
- Configuring IP addresses and FQDNs in a non-DNS environment
- 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
- Adding console FQDN to the network and accessing the GUI using the console FQDN

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 4-1 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.

A bond can be created on two, three or four public NIC; you can create a bond on eth4, eth5, eth6, and eth7.

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.

Considerations for configuration a LACP bond

Standardized protocol is used to establish and maintain link-aggregation between systems. What sets LACP bonding mode apart from the other bonding modes supported by Linux is that the systems participating in LACP continuously exchange data (LACPDUs) with each other to determine which links within the bond should actively be used for the transmission and receipt of network traffic.

While the IEEE standard defines terminology specific to LACP, some vendors have chosen to use different terminology. In this section, LAG (link-aggregation group) and MLAG are used when describing link-aggregation on the switch side and bond is used when describing link-aggregation on the appliance side.

The following list is a non-comprehensive list of equivalent vendor-proprietary terms.

- LAG
  - port-channel
  - trunk
  - multi-layer trunk (MLT)
- MLAG
  - virtual Port-channel (vPC)
  - Multi-chassis link-aggregation group (MCLAG)
  - Split multi-layer trunk (SMLT)

Interfaces in an LACP bond are either placed into an active or suspended state. In the active state an interface is allowed to collect and process incoming network traffic as well as transmit outgoing network traffic. In the suspended state an interface does not collect incoming network traffic or transmit any outgoing network traffic until the condition that caused the suspension is resolved.

For all bonded appliance interfaces to be considered active within a LACP bond, you must ensure the following:

- The switch ports which the bonded interfaces are connected to are all members of the same LAG. If the switchports are not members within the same LAG, the appliance elects an interface or a group of interfaces, which are members of the same LAG, to be active while placing all other bonded interfaces in a suspended state.
- The LAG must be configured to use LACP. If the LAG is not configured to use LACP, the appliance elects a single interface to be active while all other interfaces are placed into a suspended state.

A LAG cannot span multiple switches without the following:

- Configuration of MLAG between the switches
- Deployment of EVPN LAG multihoming in the environment
- Switches being stacked use a vendor proprietary technology.

Note: Check the switch documentation to ensure that the make and model of the switch to ensure LACP functions properly with the switches in a stacked configuration.

Stacked switches, MLAG, and EVPN LAG multihoming allow multiple switches to function as a single logical switch from the perspective of LACP. When stacking, MLAG or EVPN LAG multihoming are deployed and the LAG is properly configured. The LACP bond on the appliance places all the bonded interfaces into an active state. If the appliance's bond spans across multiple switches without MLAG, EVPN LAG multihoming, or stacking configured on those switches, LACP on the appliance detects this and elects an interface or group of interfaces, which are members of the same LAG and connected to the same switch, to be active while placing all other bonded interfaces into a suspended state. Even if the LAGs on the switches are configured to use the same LAG ID, LACP on the appliance identifies that the LACPDUs were sent from multiple different switches.

Although Access Appliance is configured as a cluster, it is important to note that Red Hat Linux does not have a mechanism to perform MLAG. Therefore, each appliance must be treated as a unique end-system. For each bond created on the Access Appliance cluster, two LAGs must be created on the switch and each LAG must be assigned a different LAG ID.

For example, when you create a two-interface bond using the Access CLI, a bond interface is created on each node. The two switch ports which connect to *node_A* should be assigned to a LAG with a LAG ID of X and the two switch ports which connect to *node_B* should be assigned to a LAG with a LAG ID of Y. In this example,

if all four switch ports are assigned the LAG ID of X, the switch or set of switches elects a set of switch ports connected to one of the appliance nodes and places it in an active state while other interfaces are placed into a suspended state. This leads to network connectivity issues as the suspended interfaces are still physically up on the appliance which allows VCS to assign an IP address to them. If this happens, when that node attempts to send traffic out the bond, the application sends the traffic down the network stack, and it is dropped by the bonding driver.

While LACP does not perform load balancing, it does have a mechanism for selecting an outgoing physical interface. This mechanism is responsible for calculating a hash based on selected attributes in the headers of the outgoing network traffic. LACP associates a hash with an interface and all traffic matching that hash egress the selected interface. It is important to note that you cannot control the outgoing interface, you can only influence the selection process. This mechanism is statically configured on each system and is not negotiated between the systems participating in LACP. In Linux this mechanism is called the *xmit_hash_policy*.

The Access Appliance supports the following configurations of the *xmit_hash_policy*:

- layer2: A hash is calculated based on the source and destination MAC addresses in the Ethernet header. This *xmit_hash_policy* provides decent results if communication occurs between the appliance and multiple systems on the same subnet.

  Imbalances occur if traffic is not distributed amongst multiple systems on the same subnet or if the traffic must be routed through a router/gateway.
- layer2+3: A hash is calculated based on the source and destination MAC and IP address in the Ethernet and IP headers. This *xmit_hash_policy* provides decent results if communication occurs between the appliance and multiple systems.

  Imbalances occur if most of the traffic is transmitted between the appliance and a single system.
- layer3+4: A hash is calculated based on the source and destination IP and port in the IP and TCP/UDP headers. A unique hash is produced for each socket established on the appliance.

During the planning stage, it is important to consider which *xmit_hash_policy* is most appropriate based on the expected usage of the bond. It is also important to include the network administrators in this planning as the *xmit_hash_policy* only controls the outgoing traffic on the appliance and has no effect on how the switch chooses an egress interface.

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 takes 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 `virtual` for the *type*.
nodename	Any node of the cluster
fqdns	Specifies a comma-separated list of Fully Qualified Domain Name (FQDN) of the IP address. The *fqdn* 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 `any` 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 *oldipaddr* 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 *fqdn* 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. *fqdns* can have NONE value. NONE is used to remove the existing FQDN entry from the `/etc/hosts` 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 IP addresses and FQDNs in a non-DNS environment

When DNS is not configured for the cluster, you can add IP addresses and the corresponding FQDN entries to the `/etc/hosts` file of both the cluster nodes. Adding these entries is required when you configure AutoSupport for a cluster where DNS is not configured.

To add an IP address and FQDN entry to the /etc/hosts file:

- Enter the following command:

  `Network>` **`host add`** ***`ipaddr fqdn`***

ipaddr	Specify an IPv4 or an IPv6 address.
fqdn	Specify the Fully Qualified Domain Name (FQDN) of the IP address. The FQDN can include the following characters: a-z|A-Z|0-9 or a hyphen (-)

To delete an IP address and FQDN entry from the /etc/hosts file:

- Enter the following command:

  `Network>` **`host del`** ***`ipaddr fqdn`***

ipaddr	Specify the IP address that you want to delete.
fqdn	Specify the corresponding FQDN.

To display the entries in the /etc/hosts file:

- Enter the following command:

  `Network>` **`host show`**

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 `lo` 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]
```

nodename	Specifies the node to whose routing table the route is to be added. To add a route path to all the nodes, use `all` in the *nodename* field. If you enter a node that is not a part of the cluster, an error message is displayed.
ipaddr	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 ipaddr field. Use a netmask value of 255.255.255.255 for the netmask to add a host route to ipaddr.
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 `dev` in the dev field. If you omit the `dev` and `device` 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
```

nodename	Specify the node from which the route entry has to be deleted. To delete the route entry from all nodes, use the `all` option in this field.
ipaddr	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.
netmask	Specifies the IP address to be used. Netmask is used for IPv4 addresses.
prefix	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 flter 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:** *sport* and *dport* 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:** *sport* and *dport* 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 flter 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
```

Adding console FQDN to the network and accessing the GUI using the console FQDN

The console FQDN must resolve to the console IP address and vice versa.

To add the console FQDN to the network:

1. Log in to the Access CLISH.
2. Run the following command:

   ```
   network ip addr modify console-ip console-ip netmask device
   console-fqdn
   ```

3. Restart the GUI:

   ```
   system guidisable

   system guienable
   ```

The appliance administrator can now log in to the Access GUI using the console FQDN.

Chapter 5

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 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.

See “Configuring LDAP server settings” on page 71.

See “Administering the Access Appliance cluster's LDAP client” on page 74.

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 *value*	Specifies the value for the users-basedn. For example: `ou=users,dc=example,dc=com (default)`
groups-basedn *value*	Specifies the value for the groups-basedn. For example: `ou=groups,dc=example,dc=com (default)`
netgroups-basedn *value*	Specifies the value for the netgroups-basedn. For example: `ou=netgroups,dc=example,dc=com (default)`

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}
```

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}
```

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)

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**.

In Access Appliance, AD domain can be used as a domain for users and groups that can have administrator role for the appliance.

Note that a user that is used to join the AD domain must have permission to search and list users/groups in the domain and to list all user's/ group's information.

The search and list permissions are needed to facilitate various other actions that are available through CLISH and UI.

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 idmapupperbound
  ```

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
idmapupperbound	idmap upper bound for AD users

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 **Setiings** > **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 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: ■ *value1* (required)- { files/nis/winbind/ldap } ■ *value2* (optional) - { files/nis/winbind/ldap } ■ *value3* (optional) - { files/nis/winbind/ldap } ■ *value4* (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:

- Digital certificate or smart card authentication can be configured for LDAP, AD, and local users.
- To authenticate LDAP users using digital certificate or smart card, 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 70.
  See "Configuring LDAP server settings" on page 71.
- To authenticate AD users using digital or smart, ensure that AD is configured. Go to **Settings > Directory Services management** and click **Configure** to configure AD, if not already done so. For details about how to configure AD:
  See "Configuring AD server settings" on page 75.
- To authenticate local users using digital or smart card, ensure that you create a local user. Go to **Settings > User management > Add**.
- 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:
     - For local user select none.
     - For an AD user, select the configured AD server.
     - For an LDAP user, select the configured LDAP server.
   - Click **Common name** to select the common mapping attribute.
   - Optionally, enter the Online Certificate Status Protocol (OCSP) URI. OSCP 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 the settings in the UI. LDAP users that do not have this role can only create S3 keys. AD users with Appliance administrator role have access to all the settings in the UI. AD users that do not have this role can only create S3 keys. Local users with Appliance administrator role have access to all the settings in the UI. Local users that do not have this role have access only to change password screen.

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.

Section 3

Managing Access Appliance storage

- Chapter 6. Configuring storage
- Chapter 7. Managing disks

Chapter 6

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
- 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
- Adding a node to the cluster using the GUI

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.

Note: When you provision storage using the GUI or the command-line interface, specify the storage pool size and the size in GiB, TiB, or PiB. You can also specify the input size in decimal values.

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 `Storage> disk list` 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
```

old_name	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.
new_name	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 85.

   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.

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 and Windows users can be configured for file systems.

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(2^53 - 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 9999999999999999999(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 91.

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 90.

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.

user_or_group_names	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 `Storage> quota fs set` 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]
```

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 `/sys/class/fc_host/${host}/supported_classes`.
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
```

Adding a node to the cluster using the GUI

You can add a new node to the cluster to increase its compute resources. When the add node operation is in progress, the backup and restore jobs continue to run on the cluster nodes. The new node must have the same hardware configuration as the existing cluster node.

Considerations before adding a new node to a cluster

- Verify that all the existing nodes in the cluster are in a healthy state. You cannot add a node if a cluster node is not healthy.
- Private links should be modified such that private NIC (eth2) of cluster node is connected to private NIC (eth2) of new node and private NIC (eth3) of cluster node should be connected to private NIC (eth3) of the new node.
- The node that you want to add must be connected to the same private and public network as that of the existing cluster. For details about the cabling information, see the *Hardware installation procedures* in the *Veritas Access 3360 Appliance Hardware Installation Guide*.
- Management interface IP (eth1) should be assigned using the node-level CLI command.
- If add node operation is in progress, you cannot perform any other cluster operation until the existing operation is completed.
- If DNS is configured on the cluster node, IP and FQDN mapping should be done correctly.
- If you have replaced a disk on an existing cluster node, check the status of the `show hardware health` command and wait until the replaced disk status is *OK*. Once the disk state is in *OK* state and it is associated with a volume group, you can trigger the add node operation.
- The network of the new node's management interface IP should be the same as the network of the cluster node's management interface.

- If you have upgraded from Access Appliance version 8.3 to 8.6 using the GUI, and if any patch, EEBs or add-ons are already installed on the cluster node, then the same patch, EEBs or add-ons are installed on the new node before adding it to the cluster. If you have upgraded from an earlier version of Access Appliance, you have to manually upgrade the new node before adding it to the cluster.
- In a single-node cluster, both adding a node and upgrading the cluster must be performed through the GUI.
- Add node operation upgrades the node only if the upgrade path is supported.
- The upgrade path between the software version of the node being added and the version running on the existing cluster node must be supported.
- The required patch must already be present on the node being added.
- Once the add node operation is complete, all the configurations are copied to the new node.
- Do not trigger the add node operation from both the GUI and CLISH simultaneously.

To add a new node the cluster

1. Using an user account with an Appliance Administrator role, log in to the Access Appliance GUI.
2. Navigate to the **Infrastructure** section. Under the **Nodes** tab, select the nodes that you want to add and click **+ Add**.
3. In the **Add node** window, select **Continue**.
4. Enter the following information for the selected nodes:
   - Host name: Enter the host name assignments for each node to be added to the cluster. The node name should have at least 3 characters and can contain a maximum of 63 characters, excluding the domain name. It should start with a lower case character. It cannot start with a number or a special character.
   - Management Network: Enter the management interface assignments for each node to be added to the cluster. FQDN is an optional parameter.

   Click **Add**.

The progress of the add node operation is displayed on the **Infrastructure** page.

5 To monitor the status of each of the tasks, click **View details** on the **Infrastructure** page. The ongoing and completed tasks for the add node operation are also displayed in **Recent activity**. After the node is added, you can view the new node and its details on the **Nodes** tab.

Note: If two instances of Veritas Data Deduplication are running on one node when add node operation is performed, then after the add node operation is complete, you should explicitly change the affinity of any one of the instances to the newly added node. Else, you may experience a higher utilization of CPU, and RAM. The affinity can be changed using the `dedupe affinity set <node_name> <config_name>` command.

Chapter 7

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 101.

See “About storage provisioning and management ” on page 84.

See “Viewing information about disks” on page 101.

See “Monitoring command activity” on page 192.

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 100.

See “About storage provisioning and management ” on page 84.

See “Monitoring command activity” on page 192.

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 84.

See “Adding disks to a storage pool” on page 100.

See “Removing disks from a storage pool” on page 101.

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 84.

See “Viewing information about disks” on page 101.

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 84.

See “Viewing information about disks” on page 101.

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 84.

See “Viewing information about disks” on page 101.

Section 4

Managing Access Appliance file access services

- Chapter 8. Configuring the NFS server
- Chapter 9. Using Access Appliance as an Object Store server

Chapter 8

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 106.

The kernel NFS server is enabled by default.

Note: The IPs used by NFS server should be different from the virtual IP used by the Veritas Data Deduplication server.

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 261.

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.

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 110.
Configure the NFS server for ID mapping	See “Configuring Access Appliance for ID mapping for NFS version 4” on page 108.
Configure the NFS client for ID mapping	See “Configuring the NFS client for ID mapping for NFS version 4” on page 108.
Exporting an NFS share for Kerberos authentication	See “Exporting an NFS share for Kerberos authentication” on page 262.
Mount the NFS share from the NFS client	See “Mounting an NFS share with Kerberos security from the NFS client” on page 263.

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 108.

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.

Chapter 9

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
- Object Access SSL certificate
- Object Access endpoints
- 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 for the following use case.

Use Case 1: 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.

Note: The `map` command should only be used to map one bucket per file system.

Note: A cluster name cannot contain any special symbols except for a hyphen. If the cluster name has special symbols other than the hyphen, then the S3 service does not work as the DNS hostname restrictions have not been followed. For valid characters for naming an Access Appliance , see:
https://technet.microsoft.com/en-us/library/cc959336.aspx

Configuring the Object Store server

This section describes how to configure the Object Store server.

Note: The IPs used by S3 server must be different from the virtual IP used by the Veritas Data Deduplication 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 or LDAP).

   See the following manual pages for more information on configuring AD or LDAP:

   - `CLISH> network man ldap`
   - `CLISH> network man ad`

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_worm         yes
retention_min   3600s
retention_max   36000s
```

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.

Note: It is recommended to have only one bucket per file system.

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 9-1 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 9-1 Object and bucket restrictions

Description	Limit
Maximum recommended parallel threads	8

Table 9-1 Object and bucket restrictions *(continued)*

Description	Limit
Maximum number of buckets per file system without MAP command	50 The maximum number of supported file systems is 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 9-1 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 9-2 describes the supported S3 mode with cluster lockdown mode for a non MSDP-C use case.

Table 9-2 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 9-3 describes the supported S3 mode with cluster lockdown mode for an MSDP-C use case.

Table 9-3 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 9-3 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.

Object Access SSL certificate

Object Access server can be started with SSL. The admin user can use the `objectaccess set ssl_enabled` command to use ObjectAccess over the HTTPS protocol.

You can set SSL for the Object Access server using the following command:

```
Objectaccess> set ssl_enabled <value>
```

where *value* can be (yes/no).

Access Appliance uses a single certificate for all the services. The CA certificate can be listed using the `system certificate show` command.

For more information on certificates, see See "About external certificates on Access Appliance" on page 170.

Note: Private key encryption for external certificates is not supported.

Object Access endpoints

The admin user can change the Object Access admin and data endpoints. You can set a maximum of 4 admin and data endpoints each. A new endpoint can be an FQDN, IP (IPv4 or IPv6) or an alias. The admin user can use the `objectaccess set data_endpoints` command to change data endpoint and the `objectaccess set admin_endpoints` command to set the admin endpoint. If you wants to use old endpoints along with the new endpoints, the command should contain all the new and old endpoints.

Note: The admin and data endpoints can be set only when the Object Access server is in stopped state.

You can set the admin endpoints using the following command:

```
Objectaccess> set admin_endpoints <endpoints>
```

where *endpoints* is a comma-separated list of endpoints.

You can set the data endpoints using the following command:

```
Objectaccess> set data_endpoints <endpoints>
```

where *endpoints* is a comma-separated list of endpoints.

If the Object Access server is started with SSL, the admin user must renew the SSL certificate. If certificate mode is set to internal, use the `system certificate renew` command. If certificate mode is set to external, use the GUI to create an external certificate.

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. 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 113.

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.
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 *Access Appliance RESTful API Guide* 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:

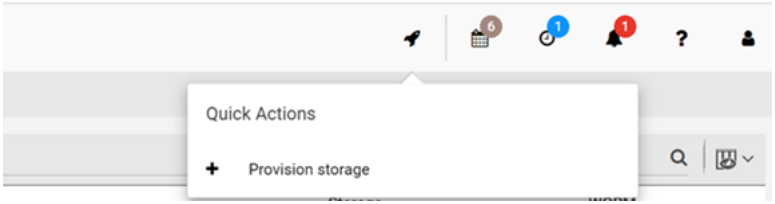
- Default parameters for creating S3 buckets
- Group-related parameters for associating a group with the S3 buckets
- S3 users and their access keys

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**.

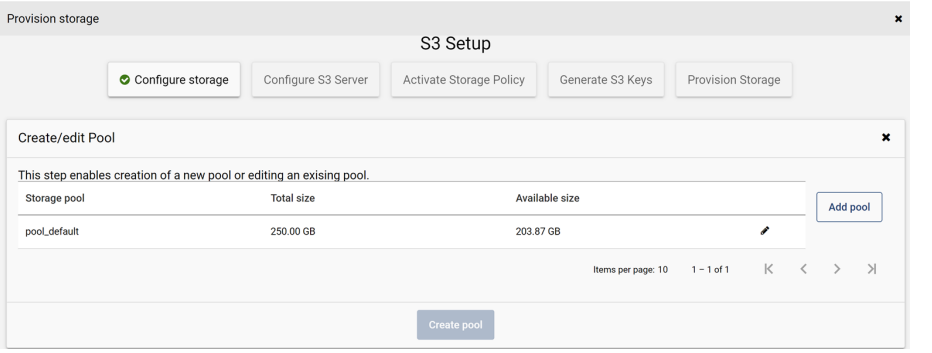


2. Select **Storage type** as **Storage for Buckets** and select **S3 without cloud** for generic S3 bucket.

Note: Select **S3 without cloud option** under **Storage for NetBackup** storage type for configuring NetBackup with S3. For more details, see the *Access Appliance Solutions Guide for NetBackup*.

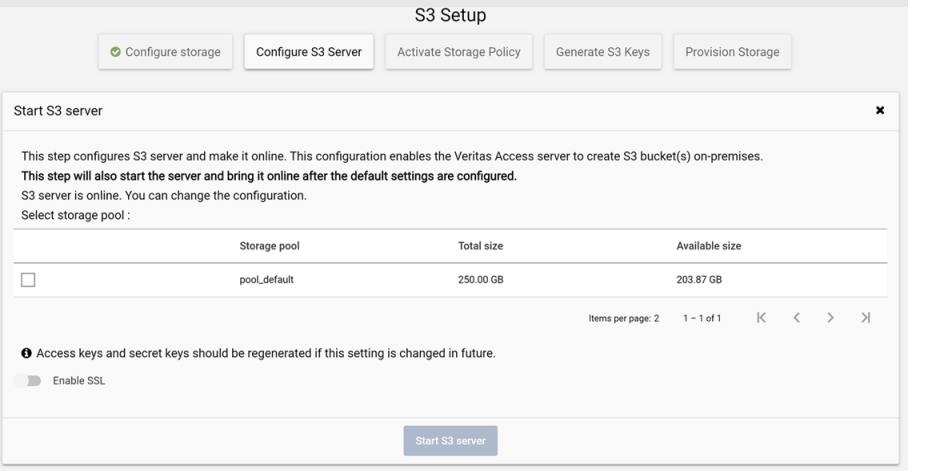
---

3 In the **Configure Storage** tab, *pool_default* is added as the default pool. You can create new pools or edit existing pools.

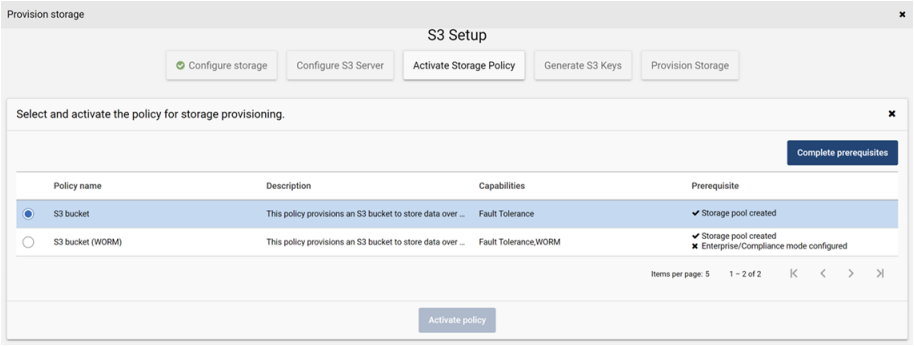


4 Select the appropriate storage pool to configure the S3 server and click **Start S3 server** to make it online.

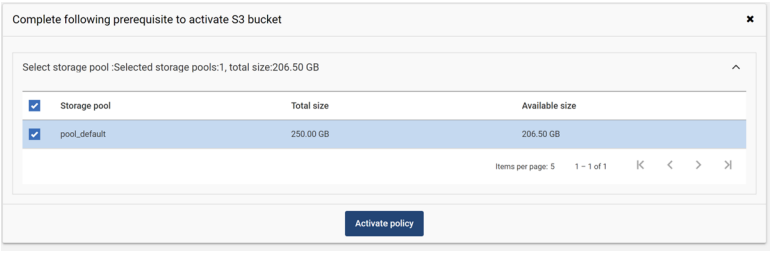
Note: SSL can be enabled, if required.



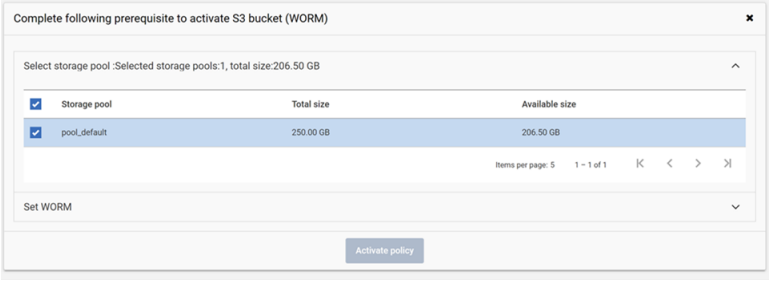
5 In the **Activate Storage Policy** page, choose the policy for storage provisioning,



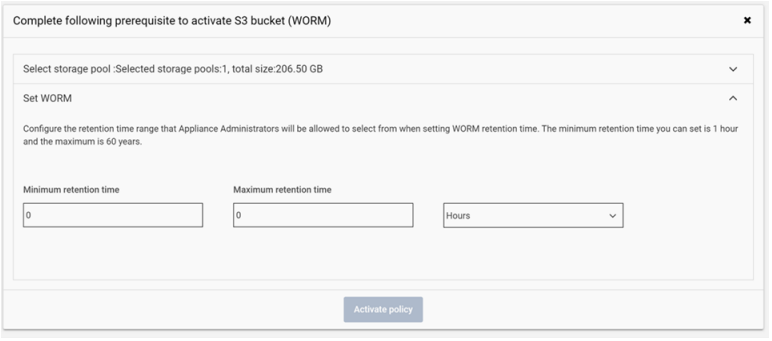
- For S3 bucket (non-WORM), click **Complete prerequisites**. Choose a storage pool from the list of pools available.



- 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.



Set the minimum and maximum retention time.



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.

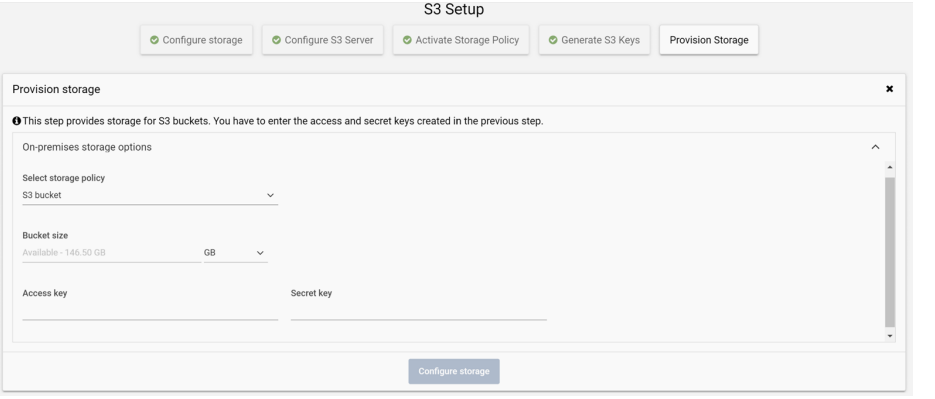
Click **Activate policy**.

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**.

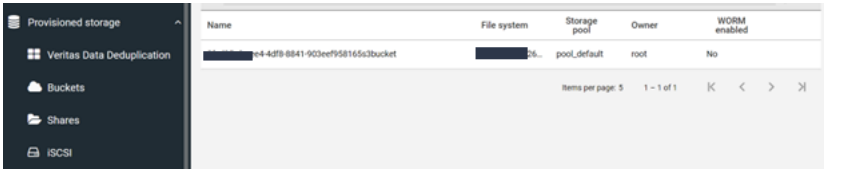


Note down the access key and secret key generated in this step for future use.

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**.



You can view the buckets by navigating to **Provisioned Storage** > **Buckets**.



Section 5

Managing Access Appliance security

- Chapter 10. Managing security
- Chapter 11. Setting up FIPS mode
- Chapter 12. Configuring STIG
- Chapter 13. Setting the banner
- Chapter 14. Setting the password policy
- Chapter 15. Immutability in Access Appliance
- Chapter 16. Deploying certificates on Access Appliance
- Chapter 17. Single Sign-On (SSO)
- Chapter 18. Configuring multifactor authentication

Chapter 10

Managing security

This chapter includes the following topics:

- Security overview
- About the security meter
- Setting the BIOS password

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.

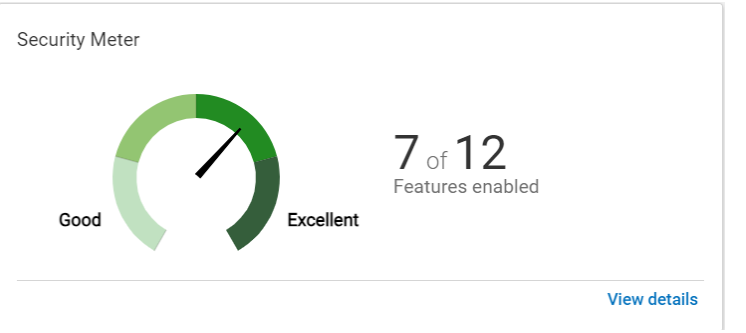
Note: IP forwarding is enabled in Access Appliance by design; it is used to facilitate network communication between application instances and external networks.

About the security meter

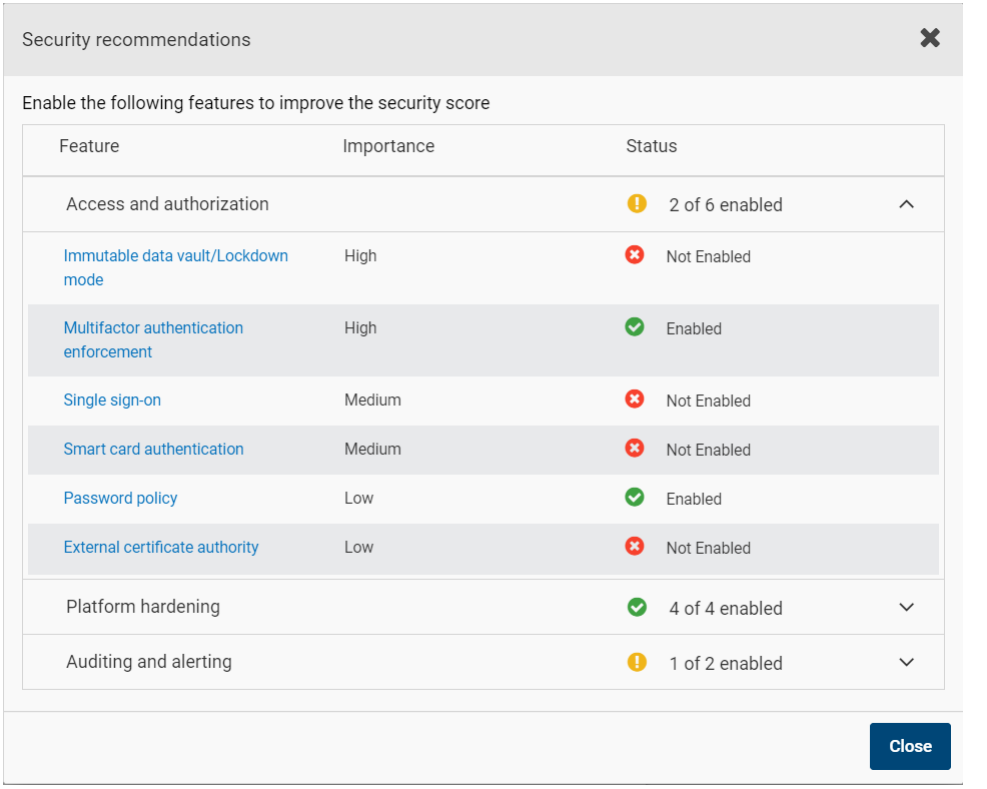
The Access Appliance includes a security meter to view and configure the appliance security settings from one location. The security meter provides security insights and recommendations to improve the appliance security. The security meter keeps a track of the security settings and shows you a list of available security features with quick links to configure them. The security meter displays the current security index from good to excellent based on how many security features are turned on. Built-in security features are turned on and shown as **Enabled** in the security meter.

The security meter can be found on the appliance dashboard. Only an administrator user can view and manage the settings from the security meter.

The following figure shows the security meter that is shown the appliance dashboard:



The following figure shows the available security settings that you can track and manage:



Setting the BIOS password

You can now set a BIOS password for the Access 3360 Appliance. It is recommended that you configure the BIOS password only through the supported Access Appliance interfaces. Supported methods include:

- During cluster configuration
- After cluster configuration using the `admin passwd bios` command

Configuring the BIOS password through these interfaces ensures consistent behavior across cluster nodes.

Considerations while setting the BIOS password:

- For single-node deployments, the BIOS password can also be set using the same supported interfaces.
- If a single-node system is later expanded into a cluster, you must run the `admin passwd bios` command and provide the existing BIOS password (configured on the first node) as both the old password and the new password to synchronize the BIOS password across the cluster.
- The `admin passwd bios` command cannot set or change the BIOS password at the cluster level if the BIOS passwords on the two nodes are different and were configured externally using the IPMI Console menu options.

Chapter 11

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
- Disabling FIPS mode for NetBackup MSDP
- Managing the FIPS mode 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.

Starting with Access Appliance version 8.2, the application layer is FIPS-compliant.

You can enable the FIPS 140-2 standard for NetBackup MSDP to increase appliance security. See "Enabling FIPS for Access Appliance" on page 139.

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 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 Appliance 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 VxOS FIPS status for both the cluster nodes and the MSDP FIPS status for each Veritas Data Deduplication server configured for the cluster is displayed:

   - **Enabled**: The FIPS mode is enabled.
   - **Disabled**: The FIPS mode is disabled.
   - **Unknown**: The FIPS status cannot be retrieved because the Veritas Data Deduplication server is stopped, shut down, or unreachable.

Enabling FIPS for Access Appliance

You can enable the FIPS mode for NetBackup MSDP for each Veritas Data Deduplication server that is configured for the cluster. A maximum of two Veritas Data Deduplication servers can be configured.

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:

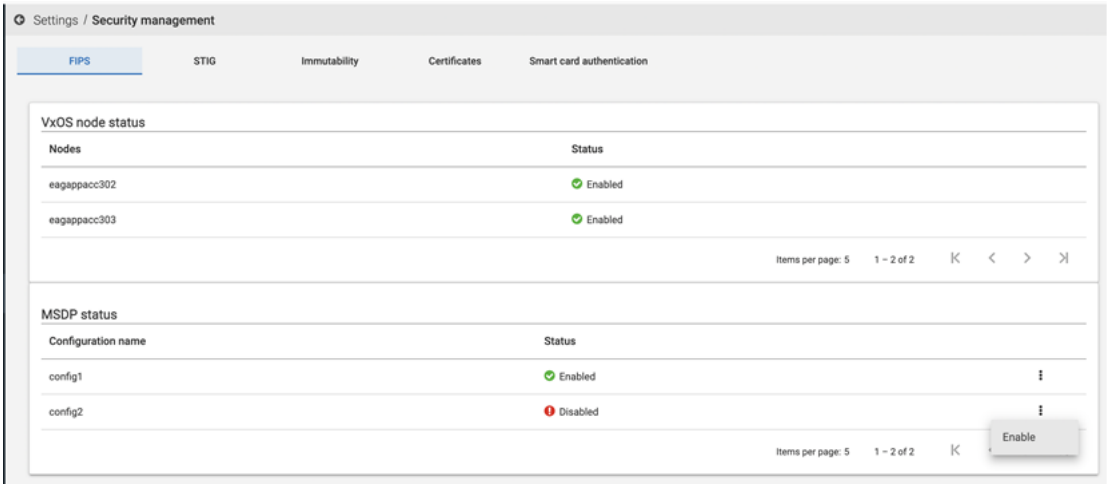
1. Log in to the Access Appliance web interface of the configured cluster by opening a supported browser and typing:

   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 under **MSDP status**, for the Veritas Data Deduplication server for which you want to enable the FIPS mode click the Actions menu (three vertical dots), and then click **Enable**.



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 the Veritas Data Deduplication servers configured for the cluster:

- **Enabled**: The FIPS mode is enabled for the Veritas Data Deduplication server.
- **Disabled**: The FIPS mode is disabled for the Veritas Data Deduplication server.
- **Unknown**: The FIPS status cannot be retrieved because the Veritas Data Deduplication server is stopped, shut down, or unreachable.
  You can enable the FIPS mode again if the status is set to **Unknown** .

Disabling FIPS mode for NetBackup MSDP

You can disable the FIPS mode for NetBackup MSDP for each Veritas Data Deduplication server that is configured for the cluster. A maximum of two Veritas Data Deduplication servers can be configured for the cluster.

To disable the FIPS mode for NetBackup MSDP, complete the following steps:

1 Log in to the Access Appliance 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 under MSDP status, for the Veritas Data Deduplication server for which you want to enable the FIPS mode click the Actions menu (three vertical dots), and then click **Disable**.

4 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**.

Managing the FIPS mode using the command-line interface

You can view the FIPS status and enable and disable the FIPS mode from the Access command-line interface.

Viewing FIPS status for the appliance

You can view the VxOS FIPS status for a cluster node and the FIPS status of NetBackup MSDP for each Veritas Data Deduplication server that is configured for the cluster.

To view the FIPS mode for VxOS and MSDP:

1. Log in to the Access command-line interface as an administrator by opening an SSH session to the console IP.
2. In the Access command-line interface, run the following command:

   ```
   system security fips show
   ```

Enabling FIPS mode

You can enable the FIPS mode for NetBackup MSDP for each Veritas Data Deduplication server that is configured for the cluster. A maximum of two Veritas Data Deduplication servers can be configured. When you enable FIPS for MSDP, all the 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 FIPS for MSDP:

1. Log in to the Access command-line interface as an administrator by opening an SSH session to the console IP.
2. In the Access command-line interface, run the following command:

   ```
   system security fips enable MSDP [config_name]
   ```

   config_name is the name of the deduplication server and is optional if only one deduplication server is configured for the cluster. By default the parameter value is set to **vxdefault**.

Disabling FIPS for MSDP:

1. Log on to the Access command-line interface as an administrator by opening an SSH session to the console IP.
2. In the Access command-line interface, run the following command:

   ```
   system security fips disable MSDP [config_name]
   ```

   config_name is the name of the deduplication server and is optional if only one deduplication server is configured for the cluster. By default the parameter value is set to **vxdefault**.

Chapter 12

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 8.6, 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 8 Security Technical Implementation Guide - Version 1, Release 10`

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.

The Access Appliance also supports DISA's Application Security and Development STIG Version 5, Release 3. The STIG rule lists are available in separate documents on the Veritas Support site. Two checklists are currently available, one for the OS and one for Application Security and Development STIG. The checklists are under the **Updates** section on the Access Appliance OS page of the Veritas Download Center.

To download the checklists:

1. Go to the Veritas Support website and click **Downloads**.
2. In the Veritas Download Center, in the **Products** list, select **Appliances**, in the **Sub product** list select **Access Appliance OS**. Select the version and click **Explore**.
3. Expand **Updates**.

   You can download the checklists from the **Updates** section.

See "Enabling OS STIG hardening for Access Appliance" on page 144.

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):

```
SIG for Red Hat Enterprise Linux 8 Security Technical Implementation
Guide - Version 1, Release 10
```

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`

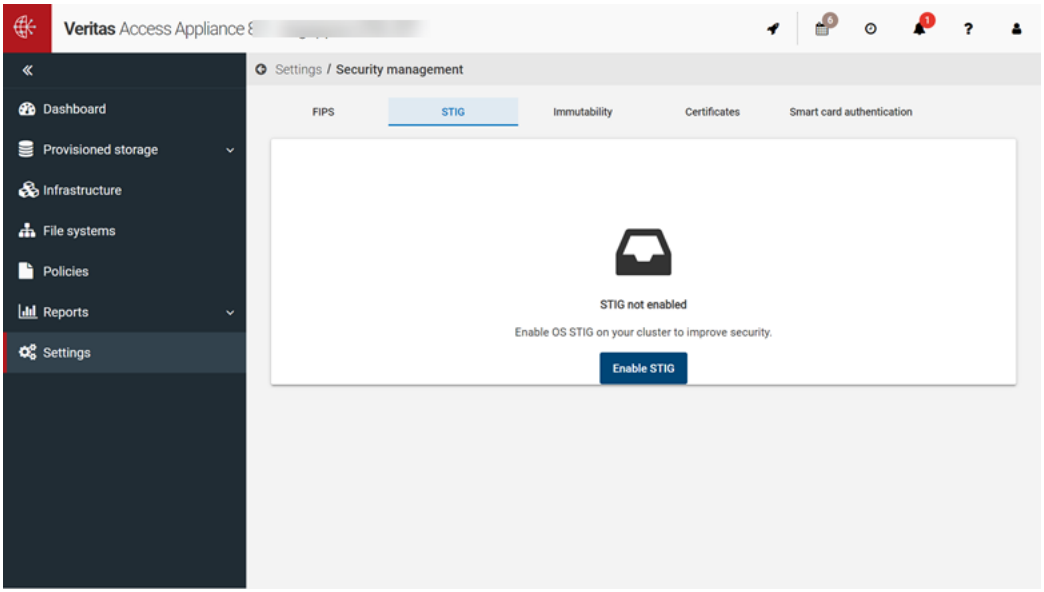
   where *console-ip* is the management console IP address where the web interface is hosted.

2. In the navigation pane, click **Settings**.

3. Click **Security management**.

4 On the **STIG** tab, click **Enable STIG**.

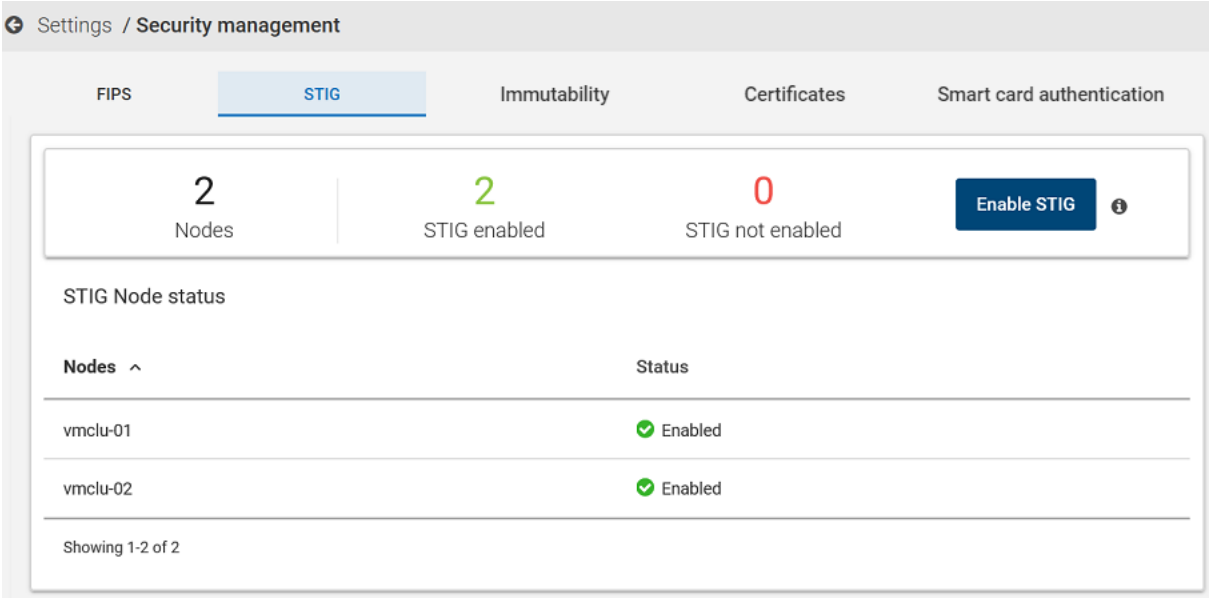
If the prerequisites are not met, you are prompted to resolve the errors. However you can choose to ignore these errors and proceed by clicking **Continue**. You can complete the prerequisites later after you enable the STIG option. If the requirements are met, review the displayed guidelines and click **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.



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 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 enable STIG without completing the required prerequisites, type the following command:

   ```
   System> security stig enable force
   ```

5 To run the precheck, enter the following command:

```
System> security stig precheck

NTP Check : PASS

Email Check : PASS

DNS Check : PASS
```

6 To enable STIG on both the nodes:

```
System> security stig enable
```

Chapter 13

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**.
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 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 banner get
```

```
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. 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 banner get
```

```
# 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.
```

Chapter 14

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.

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**.
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 14-1

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
Maximum repetitive characters	Maximum number of characters that can be repeated in a password.

Table 14-1 *(continued)*

Parameter	Description
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 From version 8.2, when you enable STIG or set the password policy, the SSH session is terminated each time you enter an incorrect password. You must open a new SSH session to log on. Previously, the SSH session was terminated only after the total number of failed attempts was reached.
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

To change the maintenance user password

1. Login to GUI with maintenance user credentials.
2. Select the **Change password** option.
3. Enter the new password. Click **Save**.

To modify the password policy irrespective of the STIG setting

1. Login to GUI with appliance administrator credentials.
2. Navigate to **Settings > User Management**. Select **Manage password policy**.
3. Edit the settings and click **Save**.

   If you do not want a password expiry date, edit the **Days after which password must be changed** field and leave it empty. Click **Save** to save your changes.

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 14-2

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.

Table 14-2 *(continued)*

Parameter	Description
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. From version 8.2, when you enable STIG or set the password policy, the SSH session is terminated each time you enter an incorrect password. You must open a new SSH session to log on. Previously, the SSH session was terminated only after the total number of failed attempts was reached.
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. 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 password-policy get
```

```
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 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 password-policy set
```

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**.

Chapter 15

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 access key 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 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 that are configured on that cluster.

Accessing the root shell in lockdown mode

You can use the following procedure to access the root shell when lockdown mode is configured.

To access the root shell when lockdown mode is configured

1. Log on to the node-level CLI on any node in the cluster.
2. Run the `show serial-number` command to get the serial number for a specific node.
3. If you have permission to generate access key on the SHI portal, then go to the portal to generate access key with the serial number that you got in step 2 . For more information, refer to the *System Health Insights User Guide*.

   If you do not have the permission to generate the access key on the SHI portal, open a ticket with Veritas Support to generate an access key. Mention the serial number in the Support case. The Veritas Support will generate an access key with the serial number provided. Set an access passphrase which will be required later to elevate to root.
4. Log on to the Access Appliance shell on the same node as in step 2.
5. Run the `support unlock` command. You are prompted to enter the maintenance password. Enter the maintenance password and press **Enter**. You are prompted to enter an access key. Enter the access key that you got in step 3. You are prompted to enter an access passphrase. Enter the access passphrase that you set in step 3. Press **Enter** to unlock the root shell access to the current node (all other nodes remain locked).
6. Run the `support elevate` command. Enter the maintenance password. You are prompted to enter an access passphrase. Enter the access passphrase set in step 3. Press **Enter** to get into the root shell.
7. Repeat steps 3 to 6 to get into the root shell of all other nodes in the cluster if you are required to unlock the nodes.
8. Run the `support lock` command on a specific node to lock that node. If no manual lock is issued, the node is locked automatically after 12 hours. All the current users are removed from the root shell in a single node.

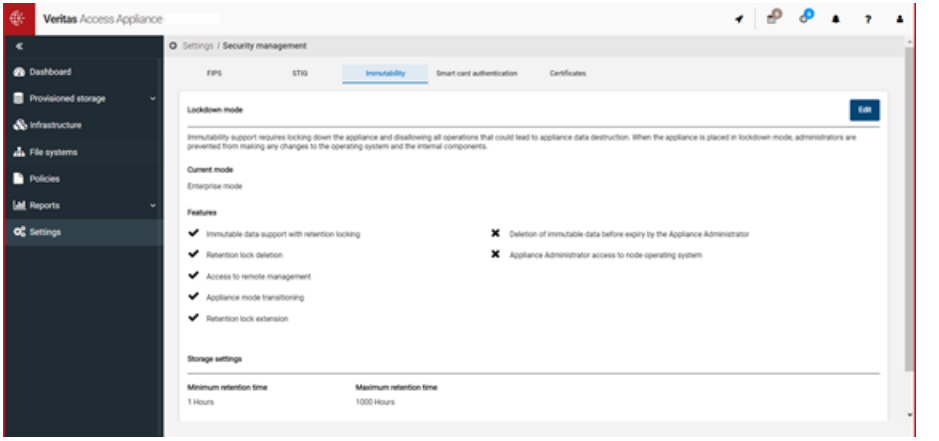
Note: The access key expires in 2 hours.

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**.

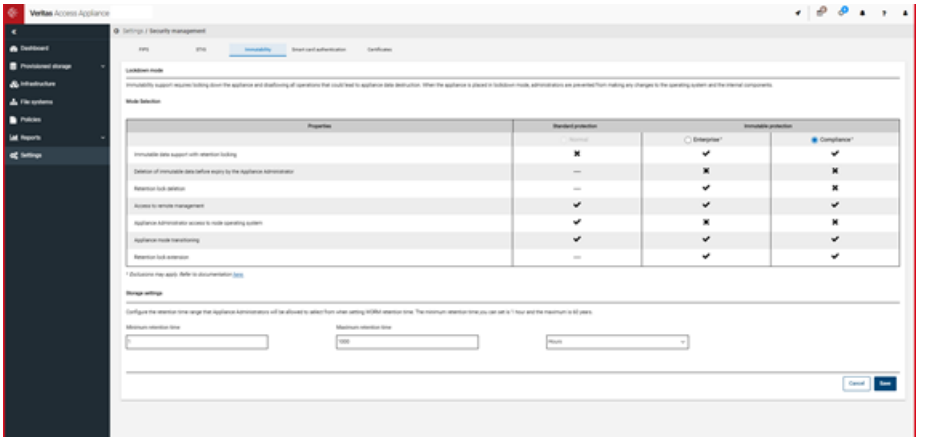


2 Choose any of the lockdown modes under **Mode Selection**. You can choose from Normal, Enterprise and Compliance. The supported features for each mode are listed.

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 (such as NFS) 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.

Chapter 16

Deploying certificates on Access Appliance

This chapter includes the following topics:

- About system certificates on Access Appliance
- About external certificates on Access Appliance
- Deploying ECA using the GUI
- Log locations

About system certificates on Access Appliance

Access Appliance supports one certificate for all services (GUI and Object Access). The certificate can be internal or external. The Appliance CA creates the internal certificate. Any CA can create the external certificate using a CSR generated using the Access Appliance GUI. The admin user can get the CA certificate, information about the certificate mode, and expiration using the `system certificate show` command. The CA certificate can also be downloaded using the GUI by navigating to **Settings > Security Management > Certificates**.

The admin user can change the certificate mode using the `system certificate mode set` command. Or you can navigate to **Settings > Security Management > Certificates** in the GUI.

After changing the certificate mode, the admin user should restart all the services to start using the new certificate. The admin user can restart the Object Access service using the following commands:

```
ObjectAccess> server stop
ObjectAccess> server start
```

The admin user can restart the GUI service by navigating to **System > guidisable** and **System > guienable**. You can use the `system certificate mode show` command to get the system certificate mode.

Internal Certificates

If the system certificate mode is set to internal, the certificate that is signed by the Access Appliance CA is used for all the Object Access and GUI services. The admin user has to renew the certificate if the Object Access endpoints get changed using `system certificate renew` command. After the certificate is renewed, the admin user must restart Object Access and GUI service.

Note: The `system certificate renew` command can be used only when the certificate mode is set to internal.

Note: You must update the clients trust-store with CA certificate to secure connection.

About external certificates on Access Appliance

Starting from the Access Appliance 8.1 release, you can generate and use external certificates instead of internal certificates. External Certificate Authority (ECA) certificates are the digital credentials that attest to the certificate owner's identity and affiliation. Once you deploy the external certificates, all the Access Appliance components use them. One certificate is deployed for all the components. These certificates are used by Access Appliance web server and S3 server for a secure client-server communication.

The external certificates also deploy a certificate bundle and (optionally) certificate revocation list. To generate an external certificate, you have to create a certificate request with proper ‘Subject Distinguished Name’ and ‘Subject Alternative Names.’ You can generate a certificate request using the GUI. The necessary FQDNs are auto-populated to generate the correct request. You can add additional information as needed. Based on the certificate request, you can create an external certificate. When deploying external certificate for the first time, you have to provide a CA certificate bundle. This is used to validate the incoming and deployed external certificate. You can also optionally provide a certification revocation list.

Some important terminologies:

- A certificate authority, also known as a certification authority, is a trusted organization that verifies websites (and other entities) so that you know who you are communicating with online. Their objective is to make the internet a

more secure place for both organizations and users. Becoming a Certificate Authority (CA) means that you (or your customers) oversee the issuing process of cryptographic pairs of private keys and public certificates.

- Certificate bundle (CA bundle) is a file that contains root and intermediate certificates. The end-entity certificate along with a CA bundle constitutes the certificate chain.
- Certificate Revocation List (CRL) is a list of digital certificates that have been revoked by the issuing Certificate Authority (CA) before their scheduled expiration date and should no longer be trusted. CRL is optional. It may be provided as a file or embedded in certificate as a URL.
- Subject Alternative Name: This field lets you specify additional host names (such as sites, IP addresses, common names, and S3 endpoints) to be protected by a single SSL certificate.

Considerations while deploying ECA:

- All certificates for communication should be obtained from a common trusted CA.
- After ECA is deployed on the cluster, you can renew or update the ECA.
- It is recommended to pause backup/restore operations before starting ECA deployment/renewal.
- The CA bundle and CRL file independent of other security artifacts.
- When you deploy security artifacts, they are validated and if inconsistencies are found, you are notified, and deployment does not proceed. If you provide an external certificate and CA certificate bundle, the EC certificate is validated against the user provided CA certificate bundle. If only one of the items is provided, it is validated against deployed artifacts.
- You receive continuous alerts and emails for 60 days before external certificates are about to expire. You must get new CA certificates and deploy it again for seamless working of the Access Appliance. If you fail to do so, the web server and S3 server stop working.

Download the Access Appliance root certificate, and add it to your web browser's list of trusted certificate authorities. This prevents your web browser from displaying security warning messages when you access the Access Appliance UI. If the appliance has been upgraded from any version before 8.1, the internal certificates are updated. You can download the certificate from the GUI by navigating to **Settings > Security management > Certificates > Download root certificate** and upload it in the client trust store for a secure client-server communication.

Deploying ECA using the GUI

You can perform all external certificates-related operations from the **Settings** > **Security management** > **Certificates** tab.

- Switch to external certificate
  - Once the external certificates are deployed, user can switch in between internal and external certificates.
  - After the external certificates are deployed, the **Switch to appliance certificates** option is enabled. The web server is restarted after certificates are switched.
  - Internal certificates or appliance certificates are the Access Appliance CA certificates which are deployed during the initial configuration.
- Upload certificate
  - An ECA that is certified to work with CA bundle and an optional certificate revocation list (CRL) is required for the Access Appliance cluster components to use an external certificate.
  - All the artifacts must be deployed in tandem for the first time. This includes external certificate, CA bundle, and CRL.
  - CRL is optional. It may be provided as a file or embedded in certificate as a URL.
  - If using your own certificate, include the certificate and provide a key. Keep the private key unencrypted.
  - Once deployed all the Access Appliance components - web server and S3 services use the external certificates.
  - The web server and S3 server are restarted after deployment of certificate artifacts is successful.
- View certificate request
  - You can view the generated certificate request and verify the subject distinguished names and subject alternative names.
  - You can also copy the generated certificate request and get the CA certificate from it.
- Generate certificate request
  - You must provide a proper subject distinguished name and subject to generate correct external certificate that can be used in Access Appliance cluster.

- You can use the GUI wizard to generate a proper request for the cluster. The wizard auto populates necessary IP/FQDN to generate correct request.

To deploy ECA using the GUI

1. Go to the **Settings** > **Security Management** > **Certificates** tab.
2. Click **Generate certificate request** and fill out the form. The SAN field is filled with default SAN entries that are mandatory for the configuration. For standard default configuration, it has the FQDN entries for all the storage servers, console IP and API gateway FQDN.
3. Click **Generate** to generate a certificate request.
4. Click **Copy** to copy the certificate request. This has to be given to a CA signing authority who uses this to generate a certificate. CA signing authority provides a certificate along with the root certificate and intermediate certificate used to generate the certificate. Authority may also provide a CRL file .
5. You also have an option to upload a private key. If you want to use your own certificate, include the certificate and provide a key. Ensure that the private key is unencrypted.
6. Go to **Certificates** > **Upload certificate** to upload the certificate artifacts. For a fresh deployment, both the certificate and CA certificate bundle are mandatory. CRL is optional. For renewal, any one of the three are required.
7. After the certificate is uploaded, **Save** gets enabled. Click **Save**. The deployment is initiated.

   After deployment is successfully completed for all components, the GUI restarts with the uploaded external certificate.

You can verify the browser certificate from the GUI to verify that it is the same one that was used during deployment.

Log locations

You can find the logs at the following locations:

- The configuration file can be found at `/shared/cluster_certs/external/config.json`. It contains information such as CSR, and whether ECA is deployed.
- The deployed artifacts can be found at `/shared/cluster_certs/external/`.
- All the logs pertaining to ECA deployment are present in the `/log/VRTSnas/log` directory.

- `/opt/VRTSnas/log/va_external_certificates.log` contains the logs of the driver script which is responsible for performing pre-checks and then calling subsequent scripts to deploy ECA on various components.

Chapter 17

Single Sign-On (SSO)

This chapter includes the following topics:

- About single sign-on (SSO) configuration
- Configuring SSO on Access Appliance
- Limitations and log locations

About single sign-on (SSO) configuration

You can configure single sign-on (SSO) with any identity provider (IDP) that uses the SAML 2.0 protocol for exchanging authentication and authorization information. Note that you can configure an IDP with more than one Veritas product. Access Appliance supports ADFS, Okta, Azure, and Ping Federate IDPs in this release.

Note the following requirements and limitations:

- To use SSO, you must have a SAML 2.0 compliant identity provider configured in your environment.
- Configuration of the IDP requires the Access Appliance GUI.
- Global logout is not supported.
- SSO is supported only for AD/LDAP users.

Configuring SSO on Access Appliance

Configuring SSO on Access Appliance involves the following steps:

Table 17-1

Task	Description
Configuring SSO on an Access Appliance cluster	See To configure SSO on an Access Appliance cluster
Adding users/group	See "Configuring AD/LDAP using the GUI" on page 78. See "Adding and removing user roles using GUI" on page 35.
Configuring an identity provider	See To configure an identity provider
Logging into Access Appliance with SSO	See Login with SSO

Configuring SSO on an Access Appliance cluster

To configure SSO on an Access Appliance cluster

1. Go to **Settings > Security management > Single sign-on (SSO)**. Click **Add**.
2. Give the IDP name and upload the IDP metadata xml and optionally provide the custom user field and group field values. The user field and group field values should be same as configured on the IDP. Click **Save**.

   The UI displays a message that confirms that the add identity provider task is triggered. You can click **View Details** to see the progress of the task. Alternatively, you can also click the **Recent Activity** icon from the top right of the UI to see the status of the most recent operations.
3. Once the configuration is complete, the SSO identify provider details are displayed on the screen. Click **Download service provider xml** to download the details and upload it on IDP server, if required.

Configuring an identity provider

To configure an identity provider

- Login with SSO works only if the configuration on the IDP side is done. Each IDP has different steps for configuration.

  Note: The IDP should always sign the assertion to ensure that the configuration is trustworthy.

  Refer to the following links for the configuration steps for each identity provider.

  - ADFS: Enrolling Access Appliance primary server as a service provider to ADFS

- Azure: Enrolling Access Appliance primary server as a service provider to Azure
- Okta: Enrolling Access Appliance primary server as a service provider to Okta
- Pingfederate: Enrolling Access Appliance primary server as a service provider to Pingfederate

Logging into Access Appliance with SSO

Login with SSO

1. Navigate to GUI login page. Click **Sign-in with single sign-on (SSO)**.
2. Enter SSO credentials and click **Sign in**.

Limitations and log locations

There are some limitations when you configure SSO on an Access Appliance cluster.

- Identity provider cannot be edited. It can be removed and added again.
- Single logout is not implemented. If SAML users log out of the application, and try to login with SSO again, the user is not asked for their login credentials unless the SSO session has expired. This applies to any other application using the same IDP.
- If after identity provider configuration, External certificate authority (ECA) is configured, then login with SSO does not work until the identity provider is updated with the latest service provider metadata xml from the Access Appliance. This can be done by downloading the service provider metadata xml from **Settings > Security > Single-Sign on > Download service provider metadata**. This metadata needs to be updated on the IDP side.
- AD/IDP server date, time, and time zone should be the same as the Access Appliance cluster. Else, the SSO login fails.

The logs can be found by logging into the Access Appliance CLISH, elevating to root and accessing the logs at:

- `/log/VRTSnas/isagui_webserver.log`
- `/log/VRTSnas/ isagui_sso_config.log`

Chapter 18

Configuring multifactor authentication

This chapter includes the following topics:

- About multifactor authentication
- Considerations when configuring multifactor authentication
- Configuring multifactor authentication for your user account
- Disabling multifactor authentication for your user account
- Enforcing multifactor authentication for all users
- Configuring multifactor authentication for your user account when it is enforced in the cluster
- Resetting multifactor authentication for a user

About multifactor authentication

Multifactor authentication is a robust security measure widely used for adding an additional layer of security to the authentication process by requiring users to provide a unique, time-limited code along with their regular login credentials. It is a multiple-step account login process that requires you to enter a 6-digit one-time password along with your password.

It is strongly recommended that you configure multifactor authentication to protect the security of your account.

See “Configuring multifactor authentication for your user account” on page 179.

If multifactor authentication is enforced in the Access Appliance cluster, all users must configure multifactor authentication for their user accounts for successful sign-in.

See "Configuring multifactor authentication for your user account when it is enforced in the cluster" on page 182.

Considerations when configuring multifactor authentication

Some considerations that you need to remember before you configure multifactor authentication:

- The Appliance administrator can see the status of all the users on the **Settings > User management** page.
- If AD/LDAP server configuration is removed from the cluster without removing the AD/LDAP user's multifactor authentication configuration, the Appliance administrator may see stale entries for AD/LDAP users.
- Multifactor authentication configuration for AD/LDAP users with no role is only possible after multifactor authentication enforcement.
- A local administrator is a non AD/LDAP user.
- Do not configure multifactor authentication if `add node` or `delete node` operations are in progress.

Perform the following steps to login to the Swagger REST APIs if multifactor authentication is configured for a particular user:

- Provide username (for which multifactor authentication is configured) and password and generate token. A token of the type, *mfa token*, is generated.
- Copy the token and paste it in the username field. Provide the OTP received in the authentication app in the password field, and enter *mfa* in the token type field.
- A bearer token is generated for that user. This token should be provided in the **Authorize** tab with the format as: *Bearer <bearer token>*.

Configuring multifactor authentication for your user account

You must first install and configure authenticator application on your smart device that provides you with the one-time password.

Supported authenticator applications

If the Access Appliance administrator has enforced multifactor authentication in the Access Appliance cluster, you must configure it for your user account for successful sign-in. You must configure multifactor authentication before the start date of enforcement. Else, you will lose access to the appliance and your automation workflow (if using login API) will also be impacted.

Even if multifactor authentication is not enforced, it is recommended that you configure it for enhanced security.

To configure multifactor authentication for your user account

1. Sign in to the Access Appliance UI.
2. On the top right, click the profile icon and click **Manage multifactor authentication**.
3. On the **Manage multifactor authentication** screen, click **Configure**.
4. On the next screen, follow the given steps.

   Install and configure authenticator application on your smart device. It generates one-time password and sends it on your smart device.
5. Scan the QR code with the authenticator application or enter the key manually.

   The manual key should be base32 encoded and can contain between 26 to 208 characters with or without padding.
6. Enter the one-time password that you see in the authenticator application.
7. Click **Configure**.

   At the time of next sign-in, you need to enter the one-time password along with the username and password.

See “Disabling multifactor authentication for your user account” on page 180.

Disabling multifactor authentication for your user account

You can disable MFA for your user account only if multifactor authentication is not enforced. However, it is strongly recommended that you configure multifactor authentication to protect the security of your account.

If multifactor authentication is enforced, and you want to reset it, See “Resetting multifactor authentication for a user” on page 183.

To disable multifactor authentication for your user account

1. Sign in to the Access Appliance UI.
2. If you are an Appliance administrator, click the profile icon on the top right, and select **Manage multifactor authentication**.

   If you are not an Appliance administrator, select **Manage multifactor authentication** in the home screen.
3. If you have already configured multifactor authentication for your user account, you can see the **Disable** button.
4. Click **Disable**.
5. Enter the one-time password and click **Submit**.

Enforcing multifactor authentication for all users

Only the Access Appliance administrator can enforce multifactor authentication for all Access Appliance users.

Before you enforce multifactor authentication:

- Multifactor authentication can be enforced only if at least two local Appliance administrator users have configured it.
- You can set a future start date for enforcement so that the users get sufficient time to configure their multifactor authentication.
- If multifactor authentication is not configured by the start date, the user will not have access to the appliance. If the user's automation workflow uses login API, then that will also be impacted.
- Once multifactor authentication is enforced, it cannot be reversed.
- It is not possible to postpone the start date of enforcement after it is set.
- You can prepone the start date for enforcement using the **Edit enforcement** button on the **Settings > Security management > Multifactor authentication** page.
- The start date for enforcement cannot be more than 90 days from the current date.

To enforce multifactor authentication for all users

1. Sign in to the Access Appliance UI.
2. Go to **Settings > Security management > Multifactor authentication**.
3. Click **Enforce** to enforce multifactor authentication for all Access Appliance users.

   Notify all users that they must configure multifactor authentication for their user accounts to be able to successfully sign in.

   See "Configuring multifactor authentication for your user account" on page 179.

Configuring multifactor authentication for your user account when it is enforced in the cluster

After multifactor authentication is enforced in the cluster, you must configure it for your user account if you have not already configured it.

If you have not configured multifactor authentication for your account after the enforcement, you cannot sign-in to the appliance if the enforcement period has expired and any automation workflow using the login API will be impacted. If the enforcement period has not expired, you can skip the multifactor authentication configuration in the login screen but it is recommended that you configure multifactor authentication to protect the security of your account.

To configure multifactor authentication after the enforcement

1. Open a web browser and go to the following URL.

   https://*console-IP*:14161/login

   The *console-IP* is the management console IP address where the web interface is hosted.
2. Enter the **Username** and **Password**.
3. Click **Sign in**. The **Configure multifactor authentication** screen is displayed.
4. On the next screen, follow the given steps.

   Install and configure an authenticator application on your smart device. It generates a one-time password and sends it to your smart device.

   Supported authenticator applications
5. Scan the QR code with the authenticator application or enter the key manually.

   The manual key should be base32 encoded and can contain between 26 to 208 characters with or without padding.

6 Enter the one-time password that you see in the authenticator application.

7 Click **Submit**.

Successful configuration takes you back to the sign-in screen.

Enter the username, password, and one-time password for successful sign-in.

Resetting multifactor authentication for a user

Only the Access Appliance administrator can reset multifactor authentication for other Access Appliance users.

Before you reset multifactor authentication:

- The logged in administrator cannot reset his own multifactor authentication. It can only be reset by another Appliance administrator.
- If multifactor authentication is not enforced, then it is possible to reset it for any user.
- If multifactor authentication is enforced, then you can reset it for a local (non AD/LDAP) administrator only if at least one other local administrator is present who has multifactor authentication configured.

To reset multifactor authentication for an Access Appliance user

1 Sign in to the Access Appliance UI.

2 Go to **Settings > User management**.

3 Navigate to the user row and click on the vertical ellipsis button from the right side of the UI and then select **Reset multifactor authentication** .

4 In the **Reset multifactor authentication** pop-up, click **Reset**.

Section 6

Monitoring and troubleshooting

- Chapter 19. Monitoring the appliance
- Chapter 20. Configuring event notifications and audit logs
- Chapter 21. Appliance log files

Chapter 19

Monitoring the appliance

This chapter includes the following topics:

- Testing the appliance hardware
- About Veritas AutoSupport on the Access Appliance
- Setting up email notifications on the appliance
- Setting up SNMP notifications on the appliance

Testing the appliance hardware

Before making any significant hardware or software configuration changes (such as initial configuration and software upgrades), you should run a hardware self-test. This test helps ensure that there are no component cable errors or disk drive errors that can cause an operation to fail.

To test the appliance hardware

1. Log on to the Access Appliance shell menu of the appliance node.
2. Type the `system self-test hardware` command.

   A **Warning** indicates a problem that can be fixed later and lets you proceed with the initial configuration. However, such problems can prevent access to the affected devices.

   An **Error** indicates a critical problem that requires immediate resolution before you can proceed with the initial configuration.
3. (Optional) If the `system self-test hardware` command output identifies any problems, check the following items:
   - Verify that all cables are connected correctly and secured.
   - Verify that all disk drives are installed and seated properly.

- Verify that all units are turned on and have started up completely.

After you have verified the previous items, reenter the command to ensure that the problems are resolved. If you cannot resolve an error after verifying all of the previous items and reentering the command, contact Veritas Technical Support.

About Veritas AutoSupport on the Access Appliance

Veritas AutoSupport is a free service that enables proactive monitoring, management, and support of the appliance's health and performance 24 hours a day, 7 days a week. The AutoSupport service identifies risks and issues with the appliance and alerts you and/or service engineers to enable proactive handling and risk mitigation.

The Veritas AutoSupport service is delivered using two components: The appliance Call Home service and the MyAppliance web portal. When Call Home is enabled, the appliance uploads diagnostic and heartbeat data over SSL-encrypted channels to a Veritas secure operations center for further processing. The MyAppliance portal then uses the Call Home data to provide a comprehensive view of appliance health and performance information, as well as support case management.

Note: Call Home is not supported in a pure IPv6 network configuration.

AutoSupport technical details

- Call Home is enabled by default and uses HTTPS (secure and encrypted protocol) with port 443 for all communication with Veritas AutoSupport servers.
- If you configured the appliance to use a proxy server to connect to the Internet, Call Home uses that proxy server to communicate with the AutoSupport servers.
- The appliance initiates all communications with the AutoSupport servers.

For more information about the data that AutoSupport collects and when it is sent to Veritas, refer to the *Veritas Appliance AutoSupport 2.0 Reference Guide*.

See “Setting up AutoSupport on the appliance” on page 186.

Setting up AutoSupport on the appliance

Table 19-1 lists the steps that you need to carry out to set up AutoSupport for the appliance.

Table 19-1 Steps to set up AutoSupport

Step	Task
Step 1	Register the appliance on the MyAppliance portal. See "To register the appliance on the MyAppliance portal" on page 187.
Step 2	Configure Call Home settings on the appliance. **Note:** See "To enable Call Home on the appliance" on page 187.Call Home is enabled by default.

Step 1

To register the appliance on the MyAppliance portal

1. Log on to the Veritas MyAppliance portal.

   https://my.veritas.com/

2. On the **Appliances** page, click **My Appliances**.
3. Follow the prompts to register the appliance.

Step 2

To enable Call Home on the appliance

1. Log on to the Access Appliance shell menu.
2. If the appliance is configured to use a proxy server, configure the proxy server.

   See "Using a proxy server with the appliance" on page 188.

3. Configure the SMTP server.

   See "Setting up email notifications on the appliance" on page 188.

4. Set up email notifications.

   See "Setting up SNMP notifications on the appliance" on page 190.

About moving an appliance to another geographic location

If you plan to move the appliance from one geographic location to another, consider the following points to ensure continuance of maintenance and support coverage:

- You should not move an appliance to another country.
- Certain locations or regions of the world may not be enabled or set up to handle field service calls for parts replacement.
- Certain locations or regions may not be able to meet the defined service level agreement (SLA) contract(s) to which the appliance may be associated with.

- If it is imperative to move the appliance, you must contact your account access team at Veritas to understand the ramifications or impact (if any) to the SLAs associated with the appliance.
- After you have moved the appliance, it is critical to update your registration details such as contact details and location information on the MyAppliance portal to ensure continuance of coverage.

Using a proxy server with the appliance

The following features require access to the Internet:

- AutoSupport (Call Home service)
- Download software updates from Veritas website.

Once you configure the appliance proxy settings, these services will travel through the proxy server.

To configure the appliance to use a proxy server

1. Log on to the Access Appliance shell menu on either node of the appliance.
2. Enter the `set proxy-server settings server= tunnel= user=` command.
   - *server*: Type the IP address or host name of the proxy server, followed by a colon (:) and then the port number.
     If the proxy server requires https, add `https://` to the address. Otherwise the appliance adds `http://` to the address by default.
   - *[tunnel]*: Type `TunnelOn` if your proxy requires tunneling (the default is `TunnelOff`).
   - *[user]*: Type the appropriate user name if the proxy server requires authentication. You are prompted for a password after you execute the command.

See “About Veritas AutoSupport on the Access Appliance” on page 186.

Setting up email notifications on the appliance

The appliance can send email alerts when hardware and software components fail or encounter errors.

To configure email notifications:

1. Log on to the Access Appliance shell menu.
2. Enter the following command to set the SMTP mail server:

   ```
   set alerts email-smtp smtp-server= smtp-account=
   ```

   where **smtp-server** is the IP address or FQDN of your SMTP mail server and **smtp-account** is name of the user account that is used to access the SMTP server.

   When prompted, enter the password for the user name if authentication is required to access the SMTP account.
3. Enter one or both of the following commands to set the email addresses that you want the appliance to send emails to:

   ```
   set alerts email-hardware email-address=
   ```

   ```
   set alerts email-software email-address=
   ```

   where **email-address** is the email address of the appliance hardware administrator and the email address of the appliance software administrator.

   Note: Separate multiple email addresses using a semi-colon (`;`). Do not add a space before or after the semi-colon.

   ---

4. (Optional) Enter the following command to set the email account that you want the emails to originate from (sender email):

   ```
   set alerts email-sender-id email-address=
   ```

   where **email-address** is the email address that you want the appliance emails to originate from.
5. Enter the following command to set the time interval between email notifications:

   ```
   set alerts email-notification-interval interval=
   ```

   where **interval** is the time interval in minutes.
6. Tto verify the appliance email notification settings, enter the following command:

   ```
   show alerts email
   ```

See “Setting up SNMP notifications on the appliance” on page 190.

Setting up SNMP notifications on the appliance

You can configure the appliance to generate and send Simple Network Management Protocol (SNMP) traps to your SNMP server for hardware monitoring purposes.

The appliance uses the SNMPv2-SMI application or the SNMPv3-SMI application protocol to monitor the appliance.

To configure SNMP notifications

1. Log on to the Access Appliance shell menu.
2. To set the SNMP server, enter the following command:

   ```
   set alerts snmp server= port=
   ```

   ```
   set alerts snmp community=
   ```

   Where *server* is the IP address or FQDN of your SNMP mail server.

   Note: The appliance uses the default community `public` and the default destination port `162` for SNMP traps. If your SNMP server uses a different community or port, use the **port** and **community** parameters.

   Your firewall must allow access from the appliance to the SNMP server through whichever port you use.

3. Enter the following command to show the appliance MIB:

   ```
   show alerts snmp mib
   ```

   Copy the Management Information Base (MIB) text and import it into your SNMP management software so that it can interpret the appliance traps.

See "Setting up email notifications on the appliance" on page 188.

Chapter 20

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 193.

See "Viewing information about alert management" on page 193.

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 193.

See "About alert management" on page 193.

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 193.

See “Viewing information about alert management” on page 193.

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 on-premises storage.
  See “Viewing cluster storage usage” on page 194.
- The File System Usage report compares the storage usage patterns for up to 5 file systems.
  See “Viewing file system usage” on page 195.

Viewing cluster storage usage

These reports show usage trends for on-premises 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 20-1 describes the valid Access Appliance severity levels in descending order of severity.

Table 20-1 Severity levels

Valid value	Description
emerg	Indicates that the system is unusable
alert	Indicates that immediate action is required
crit	Indicates a critical condition
err	Indicates an error condition

Table 20-1 Severity levels *(continued)*

Valid value	Description
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
- 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 195.

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 195.

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].

Chapter 21

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
- Collecting logs using the UI

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 21-1 describes the methods you can use to access the appliance log files.

Table 21-1 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 21-1 Methods of accessing appliance logs *(continued)*

Method	Log details
You can use the `support data-collect` commands to collect storage device logs. See "Gathering device logs with the DataCollect command" on page 200.	Appliance storage device logs

Gathering device logs with the DataCollect command

You can use the `support 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 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 adpbbucmd 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\APPLIANCE
Share can be opened using 'support collect-logs share open'
The Log files DataCollect-XXXX9000918-20180201012500.tar.gz
```

```
collected from the
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-8.2] dl365g10p-016v01.xxx.yyy.com > system software share
 open

- [Info] Created an NFS share for sharing the patches.

- [Info] You can access the NFS share at
dl365g10p-016v01.xxx.yyy.com:/inst/patch/appliance/available. To
 ensure appliance security, use the 'system software share close'
 command to remove the share after downloading the required
patches.
```

4 You can send the `DataCollect.zip` file to the Veritas Support team to resolve your issues.

See "About appliance log files" on page 199.

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
```

Data sanitization is enabled by default. To validate the current setting:

```
support data-collect sanitize-data status
```

To unmask sensitive data and disable sanitization:

```
support data-collect sanitize-data disable
```

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-IPaddrress: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

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.

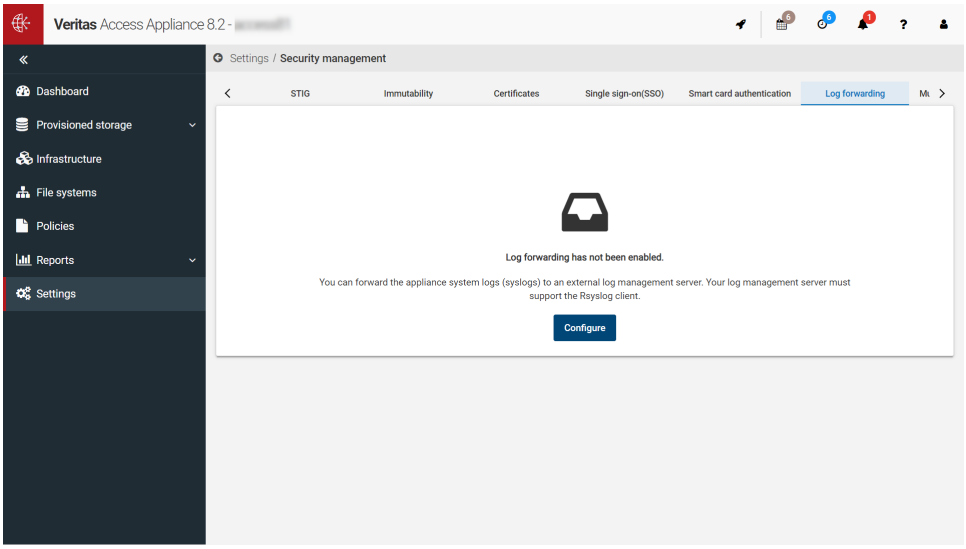
To secure the log transmission from the appliance to the log management server, you can use the TLS (Transport Layer Security) option. The Access Appliance currently supports only TLS Anonymous Authentication for log forwarding.

Configuring log forwarding using the UI

You can forward the appliance system logs (syslogs) to an external log management server. Your log management server must support the Rsyslog client.

To configure log forwarding:

1. Sign in to the Access Appliance UI by using the `http://console-ip:14161` URL, where *console-ip* is the management console IP address.
2. Do one of the following:
   - Click **Dashboard > Security Meter > View details > Auditing and alerting > Log forwarding**
   - Click **Settings > Security management > Log forwarding**
3. Click **Configure**.



4 Enter the following details:

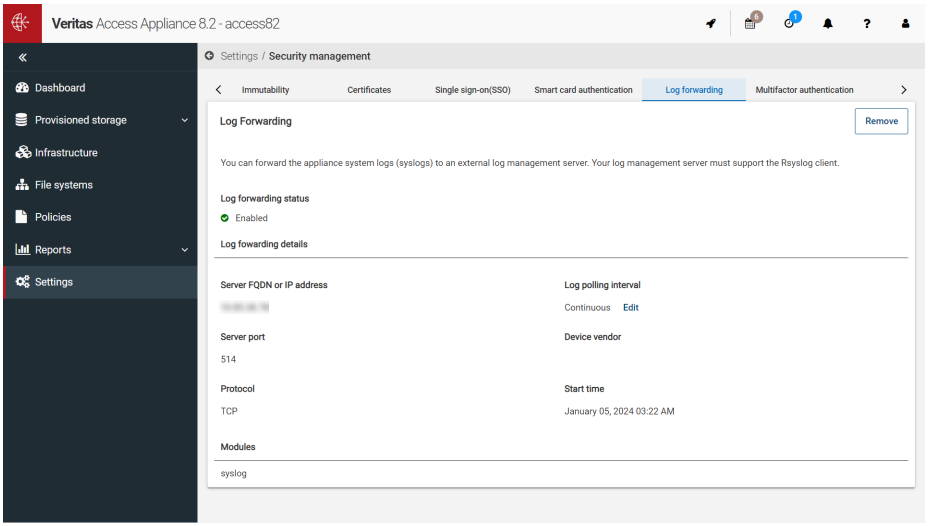
Field	Description
Server FQDN or IP address	FQDN or the IP address of the external log management server.
Server port	Port number of the external log management server. Default port is 514. You can specify a different port if the log server is configured to listen on that port.
Protocol	Select either UDP or TCP. TCP is the default protocol. With TCP protocol, you can optionally enable TLS log transmission. **Note:** Enabling TLS requires that you upload certificates obtained from CA authority and a private key to the appliance.
Log polling interval	Duration in minutes after which the log server polls for the system logs. Logs on the remote log server will be available after that duration. Set the interval in minutes. The options are **15**, **30**, **45**, **60**, **Continuous**. If you select **Continuous**, the appliance continuously forwards logs to the log server.
Device vendor	Unique name for the external log server. You can use any name to identify to which server the appliance is forwarding the logs to.

Field	Description
Enable TLS log transmission	If you want to secure the transmission of logs from the appliance to the log server, select **Enable TLS log transmission** and upload the required certificate files. Veritas recommends that you enable TLS for security purposes. This option provides end-to-end security of data sent over the network from the appliance to the log server. You need a CA certificate and the client private key to configure TLS log transmission. This option is available only if you select the TCP protocol. If you enable secure log transmission, upload CA certificate (X.509 certificate for the certificate authority in PEM format), client certificate (X.509 certificate for the appliance to communicate with the log management server, in PEM format), and client certificate key (RSA key of the client certificate) onto your log server and then upload the certificates to the appliance.

Field	Description
Modules	Types of logs that are forwarded to the log server. Only the OS logs are forwarded and the **syslog** option is selected by default.

5 Click **Enable**.

A notification about the task is displayed on the top of the page. To monitor the progress, click **View details**. After the configuration is completed successfully, a notification is displayed on top of the page. The log forwarding status is shown **Enabled** and the start time for forwarding the logs is displayed.



See “Forwarding logs to an external server” on page 205.

See “Removing log forwarding using the UI” on page 210.

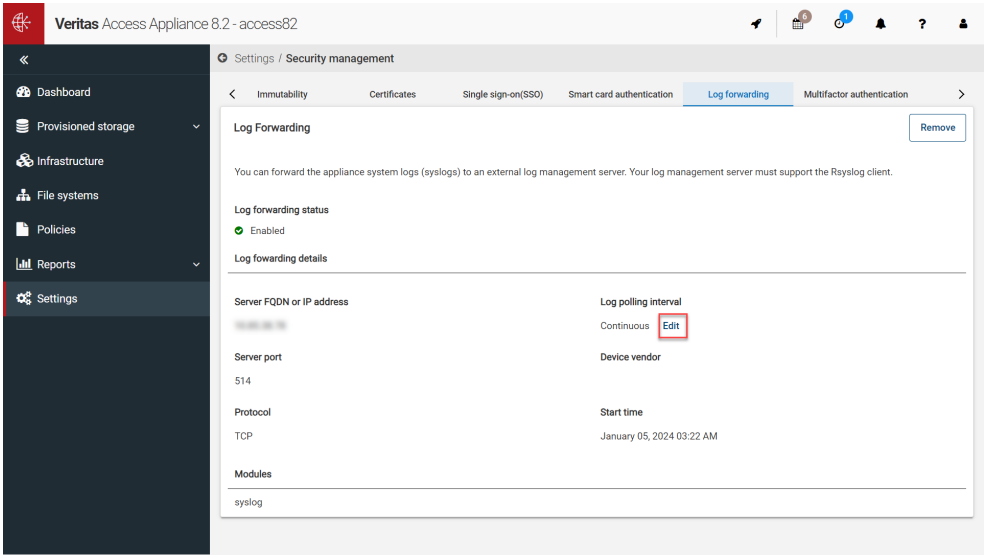
See “Modifying log forwarding settings using the UI” on page 209.

Modifying log forwarding settings using the UI

You can edit only the scheduled interval for forwarding the logs. To edit any other settings, delete the configured log server settings and reconfigure log forwarding.

To edit the settings:

1. Sign in to the Access Appliance UI by using the `http://console-ip:14161` URL, where *console-ip* is the management console IP address.
2. Do one of the following:
   - Click **Dashboard > Security Meter > View details > Auditing and alerting > Log forwarding**
   - Click **Settings > Security management > Log forwarding**
3. Under **Log polling interval**, click **Edit**.



4. Select the time interval and click **Save**.

   A notification is displayed on the top of the page. To monitor the progress, click **View details**. After the task is completed successfully a notification is displayed on the top of the page and the changed interval is displayed in the UI.

Removing log forwarding using the UI

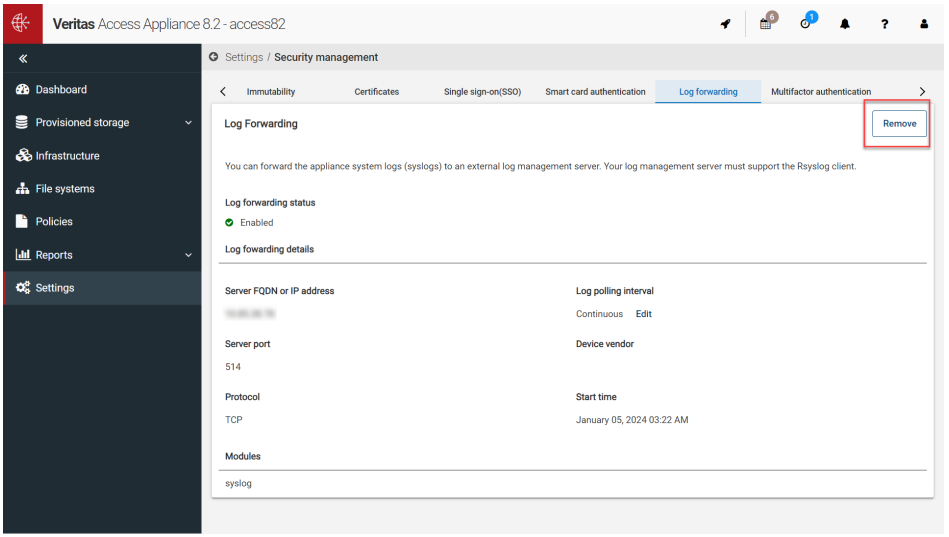
To stop forwarding logs:

1. Sign in to the Access Appliance UI by using the `http://console-ip:14161` URL, where *console-ip* is the management console IP address.
2. Do one of the following:

- Click **Dashboard > Security Meter > View details > Auditing and alerting > Log forwarding**
- Click **Settings > Security management > Log forwarding**

3 Click **Remove**.

A notification is displayed on the top of the page. To monitor the progress, click **View details**.



Setting up log forwarding using the command-line interface

You can 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 from the appliance shell menu

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 60. If the interval is set to 0, appliance continuously forwards syslogs to the log management server.
- Whether TLS is enabled for secure log transmission.

For more details, see the *Veritas Access Appliance Command Reference Guide*.

Collecting logs using the UI

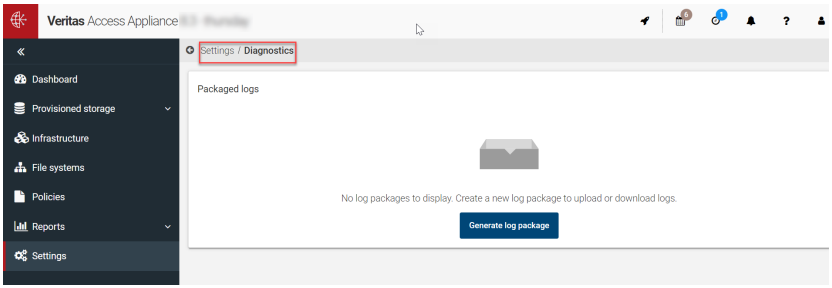
You can collect logs for an individual node in the cluster or for all the cluster nodes from the Access Appliance UI for error analysis and troubleshooting. You can choose to collect logs for all the components to get a comprehensive view of the system, or you can collect logs for specific components that have an issue.

1. Sign in to the Access Appliance UI using the following URL:

   `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 > Diagnostics**.
3. Click **Generate log package**.



   On the Generate log package page, on the **Basic** tab select the components to include in the log package.

4. To select the components for which you want to collect logs:
   - On the **Basic** tab under **Log settings**, select the components. To select all the components, click **All**.

Component	Description
Service status	Status of the services running on the cluster nodes.

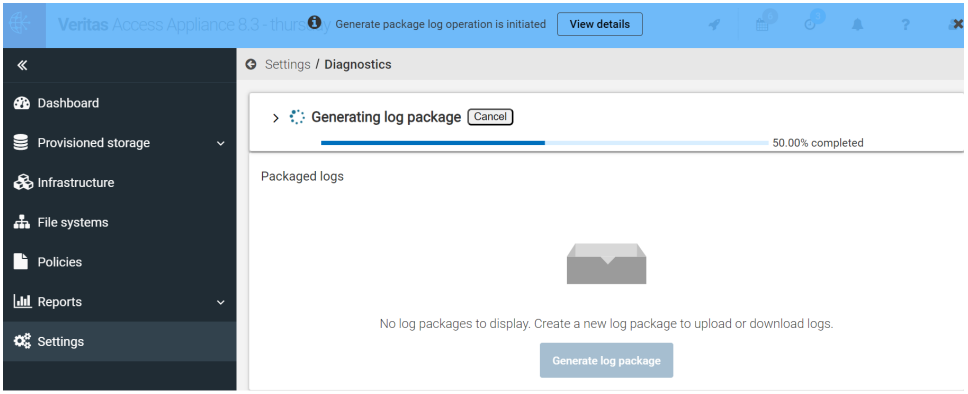
Component	Description
Performance logs	Logs for collecting performance data.
Common logs	Common logs from multiple modules.
Nas	Product information from cluster nodes.
OS	Kernel and user-level debug logs.
Sos report	SOS reports collected from all nodes using the RHEL sos utility.
Explorer	Logs and environment data collected by the `VxExplorer` utility from all the nodes.
Install	Installation logs from the `/opt/VRTS/install/logs/` log directory. If you encounter issues during the initial configuration, collect logs from this module.
NAS procstats	NAS process stacks from all the nodes. It includes the stack trace for all daemons running on the system.
Appliance	Hardware-specific logs.
SDS	Logs for the software-defined storage plug-in.
Upgrade	Logs generated during an upgrade from all the nodes
Backup	NetBackup client logs and other diagnostic information.
VDD	Veritas Data Deduplication (VDD) logs if VDD is configured on the cluster nodes.

- Under **Historical logs**, you can collect archived logs or logs for a specific duration. To include historical logs, click **Archived logs**. To include logs for a specific duration, click **Select date range** and select the duration or select **Custom range** to customize the duration by specifying the start and the end date. Including logs for a specific time period reduces the size of the generated log package and the time required to collect the logs.
- Under **Package options**, specify a name for the log package and optionally the case number if provided by Veritas Support.

- Under **Nodes**, select the nodes for which you want to collect the logs.

5 Click **Generate**.

The log package is generated in the `/log/support_logs` location. To view the progress of the operation, click **View details** on the **Diagnostics** page.



If you had specified the case number, it is prefixed to the package name. The generated log package is displayed under **Packaged logs**. You can now download the package or upload it to Veritas Support for error analysis.

See "Uploading log packages to Veritas Support" on page 214.

See "Downloading log packages" on page 215.

See "Deleting log packages" on page 215.

Uploading log packages to Veritas Support

You can share the logs with the Support personnel to help you diagnose problems. The existing log packages are displayed under **Packaged logs**.

To upload the log package:

1. Sign in to the Access Appliance UI using the following URL:

   `http://console-ip:14161`

   where *console-ip* is the management console IP address where the web interface is hosted..

2. Select the package that you want to upload to and click **Send**.

3. On the **Send packaged logs** dialog box, use the Secure File Transfer Protocol (SFTP) to upload the logs. Specify the following details, and then click **Send**.

Field	Description
Server address	Remote server name.
Port	Port number. By default the SFTP server is configured to listen to port 22.
User name	User name to authenticate with the server.
Password	Password to authenticate with the server.
Path	Location to upload the log package.

Downloading log packages

If the node is a part of the cluster, you can download log packages from the node.

To download generated log packages

1. Sign in to the Access Appliance UI using the following URL:

   `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 > Diagnostics**.

   The log packages that are already created are displayed under **Packaged logs**.

3. Select the log package that you want to download and click **Download**.

4. On the **Download packaged logs** dialog box, click **Download**.

Deleting log packages

You can delete the generated log packages from the cluster.

To delete log packages:

1. Sign in to the Access Appliance UI using the following URL:

   `http://`*`console-ip`*`:14161`

   where *console-ip* is the management console IP address where the web interface is hosted..

2. In the left pane, click **Settings > Diagnostics**.

   A list of log packages is displayed.

3. To delete a log package, select the log package and click **Delete**.

4. When prompted for confirmation click **Continue**.

   A notification is displayed on the top of the page. To monitor the progress, click **View details**. The log packages is deleted from the `/log/support_logs` directory. After the log package is deleted successfully a notification is displayed on the top of the page.

Section 7

Provisioning and managing Access Appliance file systems

- Chapter 22. Creating and maintaining file systems

Chapter 22

Creating and maintaining file systems

This chapter includes the following topics:

- About creating and maintaining file systems
- 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.

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 223.
- 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 22-1 describes the LUN throughput restrictions.

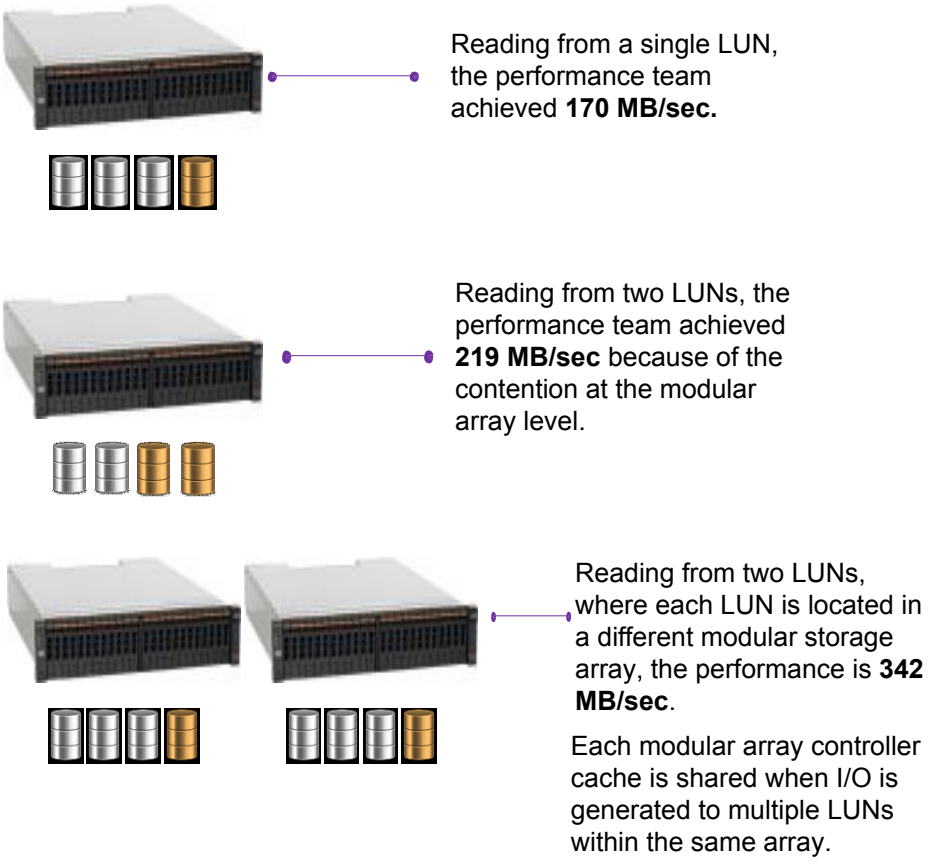
- Consider populating each storage pool with the same number of disks from each HBA. Alternatively, consider how much of the total I/O bandwidth that the disks in the storage pool can use.

  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 22-1 LUN throughput - details on the LUN throughput restrictions



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 22-1 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 tunefstab` 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:

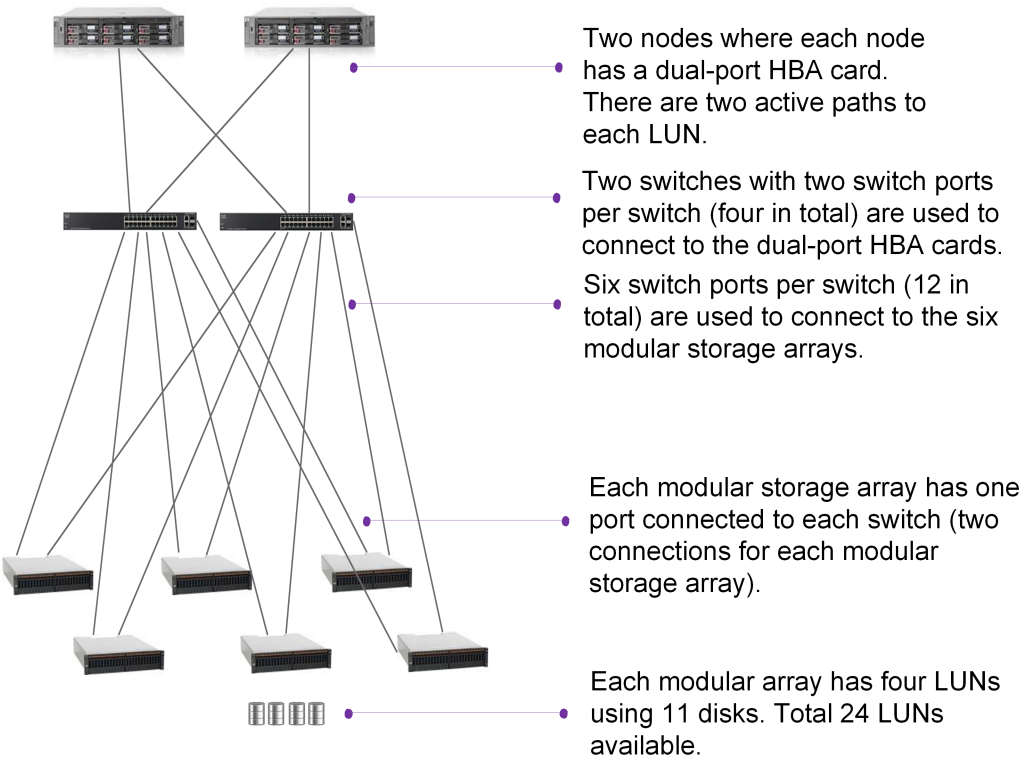
- Maximize the I/O performance.
- Proportion the I/O bandwidth available from the storage layer.
- Balance the I/O load evenly across multi-user applications running on multiple nodes in the cluster.

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)

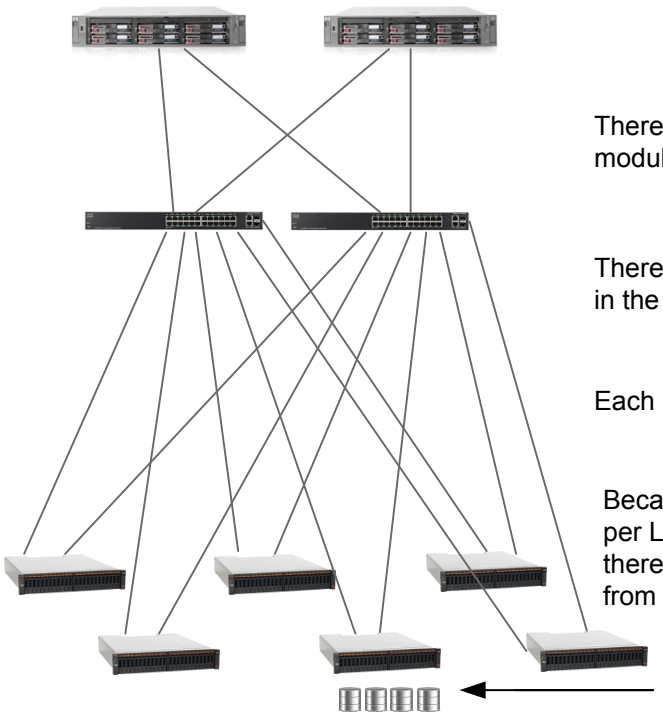
Figure 22-2 An example hardware configuration



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 22-3 LUN configuration



Each LUN is composed of 11 disks in a RAID-0 configuration.

There are 4 LUNs available per modular array controller.

There are a total of 24 LUNs available in the environment.

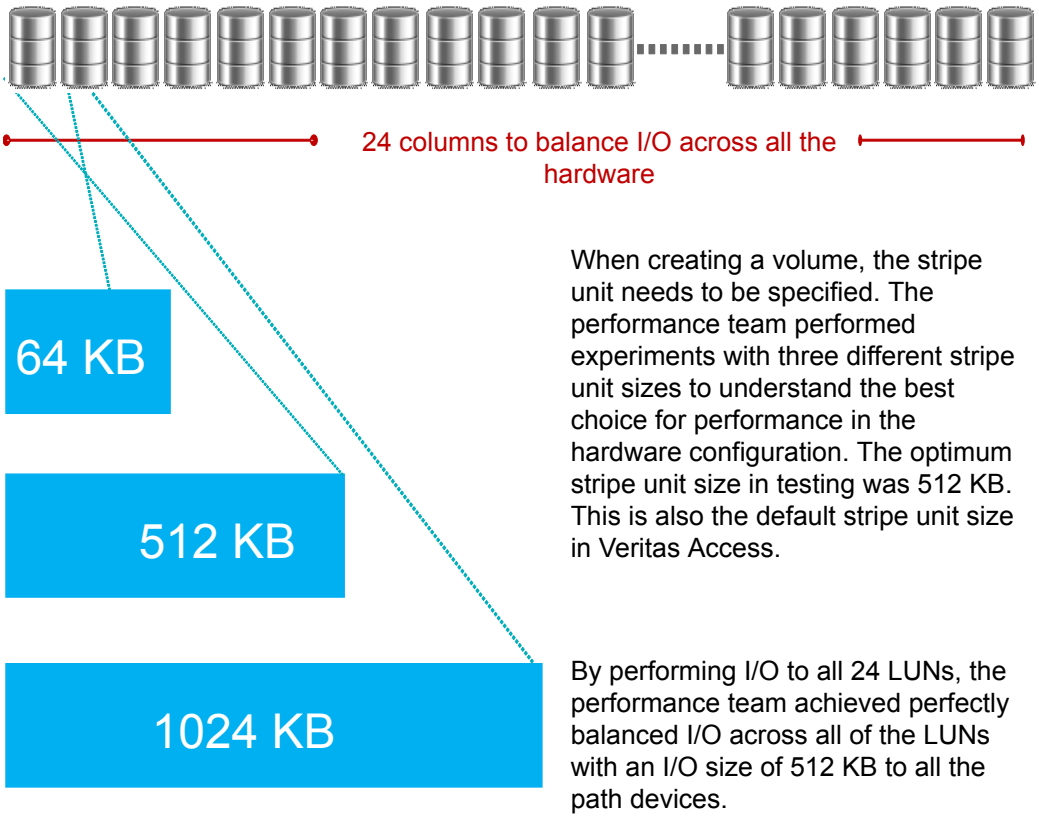
Each LUN is approximately 3 TB.

Because there are 2 active paths per LUN (dual-port HBA card), there is a total of 48 active paths from each node.

Each of the 6 modular array controllers has 4 LUNs using 11 disks. There are 24 identical LUNs available in total.

All 24 LUNs have an identical hardware configuration.

Figure 22-4 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

fs_name	Specifies the name of the file system
minret	Specifies the minimum time period of retention. It can contain: [1-9](s|S|h|H|d|D|m|M|y|Y)
maxret	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 in the Enterprise and Compliance lockdown mode but it can be decreased only in the Enterprise lockdown mode. The range of the minimum and maximum retention period should be between 1 hour and 60 years.

If the minimum retention and maximum retention values are provided with suffix as *M* or *m*, the retention period is converted to seconds using the following the calculations:

```
Days in a year: 365
Days in a month: 31
Seconds in a day: 24*60*60
```

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: 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.

When you set the retention period, remember that:

- 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.

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 84.

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 218.

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] [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]
[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]
[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]
[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]  [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 `Storage> pool free` command. See "About configuring storage pools" on page 85.
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 `Storage> pool list` command. To find a list of disks, use the `Storage> disk list` 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.
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 236.

See "Adding or removing a column from a file system" on page 238.

See "Increasing the size of a file system" on page 239.

See "Decreasing the size of a file system" on page 241.

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]
```

tier_name	Specifies which tier name. If the specified file system does not exist, an error message is displayed.
fs_name	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 `Storage> pool list` command. See "About configuring storage pools" on page 85. To find a list of the existing disks, use the `Storage> disk list` 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 `disk`. Available options are: ■ disk - if the protection is set to `disk`, then mirrorsare 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. ■ pool - if the protection is set to `pool`, 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 `Storage> fs growby` 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 `Storage> pool list` command. See "About configuring storage pools" on page 85. To find a list of the existing disks, use the `Storage> disk list` 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 `disk`. Available options are: ■ disk - if the protection is set to `disk`, 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. ■ pool - if the protection is set to `pool`, 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 `growto` and `growby` 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]
```

fs_name — 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.

time — 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.

defrag_level — Specifies the defragmentation level such as `dir`, `extent`, or `all`.

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]
```

sched_name	Specifies the the name of the schedule.
sched_duration	Specifies the duration of the defragmenatation job.
minute	Specifies the minute (0-59).
hour	Specifies the hour (0-23).
day_of_the_month	Specifies the day of the month (1-31).
month	Specifies the month of the year (1-12).
day_of_the_week	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.

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.

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 `numinodes`, *value* is the number of inodes used. The default alert value for `numinodes` is set at `0`. An alert will not be sent until you set it to a different value.

numspace — When setting the alert for `numspace`, `value` 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 — `fullspace` is the tunable for setting an alert if the file system becomes full. When file system usage is above the limit set by the fullspace tunable, all the NFS 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 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 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 — `fullinodes` is the tunable for setting an alert if the file system becomes full. When inodes on a file system reach the limit of fullinodes, the NFS 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 shares to be changed to read-write. By default, the fullinodes tunable is set to 0, which indicates that the fullinodes tunable is disabled.

fs_name — Name of the file system for which you want to set the file system alerts. *fs_name* 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 `snapshot_name`: `flags`, `ctime`, and `mtime`.

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 numinodes | numspace | fullspace | fullinodes
 [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/fullinodes 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 247.

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 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
```

Section 8

Provisioning and managing Access Appliance shares

- Chapter 23. Creating shares for applications
- Chapter 24. Creating and maintaining NFS shares
- Chapter 25. Integrating Access Appliance with Data Insight

Chapter 23

Creating shares for applications

This chapter includes the following topics:

- About file sharing protocols
- About concurrent access
- 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 253.

Table 23-1 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 112.

Table 23-1 Protocols *(continued)*

Protocol	Definition
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 105.

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 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.

The following sections describe concurrent access with multiple protocols.

See "About concurrent access with NFS and S3" on page 253.

See "About file sharing protocols" on page 252.

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.

Chapter 24

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 105.

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 256.

See “Creating an NFS share with continuous replication ” on page 255.

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 256.

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 255.

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:

`rw`	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.
`ro (Default)`	Grants read-only permission to the directory. Hosts mounting this directory will not be able to change it.
`sync (Default)`	Grants synchronous write access to the directory. Forces the server to perform a disk write before the request is considered complete.
`async`	Grants asynchronous write access to the directory. Allows the server to write data to the disk when appropriate.
`secure (Default)`	Grants secure access to the directory. Requires that clients originate from a secure port. A secure port is between 1-1024.
`insecure`	Grants insecure access to the directory. Permits client requests to originate from unprivileged ports (those above 1024).
`secure_locks (Default)`	Requires authorization of all locking requests.
`insecure_locks`	Some NFS clients do not send credentials with lock requests, and therefore work incorrectly with `secure_locks`, in which case you can only lock world-readable files. If you have such clients, either replace them with better ones, or use the `insecure_locks` option.
`root_squash (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.
`no_root_squash`	Disables the `root_squash` option. Allows root users on the NFS client to have root privileges on the NFS server.

`wdelay (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 `wdelay` option is deprecated, and is supported for backward-compatibility only.
`no_wdelay`	Disables the `wdelay` option. The `no_wdelay` option has no effect if the `async` option is also set. **Note:** The `no_wdelay` option is deprecated, and is supported for backward-compatibility only. Using the `no_wdelay` option is always effective.
`subtree_check`	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.
`no_subtree_check (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 `no_subtree_check`. One such situation might be the export of the home directory. Most other situations are best handled with `subtree_check`.
`fsid (Default)`	Allows the Access Appliance administrator to associate a specific number as fsid with the share.
`nordirplus`	Allows you to disable a readdirplus remote procedure call (RPC).
`sec`	Specifies the Kerberos security options for exporting an NFS share. The value can be `krb5`, `krb5i` , `krb5p`, or `sys`. The `sys` 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]
   ```

nfsoptions	Comma-separated list of export options from the set.
export_dir	Specifies the name of the directory you want to export. The directory name should start with `/vx`, and only `a-zA-Z0-9_/@+=.:-` characters are allowed for `export_dir`.
share_name	Specifies the share name. If the share name is not specified, then the default value is vxdefault.
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 257.

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.
   - Is a valid user on the Access Appliance cluster.
   - Is either the owner of the share directory or has the necessary permissions to access the NFS share directory.

   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 311.

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 257.

Chapter 25

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.

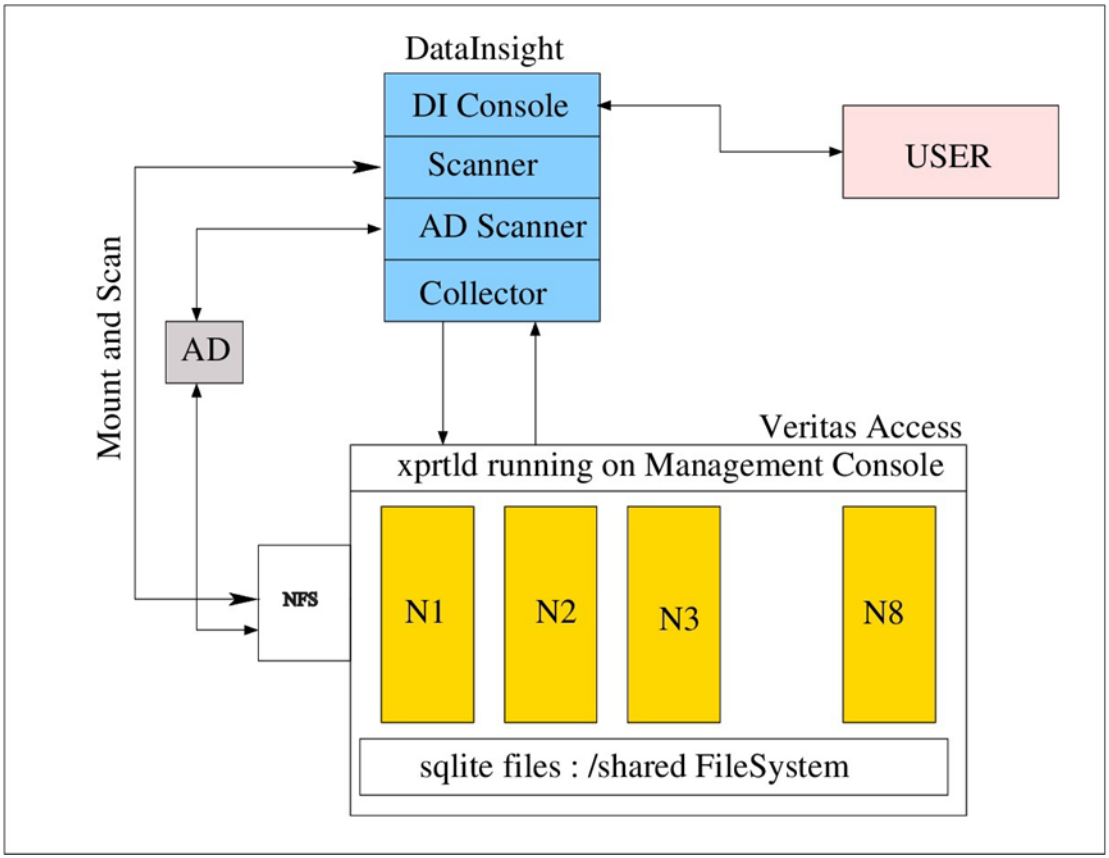
Access Appliance allows DI to display NFS share-level activities on the DI GUI and thus helps the users to make better business decisions and risk-management strategy.

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 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 25-1 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 25-2 Access Appliance integration with Data Insight

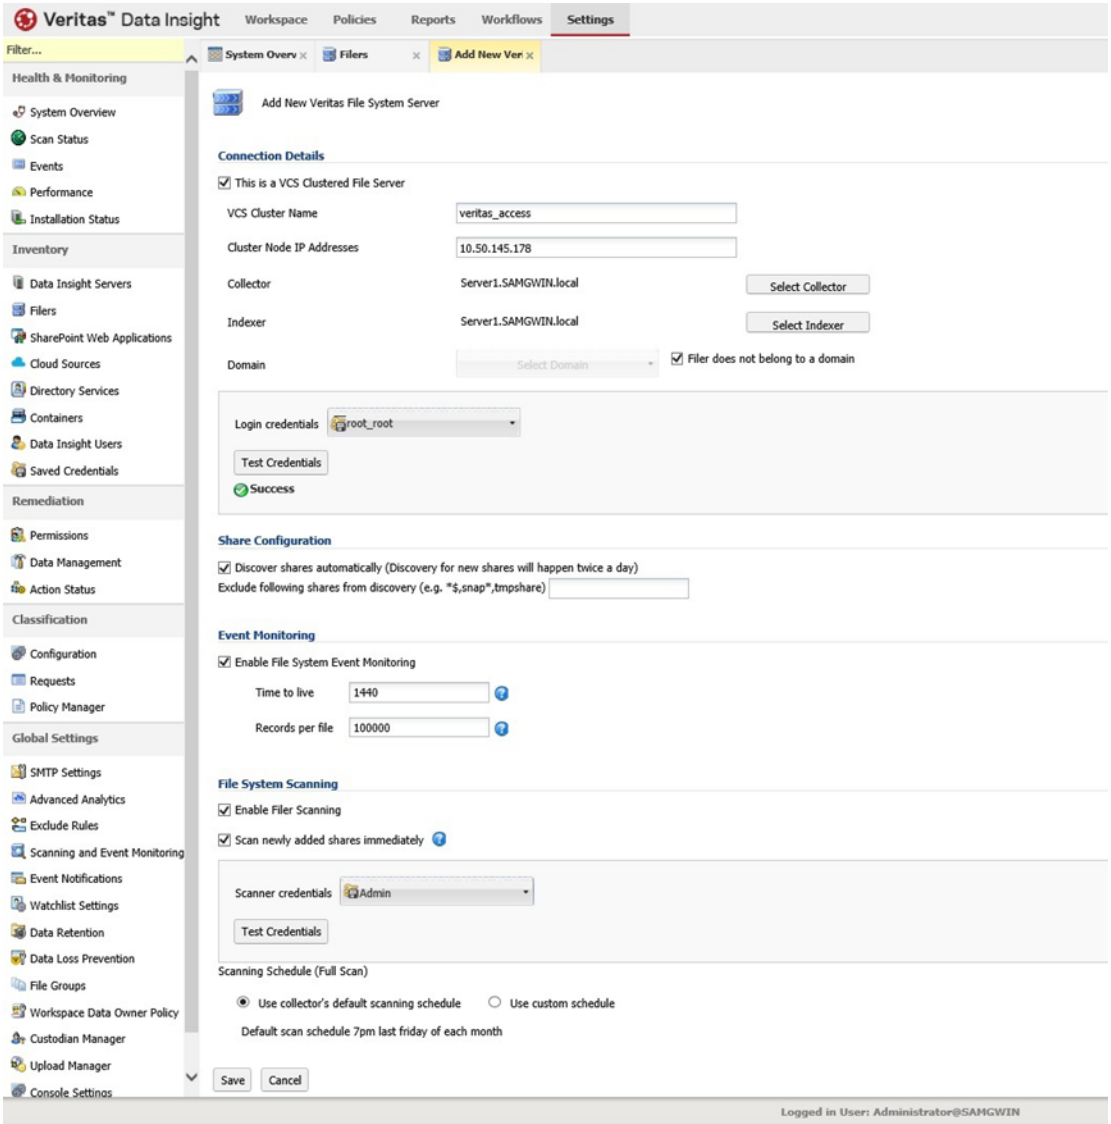
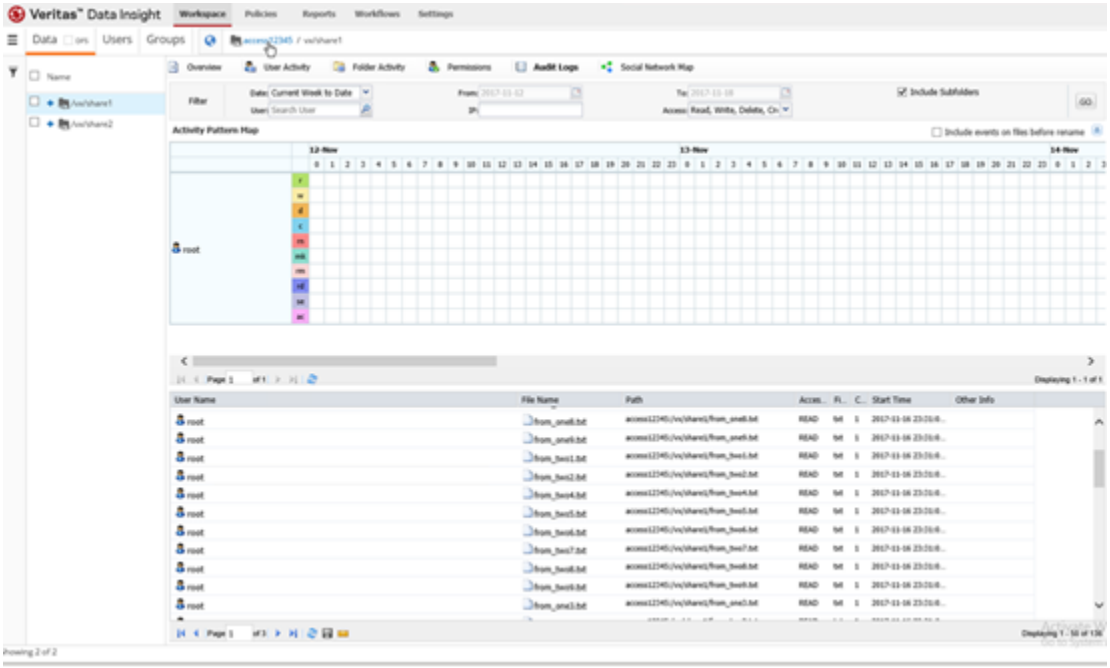


Figure 25-3 Audit information



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.

Section 9

Managing Access Appliance storage services

- Chapter 26. Configuring continuous replication
- Chapter 27. Using snapshots
- Chapter 28. Using instant rollbacks

Chapter 26

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
- Configuring a continuous replication job using the GUI
- Continuous replication failover and failback
- 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).

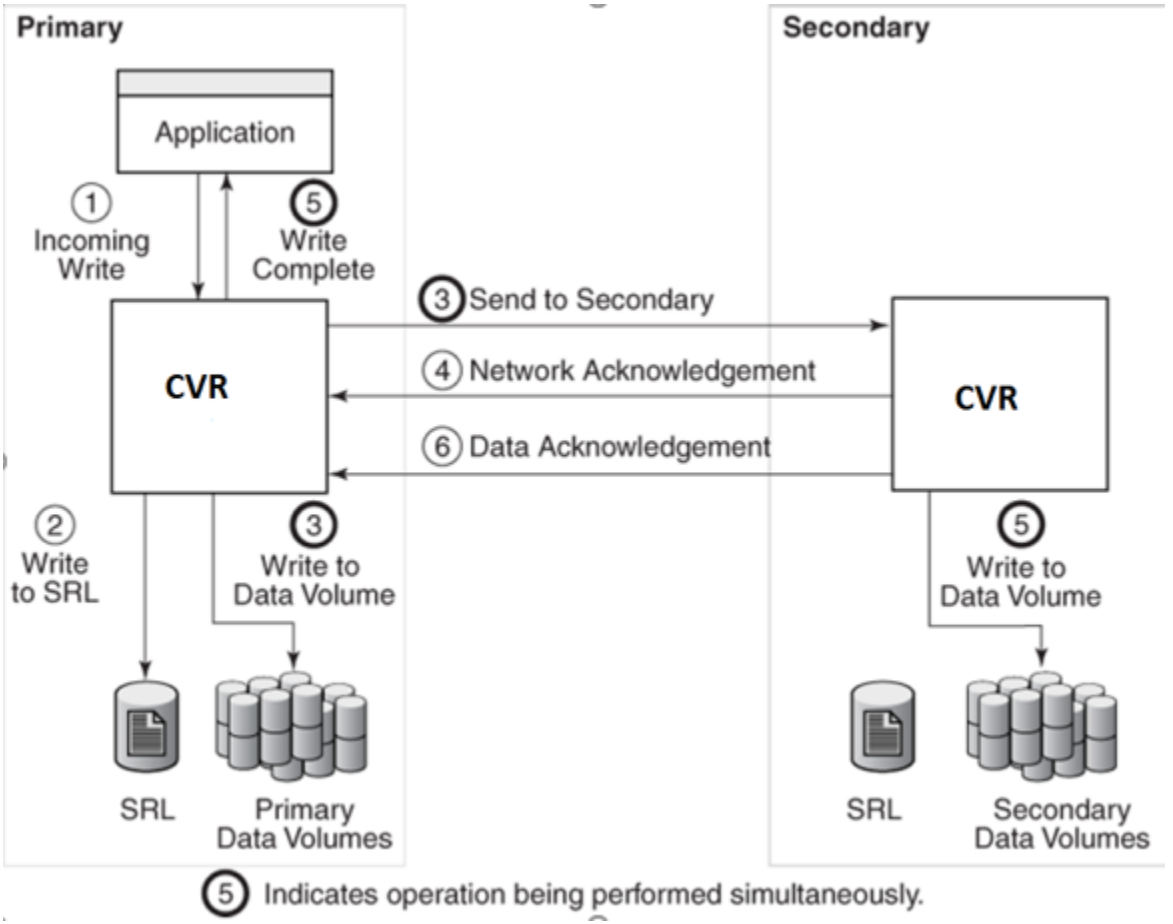
If a synchronous operation for a file system is paused, replication stops when the SRL (Storage Replicator Log) overflows and the `replication> continuous status` command displays the replication status as "`logging to DCM (needs dcm resynchronization)`". If the replicator log (SRL) overflows, the RDS (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.
- Resynchronize the DCM by running the `replication continuous resync <fs_name>` command from CLISH 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 26-1 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.

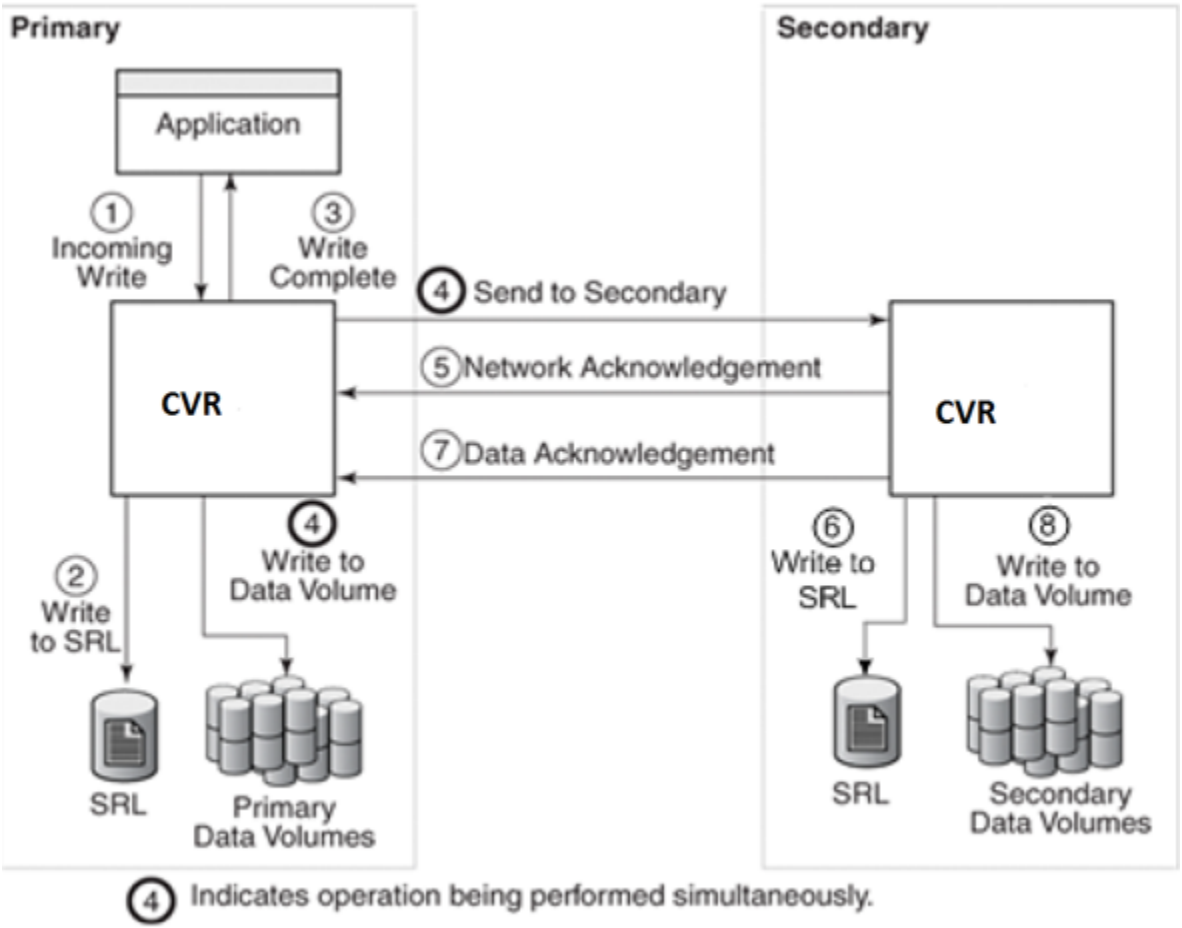
When the write is written to the data volumes on the secondary host, CVR on the secondary host sends a data acknowledgment to the primary host. CVR marks the write as complete in the SRL when the primary receives the data acknowledgment from all the secondary hosts.

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 26-2 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.
- If no physical IPs or virtual IPs have been configured, then you must provide the *device* and *netmask* parameter while binding the virtual IP. If you want to add a virtual IP from a different network, you must add a routing entry using the route add command. See “Support for multiple data subnets” on page 67.

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.

delayed=yes/delayed=no	The delayed parameter value is **yes** if you want to set up continuous asynchronous replication. By default, its value is **yes**.
rvg_name	RVG name at primary site.
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.

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 297.

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.

To modify the existing replication IP

- Before you modify the IP, ensure that you add the new IP address to the network IP address pool using the `network ip addr add` command and that the newly added IP is on the same node as the replication IP used for continuous replication.

  To modify the existing virtual IP specified in a continuous replication job, enter the following command:

  `Replication>` **`continuous config modifyip`** ***`link_name src_ip_addr tgt_ip_addr`*** **`[`*`src_netmask`*`] [`*`tgt_netmask`*`]`**

link_name	The name of the link.
src_ip_addr	The virtual IP of the source cluster.
tgt_ip_addr	The virtual IP of the destination cluster.
src_netmask	The netmask/prefix of the source replication IP address
tgt_netmask	The netmask/prefix of the destination replication IP address

  After the command is run, the old replication IP on the source and target is deleted. You can verify using the `network ip addr show` command.

To modify the replication mode

- To modify the replication mode specified in a continuous replication job, enter the following command:

  `Replication>` **`continuous config modify`** ***`rvg_name`*** **`[delayed=yes/no]`**

  This command changes the mode of replication from `delayed=yes` (asynchronous replication) to `delayed=no` (synchronous replication) and vice-versa.

rvg_name	Name of the RVG for which replication mode is to be changed.
[delayed=yes|delayed=no]	Specifies if the continuous replication should happen in delayed mode.

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 26-1 describes the important attributes displayed by the `Replication> continuous status` command.

Table 26-1

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.

Table 26-1 *(continued)*

Attribute	Description
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 `Replication> continuous config bind` command. For the secondary site, it provides the continuous replication IP which binds on the destination cluster using the `Replication> continuous config bind` command.
RVG state	Specifies the state of the primary RVG. See Table 26-2 for details on its various states.
Configured mode	Specifies the continuous replication configured mode. It may be *synchronous-override* or *asynchronous*.
Current mode	Specifies the mode of replication - *asynchronous* or *synchronous* that is used to replicate data to the secondary.
Data status	Shows the data status for the secondary. See Table 26-3 for details on its various states.
Replication status	Specifies the status of the replication to the secondary. See Table 26-4 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 26-5 for details on its various states.

Table 26-1 *(continued)*

Attribute	Description
Timestamp information	Shows the time by which secondary is lagging behind the primary.

Table 26-2 describes the values for the RVG state.

Table 26-2 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.
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 26-3 describes the values for the Data status.

Table 26-3 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.

Table 26-3 Data status *(continued)*

Value	Description
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 26-4 describes the values for the Replication status.

Table 26-4 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: ■ *needs dcm resynchronization*—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 needs_recovery state. ■ *primary needs_recovery*—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: ■ *primary paused*—primary RLINK is paused. ■ *secondary paused*—secondary RLINK is paused.

Table 26-4 Replication status *(continued)*

Value	Description
Paused due to error	Replication to secondary is paused due to the following errors: ■ *secondary config error*—secondary has some configuration error. ■ *secondary log error*—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.
Resync in progress	Resynchronization to the secondary is in progress. ■ *autosync*—Resynchronization type is autosync. ■ *dcm resynchronization*—resynchronization after an SRL overflow. ■ *failback resynchronization*-resynchronization using failback logging. ■ *smartsync*—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: ■ *primary paused*—primary RLINK is paused. ■ *secondary paused*—secondary RLINK is paused.
Resync paused due to error.	Resynchronization to Secondary is paused because of the following errors: ■ *secondary config error*—secondary has some configuration error. ■ *secondary log error*—Secondary SRL has an I/O error.
Resync paused due to network disconnection.	Resynchronization to secondary is paused due to some network problem.

Table 26-4 Replication status *(continued)*

Value	Description
Stopped	Replication to secondary is stopped due to the following: ■ *Primary detached*—Primary RLINK is detached. ■ *Secondary detached*—Secondary RLINK is detached.
N/A	The replication status cannot be determined.

Table 26-5 describes the values for the Logging to field.

Table 26-5 Logging to status

Value	Description
DCM (contains xxx KB) (log_type)	DCM is active (in use) for the replication to the secondary. The *log_type* can be *autosync*, *failback logging*, or *SRL protection logging*. 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 297.

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 evpsn ev_partition_path
```

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.

Configuring a continuous replication job using the GUI

You cannot use the GUI to enable continuous replication on an existing file system. You can use the Access Appliance CLISH to enable continuous replication.

See “Setting up the file system to replicate ” on page 285.

If file systems are not created, then you can use the **Provision storage** wizard to configure continuous replication.

Creating a replication link

Note: You cannot use the GUI to add routing entries during replication configuration. If you want to add routing entries, See “Configuring routing tables” on page 58.

To create a replication link

1 **Prerequisites:**

Create the replication VIP.

2 Create the replication link between the source and the destination clusters.

3 Enter the source and the destination cluster details.

Link Name	Enter the link name for the source and the destination clusters.
Cluster Name	Enter the cluster IP address for the destination cluster, the primary console IP where you installed Access Appliance.
Username	Enter the username for the destination cluster. To authenticate a different cluster for replication, you must provide the "primary" user credentials.
Password	Enter the password for the destination cluster.

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

Setting up a virtual IP for replication

To set up a replication virtual IP

1 Set up a dedicated virtual IP address between the source and the destination cluster.

If you have already set up a virtual IP address, this field is not enabled.

Enter a unique virtual IP address.

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

From the **Settings > Replication Management** page, you can view the following information about replication:

Replication VIP	Dedicated virtual IP address for the replication service. The replication services requires VIP addresses not already in use for the two clusters to communicate.
VIP Status	Status of the virtual IP address, ONLINE, or OFFLINE.
Replication Device	Device used by the replication service.
Online on Node	Cluster node on which the replication service is currently running.
Node DB Status	Status of the replication database, ONLINE, or OFFLINE.
Daemon Status	Status of the daemon, ONLINE, or OFFLINE.
Service Status	Status of the replication service.

Deleting a replication link

To delete a replication link

1. Select the replication link that you want to delete.
2. Click the ellipsis beside the replication link and click **Remove Link**.
3. Enter the username and password for the source and destination clusters.
4. View the **Recent Activity** panel for the status of the task to delete the replication link.

Modifying a replication unit

To modify a replication unit

1. **Prerequisites:**

   Create a replication VIP.

   Create a replication link.

   Configure a replication job.

   Set a replication schedule.
2. **File Systems**: Select the file system for which you want to modify the replication unit.
3. Click the **Replication** tab.
4. Click on the **Modify Replication Unit** button.
5. Follow the prompts and click **Add**.
6. View the **Recent Activity** panel for the status of the task.

Deleting a replication virtual IP

To delete a replication virtual IP

1. **Prerequisites:**

   Delete the replication job.

   Delete the replication link.
2. **Settings > Replication Management**: Click **Delete Replication VIP**.
3. Confirm that you want to delete the replication VIP.
4. View the **Recent Activity** panel for the status of the task.

Setting a replication schedule

Access Appliance supports periodic replication, where the data gets replicated from the source to the destination cluster at regular intervals as defined by the replication schedule.

To set a replication schedule

1. **Prerequisites:**

   Set up a virtual IP address between the source and the destination cluster.

   Create a replication link.

   Configure a replication job.

   Select a replication link.
2. Click **File Systems**, and select a file system for which you want to run replication.
3. Click **Replication > Configure Replication Job**.
4. Step through the **Configure Replication Job** wizard tasks to **Set Schedule**.
5. Specify the recurrence pattern and how often the replication job should recur.
6. Click **Set Schedule**.
7. View the **Recent Activity** panel for the status of the task.

To modify a replication schedule

1. **Prerequisites:**
2. Set up a replication schedule.
3. Click **Edit Replication Schedule** and update the appropriate options.

Viewing information about replication

From the **Settings > Replication Management** page, you can view the following information about replication:

Replication VIP	Dedicated virtual IP address for the replication service. The replication services requires VIP addresses not already in use for the two clusters to communicate.
VIP Status	Status of the virtual IP address, ONLINE, or OFFLINE.
Replication Device	Device used by the replication service.
Online on Node	Cluster node on which the replication service is currently running.
Node DB Status	Status of the replication database, ONLINE, or OFFLINE.

Daemon Status	Status of the daemon, ONLINE, or OFFLINE.
Service Status	Status of the replication service.

Stopping or starting a continuous replication job

To stop or start replication.

1. **Prerequisites:**

   Configure at least one storage pool.

   Set up a replication virtual IP address.

   Set up a replication link.

   For volume-level replication, select a file system and provide a replication link.

   For file system replication, create a collection of all the directories and all the files that need to be replicated within a file system into a replication unit. Select a replication link and schedule the replication job to be replicated at specific intervals or sync it on demand.
2. **File Systems > Replication**: At the ellipses, select **Start** or **Stop**.
3. Follow the wizard prompts.
4. View the **Recent Activity** panel for the status of the task.

See “Pausing and resuming a continuous replication job” on page 302.

Pausing and resuming a continuous replication job

To pause or resume replication

1. **Prerequisites:**

   Configure at least one storage pool.

   Set up a replication virtual IP address.

   Set up a replication link.

   For volume-level replication, select a file system, and then provide a replication link.
2. **File Systems > Replication**: At the ellipses, select **Pause** or **Resume**.
3. Follow the wizard prompts.
4. View the **Recent Activity** panel for the status of the task.

See “Stopping or starting a continuous replication job” on page 302.

Failing over or failing back a continuous replication job

To failover or failback a replication job

1 **Prerequisites:**

Configure at least one storage pool.

Set up a replication virtual IP address.

Set up a replication link.

For volume-level replication, select a file system and provide a replication link.

2 Go to **File Systems** > **Replication**: At the ellipses, select **Failover Job** or **Failback Job**.

3 Follow the wizard prompts.

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

See "Stopping or starting a continuous replication job" on page 302.

Unconfiguring a continuous replication job

You can remove replication by unconfiguring it from the Access Appliance cluster.

To unconfigure a continuous replication

1 **Prerequisites:**

Configure replication.

2 **File Systems > Replication**: At the ellipses, select **Unconfigure**.

3 Follow the wizard prompts.

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

See "Stopping or starting a continuous replication job" on page 302.

See "Pausing and resuming a continuous replication job" on page 302.

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 availble 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 share information that is associated with file system from the source cluster to the destination cluster. Share information has to be done manually.

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 uplanned 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.

While doing a planned replication continuous job failover from source to target, ensure that no NFS shares are configured with the source file system.

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.

While doing a planned replication continuous job failback from source to target, ensure that no NFS shares are configured with the source file system.

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 onlines 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 shares on the original source cluster, it is recommended that you stop the NFS 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 285.
- 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.
- If you try to add a file system to an existing RVG when the `replication continuous enable` operation is in progress, then the file system remains offline till the operation is complete. As a result, the application is not able to serve read-write requests. The time period for which the file system remains offline depends on the size of the file system and the amount of data present on it. Ensure that you consider this downtime before you run the `replication continuous enable` command.

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,
[delayed=yes|delayed=no]	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 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  Time of Key Import
                            VIP
========= ================= =================== ============================

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
```

Chapter 27

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, you must bring the snapshot online. You can then create an NFS share using the snapshot name as the path. For example: /vx/fs1/:snap1. The snapshot can only be mounted through NFS 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 `read-only` and `read-write`. Default value for mode parameter is `mode=read-only`. WORM-enabled snapshots can be online in only `read-only` 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 *schedule_name* cannot contain an underscore ('_') as part of its value. For example, `sch_1` 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 `max_snapshot_limit` is 366.
minute	This parameter may contain either an asterisk like '*/15'', which implies every 15 minutes, or a numeric value between 0-59. **Note:** If you are using the '*/xx' format, the smallest value for 'xx' is 15. You can enter */(15-59) 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 ass 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 `worm=yes/no`. Default value of this parameter is no. To create a schedule with `worm=yes`, 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]
```

fs_name	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.
schedule_name	Displays the schedule name. If no schedule name is specified, then all of the schedules created under *fs_name* 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]
```

Chapter 28

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 329.

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 329.

   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 329.

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 334.

  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.

Section 10

Reference

- Appendix A. Access Appliance documentation

Appendix A

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*
- *Veritas Access Command Reference Guide*
- *Access Appliance Release Notes*
- *Veritas Access RESTful API Guide*
- *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 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} value
       ldap get {server|port|basedn|binddn|ssl|rootbinddn|
            users-basedn|groups-basedn|netgroups-basedn}
```

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>

exit         --return to the previous menus
logout       --logout of the current CLI session
man          --display on-line reference manuals
passwd       --change the administrator password
```

```
show         --show the administrator details
supportuser  --enable or disable the support user
user         --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 253
- FastResync 226
- SNMP notifications 196

about
- Access Appliance continuous replication 273
- Active Directory (AD) 75
- bonding Ethernet interfaces 44
- buckets and objects 117
- configuring disks 85
- configuring in IPv4 and IPv6 mixed mode 64
- configuring routing tables 58
- configuring storage pools 85
- Continuous replication failover and failback 303
- creating and maintaining file systems 218
- Ethernet interfaces 52
- Multi-protocol support for NFS with S3 124
- NFS file sharing 254
- shares 252
- snapshot schedules 318
- snapshots 311
- storage provisioning and management 84
- striping file systems 223
- the IP addresses for the Ethernet interfaces 52

About Configuring
- network 44

about maximum IOPS 230

about S3
- viewing information 128

Access
- accessing the CLISH 20

Access Appliance
- about 15
- key features 15
- product documentation 336

Access appliance
- relocate 186

Access Appliance continuous replication
- about 273
- description of feature 274

Access Appliance continuous replication *(continued)*
- starting 279

Access Appliance to the Kerberos realm
- adding and configuring 110

accessing
- Access Appliance product documentation 336
- CLISH 20
- disk details 102
- man pages 337

Active Directory (AD)
- about 75

activity
- monitoring 192

adding
- a column to a file system 238
- a syslog server 197
- IP address to a cluster 52
- mirror to a file system 236
- VLAN interfaces 56

adding and configuring
- Access Appliance to the Kerberos realm 110

alerts
- email 188
- file system removing 249
- SNMP 190

Authenticate
- NFS clients 109

authenticating
- NFS clients using Kerberos 109

AutoSupport
- about 186
- configure 186
- proxy server 188

B

best practices
- creating file systems 219

bind distinguished name
- setting for LDAP server 71

bonding
- Ethernet interfaces 45

bonding Ethernet interfaces
about 44
bringing
file system online or offline 235
buckets and objects
about 117

C

cache object
destroying for an instant rollback 334
cache objects
listing 332
Call Home 186
See also AutoSupport
proxy server 188
changing
an IP address to online
on any running node 52
checking
and repairing a file system 245
for stale mirrors on file systems 246
on the status of the NFS server 106
clearing
DNS domain names 50
DNS name servers 50
LDAP configured settings 71
client configurations
displaying 74
LDAP server 74
cluster
adding an IP address to 52
changing an IP address to online for any running node 52
displaying all the IP addresses for 52
collect logs
datacollect 200
columns
adding or removing 238
command history
displaying 22
Command-Line Interface (CLI)
getting help on how to use 20
communicating
source and destination clusters 281
Concurrent access
about 253
concurrent access
NFS and S3 253
Configuring
Object Store server 113
configuring
event notifications 192, 195
IP routing 58
IPv4 and IPv6 mixed mode 64
NFS client for ID mapping 108
NIC devices 57
NSS lookup order 78
VLAN interfaces 56
Configuring Access Appliance
ID mapping for NFS version 4 108
configuring disks
about 85
configuring Ethernet interfaces
about 52
configuring routing tables
about 58
configuring storage pools
about 85
Configuring the Object Store server
use case 113
Considerations
creating a file system 219
continuous replication
display status 290
setting up file system 285
unconfiguring 295
Continuous replication failover and failback
about 303
continuous replication job
failing back 303
failing over 303
pausing 302
resuming 302
starting 302
stopping 302
converting
existing file system into a cache object 330
creating
full-sized rollback 327
mirrored file systems 231
mirrored-stripe file systems 231
NFS share with continuous replication 255
NFS share without replication and encryption 256
shared cache object 330
simple file systems 231
snapshot schedules 319
snapshots 313

creating *(continued)*
space-optimized instant rollbacks 327
storage pools 86, 100
striped file systems 231
striped-mirror file systems 231
Creating a file system
considerations 219
creating and maintaining file systems
about 218
creating file systems
best practices 219
current Ethernet interfaces and states
displaying 52

D

data flow
continuous replication synchronous mode 276
Data Insight
integrating Access Appliance 266
datacollect
device logs 200
decreasing
size of a file system 241
defrag 243
defragmenting
file systems 243
deleting
NFS options 261
route entries from routing tables of nodes in cluster 58
snapshot schedules 322
syslog server 197
VLAN interfaces 56
description of Access Appliance continuous replication 274
destroying
a file system 249
cache object of an instant rollback 334
instant rollbacks 330
snapshots 315
storage pools 86
directories
displaying exported 260
unexporting the share 261
disabling
DNS settings 50
FastResync option 246
LDAP clients
configurations 74
disabling *(continued)*
quota limits used by snapshots 315
disk layout versions
upgrading 249
Disk quotas
file systems 89
usage 89
disks
accessing details about 102
viewing information about 101
displaying
all the IP addresses for cluster 52
command history 22
current Ethernet interfaces and states 52
DNS settings 50
events on the console 198
exported directories 260
file system alert values 248
file systems that can be exported 257
information for all disk devices for the nodes in a cluster 94
LDAP client configurations 74
LDAP configured settings 71
list of syslog servers 197
NFS statistics 107
NSS configuration 78
routing tables of the nodes in the cluster 58
snapshot quotas 315
snapshot schedules 322
snapshots 314
snapshots that can be exported 257
values of the configured syslog server 197
VLAN interfaces 56
displaying WWN information 95
DNS
domain names
clearing 50
name servers
clearing 50
specifying 50
settings
disabling 50
displaying 50
enabling 50
domain name
for the DNS server
setting 50

E

enabling
DNS settings 50
FastResync for a file system 245
LDAP client configurations 74
quota limits used by snapshots 315
enabling, disabling, and displaying
file system quotas 90
Ethernet interfaces
bonding 45
event notifications
configuring 192, 195
events
displaying on the console 198
monitoring 194
exporting
an NFS share 257
NFS snapshot 265
exporting for Kerberos authentication
NFS share 262

F

failing back
continuous replication job 303
failing over
continuous replication job 303
failover and failback
about 303
FastResync
about 226
file system
converting into a cache object 330
file system alert values
displaying 248
file system alerts
removing 249
setting 247
file system quotas
enabling, disabling, and displaying 90
setting and displaying 91
file systems
adding a mirror to 236
bringing online or offline 235
checking and repairing 245
checking for stale mirrors 246
creating 231
decreasing the size of 241
defragmenting 243
destroying 249
file systems *(continued)*
disabling FastResync option 246
enabling FastResync 245
increasing the size of 239
listing with associated information 236
removing a mirror from 236
restoring from an instant rollback 328
that can be exported
displayed 257
types of layout 222
filter
about 192, 195
firewall setting 60
forcefully
importing new LUNs for new or existing pools 96
fsck
about 227

H

hardware
alerts 188, 190
notifications 188, 190
history command
using 22
hostname or IP address
setting for LDAP server 71
how to use
Command-Line Interface (CLI) 20

I

ID mapping for NFS version 4
configuring Access Appliance 108
importing
new LUNs forcefully for new or existing pools 96
increasing
size of a file system 239
initiating host discovery of LUNs 96
instant rollbacks
about 325
bringing online 329
creating a shared cache object 330
creating full-sized 327
creating space-optimized 327
destroying 330
listing 328
refreshing from a file system 329
restoring a file system from 328
taking offline 329

integrating Access Appliance with Data Insight 266
IP addresses
adding to a cluster 52
displaying for the cluster 52
modifying 52
removing from the cluster 52
IP addresses for the Ethernet interfaces
about 52
IP routing
configuring 58

K

Kerberos authentication
authenticating NFS clients 109
Kerberos realm
adding and configuring Access Appliance for 110
Kerberos share
mounting from the NFS client 263
Kernel-based
NFS server 106

L

layouts
types of file system 222
LDAP
before configuring 70
LDAP password hash algorithm
setting password for 71
LDAP server
clearing configured settings 71
disabling client configurations 74
displaying client configurations 74
displaying configured settings 71
enabling client configurations 74
setting over SSL 71
setting port number 71
setting the base distinguished name 71
setting the bind distinguished name 71
setting the hostname or IP address 71
setting the root bind DN 71
setting the users, groups, and netgroups base DN 71
listing
all file systems and associated information 236
cache objects 332
free space for storage pools 86
instant rollbacks 328
storage pools 86
log files
introduction 199
LUNs
initiating host discovery 96

M

man pages
how to access 337
managing
continuous replication 287
managing NFS shares
using netgroups 261
maximum IOPS
setting 231
mirrored file systems
creating 231
mirrored-stripe file systems
creating 231
modifying
an IP address 52
snapshot schedules 322
monitoring
activity 192
events 194
more command
using 23
mounting
NFS share from the NFS client 263
mounting snapshots 317
Multi-protocol support for NFS with S3
limitations 124

N

naming requirement for new users 33
navigating CLISH
Access 20
networking
proxy server 188
NFS client
configuring for ID mapping 108
NFS clients
authenticating 109
NFS file sharing
about 254
NFS options
deleting 261
NFS server
about 105

NFS server *(continued)*
checking on the status 106
kernel-based 106
starting 106
stopping 106
NFS share
exporting 257
exporting for Kerberos authentication 262
NFS share with continuous replication
creating 255
NFS share without replication and encryption
creating 256
NFS shares
managing using netgroups 261
NFS snapshot
exporting 265
NFS statistics
displaying 107
resetting 107
NIC devices
configuring 57
node
in a cluster
displaying information for all disk devices 94
notifications
email 188
SNMP 190
NSS
displaying configuration 78
lookup order
configuring 78

O

object server 112
object server group-specific parameters
setting up 127
Object Store server
configuring 113
objectstore buckets 120
offline
taking an instant rollback offline 329
online
bringing an instant rollback online 329

P

pausing
continuous replication job for VVR 302
planned failback
process overview 305
planned failover
process overview 304
preserving
snapshot schedules 322
Private network
configure 44
proxy server
configure 188
Public network
configure 44

Q

quota commands
setting and displaying file system quotas 91
quota limits
enabling or disabling snapshot 315
quotas
setting user quotas for users of specified groups 93

R

refreshing
instant rollbacks from a file system 329
removing
a column from a file system 238
IP address from the cluster 52
mirror from a file system 236
snapshot schedules 322
renaming
storage pools 86
replication job for VVR
unconfiguring 303
resetting
NFS statistics 107
restoring
a file system from an instant rollback 328
snapshots 318
resuming
continuous replication job 302
resynchronizing
stale mirrors on file systems 246
route entries
deleting from routing tables 58
routing tables
of the nodes in the cluster
displaying 58

S

self-test 185
setting
  base distinguished name for the LDAP server 71
  bind distinguished name for LDAP server 71
  domain name for the DNS server 50
  file system alerts 247
  filter of the syslog server 197
  LDAP server hostname or IP address 71
  LDAP server over SSL 71
  LDAP server port number 71
  LDAP users, groups, and netgroups base DN 71
  maximum IOPS 231
  prior to configuring LDAP 70
  retention 228
  root bind DN for the LDAP server 71
  severity of the syslog server 197
  user quotas for users of specified groups 93
  WORM over NFS 229
setting up
  object server group-specific parameters 127
severity levels
  about 192, 195
shared cache object
  creating 330
shares
  about 252
showing
  snapshot schedules 322
snapshot schedules
  about 318
  creating 319
  deleting 322
  displaying 322
  modifying 322
  preserving 322
  removing 322
  showing 322
snapshots
  about 311
  creating 313
  destroying 315
  displaying 314
  displaying quotas 315
  enabling or disabling quota limits 315
  mounting 317
  restoring 318
  that can be exported
    displayed 257
snapshots *(continued)*
  unmounting 317
SNMP notifications
  about 196
software updates
  proxy server 188
source and destination clusters
  communicating 281
specifying
  DNS name servers 50
SSL
  setting the LDAP server for 71
starting
  Access Appliance continuous replication 279
  continuous replication job 302
  NFS server 106
stopping
  continuous replication job 302
  NFS server 106
storage pools
  creating 86, 100
  destroying 86
  listing 86
  listing free space 86
  renaming 86
storage provisioning and management
  about 84
striped file systems
  creating 231
striped-mirror file systems
  creating 231
striping file systems
  about 223
support
  AutoSupport 186
syslog server
  adding 197
  deleting 197
  displaying the list of 197
  displaying the values of 197
  setting the filter of 197
  setting the severity of 197

T

test
  hardware 185
traps
  SNMP 190

troubleshooting
about 191

U

unconfiguring
replication job for VVR 303
unexporting
share of exported directory 261
unmounting snapshots 317
unplanned failback
process overview 306
unplanned failover
process overview 306
upgrading
disk layout versions 249
Use case
configuring the Object Store server 113
using
history command 22
more command 23
Using NFS Server 105

V

viewing
information about disks 101
viewing information
about S3 128
VLAN
adding interfaces 56
configuring interfaces 56
deleting interfaces 56
displaying interfaces 56

W

workflow
object server 112
WORM over NFS
setting 229
WWN information
displaying 95