



## **Daily Operations Guide**

Avere OS 3.2

Avere Systems, Inc.  
[www.averesystems.com](http://www.averesystems.com)

Revision C  
2014-Jun-05

## Copyright Information

Copyright © 2009-2014 Avere Systems, Inc. All rights reserved. Specifications subject to change without notice.

No part of this document covered by copyright may be reproduced in any form or by any means – graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system – without prior written permission of the copyright owner.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

## Trademark Information

Avere, FlashCloud, FlashMove, and FlashMirror are registered trademarks or trademarks of Avere Systems, Inc. in the United States and/or other jurisdictions.

Adobe and Flash are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Apple, Bonjour, and Safari are trademarks of Apple Inc., registered in the U.S. and other countries.

Google and Google Chrome are trademarks of Google Inc.

Intel is a trademark of Intel Corp. in the U.S. and other countries.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Microsoft, Active Directory, Windows, Windows NT, and Internet Explorer are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Mozilla and Firefox are registered trademarks of the Mozilla Foundation.

NetApp and Data ONTAP are registered trademarks of NetApp, Inc., in the U.S. and other countries.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

For licensing information on the third-party software used by the Avere product, see the *Third-Party Licenses Reference*.

## Revision History

Rev. A	2014-Apr-01	Initial document for version 3.2
Rev. B	2014-Apr-21	Changed pagination to remove blank pages
Rev. C	2014-Jun-05	Added information about Amplidata cloud core filers in Section 3.2 Updated cache policy screen shot in Section 6.2 Added warnings to local directory description in Section 6.5.2 Added an index

# Table of Contents

Chapter 1. Introduction to Avere OS .....	1
1.1. New in Avere OS 3.1 and 3.2 .....	1
1.1.1. Object Compression (3.2) .....	1
1.1.2. Additional Cloud Filer Options (3.2) .....	1
1.1.3. SMB (CIFS) Enhancements (3.1 & 3.2) .....	1
1.1.4. High Availability Improvements (3.1 & 3.2) .....	2
1.1.5. FlashCloud (3.1 & 3.2) .....	2
1.1.6. Hot File & Client Visualization (3.1 & 3.2) .....	2
1.1.7. Automated Drive Failure Recovery (3.1 & 3.2) .....	2
1.1.8. Maximum Core Filers per Cluster (3.1 & 3.2) .....	2
1.1.9. Multiple VLANs on the Core Filer Network (3.1 & 3.2) .....	2
1.1.10. Kerberos Authentication (3.1 & 3.2) .....	2
1.1.11. Search Functionality .....	3
1.1.12. FlashMove and FlashMirror Enhancements (3.1 & 3.2) .....	3
1.1.13. Secure Proactive Support (SPS) (3.1 & 3.2) .....	4
1.2. Avere NAS Optimization Solutions .....	4
1.3. Prerequisites for Administering an FXT Series Cluster .....	5
1.4. Adding and Removing a License .....	6
1.5. Using the Avere Control Panel .....	7
1.5.1. Logging In to the Avere Control Panel .....	7
1.5.2. Control Panel Tabs .....	8
1.5.3. Logging Out of the Avere Control Panel .....	9
1.6. Avere Documentation .....	9
Chapter 2. Configuring the Cluster (Settings Tab   Cluster) .....	10
2.1. General Setup .....	11
2.1.1. Modifying the Cluster Name .....	11
2.1.2. Allowing Unconfigured FXT Nodes Automatically Join the Cluster .....	11
2.1.3. Web Proxy Information .....	11
2.1.4. Setting Link-Aggregation Parameters (includes LACP) .....	11
2.2. Managing FXT Nodes .....	13
2.2.1. Adding FXT Nodes to the Cluster .....	14
2.2.2. Automatic Node Naming .....	14
2.2.3. Modifying a Node's Name/Node Detail Page .....	15
2.3. Administrative Network Settings .....	17
2.3.1. Management IP .....	19
2.3.2. Management Netmask .....	19
2.3.3. Management VLAN .....	19
2.3.4. Default MTU Value .....	20
2.3.5. Setting the Default Router .....	20
2.3.6. Static Routes .....	21
2.3.7. Setting DNS Parameters .....	22
2.3.8. Timezone .....	23
2.3.9. Setting NTP Parameters .....	23
2.3.10. Reserving FXT Series Ports for the Management Network .....	24
2.3.11. Node Management Addresses .....	24
2.4. Cluster Networks .....	25
2.4.1. Creating Cluster Networks .....	26
2.4.2. Cluster Network Settings .....	26
2.5. Managing Exports (Settings Tab   VServer) .....	28
2.5.1. Creating or Modifying Export Rules .....	28

2.5.2. Managing Export Policies .....	31
2.5.3. Unmounting Exports .....	32
2.5.4. Hierarchical Exports .....	33
2.6. High Availability .....	33
2.6.1. Specifying an HA Data Repository .....	34
2.6.2. Disabling High Availability .....	36
2.6.3. Unsetting the HA Export and Directory .....	36
2.7. Defining Schedules .....	36
2.7.1. Creating a Schedule .....	37
2.7.2. Updating an Existing Schedule .....	38
2.7.3. Deleting an Existing Schedule .....	38
2.8. Controlling Authentication .....	39
2.8.1. Selecting and Configuring a Directory Service .....	39
2.8.2. Configuring Netgroups .....	42
2.8.3. Enabling Kerberos Authentication .....	43
2.8.4. Specifying LDAP Login Services .....	45
2.8.5. Specifying the Source for Usernames .....	46
2.9. Optimizing the Avere Cluster for Use in a VMware® Environment .....	48
2.10. Configuring IPMI Cards .....	49
Chapter 3. Configuring a Core Filer (Settings Tab   Core Filer) .....	50
3.1. Creating an NFS Core Filer .....	51
3.2. Creating a Cloud Core Filer .....	53
3.2.1. Bucket Settings .....	55
3.2.2. Cloud Encryption Settings .....	57
3.2.3. Creating a New Key File After the Filer is Created .....	59
3.3. Maintaining Core Filers .....	60
3.3.1. Changing Core Filer Details .....	60
3.3.2. Managing Cloud Credentials .....	63
Chapter 4. Virtual Servers Used for Client Access .....	64
4.1. Virtual Servers and Namespaces .....	64
4.2. Managing Virtual Servers (Settings Tab   VServer) .....	65
4.2.1. Creating a Virtual Server .....	66
4.2.2. Configuring and Modifying a Virtual Server (VServer Details Page) .....	68
4.2.3. Client-Facing Network Settings .....	70
4.3. Creating and Maintaining a Global Namespace .....	72
4.3.1. Designing a Global Namespace .....	72
4.3.2. Elements of a Global Namespace .....	73
4.3.3. Warnings and Limitations for Global Namespaces .....	73
4.3.4. Creating and Modifying a Global Namespace .....	73
Chapter 5. Advanced Networking and VLANs .....	77
5.1. Enabling Advanced Networking .....	78
5.2. Creating a VLAN .....	78
5.3. Configuring Static Routes for a VLAN .....	80
5.4. Management Role VLANs .....	81
5.5. Client Role VLANs .....	82
5.6. Cluster Role VLANs .....	82
5.7. Core Filer Access Role VLANs .....	83
Chapter 6. Setting the Cache Policy .....	84
6.1. Understanding Avere's Cache Policies .....	84
6.2. Specifying the Cache Policy .....	85
6.2.1. Setting Read Mode .....	86
6.2.2. Setting Read/Write Mode .....	87

6.3. Setting Advanced Cache-Policy Features .....	90
6.3.1. Enabling Advanced Features .....	90
6.3.2. Setting a Custom Writeback Delay .....	90
6.3.3. Setting a Custom Core Filer Verification Time .....	90
6.3.4. Setting a Custom Core Filer Directory Verification Time .....	90
6.3.5. Enabling and Disabling Core Filer Verification in Read/Write Mode .....	91
6.4. Controlling Write Bandwidth .....	91
6.5. Controlling Cache Utilization on the FXT Series Cluster .....	92
6.5.1. Precautions While Using Cache-Utilization Controls .....	93
6.5.2. Enabling Local Directories .....	94
Chapter 7. Configuring CIFS Access .....	95
7.1. Recommendations for CIFS Configuration .....	95
7.2. Core Filer Prerequisites for CIFS .....	95
7.3. Overview: CIFS Configuration on the Cluster .....	96
7.4. CIFS Limitations .....	96
7.5. Joining an Active Directory Domain .....	97
7.5.1. Creating a Username Map .....	97
7.5.2. Configuring the Cluster to Use Active Directory .....	98
7.6. Enabling and Configuring CIFS .....	99
7.6.1. Gathering Information and Confirming Access .....	100
7.6.2. Enabling CIFS Access .....	101
7.6.3. Updating the CIFS Configuration .....	102
7.7. Selecting an Access-Control Mechanism .....	103
7.7.1. Considerations for Selecting Access Control .....	103
7.7.2. Enabling Constrained Delegation for Use With Native CIFS ACLs .....	104
7.8. Creating CIFS Shares .....	107
7.8.1. Accessing a CIFS Share .....	109
7.8.2. Deleting or Modifying a CIFS Share .....	109
Chapter 8. Moving and Mirroring Data on Core Filers (Data Management Tab) .....	110
8.1. Understanding FlashMove .....	110
8.2. Understanding FlashMirror .....	111
8.3. FlashMove and FlashMirror Prerequisites .....	111
8.4. Creating and Running a Data Management Job .....	112
Chapter 9. Monitoring the Cluster (Dashboard Tab) .....	117
9.1. Overview of the Dashboard .....	117
9.2. Viewing System Performance .....	118
9.2.1. Choosing the Statistics to be Displayed .....	119
9.2.2. Downloading Statistics Data .....	121
9.3. Monitoring Conditions and Alerts .....	121
9.3.1. Managing Conditions .....	122
9.3.2. Managing Alerts .....	124
9.3.3. Unhiding Conditions and Canceling Auto-Dismiss .....	125
9.3.4. Viewing Condition and Alert History .....	126
9.4. VServers Tab .....	127
9.5. Core Filers Tab .....	128
9.6. Nodes Tab .....	129
9.7. Clients Tab .....	129
9.8. Hot Files Tab .....	131
9.9. Monitoring the Cluster from Outside the Avere OS .....	132
9.9.1. Sending Alerts Via Email .....	133
9.9.2. Specifying a Syslog Server .....	134
9.9.3. Configuring SNMP .....	134

Chapter 10. Using Graphs (Analytics Tab) .....	135
10.1. Basic Graph Use .....	136
10.2. Sample Graphs from the Analytics Tab .....	137
10.3. Working with Graphs .....	139
10.3.1. Editing a Graph .....	139
10.3.2. Adding a User-Defined Graph .....	140
10.3.3. Adding a Predefined Graph .....	140
10.3.4. Modifying an Existing Graph .....	140
10.3.5. Working with Graph Sets .....	140
Chapter 11. Updating and System Maintenance (Settings   Administration) .....	142
11.1. Upgrading the Avere OS Software .....	142
11.2. Using the System Maintenance Page .....	144
11.2.1. Cluster Services (Restarting) .....	145
11.2.2. Monitoring (Statistics Collection) .....	145
11.2.3. Client Access .....	145
11.2.4. Data (Resetting Cache and Nodes) .....	146
11.2.5. Delete Data and Configuration .....	147
11.3. Adding and Modifying Users .....	148
Chapter 12. Troubleshooting and Getting Support .....	149
12.1. Possible Troubleshooting Scenarios .....	149
12.2. Configuring Support Settings .....	150
12.2.1. Support Package Update .....	151
12.2.2. Customer Info .....	151
12.2.3. Secure Proactive Support (SPS) .....	152
12.2.4. Custom Settings .....	152
12.3. Contacting Avere Global Services (Support Tab) .....	153
12.4. When to Contact Avere Global Services .....	154
12.5. Saving Cluster Configuration Data .....	155
Appendix A. Core Filer-Specific Configuration Notes .....	156
A.1. Configuring Export FSIDs for GNU Linux NFS Servers .....	156
A.2. NetApp Data ONTAP 7G and Data ONTAP 8.0 7-Mode .....	157
A.2.1. Using Native CIFS Volumes for CIFS ACLs .....	157
A.2.2. Using NFSv4 for CIFS ACLs .....	157
A.3. ZFS on OpenSolaris and Oracle Solaris .....	158
A.4. EMC Isilon OneFS .....	159
Appendix B. Password and Group File Formats .....	160
B.1. General Formatting Rules .....	160
B.2. Searching Rules .....	160
B.3. Password-Entry Parsing .....	161
B.4. Group-Entry Parsing .....	161
B.5. The Format of the Netgroup File .....	162
Index .....	163

# Chapter 1. Introduction to Avere OS

The *Daily Operations Guide* is written for system administrators who need to manage an Avere cluster. It assumes that you have a basic knowledge of networked storage, including network access protocols such as the Network File System (NFS) and the Common Internet File System (CIFS). *It also assumes that you have an Avere cluster installed and configured as described in the Avere Quick Start Guide.*

Refer to the [Avere OS 3.2 Release Notes](#) for the most recent information.

The following highlights are used in this document:



## Caution

A caution indicates that failure to follow directions in the statement can result in damage to equipment or loss of data.



## Important

An important statement presents clarifying information or specific instructions.



## Note

A note presents commentary, sidelights, or interesting points of information.

## 1.1. New in Avere OS 3.1 and 3.2

### 1.1.1. Object Compression (3.2)

You now have the option to compress objects before they are written to a cloud core filer. This option can use LZ4 or LZ4HC compression, and can be set when you create the cloud filer. You can change compression options at any time; objects written to the cloud bucket (or vault) after compression is enabled will be compressed. For more information, refer to Section 3.2, “Creating a Cloud Core Filer” on page 53 and Section 3.3.1.2, “Cloud Filer Details” on page 61.

### 1.1.2. Additional Cloud Filer Options (3.2)

You can now choose what type of cloud core filer to create, using a drop-down list

- Cleversafe
- Amplidata

Amazon continues to be supported. For more information, refer to Section 3.2, “Creating a Cloud Core Filer” on page 53.

### 1.1.3. SMB (CIFS) Enhancements (3.1 & 3.2)

- Oplocks support provides increased client performance. For more information, refer to Section 7.8, “Creating CIFS Shares” on page 107.
- The maximum number of connections for a node has been increased to one thousand (1000).
- CIFS ACLs are now supported by FlashMove and FlashMirror. For more information, refer to Section 8.4, “Creating and Running a Data Management Job” on page 112.
- Improved monitoring and statistics are available on the GUI.

#### 1.1.4. High Availability Improvements (3.1 & 3.2)

High Availability (HA) now uses up to three nodes to protect modified data, increasing redundancy and eliminating dependency on core filers for HA state within the cluster.

#### 1.1.5. FlashCloud (3.1 & 3.2)

In addition to traditional NAS-based core filers, Avere OS now allows you to integrate cloud storage by connecting to public and private object-based core filers. This capability requires you to obtain a license for each vendor; for example, one for Amazon and one for Cleversafe.

FlashCloud allows you to store cold data on cost-effective cloud storage, and efficiently moves active data to the edge near users.

FXT Series nodes integrate NAS and cloud storage into the same global namespace (GNS), providing scalable performance for users familiar with NAS protocols, the flexibility of storing data where it makes the most sense, and allowing adoption of cloud storage at a comfortable pace. In addition, FlashCloud enables cloud storage to be used with all NAS applications, without changes to existing applications or user access methods.

For more information about FlashCloud, refer to Chapter 3, *Configuring a Core Filer (Settings Tab / Core Filer)* on page 50.

#### 1.1.6. Hot File & Client Visualization (3.1 & 3.2)

Enhanced hot file and client information is now provided on the Dashboard, with sorting, and customer-configurable sample periods for clients. For more information, refer to Section 4.2.2, “Configuring and Modifying a Virtual Server (VServer Details Page)” on page 68.

#### 1.1.7. Automated Drive Failure Recovery (3.1 & 3.2)

When an FXT node data drive fails under usual failure circumstances, Avere OS will automatically reformat the node to use one less drive, and an alert will appear to inform you that the drive has been excluded, and should be physically replaced. Refer to the FXT Series Field Service Guide for more information.

#### 1.1.8. Maximum Core Filers per Cluster (3.1 & 3.2)

The maximum number of core filers supported in a cluster has been increased. Now up to fifty core filers of different types can be combined into one GNS.

#### 1.1.9. Multiple VLANs on the Core Filer Network (3.1 & 3.2)

VLAN support has increased so that each core filer can now be associated with a different VLAN, and each VLAN can have a different maximum transmission unit (MTU). For more information, refer to Section 2.3.4, “Default MTU Value” on page 20, Section 2.4, “Cluster Networks” on page 25, and Section 5.2, “Creating a VLAN” on page 78.

#### 1.1.10. Kerberos Authentication (3.1 & 3.2)

Kerberos authentication is now available between the cluster and the core filer, in addition to the previous cluster to client operations. Refer to Section 2.8.3, “Enabling Kerberos Authentication” on page 43 for more information.

### 1.1.11. Search Functionality

Several pages in the Avere Control Panel now include a **Search** field. These include the following:

- **Manage VServers**, described in Section 4.2 on page 65.
- **Manage Core Filers**, described in Section 3.1 on page 51 and Section 3.2 on page 53.
- **Data Management Tab**, described in Chapter 8 on page 110.

You can search for an item in any column of the associated table, below the **Search** field. This filters the table so that only rows that have items matching the search string are displayed.

In the following example, the first panel shows a list of cluster networks with no filter applied, the second panel shows the list filtered by “cluster”, which is applied to the **Network Name** column, and the third panel shows the list filtered by “6”, which is applied to the **Total Addresses** column. Note that if the search filter were, for example, “10.1”, there would be no changes, because “10.1” is present in every row.

The figure consists of three vertically stacked screenshots of the 'Cluster Networks' interface, demonstrating search functionality. Each screenshot shows a table with columns: Network Name, Addresses Per Node, Address Ranges, Total Addresses, and Actions. Buttons for 'Add', 'Modify', and 'Remove' are present above and below the table. A search field is located above the table in each panel.

**Panel 1: No filter applied**

Showing 1 to 2 of 2 entries

Network Name	Addresses Per Node	Address Ranges	Total Addresses	Actions
cluster	2	10.1.20.131-10.1.20.132 10.1.20.133-10.1.20.134	4	<input type="checkbox"/>
new	2	10.1.20.139-10.1.20.144	6	<input type="checkbox"/>

Showing 1 to 2 of 2 entries

**Panel 2: Filtered by 'cluster' (Network Name)**

Showing 1 to 1 of 1 entries (filtered from 2 total entries)

Search: cluster

Network Name	Addresses Per Node	Address Ranges	Total Addresses	Actions
cluster	2	10.1.20.131-10.1.20.132 10.1.20.133-10.1.20.134	4	<input type="checkbox"/>

Showing 1 to 1 of 1 entries (filtered from 2 total entries)

**Panel 3: Filtered by '6' (Total Addresses)**

Showing 1 to 1 of 1 entries (filtered from 2 total entries)

Search: 6

Network Name	Addresses Per Node	Address Ranges	Total Addresses	Actions
new	2	10.1.20.139-10.1.20.144	6	<input type="checkbox"/>

Showing 1 to 1 of 1 entries (filtered from 2 total entries)

### 1.1.12. FlashMove and FlashMirror Enhancements (3.1 & 3.2)

The FlashMove and FlashMirror features have been improved to allow a seamless transition between caching policies that do not support FlashMove and FlashMirror, and caching policies that **do** support these features. Previous releases of Avere OS required all clients to remount the filesystem after the cache policy transition occurred, which involved rebooting or restarting all clients. In this release, Avere OS automatically manages the transitions between caching policies, which removes the requirement to restart all clients during cache policy transition.

### 1.1.13. Secure Proactive Support (SPS) (3.1 & 3.2)

This enhancement provides features for remote troubleshooting with minimal to no customer involvement. For more information, refer to Section 12.2.3, “Secure Proactive Support (SPS)” on page 152.

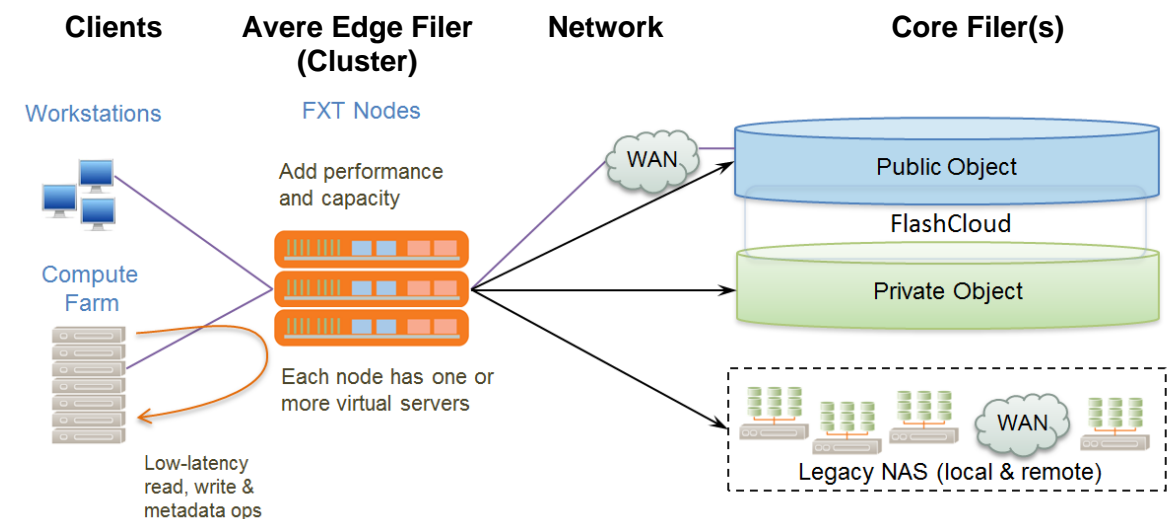
## 1.2. Avere NAS Optimization Solutions

Avere clusters and software provide a high-performance interface that accelerates client access to one or more network-attached storage (NAS) servers (*core filers*).

Avere NAS optimization consists of the Avere OS software running on one or more FXT Series *nodes* (physical servers); nodes that run together in an Avere configuration are referred to as an Avere *cluster*. The performance of the cluster scales nearly linearly with each node added to the cluster.

Avere OS caches files being actively accessed on high-performance nodes in the cluster, circumventing the overhead of sending each read and write request to the core filer’s disks. When a file becomes accessed less frequently, the cluster transfers it to the core filer and removes it from the working set, freeing up resources for files that are in higher demand.

The cluster writes updated files to the core filer at a specified rate (or faster), so that changed data is backed up regularly. In addition, FXT Series nodes are backed by NVRAM to ensure that changed data is eventually written to the core filer even in the event of a system outage. *High availability*, or HA, can be enabled to provide additional protection in the event of a system outage.



#### Avere Edge Filer

- Performance optimized
- Accelerates read, write, and metadata operations
- Transfers data to core filers (write, move, mirror)

#### Core Filer

- Capacity optimized
- FlashCloud support for public and private objects
- Supports FlashCloud object core filers and NAS core filers from different vendors in the same cluster (for example, NetApp, EMC/Isilon, Oracle, BlueArc, Nexenta)

Clients connected to a core filer through a cluster can access and modify files that are not in the working set. The cluster provides two built-in methods for handling data between clients and the core filer, with

customizations for each method. Refer to Chapter 6, *Setting the Cache Policy* on page 84 for more information.

The Avere solution can be deployed without significant disruption to your current storage or network infrastructure. Because the cluster is placed between your clients and your core filer, you can enable or disable it at will.

## 1.3. Prerequisites for Administering an FXT Series Cluster

Avere OS supports NAS access using the NFS version 3, SMB (CIFS) and SMB2.0 protocols.

The cluster is managed by using the Avere Control Panel, a secure browser-based application. The requirements and recommendations for the browser include:

- A standards-compliant browser that can access the management IP address on the cluster's network. For the following browsers, Avere supports the most recent versions of these three browsers: Mozilla Firefox®, Google® Chrome, and Apple Safari®. Microsoft Internet Explorer® versions 8 and higher are supported. The browser must have JavaScript enabled.
- If your browser supports the DNS Service Discovery (DNS-SD or Bonjour®) protocol, either through a plug-in, or in the case of Apple Safari, natively, you can use it to discover Avere clusters without needing to know machine names or IP addresses. The browser's workstation must run on the same subnet as the cluster's management interface.

This guide assumes that you have a cluster with at least one node and one virtual server (vserver) currently set up and running. Refer to the following for more information:

- The *FXT Series Installation Guide*
- The *Avere Quick Start Guide*

In addition, you will need to plan for the following:

- Setting up *exports* on your particular core filer, as described in Section 2.5, “Managing Exports (Settings Tab | VServer)” on page 28.
- Deciding on an authentication method, described in Section 2.8, “Controlling Authentication” on page 39.
- Deciding if you are going to use advanced networking, described in Chapter 5, *Advanced Networking and VLANs* on page 77.
- Planning how to set the *cache policy* (also called *write mode*) for each core filer. This will generally depend on whether you are using the core filer primarily for a high rate of data exchange, or primarily to store data long-term, as described in Chapter 6, *Setting the Cache Policy* on page 84. You will need to set the cache policy as you add core filers to the cluster.

## 1.4. Adding and Removing a License

Several Avere OS options, particularly FlashCloud, FlashMove, and FlashMirror, require a separate license. After you purchase a license, you will need to give your license ID to Avere Global Services, and they will provide you with a license key.

➤ To obtain and activate a license:

1. Navigate to the **Settings > Cluster > Licenses** page.

**Licenses**

Licensing ID: 9b01ce73-846b-11e3-8fef-000c29153b30

Licensed Features: FlashMove FlashMirror LocalDirectories  
FlashCloudAmazonS3

Licensed Cluster Size: 8

Compliance State: OK

Id	License Key	Description	Expiration	Actions
1	UpUVKQwAgAD2PYFJp6Dtg==	FlashMove FlashMirror FlashCloudAmazonS3, 8 nodes		<button>Remove</button>

**Add License**

License Key

Submit

2. Contact Avere Global Services, and provide the **Licensing ID** listed at the top of the **Licenses** page.
3. Enter the license key provided by Avere Global Services in the **License Key** field.
4. Click the **Submit** button.

The new license and its associated features appear in the **Licenses** area.

➤ To remove a license (for example, if you need to move the rights to another cluster):

1. From the **Settings > Cluster > Licenses** page, select the **Remove** button in the row of the license you want to remove.
2. A pop-up window appears asking if you are sure you want to remove the license. Click **OK** to continue.

The license is removed from the list, and you will no longer be able to create FlashMove or FlashMirror jobs. Previous jobs can still be viewed from the **Data Management** tab.

## 1.5. Using the Avere Control Panel

The cluster is managed by using the Avere Control Panel, a secure browser-based application. The Avere Control Panel presents a single system image of the cluster and all of its constituent nodes; you can manage the entire cluster or any individual node from the Avere Control Panel.

### 1.5.1. Logging In to the Avere Control Panel

➤ To log in to the cluster for management and monitoring tasks:

1. Use a Web browser to go to `https://management_IP_address/fxt`, where *management\_IP\_address* is the cluster's management IP address.

Alternatively, if you are using a browser that supports the DNS-SD protocol, you can navigate to the bookmark labeled Avere Cluster Management: *cluster\_name*.



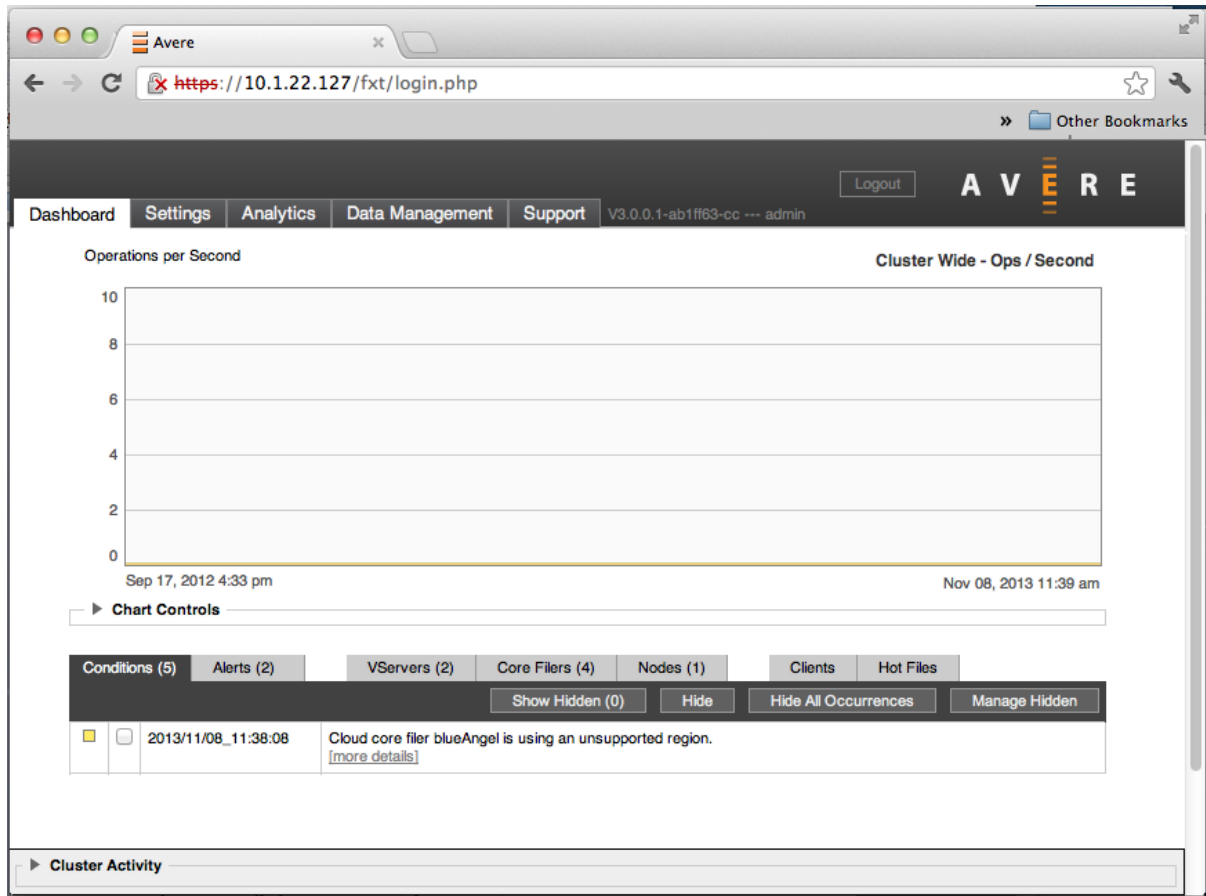
#### Note

If your browser first prompts you to accept an SSL certificate from the Avere system, accept the certificate.

2. In the **User** field, enter the cluster's administrative user name.
3. In the **Password** field, enter the administrative password.

The Avere Control Panel's **Dashboard** tab is displayed.

## 1.5.2. Control Panel Tabs



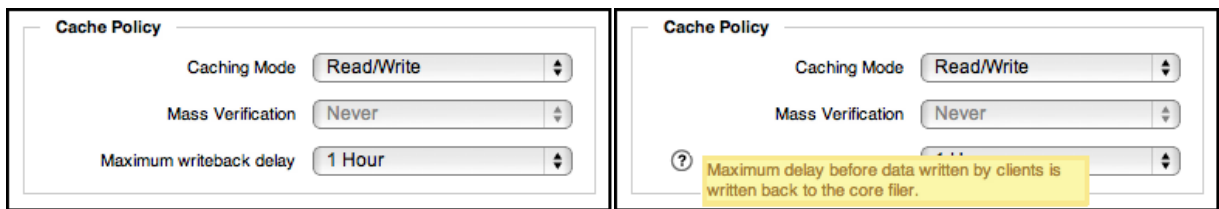
The Avere Control Panel is divided into the following tabbed pages, with the Avere OS release currently running on the cluster listed to the right of the tabs.

- **Dashboard** – This is the Avere Control Panel’s default tab. It displays performance graphs and statistics, system alerts, and information about virtual servers (vservers), FXT Series nodes, clients, and hot files. Refer to Chapter 9, *Monitoring the Cluster (Dashboard Tab)* on page 117 for more information.
- **Settings** – This tab enables you to configure and maintain the cluster.
  - Refer to Chapter 4, *Virtual Servers Used for Client Access* on page 64 for more information about managing vservers, Section 4.3, “Creating and Maintaining a Global Namespace” on page 72 for creating namespaces, Section 2.5, “Managing Exports (Settings Tab | VServer)” on page 28 and Appendix A, *Core Filer-Specific Configuration Notes* on page 156 for managing exports, and Chapter 7, *Configuring CIFS Access* on page 95 for information on CIFS.
  - Refer to Chapter 6, *Setting the Cache Policy* on page 84 and Appendix A, *Core Filer-Specific Configuration Notes* on page 156 for more information about managing core filers.
  - Refer to Chapter 2, *Configuring the Cluster (Settings Tab | Cluster)* on page 10, Section 2.2, “Managing FXT Nodes” on page 13, and Chapter 11, *Updating and System Maintenance (Settings / Administration)* on page 142 for more information about managing the cluster from the **Settings** tab.
- **Analytics** – This tab enables you to obtain detailed, in-depth information about activity on your cluster and storage network. Refer to Chapter 10, *Using Graphs (Analytics Tab)* on page 135 for more information.

- **Data Management** – This tab enables you to move and mirror data on NAS servers registered as core filers with the cluster. Data management operations can run concurrently with client access to data, and do not interrupt client access. Refer to Chapter 8, *Moving and Mirroring Data on Core Filers (Data Management Tab)* on page 110 for more information.
- **Support** – This tab enables you to work collaboratively with Avere Global Services if problems occur. Refer to Section 12.3, “Contacting Avere Global Services (Support Tab)” on page 153 for more information.

If a condition occurs that affect the operation of the cluster, the Avere Control Panel displays a system error (red) or alert (yellow) at its upper left-hand corner. Clicking on the notification takes you to the dashboard, where you can view problem details and troubleshooting information.

You can view online help text for control elements on the screen by moving your cursor to the left of the element’s label. A question mark appears; holding your cursor over the question mark displays help text for that control element.



### 1.5.3. Logging Out of the Avere Control Panel

➤ To log out of the Avere Control Panel:

1. Click the **Logout** button near the upper right-hand corner of the browser window. Your administrative session ends and the browser displays the Avere Control Panel’s Login page.
2. To ensure a complete logout, shut down the browser.

## 1.6. Avere Documentation

In addition to this *Daily Operations Guide*, the following documentation is available from the Support web site.

- *FXT Series Installation Guide* – How to install and maintain FXT Series hardware
- *Avere Quick Start Guide* – How to set up and initially configure an Avere production system
- *Release Notes* – Late-breaking information about the Avere product
- *Third-Party Licenses Reference* – Provides licensing information for the third-party software used by the Avere product

## Chapter 2. Configuring the Cluster (Settings Tab | Cluster)

You can display and adjust cluster-wide settings from the pages listed below the **Cluster** heading on the **Settings** tab. Many of the parameters listed on the Cluster pages are defined during initial cluster configuration, but you can make updates at any time.

For initial setup, refer to the *Avere Quick Start Guide*. You can then configure the cluster for your particular situation.

Dashboard Settings Analytics Data Management Support V3.0.0.1-1dea01d --- admin

VServer + General Cluster Setup

Core Filer +

Cluster -

General Setup

Administrative Network

Cluster Network

FXT Nodes

High Availability

Monitoring

Schedules

Directory Services

Kerberos

Login Services

Active Directory

Optimization

IPMI

Support

Licenses

Cloud Credentials

VLAN

Administration +

Only appears if advanced networking is enabled.

General

Cluster name LiGo\_Cluster

Allow unconfigured FXT nodes to join this cluster ☒ Only available if this setting was enabled in a previous Avere OS version.

Enable advanced networking ☒

Internet VLAN Default (tag: -, gateway: 10.1.0.76 )

Web Proxy Information (Optional)

URL http://ourproxy.company.com

User name administrator

Password .....

Revert Change cluster parameters

Network Options

Enable link aggregates for all interfaces in the cluster ☒

Enable dynamic LACP ☐

Revert Submit

Node Name Settings

? Node Prefix

First Node Number

Change Node Names

Cluster settings not covered in this chapter can be found in the following sections:

- Chapter 5, *Advanced Networking and VLANs* on page 77
- Section 7.5, “Joining an Active Directory Domain” on page 97
- Chapter 8, *Moving and Mirroring Data on Core Filers (Data Management Tab)* on page 110
- Section 9.9, “Monitoring the Cluster from Outside the Avere OS” on page 132

## 2.1. General Setup

The **General Cluster Setup** page allows you to change the cluster name, determine how nodes are added to the cluster, and enable networking options.

### 2.1.1. Modifying the Cluster Name

The cluster name can include letters, numerals, and the dash (-) and underscore (\_) characters. It cannot include other punctuation or special characters. The cluster name can be from one to 16 characters in length.

The base cluster name is used to create default names for new nodes that join the cluster. The base cluster name does not affect the names of nodes that previously joined the cluster or that were manually renamed.

➤ To change the name of a cluster:

- Enter the new name in the **Cluster name** field and click the **Change cluster parameters** button.



#### Note

If you have configured CIFS access and are using the cluster (NetBIOS) name as the name of the CIFS server, changing the cluster name does not change the CIFS server name. The Avere CIFS service continues to use the server name listed on the CIFS Configuration page.

➤ To restore the original name of the cluster before clicking the **Change cluster parameters** button:

- Click the **Revert** button.

The **Revert** button has no effect if the new name has been committed by clicking the **Change cluster parameters** button.

### 2.1.2. Allowing Unconfigured FXT Nodes Automatically Join the Cluster

If this option has previously been enabled, you can deselect the **Allow unconfigured FXT nodes to join this cluster** checkbox and click the **Change cluster parameters** button.



#### Note

This option is being removed, and the field will *only* appear if it has previously been enabled. If you disable this option, you will not be able to re-enable it.

### 2.1.3. Web Proxy Information

If your system uses a web proxy server, enter its information here to allow you to update the Avere Systems software and send statistics to Avere Global Services, if needed. HTTP, sock5, and sock4 protocols are supported.

### 2.1.4. Setting Link-Aggregation Parameters (includes LACP)

If you want to use link aggregation, either static link aggregates or dynamic Link Aggregation Control Protocol (LACP) aggregates, you will need to enable link aggregation.

In most cases, enabling link aggregation on the cluster requires a configuration change on your network switch. Refer to the documentation for your switch, or consult with your Avere Systems representative for more information.

➤ To change the cluster's link-aggregation parameters:

1. For either static or dynamic link aggregation, select the **Enable link aggregates for all interfaces in the cluster** checkbox.

In a static link aggregate, outgoing traffic is load-balanced across the active ports, and incoming traffic from all active ports is also accepted.

2. If your network uses Link Aggregation Control Protocol (LACP) as specified by IEEE 802.1AX-2008 (formerly 802.3ad), select the **Enable dynamic LACP** checkbox. Do not select this checkbox if your network uses static aggregation.
3. Choose **Revert** to restore the original networking parameters, or **Submit** to commit any changes.

A FXT cluster typically includes two 10GB ports, four 1GB ports, and two 1GB management ports (e0a and e0b). How the ports are aggregated depends on the **Administrative Network** settings, described in Section 2.3.10, “Reserving FXT Series Ports for the Management Network” on page 24.

#### 2.1.4.1. Using a Separate Management Network

If **Use separate management network** is enabled on the **Administrative Network** settings (Section 2.3.10, “Reserving FXT Series Ports for the Management Network” on page 24):

- A single 2-port link aggregate will be created, including both 10GB ports (2 x 10GB), and a separate 2-port link aggregate (2 x 1GB) reserved for the management network. This yields:
  - 2 x 10GB Client / Cluster Network link aggregate
  - 2 x 1GB Separate Management Network link aggregate
- If both ports in the 10GB aggregate fail, Avere OS transitions to a 4-port link aggregate using the four remaining 1GB interfaces (4 x 1GB). This yields:
  - 4 x 1GB Client / Cluster Network link aggregate
  - 2 x 1GB Separate Management Network link aggregate



#### Note

When **Use separate management network** is enabled, then if the optional cluster MTU and Netmask have been specified, they continue to be used, even after the checkbox is deselected.

Make sure you remove the cluster MTU and Netmask settings *before* deselecting **Use separate management network**.

When joining a node to a cluster with a physically separate management network, connect only the cluster ports on the joining node. Once the node has joined, the management network ports may be connected.

#### 2.1.4.2. Management Network is Not Separate

If **Use separate management network** is *not* enabled on the **Administrative Network** settings (Section 2.3.10, “Reserving FXT Series Ports for the Management Network” on page 24):

- A single 2-port link aggregate will be created, including both 10GB ports (2 x 10GB). This yields:
  - 2 x 10GB Client / Cluster / Management Network link aggregate
- If both ports in the 10GB aggregate fail, Avere OS transitions to a 6-port link aggregate using the six remaining 1GB interfaces (6 x 1GB, including the two that would have been used for Separate Management Network). This yields:
  - 6 x 1GB Client / Cluster / Management Network link aggregate

## 2.2. Managing FXT Nodes

The **FXT Nodes** page displays information about the selected node, and a list of unconfigured FXT nodes that are visible to the cluster on the network.

Refer to Section 2.2.1, “Adding FXT Nodes to the Cluster” on page 14 for information on adding FXT nodes to the cluster.

You can get general information about the nodes in your cluster from the **Nodes** tab on the **Dashboard**, as described in Section 9.6, “Nodes Tab” on page 129.

You can use the **FXT Nodes** page to add a node to the cluster, to restart or shut down a node, to reformat the node if necessary, or to remove it from the cluster.

Dashboard

Settings

Analytics

Data Management

Support

V3.0.0.1-1dea01d --- admin

VServer

Core Filer

Cluster

General Setup

Administrative Network

Cluster Network

FXT Nodes

High Availability

Monitoring

Schedules

Directory Services

Kerberos

Login Services

Active Directory

Optimization

IPMI

Support

Licenses

Cloud Credentials

VLAN

Administration

Cluster FXT Nodes

Name	Model	Status	Actions
LiGo	N/A	up	<div>Modify settings</div> <div>Restart services</div> <div>Reboot</div> <div>Power down</div> <div>Remove</div> <div>Offline</div> <div>Suspend</div> <div>Reformat</div>

Unconfigured FXT Nodes

Name	Address	Software Version	Status	Actions
huri03	10.1.4.112	V3.1.1.1-27f73b6	Wants to join	<div>Allow to join</div> <div>Upgrade software version</div>
cloudtest00	10.1.4.112	V3.1.1.1-27f73b6	Wants to join	<div>Allow to join</div> <div>Upgrade software version</div>
barma6	10.1.4.112	V3.1.1.1-27f73b6	Wants to join	<div>Allow to join</div> <div>Upgrade software version</div>
debug2	10.1.4.112	V3.1.1.1-27f73b6	Wants to join	<div>Allow to join</div> <div>Upgrade software version</div>
masri2	10.1.4.112	V3.1.1.1-27f73b6	Wants to join	<div>Allow to join</div> <div>Upgrade software version</div>
neem3	10.1.4.112	V3.1.1.1-27f73b6	Wants to join	<div>Allow to join</div> <div>Upgrade software version</div>



### Caution

Any node operations that affect the node’s connectivity to an active cluster, including removal, restart, and shutdown, can cause temporary service disruptions to clients.

## 2.2.1. Adding FXT Nodes to the Cluster

Node names are automatically assigned as each node joins a cluster. You can customize the node naming scheme as described in Section 2.2.2, “Automatic Node Naming” on page 14.

➤ To add specific nodes to the cluster:

1. Navigate to the **Settings > Cluster > FXT Nodes** page.
2. Locate the list of unconfigured nodes on the network in the **Unconfigured FXT Nodes** table.
3. For each node that you want to add to the cluster, click the **Allow to join** button.

When a new node is added to the cluster, all client and core filer interfaces are automatically rebalanced across all nodes, including the new node. This can result in a brief suspension of file services. No data loss or corruption results from the rebalancing operation.

You can use the **Update Software Version** button to update the software on an unjoined node, allowing you to keep unjoined nodes in reserve with the same version of the software running on the rest of the cluster.

## 2.2.2. Automatic Node Naming

Node names are automatically assigned as each node joins a cluster. The names are generally of the form `nodepartial_UUID`.

➤ To customize the node naming scheme:

1. In the **Node Prefix** field, enter a prefix for all nodes. For example, if you want all FXT nodes to be named `tiered_number`, enter **tiered\_** in the field.

You can specify trailing pound signs (#) to indicate placeholder digits. Thus, if you enter **tiered\_###** in the field, nodes would be named `tiered_001` or `tiered_020`, instead of `tiered_1` or `tiered_20`.

2. In the **First Node Number** field, optionally enter the number for the first node in the cluster. Additional nodes in the cluster are numbered sequentially from the specified number. For example, if you specify **10** as the first node number with the prefix **tiered\_###**, nodes are named `tiered_010`, `tiered_011`, and so forth.
3. Click the **Change Node Names** button. The specified naming scheme is applied to nodes that are already in the cluster and to any nodes that are subsequently added to the cluster.

You can change an individual node name from its **Node Detail** page, as described in Section 2.2.3, “Modifying a Node’s Name/Node Detail Page” on page 15).

- To restart all Avere OS services on a node, click the **Restart services** button.
- To restart a node, click the **Reboot** button.
- To shut down a node, click the **Power down** button.
- To modify a node’s settings, including its name and its optional IPMI configuration, click the **Modify settings** button to go to the **Node node\_name - Settings** page. See Section 2.10, “Configuring IPMI Cards” on page 49 for details on configuring a node’s IPMI settings.
- To remove a node from the cluster, click the **Remove** button.



### Note

Before removing a node from the cluster, ensure that the **Allow unconfigured FXT nodes to join this cluster** checkbox is disabled, as described in Section 2.1.2, “Allowing Unconfigured FXT Nodes Automatically Join the Cluster” on page 11. If you remove a

node from a cluster that unconfigured nodes automatically join, the removed node can immediately rejoin the cluster from which it was just removed.

### 2.2.3. Modifying a Node's Name/Node Detail Page

You can change a node name in one of two places:

- From the **Node Detail** page, accessed when you click on the name of a node in the head of the node's table of basic information
- From the **Node Settings** page, accessed when you click **Modify settings** from the **FXT Nodes** page.

If you click on the name of a node in the head of the node's table of basic information, the **Node Detail** page opens and displays the following information:

- A node-name drop-down list – Choose the name of the node in the cluster for which you want to view details.

Dashboard

Settings

Analytics

Data Management

Support

V3.0.0.1-1dea01d --- admin

VServer

Core Filer

Cluster

General Setup

Administrative Network

Cluster Network

FXT Nodes

High Availability

Monitoring

Schedules

Directory Services

Kerberos

Login Services

Active Directory

Optimization

IPMI

Support

Licenses

Cloud Credentials

VLAN

Administration

Node Detail

[Back to Node List](#)

Choose node

line3

Node performance

Instantaneous: 0 ops/sec; 100% cache hit rate; 0 ms latency.  
1 min average: 0 ops/sec; 100% cache hit rate; 0 ms latency.

Cluster performance

Instantaneous: 0 ops/sec; 100% cache hit rate; 0 ms latency.  
1 min average: 0 ops/sec; 100% cache hit rate; 0 ms latency.

Node Info

Node Name

line3

Reset

Submit

Primary Cluster IP:

10.1.22.195

Cluster IPs:

10.1.22.195  
10.1.22.196

Client Facing IPs:

10.1.22.199  
10.1.22.200  
10.1.22.201  
10.1.22.202

Active Image:

AvereOS\_V3.0.0.1-1dea01d

Alternate Image:

AvereOS\_V3.0.0.1-1dea01d

Hardware Summary:

Component	Details
CPU	2 (2 x Pentium(R) Pro)
Storage	4 GB data (4 x 1 GB), 21 GB system
Storage Type	Data is stored on VM disks
Memory	3 GB (1 x 1024 MB, 1 x 2048 MB)
Port e0a	00:0c:29:15:3b:30 UP, active, 1000baseT full-duplex
Port e0b	00:0c:29:15:3b:3a UP, active, 1000baseT full-duplex
Port e0c	00:0c:29:15:3b:44 UP, active, 1000baseT full-duplex
Port e0d	00:0c:29:15:3b:4e UP, active, 1000baseT full-duplex
Serial Number	VMware-56 4d de cc 67 b9 e8 0a-6d 5c 52 75 86 15 3b 30
System	sim

Refresh

Core Filer Connections:

thor.com:10.1.11.50:2049 (18 conns)  
grape:10.0.8.82:2049 (4 conns)

- **Node performance** – Performance at the current time and averages over the previous minute, including:
  - Operations per second
  - Cache hit rate percentage
  - Latency in milliseconds
- **Cluster performance** – Performance at the current time and averages over the previous minute, including:
  - Operations per second
  - Cache hit rate percentage
  - Latency in milliseconds
- **Node Name** – The editable name of the node.

➤ To change the node name:

1. In the **Node Name** field, enter the new name.
2. Choose **Reset** to restore the original IP address settings, or choose **Submit** to commit any changes.



### Note

Node names can contain alphanumeric characters, underscores (\_), and dashes (-); they cannot contain spaces or other special characters.

- **Cluster IP addresses**, including the following:
  - The primary cluster address for the node
  - List of all cluster addresses for the node
  - List of client-facing addresses assigned to the node
- **Image information** – The software available for running, including the following:
  - Active (currently running) software image
  - Alternate software image
- **Hardware Summary** – Information about the hardware. This includes the following:
  - A list of hardware components, including the component name and details for each component.
  - Data from environmental sensor ports, including:
    - Sensor name
    - Current reading
    - Sensor status
  - The serial number of the VMWare connection.
  - The FXT Node type, or whether the system is a simulation.
- **Core Filer Connections** – The core filer name or IP address to which the node is connected, and the number of connections to it.

## 2.3. Administrative Network Settings

The **Administrative Network** page, located under the **Cluster** heading on the **Settings** tab, allows you to configure administrative settings for the cluster.



### Note

The fields on this page will vary, depending on whether advanced networking is enabled or not, as described in Section 5.1, “Enabling Advanced Networking” on page 78. However, as of Avere OS 3.2, advanced networking is enabled by default for new clusters.

The screenshot shows the 'Administrative Network' settings page. The left sidebar contains a navigation menu with the following items: VServer, Core Filer, Cluster, General Setup, Administrative Network (selected), Cluster Network, FXT Nodes, High Availability, Monitoring, Schedules, Directory Services, Kerberos, Login Services, Active Directory, Optimization, IPMI, Support, Licenses, Cloud Credentials, VLAN, and Administration. The main content area is titled 'Administrative Network' and contains the following fields:

- Management IP: 10.1.2.199
- Management netmask: 255.255.224.0
- Default MTU (optional):
- Default Router: 10.0.0.0
- Static Routes: 1.1.1.0 255.255.255.0 10.0.8.1
- DNS server(s): 10.0.0.1
- DNS domain: domain.com
- DNS search:
- Timezone: America/New York
- NTP server(s): ntp.company.com
- Use multicast NTP servers: ☐
- Use separate management network: ☐

A red box highlights the 'Default Router' and 'Static Routes' fields, with a callout stating: 'Only appears if advanced networking is not enabled.'

Below the main settings is the 'Node Management Addresses (optional)' section, which contains the following fields:

- Node admin first IP: 10.12.140.1
- Node admin last IP: 10.12.140.3
- Number of IPs in range: 3

At the bottom right of each section are 'Revert' and 'Submit' buttons.

### Administrative Network Settings – With Advanced Networking *Not* Enabled

Dashboard

Settings

Analytics

Data Management

Support

V3.1.1.1-6a7c480 --- admin

LiGo Cluster

VServer

Core Filer

Cluster

General Setup

Administrative Network

Cluster Network

FXT Nodes

High Availability

Monitoring

Schedules

Directory Services

Kerberos

Login Services

Active Directory

Optimization

IPMI

Support

Licenses

Cloud Credentials

VLAN

Administration

Management IP

10.1.18.149

Management netmask

255.255.224.0

Management VLAN

Default (tag: -, gateway: 10.1.0.76 )

Default MTU (optional)

DNS server(s)

10.0.8.4

DNS domain

dns.company.com

DNS search

Timezone

America/New York

NTP server(s)

ntp.company.com

Use multicast NTP servers

Use separate management network

Revert

Submit

These fields only appear when advanced networking is enabled.

Node Management Addresses (optional)

Address Range	Subnet Mask	VLAN	Actions
Add New Range			

Add New Node Management Address Range

Node admin first IP

10.12.140.1

Node admin last IP

10.12.140.3

Number of IPs in range

3

Node admin subnet mask

255.255.224.0

Node admin VLAN

Default (tag: -, gateway: 10.1.0.76 )

Submit

**Administrative Network Settings – With Advanced Networking *Enabled***

### 2.3.1. Management IP

- To change the management interface:
  1. Navigate to the **Settings > Cluster > Administrative Network** page.
  2. Enter, verify, or change the IP address in the **Management IP** field. The address must be in IPv4 dotted notation.
  3. Choose **Revert** to restore the original management interface settings, or choose **Submit** to commit any changes.
  4. If necessary, refresh the browser window or enter the management interface's new URL in the browser's navigation bar, then log back into the Avere Control Panel.

### 2.3.2. Management Netmask

You can specify a netmask for the cluster's management network. The netmask is applied to all cluster interfaces, including client-facing addresses, cluster addresses, and management addresses.

- To set the cluster's management netmask:
  1. Navigate to the **Settings > Cluster > Administrative Network** page.
  2. Enter, verify, or change the netmask in the **Management Netmask** field. The netmask can be specified in either dotted IPv4 notation (for example, **255.255.255.0**) or in the format **/number\_of\_bits** (for example, **/16**).
  3. Choose **Revert** to restore the netmask settings, or choose **Submit** to commit any changes.

### 2.3.3. Management VLAN



#### Note

This field is *only* available if advanced networking is enabled.

There is at a minimum one VLAN in a cluster when advanced networking is enabled. If you have created additional VLANs, as described in Section 5.2, "Creating a VLAN" on page 78, you can select one of those, rather than the default VLAN, as a management VLAN.

- To set the cluster's management VLAN:
  1. Navigate to the **Settings > Cluster > Administrative Network** page.
  2. Select the VLAN you want to use from the **Management VLAN** drop-down list.
  3. Choose **Submit** to commit any changes.

### 2.3.4. Default MTU Value

You can optionally specify a default maximum transmission unit (MTU) setting for the cluster. The MTU setting is used on all cluster interfaces, including client-facing addresses, cluster addresses, and management addresses. However, the MTU can be changed for each individual VLAN, as described in Section 5.2, “Creating a VLAN” on page 78.

Typical values for the MTU setting are 1500 (the default) and 9000. Custom values are also permitted. If you specify a value greater than 1500, jumbo frames (with more than 1500 bytes of payload) are automatically enabled for the cluster.



#### Important

Before setting or changing the cluster’s MTU value, verify the proposed value with your network administrator, particularly if you are specifying a value that enables jumbo frames.

➤ To set the cluster’s default MTU value:

1. Navigate to the **Settings > Cluster > Administrative Network** page.
2. Enter, verify, or change the MTU value in the **Default MTU (optional)** field.
3. Choose **Revert** to restore the original management interface settings, or choose **Submit** to commit any changes.

### 2.3.5. Setting the Default Router

The default router will be used by all of the cluster’s interfaces.



#### Note

This field is *only* available if advanced networking is *not* enabled.

➤ To specify the default router:

1. Navigate to the **Settings > Cluster > Administrative Network** page.
2. Enter, verify, or change the default router in the **Default Router** field. The router must be specified in dotted IPv4 notation.
3. Choose **Reset** to restore the default router settings, or choose **Submit** to commit any changes.

## 2.3.6. Static Routes

If your network infrastructure uses static routes, you can configure the cluster to use them.



### Note

This field is *only* available if advanced networking is *not* enabled.

➤ To configure the static routes:

1. Navigate to the **Settings > Cluster > Administrative Network** page.

The screenshot shows the 'Static Routes' configuration page. On the left is a sidebar with navigation links: VServer, Core File, Cluster (selected), General Setup, Administrative Network, Cluster Networks, FXT Nodes, High Availability, Monitoring, Schedules, Directory Services, Kerberos, Login Services, Active Directory, Optimization, IPMI, Support, Licenses, Cloud Credentials, VLAN, and Administration. The main content area has a 'Static Routes' section with a table. The table has columns: Router, VLANs, Static Routes, and Actions. It contains two rows: one for Router 10.1.0.76 with VLAN Default (tag: -) and Static Routes destIP:102.1.0.0,netmask:255.255.255.0,gateway:192.168.1.1; and another for Router 110.1.0.2 with VLAN my\_VLAN(tag:42) and an empty Static Routes field. Below the table is a 'Modify Static Routes For 10.1.0.2' section with two route configuration blocks. 'Route 1' has Destination IP 102.1.0.1, Netmask 255.255.255.0, and Gateway 192.168.1.1. 'Route 2' has Destination IP 102.1.0.2, Netmask 255.255.255.0, and Gateway 192.168.1.1. At the bottom are links for 'Go back to VLAN configuration page', 'Add Another Route', and an 'Update Routes' button.

Router	VLANs	Static Routes	Actions
10.1.0.76	Default (tag: -)	destIP:102.1.0.0,netmask:255.255.255.0,gateway:192.168.1.1	Modify Delete
110.1.0.2	my_VLAN(tag:42)		Add

**Modify Static Routes For 10.1.0.2**

**Route 1**

Destination IP: 102.1.0.1  
Netmask: 255.255.255.0  
Gateway: 192.168.1.1  
[Remove This Route](#)

**Route 2**

Destination IP: 102.1.0.2  
Netmask: 255.255.255.0  
Gateway: 192.168.1.1  
[Remove This Route](#)

[Go back to VLAN configuration page](#) [Add Another Route](#) [Update Routes](#)

2. In the **Static Routes** field, enter the following three values as a space-separated list. All three values must be specified in dotted IPv4 notation.

- Destination IP address
- Netmask (can be specified either in dotted IPv4 notation or in */number\_of\_bits format*)
- Gateway

The following example specifies a static route with the destination IP address 1.1.1.0, the netmask 255.255.255.0, and the gateway 10.0.8.1:

**1.1.1.0 255.255.255.0 10.0.8.1**

3. Enter additional static routes as a comma-separated list of space-separated tuples. The following example adds an additional static route to the initially entered route:

**1.1.1.0 255.255.255.0 10.0.8.1, 2.2.2.0 255.255.255.0 10.0.8.1**

4. Choose **Reset** to restore the original static routes, or choose **Submit** to commit any changes.



The following **nsupdate** commands provide an example of configuring DNS correctly:

```
update add avere.example.com. 86400 A 10.0.0.10
update add avere.example.com. 86400 A 10.0.0.11
update add avere.example.com. 86400 A 10.0.0.12

update add client10.example.com. 86400 A 10.0.0.10
update add client11.example.com. 86400 A 10.0.0.11
update add client12.example.com. 86400 A 10.0.0.12

update add 10.0.0.10.in-addr.arpa. 86400 PTR client10.example.com.
update add 11.0.0.10.in-addr.arpa. 86400 PTR client11.example.com.
update add 12.0.0.10.in-addr.arpa. 86400 PTR client12.example.com.
```

### 2.3.8. Timezone

➤ To set the cluster's timezone:

1. Navigate to the **Settings > Cluster > Administrative Network** page.
2. Select the desired timezone from the **Timezone** drop-down list.
3. Choose **Reset** to restore the original timezone, or choose **Submit** to commit any changes.

### 2.3.9. Setting NTP Parameters



#### Important

- It is strongly recommended that you use NTP in your cluster. Not using NTP can result in inconsistent timestamps, causing files to appear as if they have gone forward or backward in time.
- It is recommended that you use either one NTP server or three or more NTP servers. Using two NTP servers can result in time synchronization failures, inconsistent timestamps, and out-of-order file operations.

➤ To change the cluster's NTP settings:

1. Navigate to the **Settings > Cluster > Administrative Network** page.
2. Enter one or more NTP servers (DNS names or IP addresses) in the **NTP server(s)** field. Enter multiple servers as a comma-separated list.
3. If your site uses multicast/broadcast NTP, check **Use multicast servers**.
4. Choose **Reset** to restore the original NTP settings, or choose **Submit** to commit any changes.



#### Important

CIFS access requires that NTP be enabled and use the same time source on all CIFS-related system components, including the cluster, the core filer, the Active Directory/Kerberos server, and CIFS clients, be within five (5) minutes of one another. If the clock skew between any two components is more than five minutes, CIFS access is denied.

### 2.3.10. Reserving FXT Series Ports for the Management Network

Avere network interfaces typically migrate to any available port if there is a port failure or other network interruption. If the available ports include the 1-GbE e0a and e0b ports on an FXT Series node, a failed interface (including a 10-GbE interface used for client or cluster traffic) can migrate to either of these ports.

- To prevent this situation and ensure that the e0a and e0b ports are always used for management:
  1. Navigate to the **Settings > Cluster > Administrative Network** page.
  2. Check **Use separate management network**.
  3. Choose **Reset** to restore the original NTP settings, or choose **Submit** to commit any changes.

### 2.3.11. Node Management Addresses

Optionally, you can specify a range of IP addresses that are reserved for use by management interfaces, and are distributed across each node in the cluster so that each node is guaranteed to have at least one management IP.

If you specify a range of management IP addresses, the cluster retains them until they are needed. For example, if you specify a range of twelve IP addresses in a three-node cluster, the cluster allocates one management address to each node in the cluster and retains the remaining nine addresses in case more nodes are added to the cluster. If nodes are added, the cluster assigns each new node a reserved IP address until the range is exhausted.

If you specify a reserved range of management IP addresses, ensure that the number of nodes in your cluster does not exceed the number of reserved IP addresses. You can add more IP addresses to the range if your cluster grows beyond the range number.



#### Important

It is strongly recommended that the gateway of the management VLAN should match the gateway of the (optional) node administration VLAN. Otherwise, the management interface may exhibit strange behavior.

- To specify a reserved range of management IP addresses:
  1. Navigate to the **Settings > Cluster > Administrative Network** page.
  2. Select **Add New Range**, or **Modify** for an existing address range. You can also choose **Remove** to remove an existing range, freeing the IP addresses for other uses.
  3. In the **Add New Node Management Address Range** or **Modify Node Management Address Range** panel, enter the first IP address in the range of reserved management addresses into the **Node admin first IP** field.
  4. Enter the last IP address in the range of reserved management addresses in the **Node admin last IP** field.

At this point, the number of IP addresses being allocated will appear in the **Number of IPs in range** field. If the number is 128 or more, you will be asked if you really want that many addresses allocated. If the number is what you want, you simply need to submit the changes; otherwise, you will have an opportunity to correct the entry.
  5. (Advanced networks only) Enter the subnet netmask of the range of addresses in the **Node admin subnet mask** field.

- (Advanced networks only) From the **Node admin VLAN** drop-down list, select the VLAN that the node administration IP addresses are to use, preferably with the same gateway as the (optional) management VLAN. Possible values include **Default** and any management-role VLANs that you created in Section 5.2, “Creating a VLAN” on page 78.
- Choose **Submit** to commit any changes.
- The Avere OS displays a pop-up that warns you about client disruption and asks you for confirmation to proceed. Click **OK**.

If you specify a reserved range of management IP addresses, the IP address assigned to each node is displayed on the Dashboard. These IP addresses are listed under a column named **Node Mgmt IP** on the Dashboard’s **Nodes** tab. The Avere Control Panel displays this column only if there is a reserved range of management IP addresses.

## 2.4. Cluster Networks

The operation of the cluster depends on network links among the cluster’s constituent nodes. Cluster IP addresses must be manually assigned (that is, not assigned by DHCP) and must be in a contiguous range. The **Cluster Networks** page allows you to assign these addresses.

Each cluster network includes one or more network address range. After you create a network and assign the number of IPs per node, you can create one or more address ranges from the **Address Ranges** panel to assign to the network.



### Note

The fields on this page will vary, depending on whether advanced networking is enabled or not, as described in Section 5.1, “Enabling Advanced Networking” on page 78. However, as of Avere OS 3.2, advanced networking is enabled by default for new clusters.

Dashboard
Settings
Analytics
Data Management
Support
V3.1.1-1-6a7c480 --- admin  
LiGo\_Cluster

VServer
Core Filer
Cluster
General Setup
Administrative Network
Cluster Networks
FXT Nodes
High Availability
Monitoring
Schedules
Directory Services
Kerberos
Login Services
Active Directory
Optimization
IPMI
Support
Licenses
Cloud Credentials
VLAN
Administration

### Cluster Networks

Add Modify Remove

Showing 1 to 2 of 2 entries Search:

Network Name	Addresses Per Node	Address Ranges	Total Addresses	Actions
cluster	2	10.1.20.131-10.1.20.134	4	<input type="checkbox"/>
new	2	10.1.20.139-10.1.20.144	6	<input checked="" type="checkbox"/>

Add Modify Remove

### Address Ranges

Add Modify Remove

Showing 1 to 2 of 2 entries Search:

Network Name	Address Range	Subnet Mask	VLAN	Total Addresses	Actions
cluster	10.1.20.131-10.1.20.134	255.255.224.0		4	<input type="checkbox"/>
new	10.1.20.139-10.1.20.144	255.255.255.0		6	<input type="checkbox"/>

Add Modify Remove

## Cluster Network Settings – With Advanced Networking *Enabled*

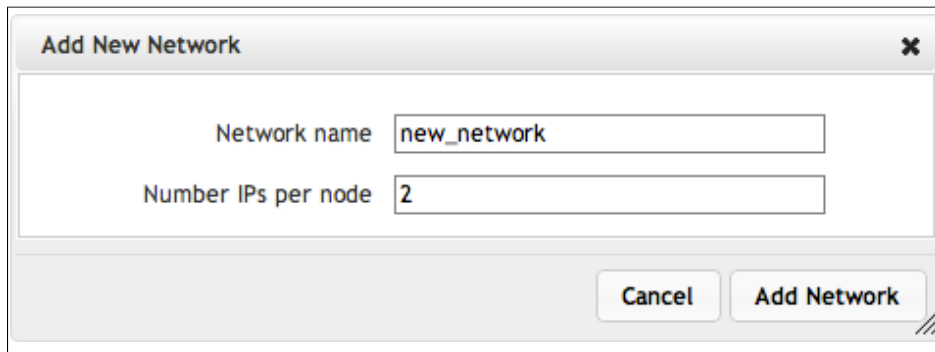
Enter a search term in the field above the **Cluster Networks** and **Address Ranges** tables to filter the display, as described in Section 1.1.11, “Search Functionality” on page 3.

### 2.4.1. Creating Cluster Networks

➤ To create a cluster network:

1. Navigate to the **Settings > Cluster > Cluster Networks** page.
2. If you are modifying or removing an existing network, select the network in the **Actions** column of the **Cluster Networks** panel.
3. From the **Cluster Networks** panel, click on one of the following:
  - **Add** (Advanced networks only) – If advanced networking is *not* enabled, you can have one and only one network.
  - **Modify** – Allows you to change the number of IPs assigned to each node.
  - **Remove** (Advanced networks only) – Removes an existing network, freeing the IP addresses used by that network for other uses.

If you have selected **Add** or **Modify**, the **Add New Network** or **Modify Existing Network** wizard starts.



#### Note

Using the browser’s back button will exit the wizard. Use the **Back** button on the wizard instead. You can choose **Cancel** at any point.

4. If you are adding a network, enter the name in the **Network name** field. The field will not be available if you are modifying an existing network.
5. To specify a minimum number of cluster addresses to be held by each node in an FXT cluster, enter the number in the **Number IPs Per Node** field. This number must be at least **1**.

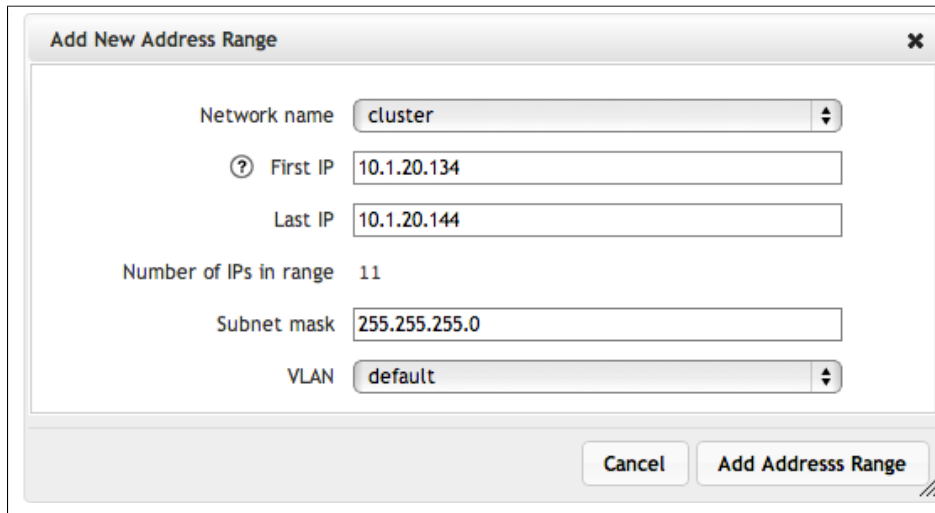
### 2.4.2. Cluster Network Settings

➤ To set the cluster IP addresses:

1. Navigate to the **Settings > Cluster > Cluster Networks** page.
2. If you are modifying or removing an existing address range, select the range in the **Actions** column of the **Address Ranges** panel.

3. From the **Address Ranges** panel, click on one of the following:
  - **Add** (Advanced networks only) – If advanced networking is *not* enabled, you must have one and only one address range.
  - **Modify** – Changes the range of IP addresses.
  - **Remove** (Advanced networks only) – Removes an existing range, freeing the IP addresses for other uses.

If you have selected **Add** or **Modify**, the **Add New Address Range** or **Modify Existing Address Range** wizard starts.



#### Note

Using the browser's back button will exit the wizard. Use the **Back** button on the wizard instead. You can choose **Cancel** at any point.

4. If you are adding a IP address range, choose the network that will use these IP addresses from the **Network name** drop-down list. The drop-down will not be available if you are modifying an existing network.
5. Enter the first IP address in the range into the **First IP** field.
6. Enter the last IP address in the range into the **Last IP** field.

At this point, the number of IP addresses being allocated will appear in the **Number of IPs in range** field. If the number is 128 or more, you will be asked if you really want that many addresses allocated. If the number is what you want, you simply need to submit the changes; otherwise, you will have an opportunity to correct the entry.
7. (Advanced networks only) Enter the cluster range's subnet mask in the **Subnet mask** field.
8. (Advanced networks only) Choose the VLAN from the **VLAN** drop-down list that cluster interfaces are to use. Possible values include **Default** and any cluster-role VLANs that you created in Section 5.2, "Creating a VLAN" on page 78.
9. Choose **Cancel** to restore the original IP address settings, or choose **Add Address Range** to commit any changes.
10. The Avere OS displays a pop-up that warns you about client disruption and asks you for confirmation to proceed. Click **OK**.

After a moment, the new address range appears in the **Address Ranges** panel.

## 2.5. Managing Exports (Settings Tab | VServer)

Clients in the Avere system access information via *exports* from each vservers. (For more information on setting up vservers, refer to Chapter 4, *Virtual Servers Used for Client Access* on page 64.) These exports basically appear as filesystems to the client. The cluster gets a list of exports for the virtual server from the core filer. You can view and manage export policies and rules from the **Settings > VServer > NFS Exports** page.

*Export policies* consist of rules that define client access granted by each policy. You can create export rules and policies, or use existing policies. The default export policy (named `default`) provides full access to the exported filesystem. Export policies can be applied to netgroups as well as exports.

**You will need to first create an export policy (or use the default policy), then create any rules you want to add to that policy.** This is because each rule requires you to apply it to an existing policy.

If you change an export policy or rule, or bind an export to a policy on the **NFS Exports** page, the cluster immediately polls for netgroup information if the required information was not previously obtained by the cluster.

For information on managing exports on a specific core filer to operate correctly with the FXT Series cluster, see Appendix A, *Core Filer-Specific Configuration Notes* on page 156.



### Note

You must set an Avere data export and directory before enabling high availability, discussed in the next section, Section 2.6, “High Availability” on page 33. The export must be at least 15 GB in size.

To view and manage export policies and rules, navigate to the **Settings > Vservers > Export Rules** page.

### 2.5.1. Creating or Modifying Export Rules

➤ To create or modify an export rule:

1. From the **Export Rules** page, select the appropriate export policy from the **Policy** drop-down list.

If you do not want to use the default policy, which allows full access to the exported filesystem, you can select **Create New Policy**, and follow the procedure in Section 2.5.2, “Managing Export Policies” on page 31 to create a policy for the rule.

2. Do one of the following:

- In the table of rules that is displayed for the export policy, locate the rule that you want to change and click **Modify**. The **Modify Rule Definition** area appears.
- Click **Add New Rule**. The **Add Rule Definition** area appears.

Dashboard Settings Analytics Data Management Support V3.1.1.1-9551a5c --- admin 3.1.Cluster

VServer Export Rules

Manage VServers Vserver Details Client Facing Network Namespace Export Policies NFS CIFS CIFS Shares Core Filer Cluster Administration

Create new policy

Policy: **cust\_limit**

Important note: for a given client, only the first matching rule is applied.

Scope	Filter	Access	Attributes				Auth	Actions
			SUID	SubDir	Squash	Anon		
default	*	R/W	yes	yes	no	-	SYS	<div>Delete</div> <div>Modify</div>

Client to test:  Filter Clear

Delete policy cust\_limit Add new rule

**Add New Rule Definition**

Rule scope: **host**

Filter:

Allowable access: **no access**

UID mapping: **map root to Anon**

Anon: **-2 (nobody)**

Allow SUID bits within this export: ☐

Allow submounts within this export: ☐

**Authentication Flavors**

UNIX/SYS ☒

Kerberos5 ☐ You have not yet enabled Kerberos. You may modify this setting on [the NFS](#) page.

Submit

3. Select a scope for the rule from the **Rule scope** drop-down list. Possible values include the following:

- **host** – The rule matches a specified host (client).
- **netgroup** – The rule matches clients in a specified netgroup.
- **network** – The rule matches clients in a specified network or subnet.
- **default** – The rule matches all clients.

4. Enter a filter for the rule. The filter format is based on the rule scope, as follows:

- **host** – The IP address or fully qualified domain name of the host to which the rule applies
- **network** – The network or subnet to which the rule applies, specified in one of the following formats:
  - *IP\_address/mask\_length*; for example, 192.168.0.1/24.
  - *IP\_address/netmask*; for example, 192.168.0.1/255.255.255.0.
  - *=NETWORK*, where *NETWORK* is the network name; for example, =EXAMPLE.COM.

- **netgroup** – The name of the netgroup to which the rule applies, preceded by the @ (at) symbol. For example, for a netgroup named BUILDHOSTS, enter **@BUILDHOSTS** into the field.
  - **default** – The asterisk (\*) character, specifying that no filtering is applied.
5. Select the appropriate access level for the rule from the **Allowable access** drop-down list, one of the following:
    - **no access**
    - **read-only**
    - **read/write**
  6. Select the type of user-ID mapping for the rule from the **UID mapping** drop-down list, one of the following:
    - **map root to Anon** – Client requests from user ID 0, typically the root user, are mapped to the anonymous user. This is commonly known as “root squashing.”
    - **map all to Anon** – Client requests from all users are mapped to the anonymous user.
    - **no UID mapping** – No user-ID mapping is performed.
  7. Select the appropriate user ID for the anonymous user from the **Anon** drop-down list, one of the following:
    - **-2 (nobody)**
    - **65534 (nobody)**
    - **-1 (no access)**
    - **65535 (no access)**
    - **0 (unprivileged root)**
  8. If you want to enable set-user and set-group ID (SUID and SGID) bits on the export, select the **Allow SUID bits within this export** checkbox.
  9. If you want to enable access to submounts on the export, select the **Allow submounts within this export** checkbox.
  10. From the **Authentication Flavors** area, select either **UNIX/SYS** (the default) or **Kerberos**. For more information on enabling Kerberos, refer to Section 2.8.3, “Enabling Kerberos Authentication” on page 43.
  11. Click the **Add rule** button.

The rule appears at the top of the **Export Rules** page.

## 2.5.2. Managing Export Policies

To view and manage export policies, navigate to the **Settings > Vservers > Export Rules** page and select the policy from the **Policy** drop-down list. The Avere Control Panel displays a page with the export policy rules, as described in Section 2.5.1, “Creating or Modifying Export Rules” on page 28.



### Caution

If you share an NFS export through CIFS, do not change the export policy or rules on the NFS export after setting up sharing. Doing so can result in unpredictable access to the export through CIFS and possible unauthorized access to data.

### 2.5.2.1. Creating and Applying a New Policy

➤ To create a new export policy:

1. Navigate to the **Settings > Vservers > Export Rules** page and select the policy from the **Policy** drop-down list.
2. Click the **Create new policy** button on the top right.
3. Enter the name of the new policy in the pop-up box that appears and click **Submit**.
4. Add any rules to the new policy as described in Section 2.5.1, “Creating or Modifying Export Rules” on page 28.

➤ To apply an export policy:

1. Navigate to the **Settings > VServers > Export Policies** page.
2. Select the vservers that will be exporting the filesystem.
3. Select the core filer that holds the files for the filesystem. The page is populated by the exports available on that core filer.

The screenshot displays the 'vservers1 - Export Policies' interface. At the top, there's a navigation bar with tabs: Dashboard, Settings, Analytics, Data Management, and Support. Below this, a sidebar on the left lists various management options: VServer, Manage VServers, vservers1, Vserver Details, Client Facing Network, Namespace, Export Policies (highlighted), Export Rules, NFS, CIFS, CIFS Shares, Core Filer, Cluster, and Administration. The main content area is titled 'vservers1 - Export Policies' and features a 'Core Filer' dropdown menu currently set to 'grape'. Below this, a section titled 'Exports available to client computers' contains a table of export rules. Each row lists a path (e.g., /vol/ker\_1, /vol/awson, /vol/adixit), a dropdown menu set to 'default', a link to 'Export Rules', and a checkbox for 'Qtrees'. At the bottom right of the table, there are 'Reset' and 'Submit' buttons.

Path	Policy	Export Rules	Qtrees
/vol/ker_1	default	<a href="#">Export Rules</a>	<input type="checkbox"/>
/vol/awson	default	<a href="#">Export Rules</a>	<input type="checkbox"/>
/vol/adixit	default	<a href="#">Export Rules</a>	<input type="checkbox"/>
/vol/loster	default	<a href="#">Export Rules</a>	<input type="checkbox"/>
/vol/aklter_nt	default	<a href="#">Export Rules</a>	<input type="checkbox"/>
/vol/asmlstanDisks	default	<a href="#">Export Rules</a>	<input type="checkbox"/>
/vol/avarney_0	default	<a href="#">Export Rules</a>	<input type="checkbox"/>
/vol/averdemo	default	<a href="#">Export Rules</a>	<input type="checkbox"/>
/vol/wntfs	default	<a href="#">Export Rules</a>	<input type="checkbox"/>
/vol/wktsas0	default	<a href="#">Export Rules</a>	<input type="checkbox"/>
/vol/wktsas1	default	<a href="#">Export Rules</a>	<input type="checkbox"/>

4. For each listed export, select the appropriate export policy from its drop-down list. The default export policy is named `default`; it provides full access to the export.
5. If applicable, click the **Qtrees** checkbox.  
  
FXT qtrees permit the Cache Utilization Control feature to limit modified data within top level directories of a root export, so that client rename and hard link operations across FXT qtree boundaries are not permitted. Note that the properties of FXT qtrees are not the same as those of NetApp qtrees.  
  
Refer to Section 6.2, “Specifying the Cache Policy” on page 85 for more information about configuring the Cache Utilization Control.
6. Choose **Reset** to restore the original export settings, or choose **Submit** to commit any changes.

### 2.5.2.2. Deleting an Export Policy

- To delete an existing export policy:
  1. Select the name of the policy that you want to delete from the **Policy** drop-down list.
  2. Click the **Delete Policy** *virtual\_server\_name.policy.policy\_name* button.
  3. A pop-up box prompts you for confirmation of the deletion operation. Click **OK** to confirm deletion.



#### Note

You can view the rules in an export policy by clicking the grey **Export Policy** link next to the export. This will take you to the **Export Rules** page for that policy, also accessible from the **Settings** tab as described in the next section.

### 2.5.2.3. Deleting a Rule from an Export Policy

- To delete a rule from an export policy:
  1. Select the appropriate export policy from the **Policy** drop-down list.
  2. In the table of rules that is displayed for the export policy, locate the rule that you want to delete and click **Delete**.
  3. A pop-up box prompts you for confirmation that you want to delete the rule. Click **OK**.

## 2.5.3. Unmounting Exports

Before unmounting any exports on your core filer, be certain that no client access to those exports is being routed through the FXT Series cluster, even if you plan to unmount and remount the clients.



#### Important

If clients attempt to access exports that the FXT Series cannot access during the export-unmount process, cluster downtime can occur.

Next, unmount the export or exports on the core filer. See the documentation for your core filer for information on unmounting exports. For example, before unmounting an export on a Data ONTAP 7G system, the volume behind the export must be marked as offline:

```
exportfs -u /vol/volX
vol offline volX
```



## Note

The **-u** parameter on the **exportfs** command unexports the specified volume without changing the Data ONTAP system's `/etc/exports` file.

If and when the volume comes back online, mark it as online by running the following commands on the Data ONTAP system:

```
vol online volX
exportfs /vol/volX
```

To ensure that the export or exports no longer appear in the list of exports on the FXT Series cluster, go to the **Export Policies** page under the **VServer** heading of the Avere Control Panel's **Settings** tab.

## 2.5.4. Hierarchical Exports

If hierarchical exports are enabled on certain core filers, the export rule of a client that changes directories into a read-only filesystem (such as a snapshot) is applied to the highest level of the export hierarchy both in the read-only filesystem and in the parent read-write filesystem (when the client changes directories back to the parent read-write filesystem from the read-only filesystem).

Depending on the permissions set in the client's export rule, the client can potentially access data at higher levels of the directory hierarchy than it is intended to access. Avere Systems recommends that you do not use hierarchical exports. Alternatively, you can work around the issue by applying identical export rules to each hierarchical export, including the highest level in the export hierarchy.

Core filers known to be affected include the following:

- Data ONTAP 7G
- Data ONTAP 8, 7-mode
- Linux
- BSD and variants
- Data Domain
- Isilon

This is not an issue for core filers such as Data ONTAP GX and Data ONTAP 8, cluster-mode, that do not use hierarchical exports.

## 2.6. High Availability

*High availability* (HA) provides continuity of data service in the event of a node outage. It protects all client-written data in addition to the interface failover provided by the cluster in both HA and non-HA configurations.

High availability requires a cluster of two or more nodes, and requires a data repository on the core filer. When you enable high availability, nodes in the cluster automatically configure and optimize themselves for HA.

In a cluster with multiple virtual servers, enabling HA for the cluster automatically enables HA for all virtual servers as well. You cannot enable HA for only a subset of a cluster's virtual servers.

## 2.6.1. Specifying an HA Data Repository

Before configuring HA, you need to specify an export and directory on the core filer where the cluster can store small configuration files. In addition, setting the cache policy, as described in Chapter 6, *Setting the Cache Policy* on page 84, will greatly aid in configuring high availability.



### Important

- You must create an Avere data export and directory to be used as a data repository, as described in Section 2.5, “Managing Exports (Settings Tab | VServer)” on page 28, before enabling high availability.
- Do not allow the export and directory for the Avere files to run out of space. If the cluster cannot read and write to the data repository, HA services can fail. The data repository on the core filer must be 15 GB or larger.
- If the cluster has a low number of nodes, it is highly recommended that you choose 2-node cluster protection, and choose a dedicated export and directory for HA use.

➤ To specify an HA data repository:

1. Optionally, create an Avere-dedicated export on your core filer. Refer to Appendix A of the documentation for your node for details.



### Note

If you create an Avere OS-dedicated export, ensure that only the FXT Series edge filers (that is, the cluster IP addresses) have read/write and root access to it. It is strongly recommended that no non-Avere data be stored on the export.

2. Navigate to the **Settings > Cluster > High Availability** page.

The screenshot shows the 'High Availability' configuration page. The left sidebar contains a menu with options: VServer, Core Filer, Cluster, General Setup, Administrative Network, Cluster Network, FXT Nodes, High Availability (selected), Monitoring, Schedules, Directory Services, Kerberos, Login Services, Active Directory, Optimization, IPMI, Support, Licenses, Cloud Credentials, and Administration. The main content area is titled 'High Availability' and includes a checkbox for 'Enable HA' which is checked. Below this is a section for 'Core Filer Data Parameters' with a checkbox for '2-node cluster protection mode (requires 15 GBytes)' which is also checked. There are three input fields: 'Core filer' with the value 'thor', 'HA data export' with the value '/vol0/lbarr', and 'HA data directory' with the value '.avere'. At the bottom right of the configuration area are buttons for 'Unset', 'Revert', and 'Submit'.

3. Select **Enable HA** to allow high-availability access.

4. Optionally, select **2-node cluster protection mode (requires 15 GBytes)**.

This option, for clusters with low numbers of nodes, allows cluster data to be stored on the core filer if only one node is available. You can select this even if the cluster has more than two nodes, and it will then be enabled even if nodes are removed and the cluster falls to two nodes.

If you do not select this option, the cluster will *not* use a core filer as its data repository, which will use more of the cluster's own total drive capacity.

5. From the **Core filer** drop-down list, choose the core filer on which the repository is to reside. This is the cluster's name for the core filer (not necessarily the network name for the core filer). In a GNS configuration, the selected vservers and core filer must be associated with each other. For more information, refer to Section 4.3, "Creating and Maintaining a Global Namespace" on page 72.



#### Note

When selecting a core filer for the HA data repository, choose a high-capacity, high-performance core filer with a reliable network connection to the Avere cluster.

6. From the **HA data export** drop-down list, select an NFS export on the core filer to which the cluster can write configuration files. If you created an Avere-dedicated export but do not see it in the list of exports, refresh the page.
7. In the **HA data directory** field, enter the name of a directory on the selected export to which the cluster can write configuration files. The default is `.avere`. It is strongly recommended that you accept the default value for the directory name.



#### Note

- You cannot enter an empty string for the directory name.
- Do not create the directory by using the **mkdir** command or a similar utility; the cluster creates the directory itself.

8. Choose **Revert** to restore the high availability parameters, or choose **Submit** to commit any changes.



#### Note

It is strongly recommended that you do not change the export and directory for the Avere files after you specify them. See Section 2.6.3, "Unsetting the HA Export and Directory" on page 36 for information on unsetting the export and directory for the data repository.

## 2.6.2. Disabling High Availability

➤ To disable high availability:

- From the **High Availability** page, deselect the **Enable HA** checkbox and choose **Submit**.

High availability is now disabled.

In addition, the **Unset** button is now activated, so you can optionally unset the export and directory for the data repository.

## 2.6.3. Unsetting the HA Export and Directory



### Note

You must disable HA before you can change or unset the data repository parameters.

➤ To unset the repository export and directory:

1. From the **High Availability** page, deselect the **Enable HA** checkbox and choose **Submit**.

The **Unset** button is now activated, so you can optionally unset the export and directory for the data repository.

2. Choose **Unset**. The Data Parameters can now be changed.
3. From the **Core Filer** drop-down list, choose **Select Core Filer** (at the top of the list).
4. Choose **Revert** to restore the high availability parameters, or choose **Submit** to commit any changes.

The core filer, export, and directory are now unset.

## 2.7. Defining Schedules

A *schedule* is the way that Avere uses to specify how caching is done for the system (that is, when data is written to the core filer). Schedules work per-cluster with clusters that have one or more virtual servers. A schedule has one or more *clauses*, with each clause consisting of the minute, hour, and day of the week on which the clause runs.

For more information about caching and writethrough periods, refer to Section 6.2.2, “Setting Read/Write Mode” on page 87.

## 2.7.1. Creating a Schedule

➤ To create a schedule:

1. Navigate to the **Settings > Cluster > Schedules** page.

The screenshot shows the 'Create New Schedule' page in the VServer settings. The sidebar on the left contains navigation links: VServer, Core Filer, Cluster, General Setup, Administrative Network, Cluster Network, FXT Nodes, High Availability, Monitoring, Schedules (highlighted), Directory Services, Kerberos, Login Services, Active Directory, Optimization, IPMI, Support, Licenses, Cloud Credentials, and Administration. The main content area is titled 'Schedules' and 'Create New Schedule'. It features a 'Schedule Name' field with the value 'weekday\_evenings'. Below this are two 'Clause' sections. 'Clause 1' has 'Minutes' set to 0, 'Hours' set to 9 pm and 11 pm, and 'Days' set to Monday, Tuesday, Wednesday, Thursday, and Friday. 'Clause 2' has 'Minutes' set to 5, 'Hours' set to 10 pm, and 'Days' set to Monday. There are buttons to 'Remove This Clause', 'Add Another Clause', and 'Create New Schedule'.

2. If any existing schedules are listed in the **Existing Schedules** area, click **Create New Schedule**. If the **Existing Schedules** area is not displayed, proceed to the next step.
3. Enter the name of a new schedule in the **Schedule Name** field.
4. From the listings provided in the **Clause 1** area, select the minute, hour, and day-of-the-week values on which the clause is to run. You can choose multiple values for each time division.  
  
Thus, the schedule in the previous figure puts the Avere system into writethrough mode every weekday at 9 p.m. and 11 p.m., and additionally, on Monday nights at 10:05 p.m.
5. Optionally, to add another clause to the schedule, click **Add Another Clause** and enter information for the new clause as described in the previous step.
6. Click the **Create New Schedule** button.

The schedule is added to the **Existing Schedules** list.

## 2.7.2. Updating an Existing Schedule

➤ To update an existing schedule:

1. From the **Existing Schedules** area on the **Schedules** page, click the name of the schedule that you want to update.
2. Edit the schedule as needed. You can perform the following actions:
  - Change existing clauses by selecting new values, deselecting current values, or both.
  - Delete a clause by clicking **Remove This Clause** in the lower right-hand corner of the clause listing.
  - Create new clauses by clicking **Add Another Clause** and specifying values for it as described in Section 2.7.1, “Creating a Schedule” on page 37.
3. Click the **Update Schedule** button.

## 2.7.3. Deleting an Existing Schedule

➤ To delete a schedule:

1. From the **Existing Schedules** area on the **Schedules** page, click the name of the schedule that you want to delete.
2. Click the **Delete Schedule** button. The browser displays an alert asking you to confirm the deletion.
3. Click **OK**.

## 2.8. Controlling Authentication

If your environment requires clients to be authenticated for file access, you can configure Avere OS to use directory services and netgroups to provide this authentication.



### Note

- You must configure a directory service, as described in Section 2.8.1, “Selecting and Configuring a Directory Service” on page 39 if you are using CIFS. See Section 7.2, “Core Filer Prerequisites for CIFS” on page 95 for more information.
- Directory-service and netgroup settings are cluster wide; they cannot be specified on a per-virtual-server basis.

Supported directory services include the Network Information Service (NIS) and the Lightweight Directory Access Protocol (LDAP).

- Refer to the [FreeBSD man pages for yp](#) or more information on NIS.
- Refer to the [FreeBSD man pages for ldap](#) or more information on LDAP.

Network groups, or netgroups, list one or more sets of hosts (clients), domains, or both that are to be given network access. See the man page for **netgroup** (`man 5 netgroup`) for general information.

The Avere cluster can obtain netgroup information either from a directory service or from an external file in `/etc/netgroup` format. Although the `/etc/netgroup` format includes user information, the Avere cluster disregards any user specifications in netgroup information. The cluster uses NFS export rules to determine access for users.

### 2.8.1. Selecting and Configuring a Directory Service

The cluster uses directory services to assist with client authentication requests that use netgroup, and to assist with CIFS user authentication.

The configuration of a directory service is required only if you are using CIFS, netgroups, or both. If you are using netgroups, you can configure the cluster to obtain netgroup information either from a directory service or from an external file in `/etc/netgroup` format on an HTTP server. If you are not using either of these Avere OS services, configuring a directory service is optional.



### Note

- If you want to allow users to log into the FXT node from Windows as an administrator, and you want to use Active Directory for directory services, you need to do the following:
  - Create a group in Active Directory called **AvereAdmins**.
  - Add these users to that group.
- All virtual servers on a Avere cluster will use the same directory-service settings.

➤ To select a directory service:

1. Use the following criteria to decide whether to use NIS or LDAP:

- Do you already have one or both services running in your environment? If both, which one is preferred?

- If you do not already have a directory service running in your environment, consider the following:
    - NIS is simpler to implement and administer and requires less planning than LDAP, but is less secure and scalable than LDAP.
    - LDAP requires up-front planning and organization and is more difficult to administer than NIS; however, it provides tighter security, finer-grained controls, and better scalability than NIS.
  - Does your network administration team prefer one over the other, or do they have experience in one but not the other?
  - Do you plan to use the directory service for services or applications other than the Avere cluster? If so, what are the requirements for use of the directory service by the other services and applications?
2. If necessary, configure the selected directory service for your network. See the documentation for the directory service and its host operating system for details.
  3. If necessary, start the directory service on the host operating system and ensure that the service is restarted in the event of a host restart. See the documentation for the directory service and its host operating system for details.
  4. Navigate to **Settings > Cluster > Directory Services**.

Dashboard Settings Analytics Data Management Support V3.1.1.1-6a7c480 --- admin  
LiGo Cluster

VServr +  
Core Filr +  
Cluster -  
General Setup  
Administrative Network  
Cluster Network  
FXT Nodes  
High Availability  
Monitoring  
Schedules  
Directory Services  
Kerberos  
Login Services  
Active Directory  
Optimization  
IPMI  
Support  
Licenses  
Cloud Credentials  
VLAN  
Administration +

### Directory Services

#### LDAP

Server

Base DN

Secure Access ☐

Require valid certificate ☐

Credentials ☒

Bind DN

Bind Password

#### NIS

NIS Server

NIS Domain

#### Netgroup

Source

Poll Period

Custom Poll Period

#### User Name

Source

Enable Dashboard Conditions ☒

NFS Domain

Poll Period

Last Poll Mon Oct 29 2012 15:28:53 GMT-0400 (EDT)  
Next Poll Mon Oct 29 2012 16:28:53 GMT-0400 (EDT)

### 2.8.1.1. Configuring an LDAP Service

You can also use an LDAP server to authenticate users that can log into the Avere OS, as described in Section 2.8.4, “Specifying LDAP Login Services” on page 45.

➤ To configure LDAP for directory services:

1. In the **LDAP Server** field, enter the fully qualified domain names or IP addresses of up to three LDAP servers. If you enter more than one LDAP server, separate the server names or addresses with single spaces.
2. In the **LDAP Base DN** field, enter the LDAP *base distinguished name*. LDAP queries are performed on the base DN, the DN of the entry, and all entries below it in the directory tree. (Active Directory requires you to use the domain component, or `dc`, method of providing a name, rather than the organization, or `o` method.)

Ask your Active Directory administrator if you don’t know your base DN.



#### Note

Base and bind DN entries use a similar format. So, for example, if the domain name is “ourdomain.server.company.com”, the DN entry is in the form

`ou=ourdomain,dc=server,dc=company,dc=com`

3. To optionally force LDAP to use TLS/SSL encryption with LDAP, select the **Secure access** checkbox. The Avere Control Panel displays the **Require valid certificate** checkbox.
  - a. To require a valid certificate from the LDAP server before a connection is established, select the **Require valid certificate** checkbox. The Avere Control Panel displays the **CA Certificate URI** field.
  - b. Do one of the following:
    - Enter the HTTP URI of the Certificate Authority (CA) that signed the LDAP server’s SSL certificate in PEM format, and click the **Download Now** button.
    - If you are using a self-signed SSL certificate, enter the URI to the certificate on the LDAP server, and click the **Download Now** button.
    - Leave this field blank to cause the Avere OS to attempt to automatically download a certificate.
4. To require login/connection credentials to the LDAP server, select the **Credentials** checkbox. The **Bind DN** and **Bind Password** fields appear.
  - a. In the **Bind DN** field, enter the distinguished name to bind to the LDAP server.
  - b. In the **Bind Password** field, enter the password for the distinguished name entered in the previous field.
5. Click the **Submit** button to save the entries.



#### Important

If you are using an OpenLDAP server, set the following values in the `slapd.conf` file:

```
# Maximum entries returned for a search
sizelimit size.soft=100 size.hard=1000 size.prtotal=unlimited
```

Not setting these values can result in incomplete data being returned for large user/group queries.

### 2.8.1.2. Configuring an NIS Service

➤ To configure NIS:

1. In the **NIS Server** field, enter the fully qualified domain names or IP addresses of up to three NIS servers. If you enter more than one NIS server, separate the server names or addresses with single spaces.
2. In the **NIS Domain** field, enter the NIS domain name. This is the single domain to which the NIS service is bound.
3. Click the **Submit** button to save the entries.

### 2.8.2. Configuring Netgroups

You can use netgroups to control access to the cluster on a host-by-host (client-by-client) basis.

The virtual server enforces export rules according to the last successful poll, if any, of netgroup information.

If a host (client) or domain listed in the netgroup file is not fully qualified, the Avere cluster uses the DNS search paths listed on the **Administrative Network** page to help resolve the names. See Section 2.3.7, “Setting DNS Parameters” on page 22 for information on setting the DNS search path.

The cluster can obtain netgroup information from either a directory service (NIS or LDAP) or an external file in `/etc/netgroup` format.

➤ To set the netgroup source:

1. Navigate to the **Settings > Cluster > Directory Services** page.
2. From the **Source** drop-down list in the **User Name** section, choose one of the following options:
  - **None** to use no netgroup source
  - **NIS** to use NIS as the netgroup source
  - **LDAP** to use LDAP as the netgroup source
  - **File** to use an external file as the netgroup source
3. If you selected **File** as the name source, enter the URL of a password file in the **File URI** field. The password file must be in standard `/etc/netgroup` format.

See Section B.5, “The Format of the Netgroup File” on page 162 for more information about the required formats for the netgroup file.

4. Select **Enable Dashboard Conditions** if you want a condition to appear with the alerts on the Dashboard when there are problems or conflicts with netgroup entries.
5. The cluster re-reads netgroup information at a specified polling interval.

From the **Poll Period** drop-down list, choose the polling period. Possible values include the following:

- **24 hours** or **12 hours**
- **1 hour** (the default)
- **Manual** – The cluster polls netgroup information only when the **Poll Now** button is clicked
- **Custom** – The cluster polls netgroup information at the interval specified, in minutes, in the **Custom Poll Period** field, which the Avere OS displays when you select the **Custom** value

6. Click the **Submit** button at the bottom of the **Directory Services** page.

The entered information is committed, and an immediate poll of netgroup information is triggered.

The **Last Netgroup Update** column on the **Manage Existing VServers** page lists the date and time of the last successful netgroup update for each virtual server. See Section 4.2, “Managing Virtual Servers (Settings Tab | VServer)” on page 65 for a list of other virtual-server information displayed on this page.

If a poll fails, the poll period is automatically reduced to approximately one minute in an attempt to handle transient outages.

You can assign export policies to netgroups the same way you assign export policies to other network entities, as described in Section 2.5, “Managing Exports (Settings Tab | VServer)” on page 28 and Section 2.5.1, “Creating or Modifying Export Rules” on page 28.

### 2.8.3. Enabling Kerberos Authentication

You can use Kerberos for cluster-client communication, or for cluster-core filer communication.

➤ To use Kerberos:

1. Configure DNS for the cluster, as described in Section 2.3.7, “Setting DNS Parameters” on page 22.
2. Configure LDAP, NIS, or File as the user name source, as described in Section 2.8.1, “Selecting and Configuring a Directory Service” on page 39.
3. **Cluster-client communication** – Navigate to **Settings > VServer > Export Rules**. Here you can select **UNIX/SYS**, **Kerberos5**, or both as the authentication flavor (type). Refer to Section 2.5, “Managing Exports (Settings Tab | VServer)” on page 28 for more information.
4. **Cluster-client communication only** – Navigate to **Settings > VServer > NFS**. Here you can upload a keytab file for each vserver, and can enable Kerberos.

**Cluster-core filer communication** – Navigate to **Settings > Cluster > Kerberos**. Here you can upload the keytab file for communication with the core filer, and can enable Kerberos, as described in step 5. Clients generally require both keytab files.



#### Note

Because a Kerberos realm is an administrative domain, all vservers in the same cluster, which have Kerberos enabled, must belong to the same realm.

The keytab must contain one or more NFS service principals for the vserver of the form:

```
nfs/Reverse_resolved_name@Realm_name
```

where *Realm\_name* is set on the Kerberos page, as described step 5, and

*Reverse\_resolved\_name* is the reverse-resolved name of each client-facing address for the vserver.

In AvereOS, each client-facing address will reverse-resolve to a different name. Thus, a principal is required for each reverse-resolved name.

5. Navigate to the **Settings > Cluster > Kerberos** page and enter the following information:



### Note

Because a Kerberos realm is an administrative domain, all vservers in the same cluster, which have Kerberos enabled, must belong to the same realm.

The screenshot shows the 'Kerberos Configuration' page in a management console. The left sidebar contains a menu with items like VServer, Core File, Cluster, General Setup, Administrative Network, Cluster Network, FXT Nodes, High Availability, Monitoring, Schedules, Directory Services, Kerberos (highlighted), Login Services, Active Directory, Optimization, IPMI, Support, Licenses, Cloud Credentials, and Administration. The main content area has two sections: 'Kerberos' and 'Kerberos Service Key Management'. The 'Kerberos' section contains fields for 'Realm' (filled with 'KERBEROS\_CO.NET'), 'KDC DNS Discovery' (checked), 'KDC' (filled with '10.0.0.23'), 'DNS Domain Discovery' (unchecked), and 'DNS Domains' (filled with '.kerberos\_co.net company.com'). A 'Submit' button is at the bottom right of this section. The 'Kerberos Service Key Management' section contains a 'Keytab File for Cluster' label, a 'Choose File' button, and the text 'No file chosen'. An 'Upload Keytab File' button is at the bottom right of this section. The top of the page has a navigation bar with 'Dashboard', 'Settings', 'Analytics', 'Data Management', and 'Support' tabs, and a user/cluster info bar showing 'V3.1.1.1-bfab49a --- admin' and '3\_1\_NoAdvancedNet\_Cluster'.

- **Realm** – Enter the Kerberos realm (domain) that contains the principal names in the Kerberos server database (for example, company.net).
- Do one of the following:
  - Select **KDC DNS Discovery** if you want Avere OS to search for the key distribution center.
  - In the **KDC** field, enter the fully-qualified domain name or IP address of the Kerberos key distribution center. This is available only if **KDC DNS Discovery** is not selected.
- Do one of the following:
  - Select **DNS Domain Discovery** if you want Avere OS to search for any available DNS domains.
  - In the **DNS Domains** field, enter a space-separated list of domains for the software to use for DNS completion, for example:  
`.kerberos_co.net kerberos_co.net .company.com company.com`  
This is available only if **DNS Domain Discovery** is not selected.
- **Keytab File for Cluster** – Browse to the cluster’s Kerberos keytab file, as described for the vserver’s keytab file in step 4

The MIT Kerberos pages provide more general information on the web about Kerberos.

## 2.8.4. Specifying LDAP Login Services

The **Login Services** page allows you to enter an LDAP source for users who are allowed to login to the Avere OS. Refer to Section 2.8.1.1, “Configuring an LDAP Service” on page 41 for information on LDAP for specific user access.

➤ To configure LDAP login services:

1. Navigate to **Settings > Cluster > Login Services**.

The screenshot shows the 'Login Services' configuration page in the LiGo Cluster interface. The top navigation bar includes 'Dashboard', 'Settings', 'Analytics', 'Data Management', and 'Support'. The left sidebar lists various settings categories, with 'Login Services' highlighted under the 'Cluster' section. The main content area is titled 'Login Services' and contains two sections: 'LDAP' and 'Login'. The 'LDAP' section includes fields for 'Server' (server.company.com), 'Base DN' (ou=ourdomain,dc=server,dc=company,dc=com), 'Secure access' (checked), 'Require valid certificate' (checked), and 'CA certificate URI' (server.company.com/path\_to\_CA). A 'Download Now' button is located below the 'CA certificate URI' field. The 'Login' section includes fields for 'Bind DN' (cn=admin,dc=ourdomain,dc=company,dc=com) and 'Bind Password' (masked with dots). Below the 'Login' section is a 'Source' dropdown menu set to 'Local/LDAP'. At the bottom right, there are 'Reset' and 'Submit' buttons.

2. In the **Server** field, enter the fully qualified domain names or IP addresses of up to three servers. If you enter more than one LDAP server, separate the server names or addresses with single spaces.
3. In the LDAP **Base DN** field, enter the LDAP *base distinguished name*. LDAP queries are performed on the base DN, the DN of the entry, and all entries below it in the directory tree. (Use the domain component, or `dc`, method of providing a name, rather than the organization, or `o` method.)

Ask your Active Directory administrator if you don't know your base DN.



### Note

Base and bind DN entries use a similar format. So, for example, if the domain name is “ourdomain.server.company.com”, the DN entry is in the form

`ou=ourdomain,dc=server,dc=company,dc=com`

4. To optionally force LDAP to use TLS/SSL encryption with LDAP, select the **Secure access** checkbox. The Avere Control Panel displays the **Require valid certificate** checkbox.
  - a. To require a valid certificate from the LDAP server before a connection is established, select the **Require valid certificate** checkbox. The Avere Control Panel displays the **CA Certificate URI** field.
  - b. Do one of the following:
    - Enter the HTTP URI of the Certificate Authority (CA) that signed the LDAP server's SSL certificate in PEM format, and click the **Download Now** button.
    - If you are using a self-signed SSL certificate, enter the URI to the certificate on the LDAP server, and click the **Download Now** button.
    - Leave this field blank to cause the Avere OS to attempt to automatically download a certificate.
5. In the **Bind DN** field, enter the distinguished name to bind to the LDAP server.
6. In the **Bind Password** field, enter the password for the distinguished name entered in the previous field.
7. Choose **Local/LDAP** from the **Source** drop-down list in the **Login** area.
8. Click the **Submit** button to save the entries.

### 2.8.5. Specifying the Source for Usernames

If you set up one or more directory services, you can specify which one will provide the UNIX usernames for the cluster.



#### Note

You must select a directory service as a UNIX username source if you are using CIFS. See Section 7.2, “Core Filer Prerequisites for CIFS” on page 95 for more information.

- To specify the source service for the usernames:
1. Navigate to the **Settings > Cluster > Directory Services** page.
  2. From the **Source** drop-down list in the **User Name** section, choose one of the following options:
    - **None** to use no name source
    - **NIS** to use NIS as the name source
    - **LDAP** to use LDAP as the name source
    - **File** to use external files as the name source
  3. If you selected **File** as the name source, perform the following steps:
    - a. Enter the URL of a password file in the **Password File URI** field. The password file must be in standard `/etc/passwd` format.
    - b. Enter the URL of a group file in the **Group File URI** field. The group file must be in standard `/etc/group` format.

See Appendix B, *Password and Group File Formats* on page 160 for more information about the required formats for the password and group files.

3. If you are using an NFSv4 domain for CIFS ACLs as described in Section 7.8, “Creating CIFS Shares” on page 107, optionally enter the name of the NFSv4 domain in the **NFS Domain** field.



### Caution

Do not clear the **NFS Domain** field if NFSv4 ACLs are in use. If it is, access to all CIFS shares that use NFSv4 ACLs is lost.

4. The cluster re-reads username information at a specified polling interval.

From the **Poll Period** drop-down list, choose the polling period. Possible values include the following:

- **24 hours** or **12 hours**
- **1 hour** (the default)
- **Manual** – The cluster polls netgroup information only when the **Poll Now** button is clicked
- **Custom** – The cluster polls netgroup information at the interval specified, in minutes, in the **Custom Poll Period** field, which the Avere Control Panel displays when you select the **Custom** value

5. Click the **Submit** button at the bottom of the **Directory Services** page.

The entered information is committed, and an immediate poll of username information is triggered.

If a poll fails, the poll period is automatically reduced to approximately one minute in an attempt to handle transient outages.

## 2.9. Optimizing the Avere Cluster for Use in a VMware® Environment

If you are using an Avere cluster to accelerate the performance of a core filer that holds VMWare images, you can optimize the cluster for a VMware-specific workload.

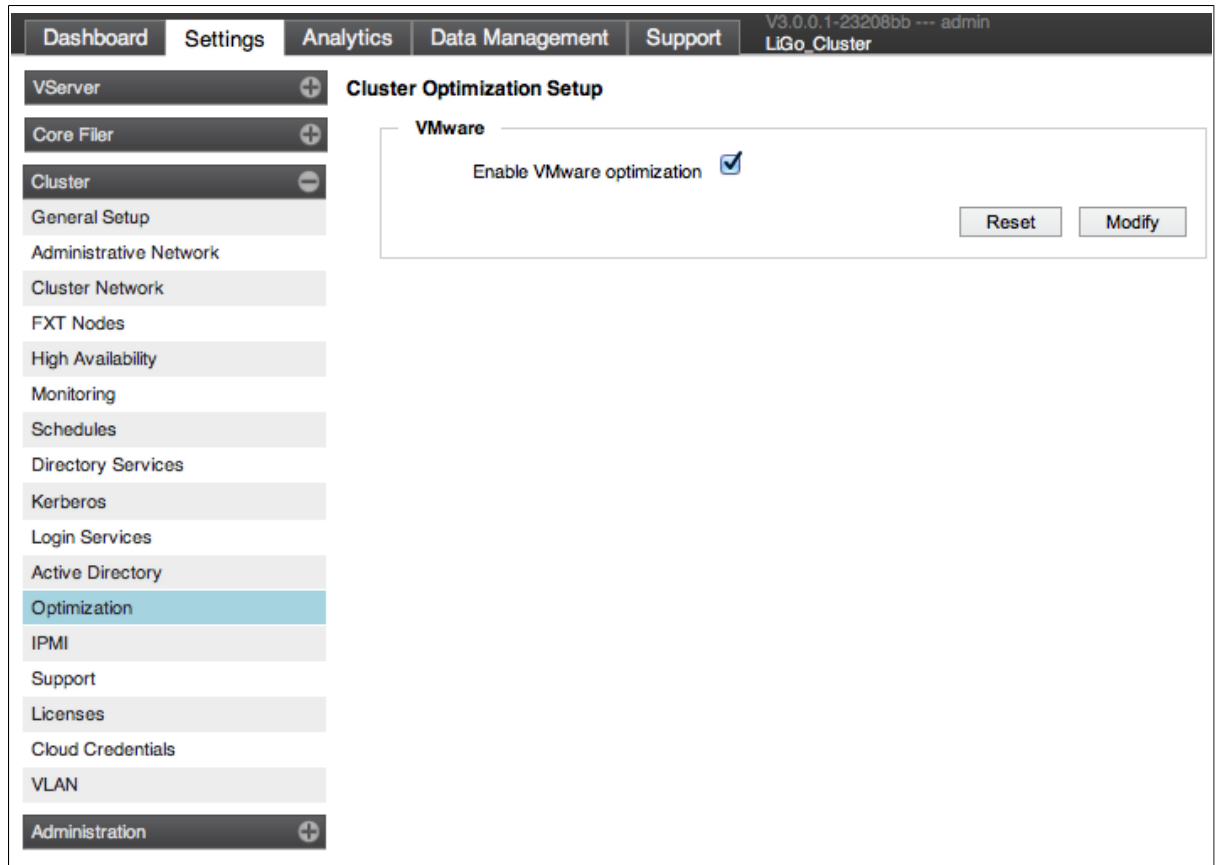


### Important

Optimizing a cluster for VMware limits the cluster's use to VMware-only workloads. Do not enable VMware optimization and then use the cluster for more general-purpose workloads.

➤ To enable VMware optimization for a cluster:

1. Navigate to the **Settings > Cluster > Optimization** page.



2. Select the **Enable VMware optimization** checkbox.
3. Click the **Modify** button.

The Avere Control Panel issues the following warning:

**WARNING:** Changing the VMware optimization setting requires restarting all nodes in the cluster and disrupts client data access. Are you sure?

4. Click the **OK** button to confirm.

All nodes in the cluster restart.

5. After all nodes in the cluster have restarted, you can mount your VMware clients to the core filer through the cluster for optimized core filer performance.

## 2.10. Configuring IPMI Cards

Each FXT Series node comes installed with an Intelligent Platform Management Interface (IPMI) card with a 100-MbE Ethernet port. This card can be used for remote operations such as power cycling. For more information on IPMI, see [www.intel.com/design/servers/ipmi/](http://www.intel.com/design/servers/ipmi/).

Avere recommends that all clusters have IPMI configured. You can configure all IPMI cards in a cluster at once or configure them individually.

- To configure the IPMI card in each node's in a cluster:

1. Navigate to the **Settings > Cluster > IPMI Configuration** page.

The screenshot displays the 'Cluster IPMI Configuration' page in the LiGo interface. The top navigation bar includes 'Dashboard', 'Settings', 'Analytics', 'Data Management', and 'Support'. The left sidebar lists various system components, with 'IPMI' highlighted under the 'Cluster' section. The main panel shows a table for configuring IPMI settings for the 'LiGo' node. The 'IPMI Settings (Optional)' section is expanded, showing fields for Mode (static), First IP (10.0.0.127), Last IP (10.0.0.132), Netmask (255.255.255.0), and Router (192.168.5.5). 'Reset' and 'Submit' buttons are located at the bottom right of this section.

2. **For a single node** From the table of nodes at the top of the page, click on the name of the node whose IPMI card you want to configure. The link takes you to the **Node Settings** page for the selected node.
3. In the **IPMI Settings (Optional)** panel.
4. From the **Mode** drop-down list, select one of the following options:
  - **static** – Allows you to enter the IP addresses manually, and they will not change (unless you change them manually).
    - **First IP** – The first IP address for the range of IPMI card addresses
    - **Last IP** – The last IP address for the range of IPMI card addresses
    - **Netmask** – The netmask for the IPMI card's IP addresses
    - **Router** – The default router for the IPMI card's IP addresses
  - **DHCP** – If you select this option, no additional information is required; the system obtains all required IP-address information from the DHCP server.
5. Click the **Submit** button.

## Chapter 3. Configuring a Core Filer (Settings Tab | Core Filer)

The second step in configuring client access through the Avere cluster, after creating a cluster, is adding a core filer for the initial vserver on the newly created cluster.

An FXT cluster is capable of serving as the front end for up to 50 core filers. The cluster sees each core filer as a single *dataset*; this is typically a standalone NAS server or a cloud object store, but it can also be a clustered filesystem or any other mechanism that serves NFS clients over NFSv3 and exposes data over NFS exports.

In the Avere OS, each core filer is identified by the cluster with an administrative name that is different from the core filer's network name (fully qualified domain name). For instance, a core filer with the network name `filer1.example.com` might be identified by the cluster as `mass1`. For ease of administration, you can change the cluster's default core filer identifier to match the name of the core filer (for example, change `mass1` to `filer1`).

You will need the following information about the core filer:

- If you are adding an NFS core filer:
  - Network name of the core filer. The fully-qualified domain name (FQDN), for example, `nas1.example.com`, not `nas1` or `127.98.12.1`, is preferable.
- If you are adding a cloud core filer:
  - A valid license for FlashCloud, activated on the Avere system. For more information, refer to Section 1.4, “Adding and Removing a License” on page 6.
  - An existing credential for your service provider, or the access and private keys to create a new one
  - Any necessary encryption details
  - The name and hostname for your existing cloud storage container

You can use the Avere Control Panel to configure, monitor, and maintain each individual virtual server and core filer, as described in Section 9.4, “VServers Tab” on page 127 and Section 9.5, “Core Filers Tab” on page 128.

Enter a search term in the field above the **Manage Core Filers** table to filter the display, as described in Section 1.1.11, “ Search Functionality” on page 3.

## 3.1. Creating an NFS Core Filer

➤ To add a new NFS core filer to the cluster:

1. Navigate to the **Settings > Core Filer > Manage Core Filers** page.

The screenshot shows the 'Manage Core Filers' page in the VServer interface. The top navigation bar includes 'Dashboard', 'Settings', 'Analytics', 'Data Management', and 'Support'. The left sidebar shows a tree view with 'VServer' expanded, containing 'Core Filer' and 'Administration'. The 'Core Filer' section is further divided into 'Manage Core Filers', 'Core Filer Details', 'Cache Policy', and 'Cloud Encryption Settings'. The main content area displays a table of core filers with columns: Name, Admin State, Cache Mode, Modified Data, Local Dirs, Vservers, and Actions. There are four entries listed, including 'blueAngel' (Amazon S3), 'earthAngel' (Amazon S3), 'grape' (NFS), and 'thor' (NFS). Buttons for 'Create', 'Invalidate', and 'Remove' are visible at the top and bottom of the table.

Name	Admin State	Cache Mode	Modified Data	Local Dirs	Vservers	Actions
blueAngel Type: Amazon S3	online	Read/Write Writeback: 30	0 Files	Enabled	None	<input type="checkbox"/>
earthAngel Type: Amazon S3	online	Read/Write Writeback: 30	0 Files	Enabled	None	<input type="checkbox"/>
grape Type: NFS	online	Read/Write Writeback: 43200	0 Files	Disabled	vserver1	<input type="checkbox"/>
thor Type: NFS	online	Read/Write Writeback: 3600	0 Files	Enabled	migrate_me	<input type="checkbox"/>

2. Click **Create**. The **Add New Core Filer** wizard starts.



### Note

Using the browser's back button will exit the wizard. Use the **Back** button on the wizard instead.

3. Select the core filer type, in this case, **NFS**.

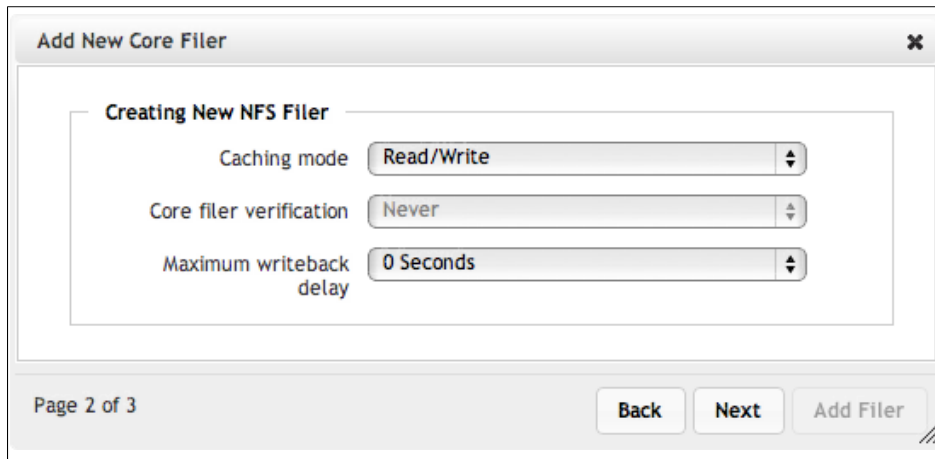
The screenshot shows the 'Add New Core Filer' wizard. It includes a title bar, a close button, and a main content area with instructions and a form. The instructions state: 'This is a step-by-step guide to adding a core filer to your cluster. Before you begin, you will need the following pieces of information: • If you're adding an NFS core filer, you will need the network name or IP (the fully-qualified domain name is preferred) of the filer you are adding. • If you're adding a cloud core filer, you will need an existing credential set for your service provider (or the access key to create a new one), and any encryption details if necessary. You will also need the name and hostname for your existing cloud storage container.' The form section, titled 'Core Filer Details', has two radio buttons for 'Core filer type': 'NFS' (selected) and 'Cloud'. Below are two text input fields: 'Core filer name' with the value 'new\_nfs' and 'Core filer network name/IP' with the value 'nfs.company.com'. At the bottom, there is a 'Page 1 of 3' indicator and three buttons: 'Back', 'Next', and 'Add Filer'.

4. Enter the desired core filer name.
5. Enter the IP address, network name, or the fully-qualified domain name (FQDN) of the core filer.

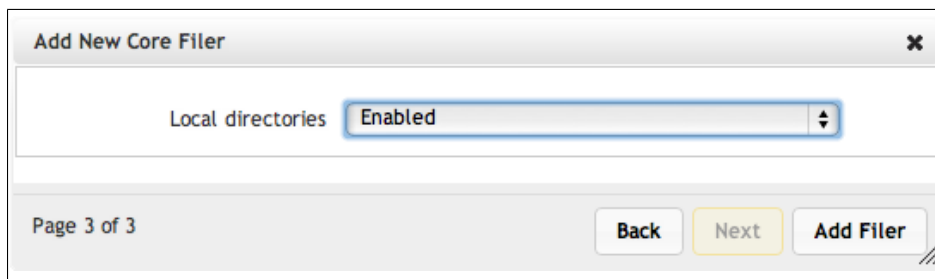
6. Click **Next**. The next wizard page appears, allowing you to set the caching mode for the filer.

The cache policy specifies how the Avere system handles file-operation requests between clients and core filers. For more information on caching, refer to Chapter 6, *Setting the Cache Policy* on page 84. You can later change the cache parameters from the **Settings > Core Filer > Cache Policy** page.

7. Select the caching mode from the **Caching mode** drop-down list.
  - If you select **Read/Write**, the **Core filer verification** drop-down list will only select **Never**. If you need core filer verification in read/write mode, you can change this later on the **Settings > Core Filer > Cache Policy** page.
  - If you select **Read** you will not be able to set a **Maximum writeback delay** time..



8. Click **Next**. The next wizard page appears, allowing you to enable local directories. For more information about local directories, refer to Section 6.5.2, “Enabling Local Directories” on page 94.



You can later change the local directory settings from the **Settings > Core Filer > Cache Policy** page.

9. Click **Add Filer**.

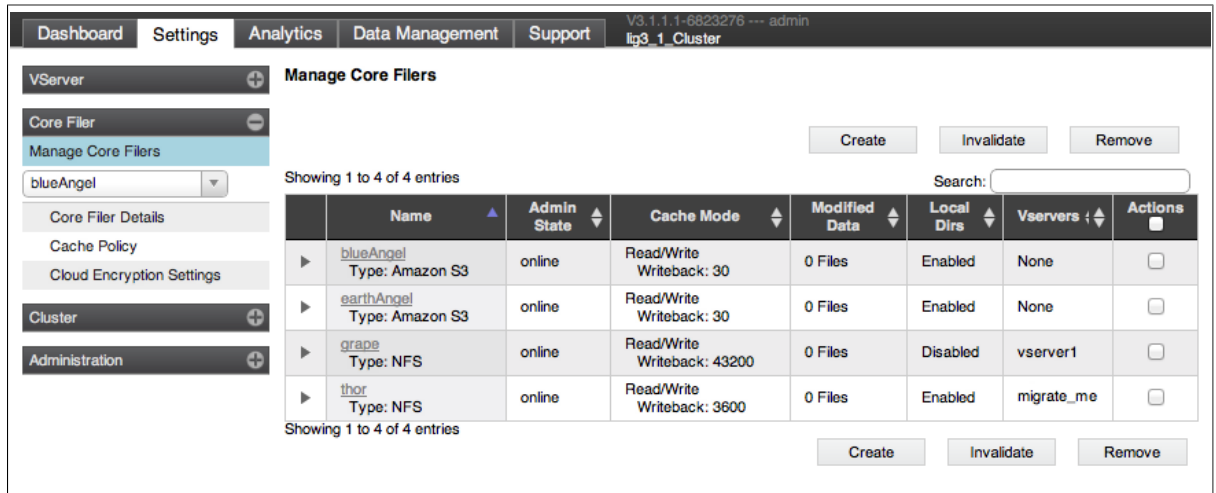
The Avere Control Panel indicates that the core filer has been added to the cluster, and the **Dashboard** tab might display alerts during the process.

## 3.2. Creating a Cloud Core Filer

If you are adding a cloud core filer, you will need either an existing credential for your service provider, or the access and private keys to create a new one. (You will be able to create a credential while configuring the cloud filer, or you can create one separately using the **Settings > Cluster > Cloud Credentials** page, as described in Section 3.3.2, “Managing Cloud Credentials” on page 63.)

➤ To add a new cloud core filer to the cluster:

1. Navigate to the **Settings > Core Filer > Manage Core Filers** page.



The screenshot shows the 'Manage Core Filers' page in the Veeam Backup & Replication console. The left sidebar contains navigation links: VServer, Core Filer (selected), Manage Core Filers, Core Filer Details, Cache Policy, Cloud Encryption Settings, Cluster, and Administration. The main area displays a table of core filers with columns: Name, Admin State, Cache Mode, Modified Data, Local Dirs, Vservers, and Actions. There are buttons for 'Create', 'Invalidate', and 'Remove' at the top and bottom of the table.

Name	Admin State	Cache Mode	Modified Data	Local Dirs	Vservers	Actions
blueAngel Type: Amazon S3	online	Read/Write Writeback: 30	0 Files	Enabled	None	<input type="checkbox"/>
earthAngel Type: Amazon S3	online	Read/Write Writeback: 30	0 Files	Enabled	None	<input type="checkbox"/>
grape Type: NFS	online	Read/Write Writeback: 43200	0 Files	Disabled	vserver1	<input type="checkbox"/>
thor Type: NFS	online	Read/Write Writeback: 3600	0 Files	Enabled	migrate_me	<input type="checkbox"/>

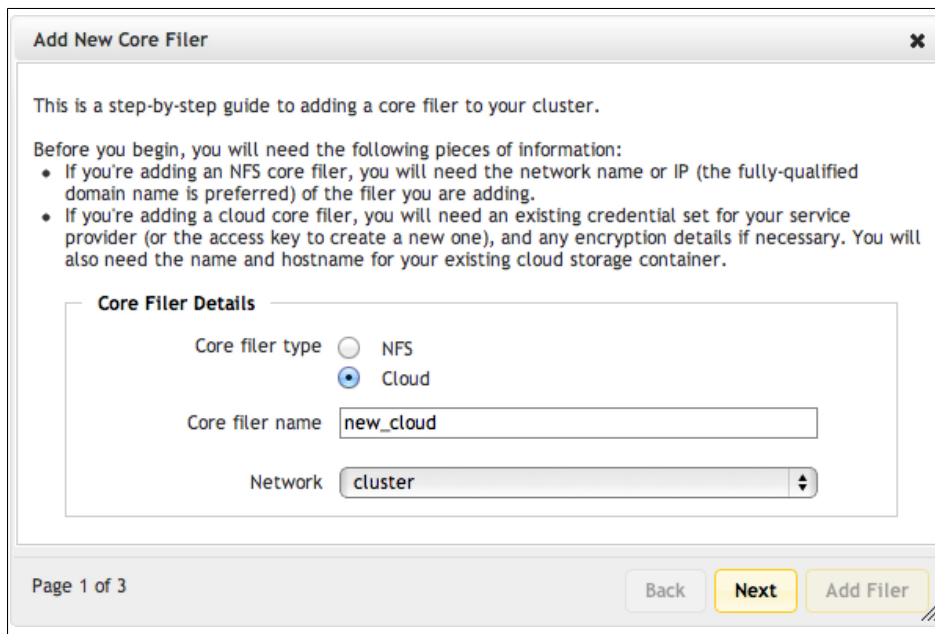
2. Click **Create**. The **Add New Core Filer** wizard starts.



### Note

Using the browser's back button will exit the wizard. Use the **Back** button on the wizard instead.

3. Select the core filer type, in this case, **Cloud**.



The screenshot shows the 'Add New Core Filer' wizard, Step 1 of 3. The title bar says 'Add New Core Filer'. The main content area contains the following text:

This is a step-by-step guide to adding a core filer to your cluster.

Before you begin, you will need the following pieces of information:

- If you're adding an NFS core filer, you will need the network name or IP (the fully-qualified domain name is preferred) of the filer you are adding.
- If you're adding a cloud core filer, you will need an existing credential set for your service provider (or the access key to create a new one), and any encryption details if necessary. You will also need the name and hostname for your existing cloud storage container.

**Core Filer Details**

Core filer type: ☐ NFS ☒ Cloud

Core filer name:

Network:

Page 1 of 3

Buttons: Back, Next, Add Filer

4. Enter the desired core filer name, and select the network to which it will belong. If you have not created any additional networks, **cluster** will be the only one available. For more information on networks, refer to Chapter 5, *Advanced Networking and VLANs* on page 77.
5. Click **Next**. The next wizard page appears, allowing you to set the service type and credential information for the filer.

#### Amplidata S3 example for selecting a service type from the drop-down list

6. From the first drop-down list, select an existing service type.
7. From the next drop-down list, either select an existing credential, or you can create a new credential.  
If you create a new credential, or select an existing credential and want to create another one later, you can enter the parameters as described in Section 3.3.2, “Managing Cloud Credentials” on page 63.
8. Every file you upload to your cloud service provider is stored in a container, called a *bucket*, (or a *vault*). From the **Bucket contents** area, select **Empty** or **Existing Avere data**. On the next wizard page, you will be able to enter the bucket information.



#### Important

Each cloud filer can only be associated with one bucket, and vice versa. Thus:

- If you intend to use an existing bucket, and choose **Empty**, you will no longer be able to access the information in the existing bucket without removing and re-creating the filer.
  - If you intend to create a new bucket, and choose **Existing Avere data**, you will be asked for an encryption key later to access the bucket. If you are trying to access a bucket that did not previously exist, you will not be able to access it (there will be no encryption key), and you will have to re-create the cloud filer.
9. The **Use HTTPS** checkbox allows you to choose the secure (HTTPS) web protocol to communicate with your cloud provider, and is selected by default, with port 443 automatically entered in the **Port** field.
    - If you do *not* want to use HTTPS, but HTTP, deselect this option. Port 80 will be automatically entered in the **Port** field.

- If your cloud provider does not use the standard ports 443, enter the port or range of ports into the **Port** field after you make your protocol choice.
10. If you want to compress objects before they are written to the cloud core filer, select either **LZ4** or **LZ4HC** compression from the **Compression mode** drop-down list.



### Note

You can change compression options at any time; objects written to the bucket (or vault) after compression is enabled will be compressed.

11. Choose **Next**.

## 3.2.1. Bucket Settings

The next wizard page appears, allowing you to enter bucket information – either for an empty (new) bucket or an existing bucket. The information you need to enter will vary according to the service you are using.

The screenshot shows a window titled "Add New Core Filer" with a close button (X) in the top right corner. Inside the window, there is a section titled "Empty Bucket Details" with a dashed border. Within this section, there are three fields: "Bucket name" with the text "consumerItems", "Endpoint/Region" with a dropdown menu showing "US West (Northern California) Region", and "Encryption type" with a dropdown menu showing "AES-256". Below the "Empty Bucket Details" section, there is a "Page 3 of 3" indicator and three buttons: "Back", "Next" (highlighted in yellow), and "Add Filer".

Selecting an Amazon S3 region from the drop-down list.  
This is only available for the Amazon S3 service.

The screenshot shows a window titled "Add New Core Filer" with a close button (X) in the top right corner. Inside the window, there is a section titled "Empty Bucket Details" with a dashed border. Within this section, there are three fields: "Bucket name" with the text "consumerItems", "Network hostname" with the text "cleversafe-company.com", and "Encryption type" with a dropdown menu showing "AES-256". Below the "Empty Bucket Details" section, there is a "Page 3 of 3" indicator and three buttons: "Back", "Next" (highlighted in yellow), and "Add Filer".

Entering a Cleversafe S3 network hostname.  
This is currently available for non-Amaon services.

➤ To configure the bucket:

1. Do one of the following:

- For an existing bucket, enter the bucket name.
- For an empty bucket, enter the name of the bucket that will be created on the cloud provider's network in the **Bucket name** field.



### Important

The bucket name cannot be changed later.

2. Do one of the following:

- For Amazon AWS services, select the endpoint from the **Endpoint/Region** drop-down list. Refer to the [Amazon Web Site](#) for appropriate Amazon endpoints.

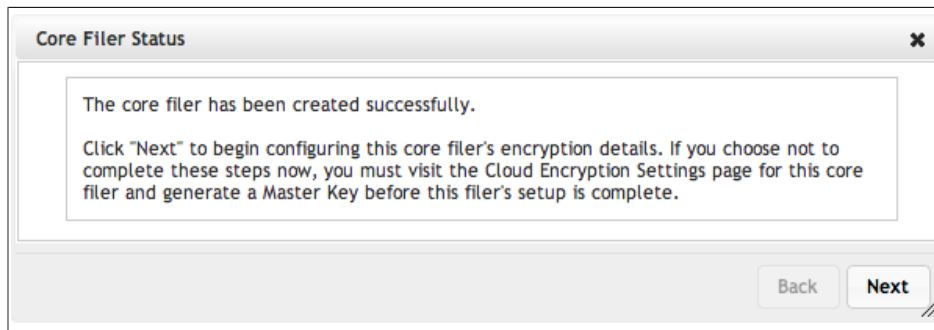


### Note

The Amazon **US East (Northern Virginia)**, also known as **US Standard**, region is not supported in Avere OS.

- For other services, enter the fully-qualified domain name (FQDN) of the bucket's endpoint in the **Network hostname** field.
3. Select the **Encryption type** from the drop-down menu. Currently, only AES-256 is supported.
- *For an empty (new) bucket, you cannot change the encryption settings once the cloud core filer has been created.*
  - For a bucket with existing Avere data, choose the same encryption type that was previously used.
4. Click **Add Filer**. If you are *not* encrypting the data in your bucket, you are now finished.

A window appears informing you that the core filer has been created. If you have chosen an encryption setting, choose **Next** to configure the core filer's encryption details.



### 3.2.2. Cloud Encryption Settings

The process of configuring bucket encryption will vary depending on whether the bucket is empty or not:

- **For an existing bucket** – you will first need to upload the old key file.
- **For all buckets** – you will need to generate a new (master) recovery key, download the new recovery key file, and then upload it for activation.

You must configure the encryption details before you can use this core filer, either at this point in the cloud filer setup, or later from the **Settings > Core Filer > Cloud Encryption Settings** page, as shown on Section 3.2.3, “Creating a New Key File After the Filer is Created” on page 59.

➤ To create and activate a new key recovery file:

1. From the **Core Filer Status** window, click **Next**.



#### Caution

For a bucket with existing data, if you do not upload the old key recovery file and generate and create a new key file (for example, if you close the window), you will not be able to access that data later unless you delete and re-create the core filer. Any data changes made to the bucket between those times will be lost.



#### Note

If you are configuring a bucket with existing data, **rather than an empty bucket**, the following window appears; otherwise, the process skips to Step 3.

2. Select the currently used key recovery file, enter the passphrase for the file, and click **Next**.

3. Enter a passphrase for the new key file, and click **Next** to save the file.



### Important

Both the passphrase and the file must be kept somewhere safe.

The screenshot shows a dialog box titled "Generate Cloud Key Recovery File" with a close button (X) in the top right corner. The main content area is titled "Generate and Download Key Recovery File". It contains two text input fields: "Key recovery passphrase" and "Confirm key recovery passphrase", both filled with dots. Below the fields is a paragraph of text: "When you click 'Next', you will be prompted to download a key recovery file. Save this file in a secure location; you will need it in the next step." At the bottom left, it says "Page 2 of 4". At the bottom right, there is a "Next" button and a double-slash icon (//).



### Caution

If you lose the key file, or the passphrase, you will *never* be able to access the data later.

4. Select the key recovery file from your saved location and click **Next**.

The screenshot shows a dialog box titled "Generate Cloud Key Recovery File" with a close button (X) in the top right corner. The main content area is titled "Upload and Activate Key Recovery File". It contains a paragraph of text: "Select the key recovery file that you saved in the last step and upload it to activate the master encryption key." Below this is a label "Upload key recovery file" followed by a "Choose File" button and a text field containing the filename "cloud\_test\_CloudM...1000000000000000002". At the bottom left, it says "Page 3 of 4". At the bottom right, there is a "Next" button and a double-slash icon (//).

5. The Avere OS displays a message that the key file has been uploaded. Choose **Next** to complete the process.

### 3.2.3. Creating a New Key File After the Filer is Created

If you want to re-download a recovery file for the cloud filer (for example, if you would like to have a copy for backup), select **Redownload Recovery File**.

The screenshot shows the 'Encryption Key Settings – emptyAngel' page. The left sidebar contains navigation links: VServer, Core Filer, Manage Core Filers (with a dropdown showing 'emptyAngel'), Core Filer Details, Cache Policy, Cloud Encryption Settings (highlighted), Cluster, and Administration. The main content area has a breadcrumb trail: 'Core Filer List > Filer Details > Choose Filer: emptyAngel'. Below this, key information is displayed: Access key ID: N/A, Encryption key activation date: N/A, Pending Access Key ID: 10000000000000002, and Recovery file signature: 339328a52c£733b0541a2aae888fd2b899741771. A 'Redownload Recovery File' button is present. The 'Generate a New Master key' section includes 'Step one' instructions, a warning about losing the recovery file, and input fields for 'Key recovery passphrase' and 'Confirm key recovery passphrase'. A 'Generate Key and Download File' button is at the bottom of step one. 'Step two' instructions follow, with an 'Upload key recovery file' section containing a 'Choose File' button and a text field showing 'emptyAnge...00000002'. An 'Activate Key' button is at the bottom right of step two.

➤ If you want to create a **new** key recovery file:

1. Enter a passphrase for the new key file that you will be able to remember.
2. Select **Generate Key and Download File**. The key recovery file will be downloaded to your browser's default location, depending on your browser's settings.
3. Select the key recovery file from your saved location.
4. Select **Activate Key**.

## 3.3. Maintaining Core Filers

From the **Settings > Core Filer > Manage Core Filers** page, you can click on the triangle to the left of each listed core filer to see its details.

You can change some of these details from the **Settings > Core Filer > Core Filer Details** page.

The screenshot shows the 'Manage Core Filers' page in the Avere OS interface. The page has a sidebar with navigation options: VServer, Core Filer, Manage Core Filers, Core Filer Details, Cache Policy, Cloud Encryption Settings, Cluster, and Administration. The main content area displays a table of core filers. The first core filer, 'blueAngel' (Type: Amazon S3), is expanded to show its details. The details include: Name (blueAngel), State (online), Network (cluster), Bucket Name (ccork0), Cloud credential (encryptionCred), Encryption mode (DISABLED), HTTPS (yes), Network Name (s3-us-west-1.amazonaws.com), and Vservers (None). The other three core filers are 'earthAngel' (Type: Amazon S3), 'grapnel' (Type: NFS), and 'thor' (Type: NFS). The 'earthAngel' core filer has a checkmark in the Actions column. The page also includes a search bar and buttons for Create, Invalidate, and Remove.

Name	Admin State	Cache Mode	Modified Data	Local Dirs	Vservers	Actions
blueAngel Type: Amazon S3	online	Read/Write Writeback: 30	0 Files	Enabled	None	<input type="checkbox"/>
earthAngel Type: Amazon S3	online	Read/Write Writeback: 30	0 Files	Enabled	None	<input checked="" type="checkbox"/>
grapnel Type: NFS	online	Read/Write Writeback: 43200	0 Files	Disabled	vserver1	<input type="checkbox"/>
thor Type: NFS	online	Read/Write Writeback: 3600	0 Files	Enabled	migrate_me	<input type="checkbox"/>

In addition to creating a core filer, you can invalidate or remove each listed core filer.

- **Remove** – Removes the core filer and all of its settings from the cluster. The Avere OS prompts you for confirmation before completing the operation.
- **Invalidate** – Invalidates all client data held by the Avere cluster for a given core filer. The Avere Control Panel prompts you for confirmation before completing the operation.

### 3.3.1. Changing Core Filer Details

Use the **Settings > Core Filer > Core Filer Details** page, or click on the name of core filer, to view the details of the core filer, and to change some of the details.

You can select the core filer using the drop-down list under the **Core Filer** menu, or the drop-down list at the top of the **Core Filer Details** page.

The core filer details will vary according to whether it is an NFS core filer or a cloud core filer.

#### 3.3.1.1. NFS Filer Details

- **Core filer name** – You can change the visible name of the core filer here.
- **Core filer network name/IP** – You can change the information about the network name or IP address, if, for example, you have moved the core filer to a different network.

- **Snapshot directory** – Many NFS core filers provide “snapshots”; that is, periodic, read-only copies of data, which are accessible to clients through a specially named directory. The Avere OS always treats snapshot directories as read-only, no matter what the caching policy for the core filer is set to. If your core filer supports snapshots, specify the name of the snapshot directory that it uses. The default is `.snapshot`.
- **Network** – The network that the cluster uses to communicate with the core filer.
- **Enable Bandwidth Control** – Select this option to limit the bandwidth consumed between the cluster and the core filer, or to allow loads to be balanced more evenly across the cluster. For more information, refer to Section 6.4, “Controlling Write Bandwidth” on page 91.
- **Enable WAN Optimization** – If the core filer is in a wide-area network (WAN) environment, select this option from the **NFS Core Filer Details** to optimize communication. This does not help, and may degrade, performance in a local-area network (LAN) environment.
- **Enable Kerberos** – Select this option to enable Kerberos authentication. If you have not yet installed a keytab file, you can still select this option, but will need to access the **Kerberos Configuration** page to install one. For more information, refer to Section 2.8.3, “Enabling Kerberos Authentication” on page 43.
- **Hardware Details (Optional)** – Information about the core filer’s hardware, if it has been entered.
  - Manufacturer
  - Model
  - Description
- **Details**
  - Administrative state – The state of the core filer; for example, “online”, “removed”, or “flushing”.
  - Associated vservers – Any vservers associated with the core filer. For more information, refer to Chapter 4, *Virtual Servers Used for Client Access* on page 64.
  - Core filer management IPs – Any management IP addresses that the core filer uses to communicate with the cluster. These cannot be changed.
- **Cache Policy Details** – You need to go to the **Cache Policy** page to change these details; a link to this page is provided in this area.
  - Caching mode – Either read mode or read/write mode.
  - Core filer verification – How often the FXT Series system verifies cached file and directory attributes against the core filer’s file and directory attributes.
  - Maximum writeback delay – the maximum amount of time that changed data is stored on the cluster before being committed to the core filer.
  - Local directories – Whether local directories are enabled or not.

For more information, refer to Chapter 6, *Setting the Cache Policy* on page 84.

### 3.3.1.2. Cloud Filer Details

- **Core filer name** – You can change the visible name of the core filer here.
- **Core filer network name/IP** – You can change the information about the network name here. For a cloud filer, the IP address is not applicable.
- **Cloud credential** – Which cloud credential the core filer is using. Use the drop-down list to choose from any other existing credentials. For more information, and to create a new credential, refer to Section 3.3.2, “Managing Cloud Credentials” on page 63.
- **Network** – The network that the cluster uses to communicate with the core filer.

- **Enable Bandwidth Control** – Select this option to limit the bandwidth consumed between the cluster and the core filer, or to allow loads to be balanced more evenly across the cluster. For more information, refer to Section 6.4, “Controlling Write Bandwidth” on page 91.
- **Compression mode** – Any compression mode (LZW, LZW4, or None) can be selected here.
- **Enable HTTPS** – Select this option if you want to use the secure (HTTPS) web protocol to communicate with your cloud provider. If your cloud provider does not use the standard port 443, this option will be disabled.
- **Details**
  - Administrative state – The state of the core filer; for example, “removed”, or “flushing”.
  - Associated vservers – Any vservers associated with the core filer. For more information, refer to Chapter 4, *Virtual Servers Used for Client Access* on page 64.
  - Core filer management IPs – Any management IP addresses that the core filer uses to communicate with the cluster. These cannot be changed.
- **Cloud Details**
  - Cloud type – Currently, the only available cloud type is S3.
  - Bucket name – The name of the bucket used by the cloud filer. This cannot be changed.
  - Encryption type – The encryption type, if any, used on the cloud filer. Currently, only AES-256 is supported. This cannot be changed.
- **Cache Policy Details** – You need to go to the **Cache Policy** page to change these details; a link to this page is provided in this area.
  - Caching mode – Either read read/write mode.
  - Maximum writeback delay – the maximum amount of time that changed data is stored on the cluster before being committed to the core filer.

For more information, refer to Chapter 6, *Setting the Cache Policy* on page 84.

### 3.3.2. Managing Cloud Credentials

From the **Settings > Cluster > Cloud Credentials** page, you can add, remove, and modify credentials. (You can also add a credential when you create a cloud core filer.)

Name	Type	Access Key	Associated Filers	Note	Actions
updatedCred	s3	AKIAI6IWCKSFC73REI7Q		Updated information	<button>Remove</button> <button>Modify</button>
encryptionCred	s3	AKIAI6IWCKSFC73REI7Q	blueAngel, earthAngel	Credential for help	<button>Remove</button> <button>Modify</button>

Add Credential

**Create Credential**  
Credential name:   
Service type:   
Access key:   
Private key:   
Note:   
Submit

#### 3.3.2.1. Adding a Credential

- To add a new credential:
  1. Click **Add Credential**. The **Create Credential** panel appears.
  2. Enter a descriptive name in the **Credential name** field.
  3. Enter the access key and the private key generated or sent to you by your cloud provider.
  4. Enter any descriptive notes you want for the credential.
  5. Click **Submit**.

The credential appears in the list on the **Cloud Credential** page.

#### 3.3.2.2. Removing a Credential

- Click **Remove** next to an unused credential to remove it and all of its information from the cluster. You can only remove a credential if it is not being used.  
*The Avere OS does not prompt you for confirmation before completing the operation.*

#### 3.3.2.3. Modifying a Credential

- Click **Modify** next to a credential to change any of the following information:
  - **Credential name**
  - **Access key** – If you change this, you must also change the private key.
  - **Private key**
  - **Note**

## Chapter 4. Virtual Servers Used for Client Access

The Avere system provides an abstraction layer between your clients and one or more NAS servers (core filers). Clients do not connect directly to the core filers, but connect to one or more virtual servers (*vservers*) on the cluster; each virtual server provides access to data residing on one or more core filers. This architecture enables many combinations of vservers and core filers, including a global namespace (GNS) that can be used across several core filers.

You can use the Avere Control Panel to configure, monitor, and maintain each individual virtual server and core filer, as described in Section 9.4, “VServers Tab” on page 127 and Section 9.5, “Core Filers Tab” on page 128.

### 4.1. Virtual Servers and Namespaces

A vserver is a virtual entity on the FXT Series cluster that connects clients to data served by the cluster and the core filers. Each vserver interacts with one or more core filers, allowing clients to access data from multiple core filers through a single vserver; allowing different sets of clients to access data from a single core filer through multiple vservers; and allowing one or more sets of clients to access data from one or more core filers through one or more vservers. The one-to-many mapping of vservers to core filers, instead of one vserver to a core filer, is what enables you to create a global namespace.

A *namespace* is a set of data that clients can access through the single mountpoint of a vserver, allowing information from several filesystems to look like one file structure. You can create two types of namespace: *simple* and *global*. In a simple namespace, one vserver provides data access for one core filer, and the namespace seen by clients is a replica of the namespace on the core filer. In a global namespace, a vserver provides data access for one or more core filers, and the top-level namespace seen by clients is a logical construction in the vserver. For more information, refer to Section 4.3, “Creating and Maintaining a Global Namespace” on page 72.

All virtual servers on a cluster share the cluster’s resources, including storage, network, memory, and CPU. The client-facing IP addresses of each virtual server are distributed evenly across all nodes in the cluster.



#### Note

There are no tunable quality-of-service (QoS) parameters for virtual servers.

If the vserver associated with a particular core filer or core filers becomes unhealthy (for example, if a core filer fails or the network connection between a core filer and the cluster is dropped), the cluster continues to provide service for other vservers, and automatically constrains resource allocation to the unhealthy vserver until it returns to a healthy state.

The maximum number of vservers permitted on an FXT cluster is 24. However, if the cluster includes close to the maximum number of core filers (50), keep the number of vservers at 15 or fewer.

## 4.2. Managing Virtual Servers (Settings Tab | VServer)



### Note

As described in Section 9.1, “Overview of the Dashboard” on page 117, the Dashboard tab’s status bar provides a high-level overview of virtual servers.

You can perform the following actions by navigating to the **Settings > VServer > Manage VServers** page:

The screenshot shows the 'Manage VServers' page. On the left is a sidebar with navigation links: VServer, Manage VServers (selected), migrate\_me, Vserver Details, Client Facing Network, Namespace, Export Policies, Export Rules, NFS, CIFS, CIFS Shares, Core File, Cluster, and Administration. The main area has a 'Manage VServers' header with 'Create', 'Suspend', and 'Remove' buttons. Below the header is a table with columns: Name, Admin State, Namespace Type, CIFS, Netgroup Update, Filers, Hot Clients, and Actions. The table shows two entries: 'migrate\_me' (online, Global, No, 12/18/2013 11:35:01 AM, thor, No) and 'vserver1' (online, Simple, No, 12/18/2013 11:02:58 AM, grapnel, No). The 'vserver1' entry is expanded, showing details: Name vserver1, State Suspended, Client Facing IPs 10.1.20.135 - 10.1.20.138 (4), and buttons for 'Unsuspend', 'Create', 'Suspend', and 'Remove'.

Click on the ► symbol next to a vserver to display some of its details.

Enter a search term in the field above the **Manage VServers** table to filter the display, as described in Section 1.1.11, “ Search Functionality” on page 3.

- **Create** – Allows you to enter the information for a new vserver, as described in Section 4.2.1, “Creating a Virtual Server” on page 66.
- **Suspend** – Suspends client access to the virtual server, which allows all current operations to complete and blocks all new client operations. You can suspend more than one vserver at one time.

➤ To suspend a vserver:

- In the **Actions** column, check the box for the vserver.
- Click the **Suspend** button. The Avere OS asks if you are sure you want to suspend the vserver(s).
- Choose **Yes** (or choose **No** to cancel).

The vserver is suspended, and the **Unsuspend** button appears in the details under that vserver on the **Manage VServers** page.

- **Remove** – Removes the virtual server and all of its settings. This is a safe operation only if all clients have been disconnected from the virtual server and do not need to access the virtual server’s associated core filer.

➤ To remove a vserver:

- In the **Actions** column, check the box for the vserver.
- Click the **Remove** button. The Avere OS asks if you are sure you want to remove the vserver.
- Choose **Yes** (or choose **No** to cancel).

The vserver is removed from the list.



### Note

The **Remove** operation does not necessarily work on vservers with active core filers. In such cases, contact Avere Global Services for assistance.

- **Unsuspend** – Unsuspends a suspended virtual server. This option only appears if you have previously suspended the vserver. You can unsuspend more than one vserver at one time.

➤ To unsuspend a vserver:

- In the **Actions** column, check the box for the vserver.
- Click the **Unsuspend** button.

The vserver is returned to an online state, and the **Unsuspend** button for that vserver disappears.

## 4.2.1. Creating a Virtual Server

The first step in configuring client access through the Avere cluster is creating a vserver. You will need the following information:

- For a simple vserver, the network name or IP address of a core filer
- The name you want to use for the vserver
- The range of client-facing IP addresses



### Important

The set of client-facing IP addresses for each virtual server must have static IP addresses manually assigned (that is, not assigned by DHCP) and must be in a contiguous range.

➤ To create a new virtual server:

1. Navigate to the **Settings > Vserver > Manage VServers** page.
2. Choose **Create**. The **Add New VServer** window appears.

3. Select the type of namespace for the vserver, one of the following:
  - **Global Namespace** for a global namespace involving one or more core filers that are or are not configured on the cluster.
  - **Simple Namespace** for a core filer that is already configured on the cluster.

4. In the **VServer Name** field, enter the name for the new virtual server.
  - For a simple namespace, this is typically the abbreviated name of the core filer (for example, `nas1`).
  - For a global namespace, it is typically a name that indicates that the vservers exports a global namespace (for example, `GNS` or `netfiles`).

In either case, the name must consist of only alphanumeric characters (no punctuation, spaces, or special characters) and have a maximum length of 255 characters.

5. *If you chose **Simple Namespace***, the next field will be labeled **Core filer name**. Enter the fully qualified domain name of the new virtual server's core filer; for example, `nas1.example.com`, not `nas1` or `127.98.12.1`.

#### 4.2.1.1. Adding an IP Address Range

You can specify the range of client-facing IP addresses that clients use to access the virtual server. Client-facing IP addresses must be manually assigned (that is, not assigned by DHCP) and must be in a contiguous range.



#### Note

You can specify more than one range of client-facing IP addresses by enabling advanced networking, as described in Section 5.1, “Enabling Advanced Networking” on page 78.

These IP addresses will be associated with either the default VLAN, or any client-role VLANs that you created in Section 5.2, “Creating a VLAN” on page 78.

6. Enter the VLAN's subnet mask in the **Subnet mask** field.
7. In the **First IP** field, enter the first address in the assigned range of client-facing IP addresses.
8. In the **Last IP** field, enter the last address in the assigned range of client-facing IP addresses.
9. If advanced networking is enabled, there is a drop-down list named **VLAN** at the bottom of the panel. Choose the VLAN that you want the vservers to use for client-facing interfaces; that is, that will connect clients and vservers. Possible values include **Default** and any client-role VLANs that you have created.

#### 4.2.1.2. Add the VServer

10. Click the **Add VServer** button.

The Avere OS indicates that the virtual server has been created, and the **Dashboard** tab may display alerts as the new virtual server is added to the cluster. You will now need to configure the vservers.

## 4.2.2. Configuring and Modifying a Virtual Server (VServer Details Page)

After you create a vserver and an associated core filer, as described in the previous section and in Chapter 3, *Configuring a Core Filer (Settings Tab | Core Filer)* on page 50, you need to configure client access on the vserver. Additional configuration information includes the following:

- For a global-namespace vserver, the logical structure of the namespace, as described in Section 4.3, “Creating and Maintaining a Global Namespace” on page 72.
- Optionally, NFS export policies and rules, as described in Section 2.5, “Managing Exports (Settings Tab | VServer)” on page 28.
- Optionally, CIFS configuration information, as described in Chapter 7, *Configuring CIFS Access* on page 95.

You can change and view the name and hot client collection for a vserver using the **Settings > Vserver > Vserver Details** page. You can also get to this page by clicking on the vserver name on the **Manage Vservers** page.

The screenshot shows the 'VServer Details – vserver1' page. The top navigation bar includes 'Dashboard', 'Settings', 'Analytics', 'Data Management', and 'Support'. The 'Settings' tab is active. On the left, a sidebar menu shows 'VServer' (selected), 'Manage Vservers', 'Vserver Details', 'Client Facing Network', 'Namespace', 'Export Policies', 'Export Rules', 'NFS', 'CIFS', 'CIFS Shares', 'Core Filer', 'Cluster', and 'Administration'. The main content area has a title 'VServer Details – vserver1' and a dropdown menu 'Vserver List > Choose VServer: vserver1'. Below this is the 'Edit Vserver vserver1' section with fields for 'VServer name' (vserver1), 'Collect hot client information' (checkbox), 'Hot client limit' (10), and 'Hot client period' (60). There are 'Revert' and 'Submit' buttons. At the bottom is a 'Details' section showing 'Administrative state: online', 'Operational state: up', and 'Core filers: grape'.

Select the vserver using the drop-down list under the **Vserver** menu, or the drop-down list at the top of the **Vserver Details** page.

### 4.2.2.1. Renaming a VServer

➤ To rename a vserver:

1. In the **Name** field, enter the new administrative name for the vserver.
2. Click the **Rename** button.

#### 4.2.2.2. Enabling Hot Client Collection

A *hot client* is a client that generates a disproportionately high amount of demand on the cluster relative to other clients.

Enabling hot-client statistic collection reduces FXT node performance.

➤ To enable hot client collection:

1. Select the **Collect hot client information** button.
2. Enter the **Hot client limit**, the maximum number of clients on the vserver for which you want hot client information. The default is 10. Information will be collected for this number of most active clients on the vserver.
3. Enter the **Hot client period**, the polling period for collecting information, in seconds. The default is 60 (1 minute). The Avere OS will check for hot client information once each polling period.
4. Choose **Submit** to accept the changes, or **Revert** to cancel out of making changes.



#### Note

You can also change hot file collection from the Avere Control Panel, as described in Section 9.7, “Clients Tab” on page 129.

### 4.2.3. Client-Facing Network Settings

You can add or change the range of client-facing IP addresses, and set a home node for any client-facing IP addresses, on the **Settings > VServer > Client Facing Network** page.

#### 4.2.3.1. Adding or Editing an IP Address Range

1. Click either the **Add New Range** button or the **Modify** button. The Avere OS displays a set of fields for the range.

The screenshot shows the Avere OS interface for managing client-facing network settings. The top navigation bar includes Dashboard, Settings, Analytics, Data Management, and Support. The left sidebar shows the VServer management menu, with 'Client Facing Network' selected. The main content area displays the 'global - Client Facing Network' settings. A table lists the current IP range (10.1.22.127 - 10.1.22.128) with a 'Modify' button. Below the table is the 'Modify IP Range' form, which contains the following fields:

Address Range	Subnet Mask	VLAN	Actions
10.1.22.127 - 10.1.22.128	255.255.224.0	default	<button>Remove</button> <button>Modify</button>

Add New Range

**Modify IP Range**

First IP:

Last IP:

Number of IPs in range:

Subnet Mask:

VLAN:

Submit

2. Enter or change the range information, described in Section 4.2.1.1, “Adding an IP Address Range” on page 67
3. Click the **Submit** button. The Avere OS warns you about possible client disruption and asks you for confirmation to proceed.
4. Click **OK**.

#### 4.2.3.2. Setting a Home Node

Client-facing IP addresses can move from node to node as a result of node failover or network conditions. You can assign home nodes to client-facing IP addresses so that, after a failover or outage, the IP addresses return to their home nodes to ensure optimal load distribution. The cluster does not assign home nodes to client-facing IP addresses by default; you must manually specify a home node for each client-facing IP address if you want to use this feature.



#### Note

If home nodes for the client-facing IP addresses are not configured, the cluster automatically balances the client-facing IP addresses across all healthy nodes.

➤ To set a home node:

1. Navigate to the **Settings > VServer Settings > Client Facing Network**.

The screenshot shows the 'global - Client Facing Network' configuration page. It includes a table of IP ranges and a 'Modify IP Range' form. Below these is the 'IP Address Home Nodes' table.

Address Range	Subnet Mask	VLAN	Actions
10.1.22.127 - 10.1.22.128	255.255.224.0	default	<button>Remove</button> <button>Modify</button>

Add New Range

**Modify IP Range**

First IP:

Last IP:

Number of IPs in range:

Subnet Mask:

**IP Address Home Nodes**

IP Address	Current Node	Home Node
10.1.22.127	LiGo	<input type="text" value="None"/>
10.1.22.128	Liz-Golf1	<input type="text" value="None"/>

Revert Unhome All Make Current Home Submit

2. If you have more than one vservers, choose the appropriate vservers from the drop-down list immediately under the **VServer** heading.
3. The **IP Address Home Nodes** panel lists each client-facing IP address, the node on which it is currently located, and its home node if one has been assigned.
4. For each IP address to which you want to assign a home node, select a node from the drop-down list in the table's **Home Node** column. Selections include each cluster node and **None**.
  - To make each IP address's current node its home node, click the **Make Current Home** button.
  - To reset all IP addresses to have no home node, click the **Unhome All** button.
5. After you have made the appropriate home-node assignments for all IP addresses, click **Submit**.

The Avere OS pops up a warning indicating that changing home-node assignments for client-facing IP addresses can affect the interfaces as the changes are made.

6. Click **OK** to proceed.

## 4.3. Creating and Maintaining a Global Namespace

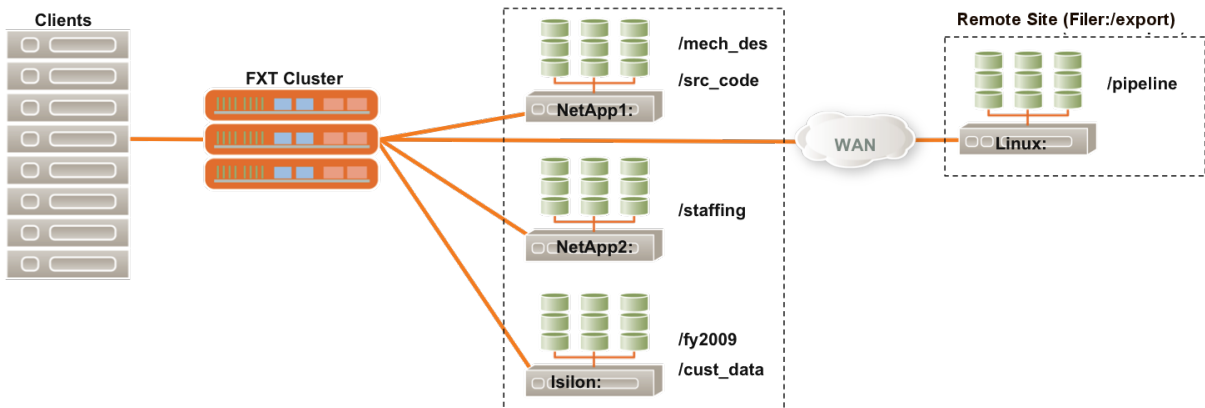
The vserver/core filer architecture gives you the ability to create a *global namespace*, also called a GNS; that is, client access through a single vserver to the data on multiple core filers. This section describes how to create and maintain a global namespace; it does not apply to simple namespaces (that is, client access through a single vserver to a single core filer).

### 4.3.1. Designing a Global Namespace

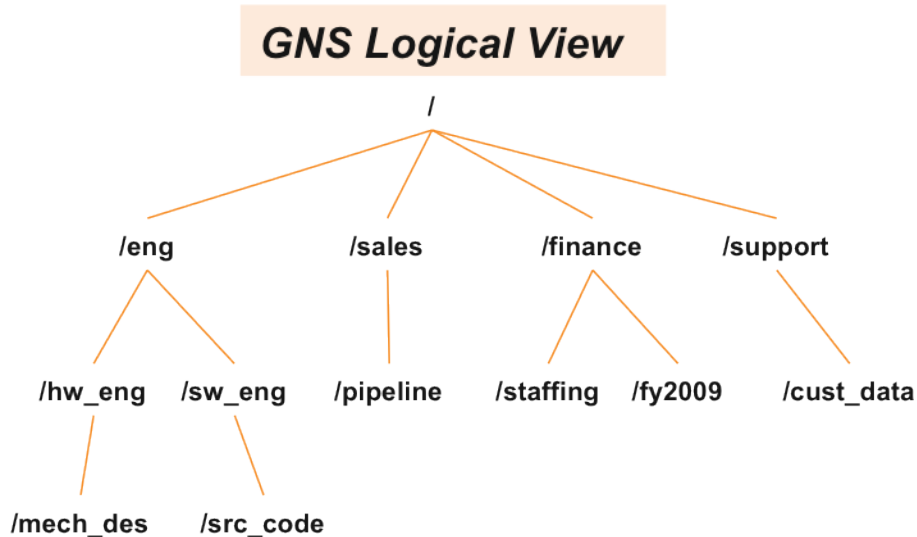
Before implementing a global namespace, it is important to consider the file and directory layout you want the global-namespace vserver to present to clients. Although the namespace can be modified after its creation, a best practice is to present clients with a finalized namespace at its introduction to prevent future confusion.

Physical and logical views of a sample global namespace are shown in the following figures.

#### Physical view of a sample global namespace



#### Logical view of a sample global namespace



### 4.3.2. Elements of a Global Namespace

The following terms are used to describe a global namespace on an Avere cluster:

- *Junction*—A link from one filesystem to a directory in another filesystem
- *Pseudo-filesystem*—A read-only filesystem contained by the global-namespace vserver that contains directory listings and junctions

### 4.3.3. Warnings and Limitations for Global Namespaces

The following warnings and limitations apply to the global-namespace implementation in Avere OS.

- Global namespaces can be constructed from core filers of the following types:
  - EMC Isilon OneFS
  - EMC VMXe
  - Hitachi Data Systems (HNAS) / BlueArc
  - Linux
  - NetApp Data ONTAP
  - Nexenta
  - Oracle Solaris ZFS

Other core filers might or might not work in a global namespace but have not yet been tested.

- A global namespace can be implemented only on a newly created vserver; an existing vserver with a simple namespace cannot be modified to have a global namespace.
- A global-namespace vserver does not have the same level of fault isolation as a single-namespace vserver. Clients requests that access a junction in the top-level pseudo-filesystem can return cached attributes without notification if the core filer is not reachable.
- Removing a global-namespace vserver does not remove any of its associated core filers from the cluster.

### 4.3.4. Creating and Modifying a Global Namespace



#### Note

Remember that you will need to log into your core filer and make sure the directories you want to use in your namespace exist in the correct locations, and that access permissions are appropriately set.

#### Creating a GNS

➤ To create a global namespace:

1. Create one or more core filers as described in Chapter 3, *Configuring a Core Filer (Settings Tab | Core Filer)* on page 50.
2. Create a virtual server with the type **Global Namespace**, as described in Section 4.2, “Managing Virtual Servers (Settings Tab | VServer)” on page 65.
3. Navigate to the **Settings > VServer > Namespace** page.
4. Add one or more junctions, as described in Section 4.3.4.1, “Adding a Junction to a Global Namespace” on page 74.

## Modifying a GNS

- To modify a global namespace:
  1. Navigate to **Settings > VServers > Namespace**, and select the vserver from the drop-down menu.
  2. Do one of the following:
    - Add an additional junction, as described in Section 4.3.4.1, “Adding a Junction to a Global Namespace” on page 74.
    - Modifying an existing junction, as described in Section 4.3.4.2, “Modifying an Existing Junction” on page 76. (Junction names can only be changed by re-creating the junction.)
    - Deleting an existing junction, as described in Section 4.3.4.3, “Deleting an Existing Junction” on page 76.

### 4.3.4.1. Adding a Junction to a Global Namespace

- To add a junction to a new or existing GNS:
  1. Navigate to **Settings > VServers > Namespace**, and select the vserver from the drop-down menu.
  2. Click on the **Add New Junction** button.

The screenshot shows the LIGO Cluster management interface. The top navigation bar includes Dashboard, Settings, Analytics, Data Management, and Support. The left sidebar lists various configuration options: VServer, Core File, Cluster, and Administration. The main content area is titled 'Namespace' and shows a table of existing namespaces. Below the table is a form to 'Add A New Junction'.

Namespace Path	Core Filer	Export	Subdir	Actions
/lbarrow	thor	/vol0/lbarrow	source	Delete Modify

**Add A New Junction**

Namespace path:

Core filer admin name:

Core filer export:

Export subdirectory:

Advanced: ☒

CIFS access control:

Core filer share name:

Core filer share subdir:

3. In the **Namespace Path** field, enter the path for the new junction. This is the path that clients see when they access the global-namespace vserver.



#### Important

The value entered into this field must be an absolute path; that is, it must start at root (/) of the global namespace (pseudo-filesystem) and work its way down. For example, if you want to map a user directory named `seamus` to a top-level directory named `usr`, you must enter the path `/usr/seamus` into this field.

4. From the **Core filer admin name** drop-down list, choose the core filer containing the export to be exposed to the namespace.
5. From the **Core filer export** drop-down list, choose the NFS export to be exposed to clients.



### Important

Ensure that the NFS export has all appropriate policies and rules in place to provide or deny access as necessary to various clients and users. Refer to Section 2.5, “Managing Exports (Settings Tab | VServer)” on page 28 for details.

6. Optionally, in the **Export subdirectory** field, enter the name of a subdirectory of the **Core filer export**.

If you enter a value in this field, the namespace path will point to this subdirectory instead of to the core filer export directory. Because the export subdirectory is relative to the namespace path, do not enter a leading backslash (/).

7. Optionally, configure CIFS access for the junction.



### Note

To *configure* CIFS access, you must first *enable* CIFS access on the vservers, as described in Section 7.6, “Enabling and Configuring CIFS” on page 99.

➤ To configure CIFS access:

- a. Select the **Advanced** checkbox. The window expands to display fields for CIFS access.
- b. From the **CIFS access control** drop-down list, choose the access-control mechanism for CIFS clients on the junction. Possible values include **POSIX Mode Bits**, **NFSv4 ACLs**, and **CIFS ACLs**. Refer to Section 7.7, “Selecting an Access-Control Mechanism” on page 103 for information about CIFS access-control mechanisms.
- c. In the **Core filer share name** field, enter the name of a CIFS share through which CIFS clients can access the junction.
- d. Optionally, in the **Core filer share subdir** field, enter the name of a subdirectory of the core filer share.

If you enter a value in this field, the namespace path will point CIFS clients to this subdirectory instead of to the core filer share directory. Because the CIFS subdirectory is relative to the namespace path, do not enter a leading backslash (/).

8. Click the **Add Junction** button.

Repeat this procedure until all required NFS exports on all core filers are exposed on the desired namespace paths.

#### 4.3.4.2. Modifying an Existing Junction



##### Note

You cannot modify an existing junction's name. If you need to change the name of an existing junction, delete the original junction as described in Section 4.3.4.3, "Deleting an Existing Junction" on page 76 and recreate it with the new name and the same core filer and core filer export.

➤ To modify an existing junction:

1. Navigate to **Settings > VServers > Namespace**, and select the vserver from the drop-down menu.
2. In the table at the top of the page, locate the junction that you want to modify and click the **Modify** button in the **Actions** column of the junction's row.

The **Modify Junction** panel appears.

3. If you want to modify the junction's CIFS configuration, select the **Advanced** checkbox to display fields for CIFS access.
4. Make changes in the fields and drop-down lists in the same way as adding a junction, as described in Section 4.3.4.1, "Adding a Junction to a Global Namespace" on page 74.

Modifications can only be made on one junction at a time.

5. When you have finished the changes, click **Submit**.

#### 4.3.4.3. Deleting an Existing Junction

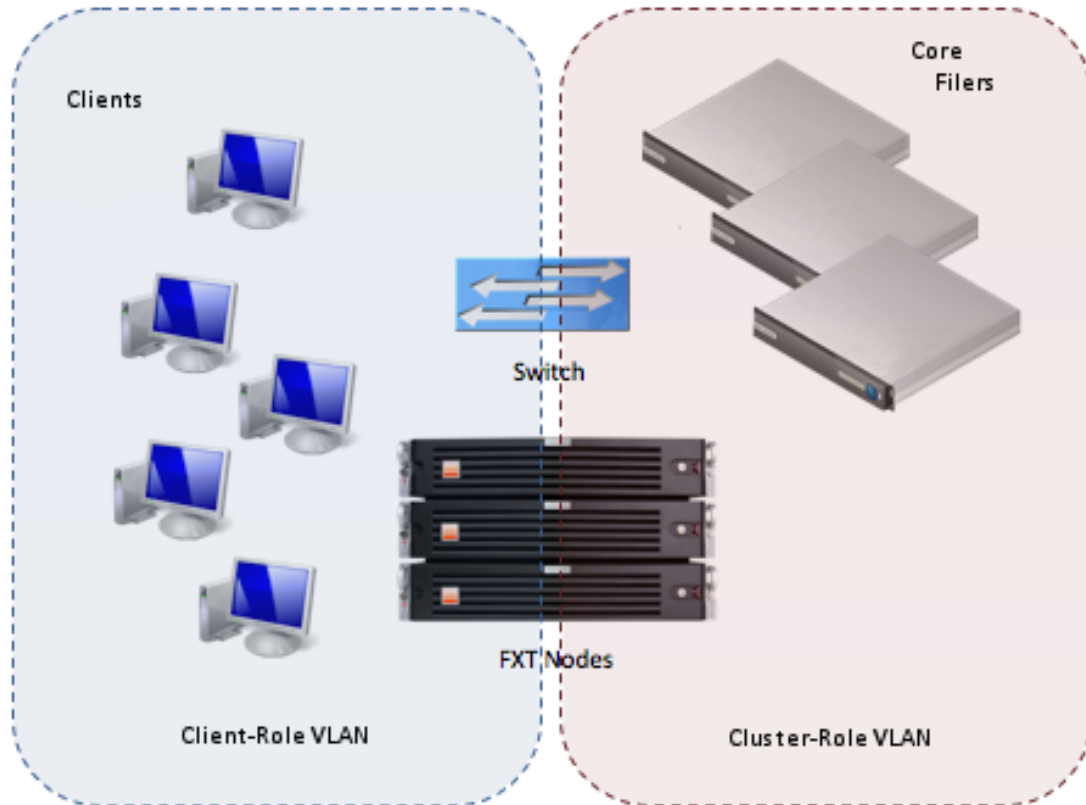
➤ To delete a junction:

1. Navigate to **Settings > VServers > Namespace**, and select the vserver from the drop-down menu.
2. In the table at the top of the page, locate the junction that you want to modify and click the **Delete** button in the **Actions** column of the junction's row.
3. A pop-up window appears, asking if you are sure you want to delete the junction. Choose **OK** to continue, or **Cancel** to, well, cancel.

The Avere OS indicates that the junction has been deleted.

## Chapter 5. Advanced Networking and VLANs

You can use *virtual local area networks* (VLANs) to connect components of the Avere cluster (core filers, vservers, and nodes). Note that VLAN access is *enabled* on a cluster-wide basis, not a per-vserver or per-core filer basis, although they can be *selected* on a per-server basis.



### Note

Advanced networking must be enabled to set up a VLAN configuration. This is done by default on clusters created with Avere OS 3.2 or later.

When advanced networking is enabled, whether by default or manually, a default VLAN is automatically created. You cannot remove the default VLAN, and can only modify its default router.

If the IP address range or VLAN number for any network interface is changed, the interface is temporarily unavailable while the changes are committed.

## 5.1. Enabling Advanced Networking

Advanced networking is enabled by default on clusters set up with Avere OS 3.2 or later. You cannot disable advanced networking after it has been enabled.

- To enable advanced networking on the cluster:
  1. Navigate to the **Settings > Cluster > General Setup** page.
  2. Select the **Enable advanced networking** checkbox.
  3. Click the **Change cluster parameters** button.

The Avere OS notifies you that advanced networking cannot be disabled after it is enabled and asks you for confirmation to proceed.

4. Click **OK** to proceed.

There might be a brief interruption to the Avere OS as advanced networking is enabled. When the enablement operation completes, the Avere OS displays a new link named **VLAN** under the **Cluster** heading on the **Settings** tab. (If necessary, refresh the browser page to allow you to see the link under the **Cluster** heading.)

## 5.2. Creating a VLAN

- To create a VLAN:
  1. Navigate to the **Settings > Cluster > VLAN Configuration** page.

The screenshot displays the 'VLAN Configuration' page in the Avere OS interface. The left sidebar shows the navigation menu with 'VLAN' selected under the 'Cluster' heading. The main content area features a table of existing VLANs and a form to add a new VLAN.

Name	VLAN Tag	Default Router	Roles	Actions
default	-	10.1.0.76	client,mgmt,cluster,core_access	<button>Remove</button> <button>Modify</button>

[Add New VLAN](#)

**Add New VLAN**

Name:

VLAN tag:

Default router:

MTU (optional):

Roles: ☒ Client ☐ Cluster ☐ Core filer access ☐ Management

[Configure Static Routes](#)

Submit

2. Click on the **Add New VLAN** button to display the panel.

3. In the **Name** field, enter the name of the new VLAN.

VLAN names can contain alphanumeric characters, underscores (\_), and dashes (-); they cannot contain spaces, periods, or other special characters.

4. In the **VLAN Tag** field, enter the tag number for the new VLAN. The tag number must be a unique integer between 1 and 4094.
5. Optionally, enter the IP address of the default router for the VLAN in the **Default Router** field.
6. Optionally, enter the VLAN's maximum transmission unit (MTU) in the **MTU (optional)** field. Otherwise, the default MTU for the cluster will be used, as described in Section 2.3.4, "Default MTU Value" on page 20.
7. In the **Roles** field, select the role or roles for the new VLAN, any combination of the following:
  - **Client** – The VLAN is used to communicate between the FXT nodes cluster and the clients.
  - **Cluster** – The VLAN is used to communicate between the FXT nodes in the cluster, and also between the cluster and core filers.
  - **Core filer access** – The VLAN is used to communicate between the FXT node cluster and the core filers.
  - **Management** – The VLAN is used to communicate between a management machine and the FXT nodes in the cluster. The IP addresses used by the VLAN are not locked to a single node.



#### Note

VLAN roles only determine whether or not the VLAN can be selected when later configuring particular networks, as described in Section 5.4, "Management Role VLANs" on page 81, Section 5.5, "Client Role VLANs" on page 82, Section 5.6, "Cluster Role VLANs" on page 82, and Section 5.7, "Core Filer Access Role VLANs" on page 83.

A cluster VLAN is used for communication with the core filer as well as among nodes. Cluster VLANs must use the default gateway assigned to the cluster.

8. Verify the values you entered into the fields, then choose **Submit**.

You can now do any of the following:

- Select **Add New VLAN** to add an additional VLAN.
- Select the **Modify** button to change a VLAN. **You can only modify the default router and the MTU for the default VLAN.**
- Select the **Remove** button to remove a VLAN from the cluster. **You cannot remove the default VLAN.**
- Select **Configure Static Routes** to configure static routes for any VLAN, as described in the following section, Section 5.3, "Configuring Static Routes for a VLAN" on page 80.

## 5.3. Configuring Static Routes for a VLAN

You can configure static routes for a VLAN in order to direct how the vservers communicate with the cluster, and how the cluster communicates with the core filers. This is particularly necessary if your core filers are located on different VLANs, and the cluster needs a gateway other than the default to reach them. You will need to know what VLANs your clients are on, what VLANs your core filers are on, and the client-facing and core filer-facing IP addresses.

- To add, modify, or delete static routes in a VLAN configuration:
1. Navigate to the **Settings > Cluster > VLAN Configuration** page.
  2. Click **Configure Static Routes**. The Avere Control Panel displays the **Static Routes** page.

The screenshot shows the 'Static Routes' configuration page in the Avere Control Panel. The top navigation bar includes 'Dashboard', 'Settings', 'Analytics', 'Data Management', and 'Support'. The left sidebar lists various configuration categories: VServer, Core Filer, Cluster, General Setup, Administrative Network, Cluster Networks, FXT Nodes, High Availability, Monitoring, Schedules, Directory Services, Kerberos, Login Services, Active Directory, Optimization, IPMI, Support, Licenses, Cloud Credentials, and VLAN. The main content area is titled 'Static Routes' and contains a table with columns: Router, VLANs, Static Routes, and Actions. The table lists two routes: one for Router 10.1.0.76 and VLAN Default (tag: -) with a static route of destIP:102.1.0.0,netmask:255.255.255.0,gateway:192.168.1.1, and another for Router 110.1.0.2 and VLAN my\_VLAN(tag:42) with an 'Add' button. Below the table, there are sections for 'Route 1' and 'Route 2' with input fields for Destination IP, Netmask, and Gateway. At the bottom, there are buttons for 'Remove This Route', 'Add Another Route', and 'Update Routes'.

Router	VLANs	Static Routes	Actions
10.1.0.76	Default (tag: -)	destIP:102.1.0.0,netmask:255.255.255.0,gateway:192.168.1.1	<button>Modify</button> <button>Delete</button>
110.1.0.2	my_VLAN(tag:42)		<button>Add</button>

**Modify Static Routes For 10.1.0.2**

**Route 1**

Destination IP:

Netmask:

Gateway:

[Remove This Route](#)

**Route 2**

Destination IP:

Netmask:

Gateway:

[Remove This Route](#)

[Add Another Route](#)

[Update Routes](#)

[Go back to VLAN configuration page](#)

3. Choose one of the following actions from the VLAN's row:
  - Choose **Delete** to remove *all* existing static routes from the VLAN.  
*The Avere OS will not give you a warning if you choose this option.*
  - Either choose **Add** to add a route to a VLAN that does not currently have any static routes, or choose **Modify** to add or change routes on a VLAN that has existing routes, to display the static route fields:
    - a. In the **Destination IP** field, enter or change the destination IP address.
    - b. In the **Netmask** field, enter or change the netmask, in IPv4 dotted format or */number\_of\_bits* format.
    - c. In the **Gateway** field, enter or change the gateway, using IPv4 dotted format.
    - d. If you want to enter additional routes, click **Add Another Route**. A set of fields for a new route is displayed.
    - e. When you have finished adding or changing the static routes, click the **Update Routes** button.

## 5.4. Management Role VLANs

Management role VLANs are used to communicate between a management machine and the FXT nodes in the cluster. The IP addresses used by the VLAN are not locked to a single node. Optionally, you can reserve a range of administrative IP addresses that are available exclusively for the use of management role VLANs.

Management role VLANs can be selected from the **Settings > Cluster > Administrative Network** page. Refer to Section 2.3, “Administrative Network Settings” on page 17 for more information.

➤ To use a VLAN for the cluster’s management interface:

1. Navigate to the **Settings > Cluster > Administrative Network** page. The **Default Router** and **Static Routes** fields that are present when advanced networking is not enabled have been replaced by a **Management VLAN** drop-down list.

The screenshot shows the 'Administrative Network' configuration page. The left sidebar contains a navigation menu with options like VServer, Core Filer, Cluster, General Setup, Administrative Network (selected), Cluster Networks, FXT Nodes, High Availability, Monitoring, Schedules, Directory Services, Kerberos, Login Services, Active Directory, Optimization, IPMI, Support, Licenses, Cloud Credentials, VLAN, and Administration. The main content area is titled 'Administrative Network' and contains several configuration fields: Management IP (10.1.22.194), Management Netmask (255.255.224.0), Management VLAN (Default (tag: -, gateway: 10.1.0.1)), Default MTU (optional), DNS server(s) (10.0.8.4), DNS domain (dns.company.com), DNS search, Timezone (America/New York), and NTP servers (ntp.mycompany.com). There are checkboxes for 'Use multicast NTP servers' and 'Use separate management network'. At the bottom right are 'Revert' and 'Submit' buttons. Below this is a section for 'Node Management Addresses (optional)' which includes a table with columns 'Address Range', 'Subnet Mask', 'VLAN', and 'Actions', and an 'Add New Range' button. The 'Add New Node Management Address Range' section contains fields for Node admin first IP (10.12.140.1), Node admin last IP (10.12.140.3), Number of IPs in range (3), Node admin subnet mask (255.255.224.0), and Node Admin VLAN (Default (tag: -, gateway: 10.1.0.1)), with 'Revert' and 'Submit' buttons at the bottom right.

Address Range	Subnet Mask	VLAN	Actions
---------------	-------------	------	---------

Add New Range

**Add New Node Management Address Range**

Node admin first IP: 10.12.140.1

Node admin last IP: 10.12.140.3

Number of IPs in range: 3

Node admin subnet mask: 255.255.224.0

Node Admin VLAN: Default (tag: -, gateway: 10.1.0.1)

Revert Submit

2. From the **VLAN** drop-down list, choose the VLAN that the management interface should use. Possible values include **Default** or any other management-role VLANs that you have created.



### Important

It is strongly recommended that the gateway of the management VLAN should match the gateway of the (optional) node administration VLAN, which can be set at the bottom of the page. Otherwise, the management interface may exhibit strange behavior.

3. Choose **Revert** to restore the original settings, or choose **Submit** to commit any changes. The Avere OS displays a pop-up that warns you about cluster session disruption and asks you for confirmation to proceed.
4. Click **OK**.

You can reserve a range of IP addresses for a management VLAN to use for node management. Refer to Section 2.3.11, “Node Management Addresses” on page 24 for more information.

## 5.5. Client Role VLANs

You can specify a VLAN for client-facing interfaces either during or after the creation of a vservers, as described in the following sections:

- Section 4.2, “Managing Virtual Servers (Settings Tab | VServer)” on page 65.
- Section 4.2.3, “Client-Facing Network Settings” on page 70.

## 5.6. Cluster Role VLANs

Cluster role VLANs are used to communicate within the FXT nodes in the cluster, and also between the cluster and the core filers. Optionally, you can reserve a range of IP addresses that are available exclusively for the use of cluster role VLANs.

Cluster role VLANs can be specified from the **Settings > Cluster > Cluster Networks** page. Refer to Section 2.4, “Cluster Networks” on page 25 for more information.



### Important

Before you modify or initially enable a VLAN for cluster interfaces, ensure that all components of both the current cluster network and the new cluster network are functioning correctly. Check that cluster traffic is already being transmitted across the existing cluster network and interfaces. If any component of either the initial cluster network or the new cluster network is unavailable when this process occurs, the cluster can become unhealthy, leading to file-service disruptions or other failures.

When you modify or enable a VLAN for cluster interfaces, the change is propagated through the existing cluster network. Services and internal processes restart on each node in the cluster; after the restart, the nodes use the new VLAN.



### Note

A cluster VLAN is used for communication between the cluster and the core filers as well as between nodes. Cluster VLANs must use the default gateway assigned to the cluster.

## 5.7. Core Filer Access Role VLANs

Core filer access role VLANs are used to communicate between FXT nodes in the cluster and the core filers. (You can also use cluster role VLANs to access the core filer.) Refer to Section 5.6, “Cluster Role VLANs” on page 82 and Section 2.4, “Cluster Networks” on page 25 for more information.

## Chapter 6. Setting the Cache Policy

The cache policy specifies how the Avere system handles file-operation requests between clients and core filers. For example, you could set two core filers in the following way:

- A core filer that has a high rate of data change (for example, an active source-code repository) can be set to provide clients with full read/write access (also called *writeback* mode), in which the cluster controls all read, write, and metadata operations between the core filer and clients, for maximum speed.
- On the other core filer (for example, one that holds binary files that clients access primarily only to read), you can specify a read-only policy (also called *writearound* mode) in which any write operations from clients go directly through to the core filer.



### Note

The terminology and implementation of cache policies, known in previous versions of Avere OS as write modes, have changed significantly as of Avere OS 2.0. The term “write mode” is being replaced; however, it is still rarely used in the Avere Control Panel and documentation.

### 6.1. Understanding Avere’s Cache Policies

The Avere system has two primary modes of handling data that clients have changed (written).

- In *read mode*, the cluster updates the core filer after each write operation and checks the attributes on files at an administratively specified interval. This enables the use of the cluster while other systems both read and write data directly to the core filer. This mode accelerates data-read operations, but the performance of data-write operations is limited to the performance of the core filer and the performance of attribute-read and attribute-write operations is constrained by the specified core filer verification interval.

Read mode is the expected configuration for initial installation and setup of the Avere system, as clients are migrated from connecting directly to the core filer to connecting through the cluster. Some clients can still mount the core filer directly. Read mode is the default caching mode for clusters and new core filers.

- In *read/write mode*, the cluster controls all read, write, and metadata (attribute) operations, and writes changed data to the core filer within an interval called the *maximum writeback delay*. Read and write performance scales with the size of the cluster independently of the performance capabilities of the core filer. This mode accelerates read, write, and metadata operations.

The caching mode is set on a per-core filer basis. For a simple-namespace vserver, there is still a single caching policy because the vserver is associated with a single core filer. For a global-namespace vserver, you can have different caching modes set for different core filers.



### Caution

If you are using read/write mode, all clients that are accessing a given dataset *must* mount the FXT Series cluster to access the core filer. If some clients are mounted on the FXT cluster and others are mounted directly on the core filer, with all clients accessing the same dataset, data corruption and inconsistency can result.



### Important

If your cluster is in read mode with core filer verification set to **Always** (previously known as write-around mode), do not use CIFS clients to access the cluster. See Section 7.4, “CIFS Limitations” on page 96 for details.

## 6.2. Specifying the Cache Policy

The following sections describe how to specify each cache policy and discuss considerations for each mode.

➤ To set a cache policy:

1. Navigate to the **Settings > Core Filer > Cache Policy** page.

The screenshot shows the 'Cache Policy - cinder' configuration page. The left sidebar contains navigation links: Dashboard, Settings, Analytics, Data Management, Support, VServer, Core Filer, Manage Core Filers, Cluster, and Administration. The 'Core Filer' section is expanded, showing 'cinder' as the selected filer. The 'Cache Policy' section is active, displaying settings for 'Caching mode' (Read/Write), 'Core filer verification' (Never), 'Maximum writeback delay' (Custom), and 'Custom writeback delay' (5). The 'Advanced Features' section is expanded, showing 'Advanced Cache Policy' (Allow core filer verification when Read/Write caching: ☐, Custom core filer directory verification time: ) and 'Cache Utilization Control' (Limit modified data by: Core filer). The 'Writethrough Scheduling' section is also expanded, showing 'Writethrough schedule' (No Schedule Selected), 'On-sync action' (URL Polling), and 'Limit modified data by' (Core filer). The 'Local Directory Control' section is expanded, showing 'Local directories' (Enabled). A red box highlights the 'Advanced Features' section, and a red arrow points from the text 'Only for NFS core filers, not cloud filers.' to this section.

Only for NFS core filers, not cloud filers.

2. If necessary, choose the appropriate core filer from the drop-down list under the **Core Filer** heading.
3. From the **Caching Mode** drop-down list, choose either **Read** for read mode or **Read/Write** for read/write mode.

4. Depending on the selected mode, choose appropriate values from the **Core filer verification** or **Maximum writeback delay** drop-down lists. These settings are discussed in Section 6.2.1, “Setting Read Mode” and Section 6.2.2, “Setting Read/Write Mode”.
5. If desired, open the **Advanced Features** panel and set the appropriate values.



### Caution

Setting values in the **Advanced Features** panel can lead to problems ranging from performance issues to data inconsistency. Contact Avere Global Services before setting any values on this panel.

6. Click the **Save Settings** button.

## 6.2.1. Setting Read Mode

You can set read mode, the default cache policy, by selecting **Read** from the **Caching Mode** drop-down list. Also select a value for the core filer verification interval from the **Core Filer Verification** drop-down list; the verification interval determines how often the FXT Series system verifies the cached file and directory attributes against the core filer’s file and directory attributes.

### Determining the Core Filer Verification Interval

When deciding on a core filer verification interval, note the following general rules:

- The longer the interval, the higher the performance of read operations.
- The shorter the interval, the higher the guarantee of file- and directory-attribute synchronization.

Choose a value that balances your need for performance against your need for attribute synchronization. This value typically varies from application to application. Consult with your Avere Systems representative to select an initial value and, if necessary, adjust it later.

Possible values for the core filer verification interval include the following:

- A range between 5 seconds and 1 day.
- **Always**—The FXT Series cluster verifies the cached attributes against the core filer’s attributes with each read operation. This value guarantees attribute synchronization but reduces the performance of the Avere system. This is the default value; most customers start with this value and later move to longer core filer verification intervals or read/write mode.
- **Never**—The FXT Series cluster never verifies the cached attributes against the core filer’s attributes.

## 6.2.2. Setting Read/Write Mode

You can set read/write mode for a core filer by selecting **Read/Write** from the **Caching Mode** drop-down list. Also select a value from the **Maximum writeback delay** drop-down list; this value specifies the maximum amount of time that changed user data and metadata (attributes) is stored on the cluster before being committed (written) from the cache to the core filer.

### 6.2.2.1. Determining the Maximum Writeback Delay

When determining the value to set for the maximum writeback delay, consider the following:

- The usage period of the average client. For example, if a standard client is a desk workstation that is used for eight hours per day, you might consider eight hours as a reasonable amount of time.
- The importance of the data being worked on by clients. For example, if clients are being used to input mission-critical data, you might consider 30 minutes, one hour, or two hours as the maximum writeback delay. However, if clients are generating scratch data, a maximum writeback delay of one day or two weeks might be reasonable.
- The amount of data being generated by the client as compared to the capacity of the core filer. For example, if your core filer is notably slow but your clients are generating data rapidly, you might consider increasing the maximum writeback delay to ensure that the core filer is able to accept data at a steady rate that does not overload it. In this example, a maximum writeback delay value of one (1) minute is far too low for realistic performance gains.
- Because enabling and disabling local directories on a core filer can take a large amount of time (as much as, or more than, 30 minutes), you can reduce this time by reducing the writeback delay. This depends on how much data is already in the cache, and what attributes need to be transferred to the core filer. After the data is flushed back to the core filer, you can then enable or disable local directories with less delay.
- The size of the working set. A larger working set typically requires a longer maximum writeback delay value for optimal performance.
- The schedule on which your core filer performs backup operations. For example, if your core filer performs a full backup every 24 hours, the maximum writeback delay needs to be less than 24 hours to ensure that each day's current or new data is backed up.

Consult with your Avere Systems representative to estimate a good initial value for the maximum writeback delay. As you become more familiar with the Avere system and the workload patterns generated by your clients and application servers, you can adjust the maximum writeback delay until you find the optimal value for your environment.

Possible values for the maximum writeback delay vary from zero seconds through one year.

### 6.2.2.2. Scheduling Read-Mode Periods

Most core filers use backup services such as snapshots, mirrors, and Network Data Management Protocol (NDMP) sessions to a tape or disk library, typically in combinations to guarantee data retention in the event of data loss on the core filer. If your cluster is using read/write mode to serve a core filer's data to clients, you can ensure that backups of the core filer's data are synchronized with the latest modifications from clients by scheduling a read-mode (or write-through) period on the FXT cluster to coincide with the core filer's backup schedule. The length of the read-mode period can be a specified length of time (for example, five minutes) or can be controlled by a simple external polling mechanism.

The FXT cluster performs the following actions when you set up read-mode periods:

1. Approximately every ten seconds, an internal cluster process checks the current time against the next scheduled read-mode period (also called the *target time*).
2. If a read-only period is approaching, the cluster automatically and gradually lowers the value of the maximum writeback delay for read-write mode, causing the cluster to write changed data from clients back to the core filer more aggressively.
3. The cluster continues to lower the value as necessary as the target time gets closer. As a result, the cluster has written all changed data to the core filer when the scheduled read-only period starts.
4. When the read-only period starts, the cluster automatically switches all client access to read-only (write-through) mode. Any changes made by clients during this period are written directly to the core filer, with the FXT cluster retaining information about changed data and metadata. The FXT cluster continues to accelerate client read-only requests.
5. The cluster remains in read-only mode until one of the following events occurs:
  - The polling URL returns the string `RELEASE Core Filer SYNC` to the FXT cluster
  - The specified waiting period expires
6. When the read-only period has ended, the cluster switches back to read/write mode at the originally specified maximum writeback delay for the core filer and stays in that mode until the next scheduled read-only period approaches.

➤ To schedule a read-mode period:

1. Create a schedule on the cluster that corresponds to the core filer's backup schedule, as described in Section 2.7.1, "Creating a Schedule" on page 37.
2. Navigate to the **Settings > Core Filer > Cache Policy** page.
3. Choose the appropriate core filer from the drop-down list under **Settings > Core Filer**.
4. Expand the **Advanced Features** panel by clicking the arrow on the left of the panel label.
5. Select **Enable advanced caching features**. The Avere Control Panel displays a warning pop-up; click **OK** to dismiss it.
6. From the **Writethrough Scheduling** panel, select the cluster schedule that corresponds to the core filer's backup schedule from the **Writethrough schedule** drop-down list.
7. From the **On-sync action** drop-down list, choose **URL Polling** if you plan to use URL polling, or **Wait Period** if you want to specify a fixed length of time for the read-mode period.
8. Depending on the selection you made in the previous step, do one of the following:
  - a. If you selected **URL Polling**, enter the URL of the external polling agent in the **Poll URL** field. External polling agents are discussed in Section 6.2.2.3, "Implementing a URL Polling Agent" on page 89.
  - b. If you selected **Wait Period**, choose a value for the waiting period from the **Wait Period** drop-down list. The value needs to be as long as or longer than the core filer backup process takes to complete. For example, if backups typically take 30 minutes but sometimes take 40 to 45 minutes, select **1 Hour**.  
  
Possible wait periods vary from 5 minutes to 6 hours.
9. Click the **Save Settings** button.

### 6.2.2.3. Implementing a URL Polling Agent

A polling agent is a CGI script external to the FXT cluster that can monitor and report on the progress of the core filer's backup job. The agent can be written in any language that supports CGI, including scripting languages such as Perl and Python and compiled languages such as C, C++, and Java. See [RFC 3875](#) for information about the components of a CGI script.

When URL polling is set, the FXT cluster checks the URL of the CGI script approximately every 10 seconds during the time period specified by the cluster schedule. The CGI script must accept the arguments `corefiler` (or `mass`) and `targetTime`, where `corefiler` is the name of the core filer and `targetTime` is a UNIX timestamp representing the next scheduled writethrough period. It must also return the string `RELEASE Core Filer SYNC` when the core filer's backup process has completed.

A simple example CGI script written in Python follows. Note that the script does not include any logic for monitoring the status of the core filer; see the documentation for your core filer, particularly any API documentation, for that information.

#### CGI Script Example

```
#!/usr/bin/env python

import os
import sys
import cgi

cgiArgs = cgi.FieldStorage(keep_blank_values=True)
corefiler = cgiArgs.getfirst('corefiler')
targetTime = cgiArgs.getfirst('targetTime')

response = "must supply corefiler and targetTime arguments"

if corefiler and targetTime:
    statefile = "/tmp/STATE.%s_%s"%(corefiler,targetTime)
    os.close(os.open(statefile, os.O_CREAT|os.O_RDWR))
    f = open(statefile, "r+")
    l = f.readline()
    if not l:
        response = "starting backup operation"
        f.write("start\n")
    elif l == "start\n":
        response = "waiting for backup operation to complete (1)"
        f.truncate(0)
        f.seek(0)
        f.write("wait1\n")
    elif l == "wait1\n":
        response = "waiting for backup operation to complete (2)"
        f.truncate(0)
        f.seek(0)
        f.write("wait2\n")
    elif l == "wait2\n":
        response = "waiting for backup operation to complete (3)"
        f.truncate(0)
        f.seek(0)
        f.write("wait3\n")
```

```

else:
    response = "backup complete.\nIt's time to RELEASE Core Filer SYNC\n"
    os.unlink(statefile)
f.close()

print "Content-Type: text/plain"
print "Pragma: no-cache"
print "Cache-Control: no-cache, must-revalidate"
print "Expires: Sat, 26 Jul 1997 05:00:00 GMT"
print
print response
sys.exit(0)

```

## 6.3. Setting Advanced Cache-Policy Features

The **Advanced Features** panel enables you to customize the cache policy for a given core filer. Some of these features are described in previous sections.

### 6.3.1. Enabling Advanced Features

- To enable and configure advanced cache-policy procedures:
  1. If necessary, choose the appropriate core filer from the drop-down list under the **Core Filer** heading on the Avere OS.
  2. Navigate to the **Settings > Core Filer > Cache Policy** page.
  3. Open the **Advanced Features** panel by clicking the triangle to the left of the panel's heading.
  4. Set advanced features as required. See the following sections for details.
  5. Click the **Save Settings** button.

### 6.3.2. Setting a Custom Writeback Delay

If none of the predefined values available for read/write mode's maximum writeback delay works for your environment, you can specify a custom value for the maximum writeback delay. Enter the value, expressed in seconds, in the **Custom Writeback Delay** field. The value must be a nonnegative integer.

### 6.3.3. Setting a Custom Core Filer Verification Time

If none of the predefined values available for read mode's core filer verification interval works for your environment, you can specify a custom value for core filer verification time. Enter the value, expressed in seconds, in the **Custom Core Filer Verification Time** field. The value must be a nonnegative integer.

### 6.3.4. Setting a Custom Core Filer Directory Verification Time

You can specify the interval at which the FXT Series system verifies the cached directory attributes against the attributes of the same directories on the core filer. Enter the value, expressed in seconds, in the **Custom Core Filer Directory Verification Time** field. The value must be a nonnegative integer.

### 6.3.5. Enabling and Disabling Core Filer Verification in Read/Write Mode

The default behavior of the FXT Series cluster in read/write mode is to write file and directory attributes to the core filer as it writes changed data to the core filer. You can override this default behavior if necessary – for example, if your network has a long average latency or if the characteristics of your dataset require it. If you override the default behavior, the cluster checks the attributes of files and directories currently retained on the cluster against the attributes of the files and directories on the core filer and, if the attributes on the core filer appear to be more recent than the attributes on the cluster, adjusts the attributes of files and directories on the cluster to match those on the core filer.



#### Caution

Enabling this feature can cause data inconsistencies between the FXT Series cluster and the core filer. Do not enable it without first consulting with Avere Global Services.

- To enable core filer verification:
  1. Select the **Allow Core Filer verification when Read/Write caching** checkbox.
  2. Specify core filer verification values using
    - The **Core filer verification** drop-down list or the **Custom Core Filer Verification Time** field
    - Optionally, the **Custom Core Filer Directory Verification Time** field
- To disable core filer verification in any write mode, do one of the following:
  - Choose **Never** as the value of the **Core Filer Verification** drop-down list
  - Enter 0 (zero) in the **Custom Core Filer Verification Time** and **Custom Core Filer Directory Verification Time** fields.

## 6.4. Controlling Write Bandwidth

By default, the Avere cluster uses bandwidth to the core filer as follows:

- When the cluster writes changed data back to the core filer, it uses all available network bandwidth between the cluster and the core filer.
- The cluster does not write changed data back to the core filer until the writeback-delay value is reached.

You can change either of these default behaviors on a per-core filer basis from the **Settings > Core Filer > Core Filer Details** page.

Reasons you might want to change the default behaviors include the following:

- You want to limit the bandwidth consumed between the cluster and the core filer, to make additional bandwidth available for other network traffic.
- You want the cluster to write changed data back to the core filer before the writeback-delay value is reached (that is, enable *early writeback*). This increases the performance demands on the cluster but can also result in more uniform load balancing across the storage infrastructure.



#### Note

- If you want the cluster to use early writeback, you must also specify a maximum bandwidth.
- Bandwidth throttling is not honored if a high-priority writeback to a particular core filer is in progress. The specified throttling settings resume after the high-priority writeback completes.

- To change the default bandwidth use or writeback timing for a core filer:
1. If you have more than one core filer configured, choose the appropriate one from the drop-down list.
  2. Navigate to the **Settings > Core Filer > Core Filer Details** page.
  3. Select the **Enable Bandwidth Control** checkbox. The Avere Control Panel displays the **Maximum Bandwidth** and **Early Writeback** controls.
  4. To select a maximum bandwidth for writebacks, choose a value from the **Maximum Bandwidth** drop-down list. You can choose from a range of values, with a **Custom** option.  
  
If you select **Custom**, the Avere Control Panel displays the **Custom Maximum Bandwidth (MB/s)** control. Enter a custom value into the field. The value must be an integer between 1 and 800.
  5. If you also want to set early writeback on the core filer, select the **Early Writeback** checkbox.



#### Note

A maximum-writeback value *must* be set if early writeback is to be enabled.

6. Click **Submit** to submit any changes you made.

## 6.5. Controlling Cache Utilization on the FXT Series Cluster

Cache-utilization controls enable you to limit the amount of modified data retained on the FXT Series cluster from clients according to core filer, user, NFS export, or file-system ID (FSID). This feature enables you to avoid a situation in which a single client or small number of clients does not fill the entire cache policy with changed data, thus blocking other clients from writing changed data to the cluster in the meantime.



#### Important

If the allocated amount of cache fills completely with modified data, clients attempting write operations receive DQUOT or EJUKEBOX errors until all modified data in the cache is written back to the core filer. If this occurs, you can force all changed data to be written back to the core filer as soon as possible by setting the core filer's maximum writeback delay to 0 (zero) seconds, as described in Section 6.2.2.1, "Determining the Maximum Writeback Delay" on page 87. As a followup step, it is recommended that you redistribute the workload across your clients.



#### Note

Cache-utilization controls limit modified data on the cache on the *cluster*, and has no connection with storage-limiting containers defined on the core filer, such as Data ONTAP qtrees.

- To enable and specify cache-utilization controls:
1. If necessary, choose the appropriate core filer from the drop-down list under the **Core Filer** heading on the Avere Control Panel.
  2. Navigate to the **Settings > Core Filer > Cache Policy** page.
  3. If necessary, click the triangle next to the **Advanced Features** label to open the advanced-features panel.

4. If advanced features are not already enabled, select the **Enable advanced caching features** checkbox. The Avere OS prompts you for confirmation.
5. From the **Limit modified data by** drop-down list in the **Cache Utilization Control** panel, choose one of the following policies:
  - **Core Filer** (the default): Cache-utilization controls will be applied to the core filer.
  - **User**: Cache-utilization controls will be applied on a per-user basis.
  - **Export**: Cache-utilization controls will be applied on a per-export basis.
  - **Fsid**: Cache-utilization controls will be applied on a per-FSID basis.
  - **Export and User**: Cache-utilization controls will be applied on a per-export and per-user basis.
  - **Fsid and User**: Cache-utilization controls will be applied on a per-FSID and per-user basis.
6. Click the **Save Settings** button to apply the specified changes.

### 6.5.1. Precautions While Using Cache-Utilization Controls

Observe the following warnings when using cache-utilization controls:

- Cache-utilization controls enable you to limit the amount of modified data retained on the FXT Series cluster according to core filer, user, NFS export, or FSID object.
- Cache-utilization controls are available only when the caching mode is set to **Read/Write** and advanced features are enabled for the cache policy.
- If you change the caching mode to **Read** when cache-utilization controls are enabled, the cluster writes any remaining modified data in the specified cache to the core filer as an asynchronous background job.
- If cache-utilization controls are enabled and you change the caching mode from **Read/Write** to **Read** and then back to **Read/Write** on the same core filer, the cache-utilization setting that was specified when the caching mode was first changed to **Read** is reenabled.

For example, if you have a core filer's cache-utilization control set to **User**, then switch the core filer's caching mode to **Read**, cache-utilization controls are disabled. If you then switch the same core filer's caching mode back to **Read/Write**, the cache-utilization controls are automatically reactivated with a setting of **User** (not, for example, the default value of **Core Filer**). If you want to use a different cache-utilization setting (for example, **Export**) after reenabling Read/Write caching mode, you must specify the desired setting in the Avere Control Panel.

- When enabled, the cache-utilization control mechanism automatically calculates the optimal amount of the cache policy allocated to each user, based on the size of the cache policy and other variables. You cannot specify the maximum cache-utilization value for a given client.

## 6.5.2. Enabling Local Directories



### Caution

If you are using read/write mode, all clients that are accessing a given dataset *must* mount the FXT Series cluster to access the core filer. If some clients are mounted on the FXT cluster and others are mounted directly on the core filer, with all clients accessing the same dataset, data corruption and inconsistency can result.



### Note

Because enabling and disabling local directories on a core filer can take a large amount of time, you can reduce this time by reducing the writeback delay. After the data is flushed back to the core filer, you can then enable or disable local directories.

To allow directory structures to be written to the cluster, and only be transferred to a core filer at specified intervals, select **Enabled** from the **Local Directories** drop-down list.

This capability significantly boosts the performance of workloads that have a high number of creation operations.

Local directory creation is set *per core filer*, not per vservers, per export, or on any other object in the Avere system.

If you are enabling or disabling local directories, a condition will appear on the dashboard as the process starts, and will be removed when the change is complete. At that point, an alert will appear, reminding you to remount your NFS clients, and to remap any CIFS clients (as described in Section 7.5.2, “Configuring the Cluster to Use Active Directory” on page 98).

## Chapter 7. Configuring CIFS Access

Short for *Common Internet File System*, CIFS is a protocol that allows users with different platforms and computers to share files without having to install new software.

A CIFS server running on each node, as opposed to file sharing between computers on a network, ensures that CIFS performance is as fast as possible. Avere vservers can be configured to serve CIFS clients, such as Microsoft Windows computers.

### 7.1. Recommendations for CIFS Configuration

For optimal CIFS performance, follow these recommendations.

- As noted in Section 2.3.7, “Setting DNS Parameters” on page 22, you should configure your DNS domain to use round-robin load distribution for client-facing IP addresses.
- NetBIOS advertises only one IP address per namespace. If CIFS clients browse to the namespace using the Windows® Network Neighborhood feature, all client requests get directed to a single client-facing IP address on a single node, resulting in poor performance. To avoid this issue, use one of the following strategies:
  - Create an entry or entries in the Microsoft® Dynamic DNS server for the namespace or the client-facing IP addresses of the namespace server.
  - Use the Windows Internet Name Service (WINS) to create an entry for client-facing IP addresses. See the Windows documentation for information.
  - Manually distribute the CIFS client load by directing individual clients to individual client-facing IP addresses.
- It is strongly recommended that core filer exports be used for access by only one file protocol, NFS or CIFS. Using both file protocols to access a core filer export can result in conflicting file locks and potential file corruption. If you have data on an export that must be accessed by both NFS and CIFS clients, contact Avere Global Services for an analysis of your environment and access patterns before enabling multi-protocol access on the export.

### 7.2. Core Filer Prerequisites for CIFS

In order to use CIFS access, the core filer must be one of the following:

- NetApp Data ONTAP 7G or NetApp Data ONTAP 8 in 7-mode, with the volume or qtree security style set to either `unix` or `ntfs`. Avere OS does not currently support the `mixed` security style. For more information, refer to Section A.2, “NetApp Data ONTAP 7G and Data ONTAP 8.0 7-Mode” on page 157.
- OpenSolaris or Oracle® Solaris 9 or 10 with ZFS volumes. For more information, refer to Section A.3, “ZFS on OpenSolaris and Oracle Solaris” on page 158.
- EMC Isilon OneFS 5.x and 6.0.x. For more information, refer to Section A.4, “EMC Isilon OneFS” on page 159.

## 7.3. Overview: CIFS Configuration on the Cluster

➤ To configure a namespace for CIFS, use the following general procedure:

1. Make sure the cluster, the core filer, the Active Directory/Kerberos server, and all CIFS clients are using a common time source; if the time skew between any two components is more than five minutes, CIFS access is denied. Refer to Section 2.3.8, “Timezone” on page 23 and Section 2.3.9, “Setting NTP Parameters” on page 23 for more information.
2. Configure the DNS settings, as described in Section 2.3.7, “Setting DNS Parameters” on page 22.
3. Configure the cluster with a directory service: NIS, LDAP, or both, as described in Section 2.8.1, “Selecting and Configuring a Directory Service” on page 39.

NIS and LDAP can both be configured, but the Avere CIFS service uses only one for username resolution. Make sure that you select the correct source service from the **Directory Services** page (at the bottom of the page where you configure NIS and LDAP).

If you are using NIS, then the UNIX and Windows usernames must be identical. For example, a UNIX user with the username “jsmith” must also have the Windows username “jsmith”, not the Windows username “Jodi Smith”.

4. Configure the cluster with an Active Directory domain with Dynamic DNS. If more than one vserver on the cluster is enabled for CIFS, all vservers must join the same Active Directory domain (that is, the entire cluster must use a single Active Directory domain). If CIFS and UNIX usernames differ in your environment, set up a username mapping file. Refer to Section 7.5, “Joining an Active Directory Domain” on page 97 for more information.
5. Enable CIFS on one or more vservers. This requires entering the username and password of an Active Directory user that is authorized to join the vserver to the Active Directory domain. These tasks are described in Section 7.6, “Enabling and Configuring CIFS” on page 99.
6. Configure each CIFS share with the appropriate access mode, as described in Section 7.8, “Creating CIFS Shares” on page 107. The access mode depends in part on what your core filer supports and exports.



### Note

Although NFSv4 ACLs can be used as an access mode for CIFS clients, the Avere product does not support serving data from core filers that export pure NFSv4 volumes.

## 7.4. CIFS Limitations

The following is a known limitation in Avere OS’s current CIFS implementation:

- The Avere system does not correctly observe the append bit on CIFS ACLs for NetApp core filers. (This behavior matches that of most other commercial CIFS servers.)

## 7.5. Joining an Active Directory Domain

To use CIFS, the cluster must use an Active Directory domain. In addition, if your CIFS usernames are not identical to your UNIX usernames (that is, if Jane Smith's CIFS username is `Jane Smith` but her UNIX username is `janes`), you need to create a mapping file for the cluster to use to match usernames. If your CIFS and UNIX usernames are identical, you do not need a mapping file.

### 7.5.1. Creating a Username Map

If you use username mapping:

- The UNIX username is the name used by the vserver to access data on the core filer.
- A CIFS user is authenticated according to his or her CIFS username but authorized according to his or her UNIX username; that is, file access is determined by the UNIX username.
- Notifications of changes to files or directories will not be sent if any of the following conditions exist:
  - Any CIFS clients are directly connected to the core filer.
  - Any NFS clients are directly mounted to the core filer.
  - Any NFS clients are mounted to the cluster.

#### Username Map File Example

The format of the username map is as follows:

- Each line maps one username in the format **`UNIX_username=AD_DOMAIN "CIFS_username"`** where:
  - **`UNIX_username`** is the user's UNIX username.
  - **`AD_DOMAIN`** is the Active Directory domain used by the cluster.
  - **`CIFS_username`** is the user's CIFS username. If the CIFS username includes spaces, enclose it in double quotation marks.

A line that begins with a pound sign (`#`) is treated as a comment.

```
# The next line maps the UNIX user timmy to the CIFS user
# "Timothy Primate" in the AD domain MYDOMAIN:
timmy=MYDOMAIN "\"Timothy Primate\""
#
# The next lines map more UNIX usernames to CIFS usernames
# in the same AD domain:
cgflynn=MYDOMAIN "\"C. Gertrude Flynn\""
fletch=MYDOMAIN "\"Lawrence Fletcher\""
jeremyj=MYDOMAIN "\"Jeremy Johnson\""
perry=MYDOMAIN "\"Perry Monotreme\""
vanessa=MYDOMAIN "\"Vanessa Doofenshmirtz\""
```

Place the file in a location where you can access it with a URI.

## 7.5.2. Configuring the Cluster to Use Active Directory

When you have the username mapping file set (if needed) you can continue with the cluster configuration.



### Note

- All CIFS-enabled vservers on a cluster must use the same Active Directory domain.
- This procedure does not actually join the Active Directory domain, it simply sets up the association with the domain for vservers to join later.

➤ To configure the cluster to use Active Directory:

1. Navigate to the **Settings > Cluster > Active Directory** page.

The screenshot shows the 'Active Directory Configuration' page in the LiGo Cluster management interface. The page has a top navigation bar with tabs: Dashboard, Settings, Analytics, Data Management, and Support. The 'Settings' tab is active. On the left is a sidebar menu with options: VServer, Core Filer, Cluster, General Setup, Administrative Network, Cluster Networks, FXT Nodes, High Availability, Monitoring, Schedules, Directory Services, Kerberos, Login Services, Active Directory (highlighted), Optimization, IPMI, Support, Licenses, Cloud Credentials, VLAN, and Administration. The main content area is titled 'Active Directory Configuration' and contains three sections:

- Active Directory**: Fields for AD domain (FQDN) set to 'actif-dir.host.com', AD domain (NetBIOS) set to 'ACTIF-DIR', Domain controller set to 'cifs.actif-dir.host.com', DC site set to 'Default-First-Site-Name', and FXT site set to 'Default-First-Site-Name'.
- Domain Controller Address Overrides**: A table with columns 'Domain (FQDN)', 'Domain (Netbios)', 'DC Addresses', and 'Actions'. Below the table are input fields for 'Domain (FQDN)', 'Domain (Netbios)', and 'DC address(es)', followed by an 'Add Override' button.
- CIFS User to NFS User Overrides**: Fields for 'Source' (set to 'File'), 'File URI' (set to 'http://hostname.com/cifsusermap'), and 'Poll Period' (set to '1 hour'), followed by a 'Poll Now' button.

At the bottom right of the configuration area are 'Reset' and 'Submit' buttons.

2. In the **Active Directory domain to join** field, enter the name of the Active Directory server. Specify either its fully qualified domain name (for example, `adserver.example.com`) or its IP address.

3. If you are associating a username map with a cluster, do the following:
  - a. In the **CIFS Users to NFS User Overrides** panel, choose **File** from the **Source** drop-down list. The page displays the **File URI** field below the **Source** drop-down list.
  - b. Enter the URL of the username map in the **File URI** field.
  - c. From the **Poll Period** drop-down list, choose an interval at which the cluster checks the username map file for changes. Possible values include **24 hours**, **12 hours**, **1 hour** (the default), **Manual**, and **Custom**.
  - d. If you selected the **Custom** value from the **Poll Period** drop-down list, the page displays the **Custom Poll Period** field below the **Poll Period** drop-down list. Enter a custom poll period in minutes.
  - e. If you want the cluster to read the file immediately, click the **Poll Now** button.



#### Note

If you selected the **Manual** value from the **Poll Period** drop-down list, the cluster reads the file only when the **Poll Now** button is clicked.

4. Click the **Submit** button at the bottom of the page.

## 7.6. Enabling and Configuring CIFS

After the cluster has been joined to an Active Directory domain, you can enable CIFS access on any of the vservers of that cluster.

Ensure that the export settings for all listed exports are correct, as described in Section 2.5, “Managing Exports (Settings Tab | VServer)” on page 28. However, do not change the export policy or rules after setting up CIFS sharing.

Ensure that directory services are configured as described in Section 2.8.1, “Selecting and Configuring a Directory Service” on page 39.



#### Important

Regardless of the directory service used (NIS or LDAP), if UNIX and Windows usernames are not identical, you must set up a username map as described in Section 7.5.1, “Creating a Username Map” on page 97.

### 7.6.1. Gathering Information and Confirming Access

Consult with your Active Directory domain administrator before attempting to join the vserver to the Active Directory domain. You will need the following access and information:

- The user name and password of a Windows user with administrative permissions to join the Active Directory® domain (the *administrative user*).



#### **Note**

The administrative password is not retained by the cluster; it is discarded immediately after being used to join the Active Directory domain.

- If needed, the Organizational Unit (OU) name.
- The administrative user must have `Owner` permissions for the following access control entries in the Organizational Unit's access control list (ACL):
  - List Contents
  - Read All Properties
  - Create Computer Objects
  - Delete Computer Objects
- If the vserver's machine account is precreated in an OU directly on the Active Directory server, the user that creates the machine account must be the first user listed on the AD server's initial computer account creation screen, a member of the group named Domain Admins, or both.

## 7.6.2. Enabling CIFS Access

➤ To enable CIFS access on a virtual server:

1. Navigate to the **Settings > VServer > CIFS** page.

The screenshot shows the 'global - CIFS Configuration' page. On the left, a sidebar lists navigation options: 'VServer' (selected), 'Manage VServers', 'global' (dropdown), 'Name', 'Client Facing Network', 'Namespace', 'Export Policies', 'Export Rules', 'CIFS' (highlighted), 'CIFS Shares', 'MASS', 'Cluster', and 'Administration'. The main content area is titled 'global - CIFS Configuration'. It has two sections: 'Enable / Disable' and 'Machine Account'. In the 'Enable / Disable' section, 'Enable CIFS on global' is checked, and 'Enable SMB2' is unchecked. There are 'Reset' and 'Update CIFS Status' buttons. The 'Machine Account' section shows 'Current Join Status' as 'JOINED' and 'Distinguished Name' as 'CN=ligo2-,CN=Computers,DC=cifs,DC=amput,DC=com'. It has input fields for 'Name for the CIFS server (the NetBIOS name)' (containing 'LiGo2-'), 'Admin username' (containing 'Administrator'), 'Admin password' (masked with dots), and 'Organizational Unit' (containing 'infrastructure/masses'). There are 'Reset' and 'Update CIFS configuration' buttons at the bottom.

2. If you have more than one virtual server, choose the appropriate virtual server from the drop-down list immediately below the **Manage VServers** heading on the **Settings** tab.
3. Make the following changes in the **Machine Account** panel:
  - a. In the **Name for the CIFS server (the NetBIOS name)** field, enter the name of the CIFS server as it is to appear to clients. The default is the name of the cluster. The name must be between one (1) and 15 characters in length and consist only of the characters A-Z, a-z, 0-9, and - (hyphen).
    - If your Active Directory server is Windows 2000 or later, the name cannot include the period (.) character. If your Active Directory server is Windows NT®, the name can include a period but cannot start with a period. For possible future compatibility issues, it is recommended that you do not use the period character in the NetBIOS name.
    - If you want to use the NetBIOS name to access the vserver, then add a DNS server Host(A) record for each IP address assigned to the vserver.
  - b. In the **Admin username** field, enter the name of the administrative user.
  - c. In the **Admin password** field, enter the password of the administrative user.
  - d. Optionally, enter the name of an Organizational Unit (OU) for the vserver's machine account:
    - Select the **Advanced** checkbox. The Avere OS displays the **Organizational Unit** field.
    - Enter the name of an Organizational Unit for the vserver's machine account.
  - e. Click **Reset** to return to the original settings, or click **Update CIFS configuration** to accept the changes.

6. In the **Enable/Disable** panel, select the **Enable CIFS on *vserver-name*** checkbox, and click the **Update CIFS Status** button.



#### Note

You cannot change the CIFS configuration while the CIFS service is disabled or when the CIFS service is transitioning from the enabled state to the disabled state. If you attempt this, the Avere OS alerts you that the changes are not accepted. You can change the CIFS configuration only when the CIFS service is enabled and running.

After some cluster activity, the **Current Joined Status** will change to JOINED, and the **Distinguished Name** field will be filled in.

### 7.6.3. Updating the CIFS Configuration

After CIFS is configured and enabled on a virtual server, you can update the following parameters:

- NetBIOS name
- Active Directory administrative user name. Sometimes an update is needed to reauthenticate with the Active Directory domain; this requires the administrative user to re-enter their password.

You can also disable CIFS service on the virtual server, for example, if a client only needs brief CIFS access, and you do not want to risk simultaneous NFS/CIFS access to the same data.

➤ To update CIFS configuration parameters:

1. Navigate to the **Settings > VServer > CIFS** page.
2. If you have more than one virtual server, choose the CIFS-enabled virtual server from the drop-down list immediately below the **Manage VServers** heading on the **Settings** tab.
3. Enter new values in the appropriate fields as required. See Section 7.6, “Enabling and Configuring CIFS” on page 99 for information on the fields.
4. Choose **Reset** to restore the original configuration values, or choose **Update CIFS configuration** to commit any changes.

## 7.7. Selecting an Access-Control Mechanism

CIFS uses access-control lists (ACLs) to control access to files. Depending on the type of filesystem the vserver's core filer supports, the Avere cluster can communicate ACLs to CIFS clients in one of three ways:

- **NFSv3** – If the core filer supports only NFSv3, the Avere cluster translates between POSIX mode bits and CIFS ACLs.

core filers that support only NFSv3 and use POSIX mode bits do not require any special configuration.

- **NFSv4** – If the core filer supports NFSv4, the Avere cluster can use the NFSv4 domain to translate between NFSv4 ACLs and CIFS ACLs.

The Avere cluster does *not* provide file service to NFSv4 clients or from NFSv4 servers; however, it can use an NFSv4 domain to provide access control to CIFS clients.

To use NFSv4 ACLs, the cluster must be configured with the name of a valid NFSv4 domain in the **NFS Domain** field in the **Directory Services** settings, as described in Section 2.8.5, “Specifying the Source for Usernames” on page 46 for details.

If you are using NFSv4 ACLs, the directory associated with the root of the CIFS share must have an appropriate ACL defined with the correct inheritance flags before ACLs can be set by CIFS clients. This is most easily achieved by setting the directory's POSIX mode bits to 0777.

- **CIFS** – If the core filer supports CIFS natively, the Avere cluster can pass CIFS ACLs directly between the client and the core filer.

In this case, the CIFS share name on the cluster *must* match the name of the matching share on the core filer.

### 7.7.1. Considerations for Selecting Access Control

When selecting an access-control mechanism for a particular CIFS share, consider the following criteria:



#### Note

CIFS ACLs are always preferred over NFSv4 ACLs.

- Does the core filer support CIFS natively (for instance, a Data ONTAP qtree with the security style `ntfs`)? If so, you *must* use CIFS ACLs.
- Does the core filer support only NFSv4? If so, use NFSv4 ACLs.
- Does the core filer support only NFSv3? If so, use POSIX mode bits.
- Does the core filer support both NFSv3 and NFSv4 AND does the core filer have a full NFSv4 implementation? If so, then you can use either NFSv4 ACLs or POSIX mode bits.

In this case, the use of NFSv4 ACLs is recommended only if your storage-administration team is familiar and comfortable with mixed NFSv3/v4 environments.

If you are using NFSv4 ACLs or native CIFS ACLs, see Appendix A, *Core Filer-Specific Configuration Notes* on page 156 for information about configuring CIFS-tested core filers for correct operation with the Avere system.

## 7.7.2. Enabling Constrained Delegation for Use With Native CIFS ACLs

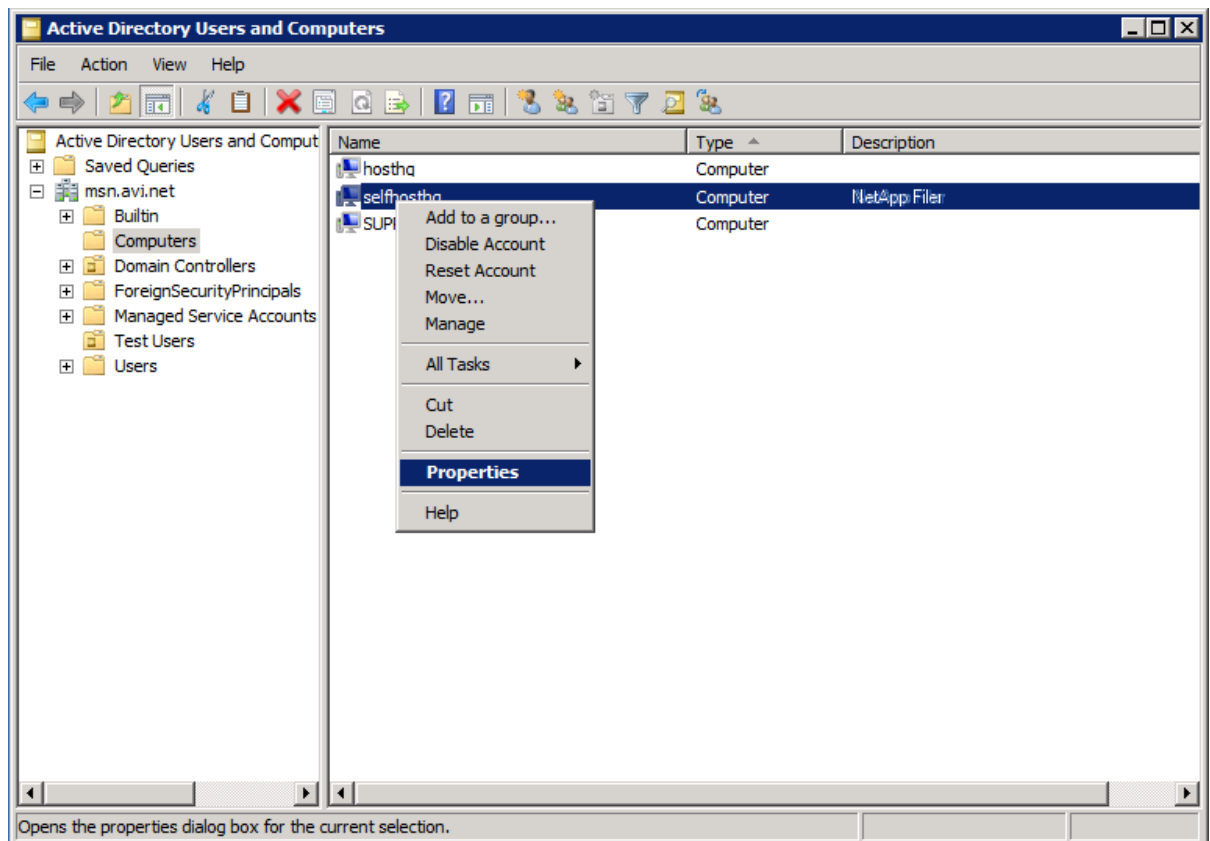
If your core filer is a native CIFS server (for example, a Data ONTAP volume with the security style `ntfs`), or if you are using any filesystem with native CIFS ACLs, you must enable constrained delegation on the Active Directory server between the core filer and the Avere cluster.

Constrained delegation configures the Microsoft Active Directory service to restrict the services and servers that your application can access. Constrained delegation requires Kerberos authentication on some servers. For more information on enabling Kerberos, refer to Section 2.8.3, “Enabling Kerberos Authentication” on page 43.

➤ To enable constrained delegation:

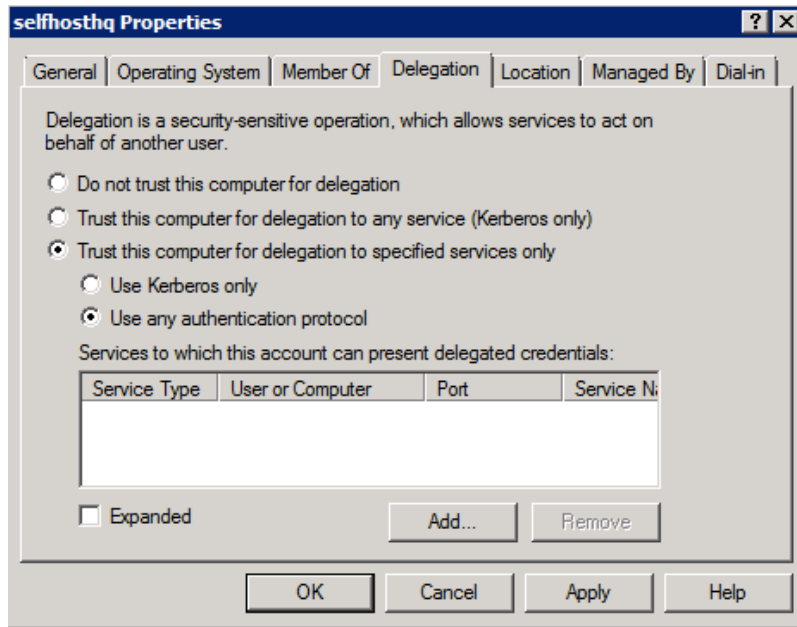
### On the Active Directory server:

1. Log in to the Active Directory server with an account that has administrative privileges.
2. From the **Start Menu**, under **Administrative Tools**, choose **Active Directory Computers and Users**.

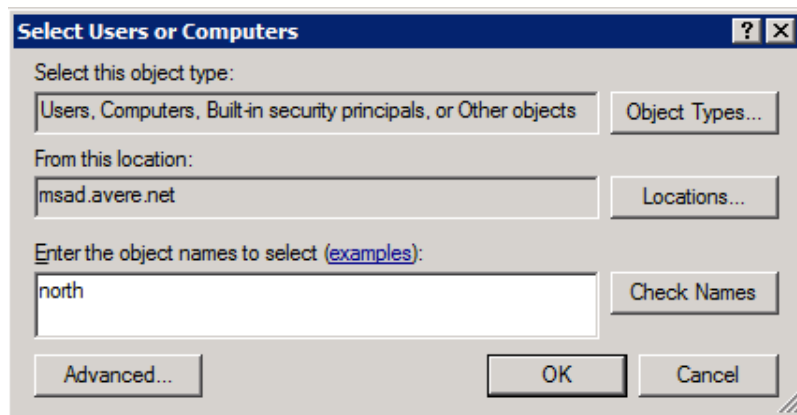


3. In the left-hand pane, expand your network item, and choose **Computers**.
4. In the right-hand pane, right click on the vserver's NetBIOS name and choose **Properties**.

5. Click on the **Delegation** tab.

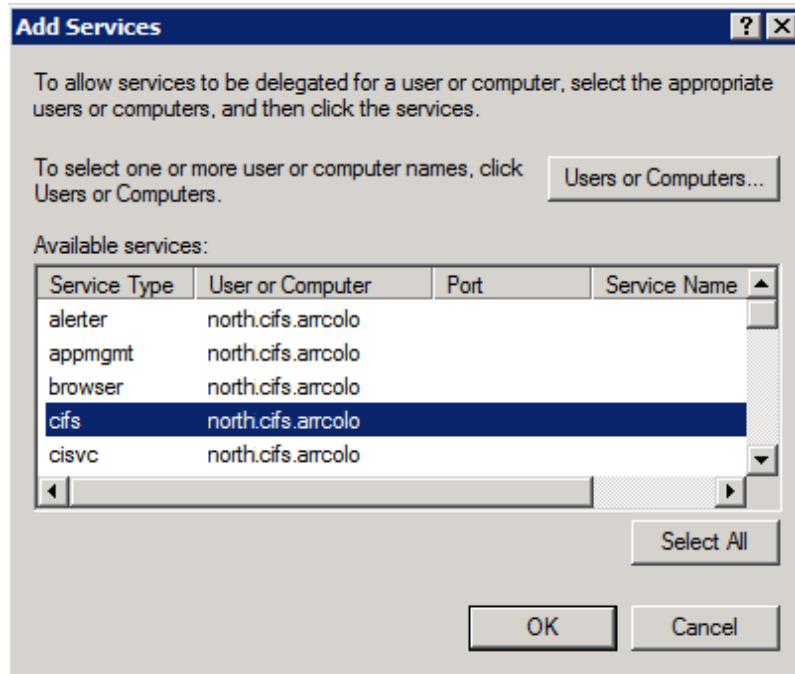


6. Select **Trust this computer for delegation to specified services only**, and then select **Use any authentication protocol**.
7. Choose **Add**. The **Add Services** window appears.
8. Find the core filer:
- Click on **Users or Computers**
  - In the bottom field of the window that appears, type part of the server name and click **Check Names**.



- Select the name of the core filer from the list that appears, and click **OK**.

You return to the **Add Services** window.



9. Select **cifs** and choose **OK**.

#### On the core filer:

1. Create a volume or filesystem (an export) that has native CIFS permissions. See the documentation for your core filer for instructions.
2. Ensure that the volume or filesystem enables UNIX root access.
3. Ensure that the volume's or filesystem's initial privileges enable full control for all users (equivalent to the POSIX mode bits 0777). Set ACLs on individual files and directories after setup is complete.

#### On the Avere system:

- Create a share for the volume or filesystem as described in Section 7.8, "Creating CIFS Shares" on page 107.



#### Important

The name of the CIFS share must *exactly* match the name of the source NFS export.

## 7.8. Creating CIFS Shares

Before CIFS clients can access data through the virtual server, you must set up CIFS shares that map to NFS exports on the Avere core filers (or to subdirectories of the exports). Perform the following steps to create shares.

CIFS shares can be either regular shares or home shares. There is no limit to the number of either type of share that can be created on an Avere cluster, and an Avere cluster can have both types.



### Caution

If you share an NFS export through CIFS, do not change the export policy or rules on the NFS export after setting up sharing. Doing so can result in unpredictable access to the export through CIFS and possible unauthorized access to data.

➤ To create a CIFS share:

1. Navigate to the **Settings > VServer > CIFS Shares** page.

The screenshot shows the 'CIFS Shares' configuration page. On the left is a sidebar with navigation links: VServer, Manage VServers, Vserver Details, Client Facing Network, Namespace, Export Policies, Export Rules, NFS, CIFS, CIFS Shares (selected), Core Filer, Cluster, and Administration. The main content area has a top navigation bar with 'Dashboard', 'Settings', 'Analytics', 'Data Management', and 'Support'. Below this is a header for 'CIFS Shares' with a dropdown for 'VServer'. A table lists existing shares:

Share name	Namespace Path	Access Control	Home	Advanced	
newg3.1share	/vol/cifshomes	posix	No	No	<button>Delete</button> <button>Modify</button>

Below the table is the 'Share definition' form:

- Share Type: Regular Share (dropdown)
- CIFS share name: arrow\_share (text input)
- NFS export: / (dropdown)
- Subdirectory: thor (text input)
- Access Control: POSIX Mode Bits (dropdown)
- Advanced: ☐
- Create share

2. If you have more than one virtual server, choose the CIFS-enabled virtual server from the drop-down list immediately below the **Manage VServers** heading on the **Settings** tab.
3. From the **Share Type** drop-down list, choose the type of CIFS share that you want to create.
  - **Regular Share** (the default) – The path to the share is determined by the NFS export path and the subdirectory path, combined.
  - **Home Share** – With a Home Share, the user's name will be substituted into the namespace path. For example, if the namespace path for the Home Share was `grap/users/%U`, and that user was mapped to the share as `juser`, the namespace path would become `grap/users/juser`. This allows an administrator to create just one share, with each user having a home directory on that share.

4. Enter a name for the new share in the **CIFS share name** field. Share names must have the following properties:
  - Share names must be unique. Note that they are case-insensitive.
  - Share names cannot contain control characters (0x00 - 0x1F)
  - Share names cannot contain any of the following characters:  
 " % \* + , / : ; < = > ? [ \ ] |
  - Share names must be at least 1 character long but no more than 242 characters long.



#### Note

The maximum length for a CIFS share name is 255 characters; however, the share name is internally used in combination with the NetBIOS name (maximum 12 characters) of the server on which the share is located. A maximum share-name length of 242 characters is therefore recommended.

5. Select the NFS export that is to be used for the share from the **NFS export** drop-down list. Multiple CIFS shares can use the same NFS export.
6. If needed, enter a subdirectory for the export in the **Subdirectory** field. The subdirectory name can include additional directory path information.

For example, if you are creating a share named **payroll** that needs to resolve to the NFS directory /vol/vol0/dept/finance/payroll and the share's NFS export is /vol/vol0/dept, enter the subdirectory as **/finance/payroll**.

If you selected **Home Share** as the share type, the **Subdirectory** field is automatically populated with **%U**, which will be translated to the user name of a user on the CIFS client.

7. Choose an access-control mechanism from the **Access Control** drop-down list. See Section 7.7, "Selecting an Access-Control Mechanism" on page 103 for information about selecting an access-control mechanism.
8. Optionally, set advanced properties for the share by selecting the **Advanced** checkbox. The Avere OS displays a list of advanced properties.
  - **Browseable** – Choose whether the share is browseable. If the share is browsable, it will be visible when a user browses to that vserver.
  - **Inherit Permissions** – Specifies whether new directories created under the share will inherit the permissions of their parent directory. The default is No.
  - **Read Only** – Choose whether the share is read-only. Setting this option to **Yes** can lead to faster performance if the data on the core filer will not need to be changed.
  - **Strict locking** – Select **Enabled** from the drop-down list so that a byte range lock check (a check on locks for a section of a file) is made each time data is read from or written to the CIFS share. You can select **Disabled** (the default) to improve performance.
  - **Oplocks** – Select **Disabled** to disable all oplock support. Select **Enabled** (the default) to allow read/write oplocks.
  - **Level 2 oplocks** – Select **Disabled** to disable all read-only oplocks. Select **Enabled** to enable read-only oplocks, *when the Oplocks option is enabled*.
  - **Read-only optimized** – Select **Yes** to enable performance-related options that are applicable to read-only shares. The default value is **No**.
  - **Create Mask** – Enter an octal value representing the UNIX permissions for newly created files. The default is **0744** (read-write-execute by owner; read-execute by group; read-execute by any user).

- **Security Mask** – Enter an octal value representing the UNIX permissions that are set on a file whose permissions are changed by a Windows NT client from the native Windows NT security dialog box. The default is **0777** (read-write-execute for all users).
  - **Directory Mask** – Enter an octal value representing the UNIX permissions of a directory that is created with DOS permissions. The default is **0755** (read-write-execute by owner; read-execute by group; read-execute by any user).
  - **Directory Security Mask** – Enter an octal value representing the UNIX permissions for a directory whose permissions are changed by a Windows NT client from the native Windows NT security dialog box. The default is **0777** (read-write-execute for all users).
  - **Force Create Mode** – Enter an octal value representing the minimum set of UNIX permissions for any file created by the Avere OS CIFS server. The default is **0000** (no permissions).
  - **Force Security Mode** – Enter an octal value representing the minimum set of UNIX permissions that can be modified on a file whose permissions are changed by a Windows NT client from the native Windows NT security dialog box. The default is **0700** (read-write-execute by owner; no permissions for others).
  - **Force Directory Mode** – Enter an octal value representing the minimum set of UNIX permissions for any directory created by the Avere OS CIFS server. The default is **0000** (no permissions).
  - **Force Directory Security Mode** – Enter an octal value representing the minimum set of UNIX permissions that can be modified on a directory whose permissions are changed by a Windows NT client from the native Windows NT security dialog box. The default is **0000** (no permissions).
  - **Force User** – Enter a UNIX username that is assigned as the default user for all users of the Avere OS CIFS server. This is useful for sharing files. Specifying an incorrect username or a username without adequate permissions can cause security (access) problems. There is no default value.
  - **Force Group** – Enter a UNIX group name that is assigned as the default group for all users of the Avere OS CIFS server. This is useful for sharing files. Specifying an incorrect group name or the name of a group without adequate permissions can cause security (access) problems. There is no default value.
9. Click the **Create share** or **Set home share** button; the latter is displayed if the share type is **Home Share**. The share appears in the table listing of shares.

### 7.8.1. Accessing a CIFS Share

After a share is created, it can be accessed from a CIFS client by using the syntax `\\NetBIOS_name\share_name`, where *NetBIOS name* is the name of the CIFS server specified in the primary CIFS configuration and *share\_name* is the name assigned to the share.

### 7.8.2. Deleting or Modifying a CIFS Share

- To delete a share, click the **Delete** button in the share's row in the table listing of shares.
- To change a share's configuration:
  1. Click the **Modify** button.  
You can change the access control or any of the advanced options.
  2. When you are finished making changes, click **Reset** to return to the original settings, or click **Modify share** to accept the changes.

## Chapter 8. Moving and Mirroring Data on Core Filers (Data Management Tab)

The **Data Management** tab allows you to migrate (FlashMove) data from one core filer to another, and to have data mirrored (FlashMirror). This capability requires you to purchase a license; for more information, refer to Section 1.4, “Adding and Removing a License” on page 6.

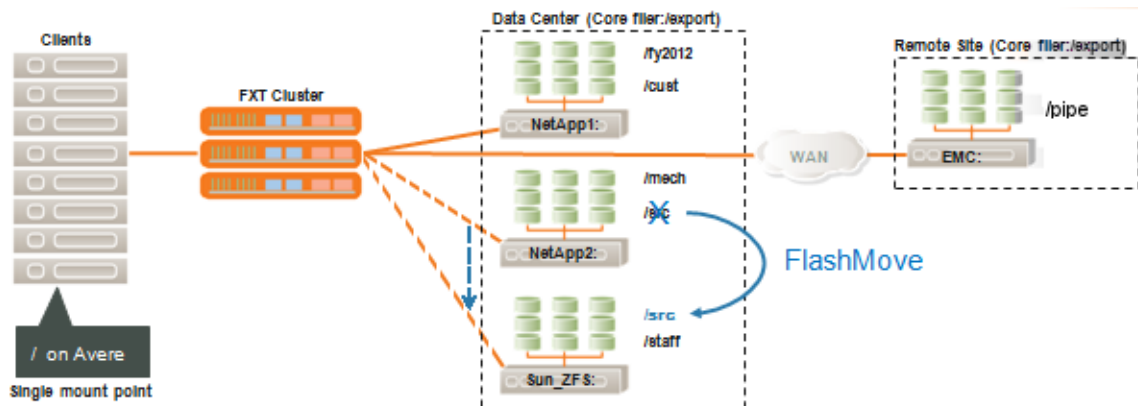
### 8.1. Understanding FlashMove

FlashMove enables you to move (copy data and change the junction between core filers) data transparently from one core filer to another, provided the two core filers are in the same GNS-enabled vserver. During the operation, data remains fully available to clients for read and write operations. For more information on global namespace (GNS) vservers, refer to Section 4.3, “Creating and Maintaining a Global Namespace” on page 72.



#### Note

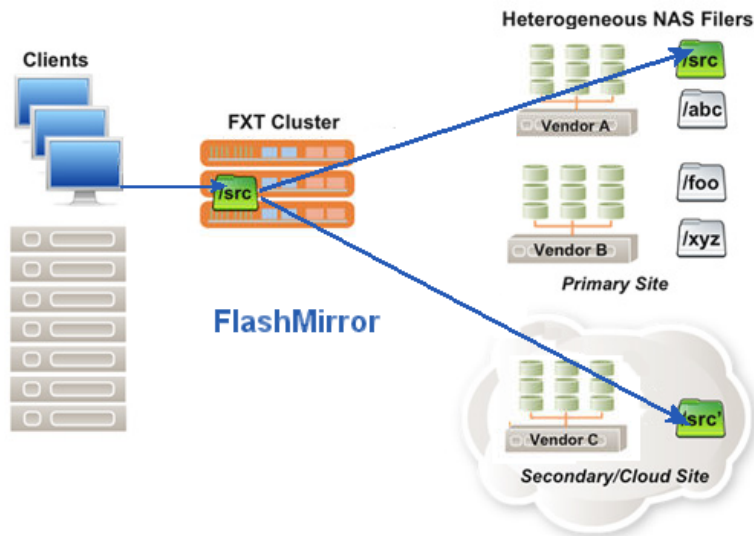
The FlashMove operation does *not* delete your original data; however, it also does not continue updating data after the operation, nor does it allow the same access to that data as before, since (from the cluster’s point of view) the data has actually “moved” from its original location.



One way to visualize FlashMove is as a copy operation from one core filer to another. After the copy is complete, the GNS namespace is automatically adjusted, so that in the future, clients see the new copy of the data, and the original copy is removed. FlashMove is useful, for instance, to capacity-balance data between several core filers, or when a server is to be decommissioned and the administrator needs to move active data elsewhere without interrupting client access to that data.

## 8.2. Understanding FlashMirror

FlashMirror is an extension of the FlashMove feature; after the initial move of the dataset from the source core filer to the destination core filer, a FlashMirror operation continues to synchronize any changes (writes) made on the source to the destination.



One way to visualize FlashMirror is as a continuing copy of the data, again, without interrupting client access to the data.

## 8.3. FlashMove and FlashMirror Prerequisites

Before you can start a data management job, your cluster must meet the following criteria:

- A valid license for FlashMove, FlashMirror, or both, activated on the Avere system.
- Because you are moving data from one core filer to another:
  - You need more than one core filer.
  - Each core filer must have at least one export available.
- As described in Section 6.2.2, “Setting Read/Write Mode” on page 87:
  - The **Core filer verification** interval must be set to **Never**.
  - A writeback delay must be set.
- Local directories must be enabled, as described in Section 6.5.2, “Enabling Local Directories” on page 94.



### Note

Because enabling and disabling local directories on a core filer can take a large amount of time, reducing this time by reducing the writeback delay is recommended, as described in Section 6.2.2.1, “Determining the Maximum Writeback Delay” on page 87. After the data is flushed back to the core filer, you can then enable or disable local directories.

- Neither the source nor the destination core filer can be associated with a simple vserver.
- The vserver must either use a global namespace, as described in Section 4.3, “Creating and Maintaining a Global Namespace” on page 72, or the destination core filer must not be associated with a vserver (for FlashMove operations).



## Important

Simple vserver (one vserver to one core filer) configurations cannot be used with data management.

- Both the source and destination users need to have domain administrator permissions, at a minimum.
- Both the source and destination core filers need to be joined to the same domain, or to a trusted domain.

## 8.4. Creating and Running a Data Management Job

You can enter a search term in the field above the table of job listings to filter the display, as described in Section 1.1.11, “Search Functionality” on page 3.

Click on the **Data Management** tab from the Avere Control Panel to configure and run a migration.

The screenshot displays the Avere Control Panel interface with the 'Data Management' tab selected. The top navigation bar includes 'Dashboard', 'Settings', 'Analytics', 'Data Management', and 'Support'. The user is logged in as 'admin' on 'V3.1.1-6a7c480' of a 'LiGo Cluster'. Below the navigation bar, there are buttons for 'Create', 'Start', 'Dismiss', and an 'Actions' dropdown. A search field is present with the text 'Showing 1 to 1 of 1 entries'.

Job ID	Source	Destination	Job Type	Status	Throughput Bytes/Files/Dirs	Create Time	Run Time Days:Hours:Minutes:Seconds
2	grape:/vol/haus_src1	thor:/vol0/arrow	move	move completed	0B / 0 / 0	1/29/14 12:35:45	00:00:00:55

Below the table, the details for the selected job are shown:

**Move Job:** grape:/vol/haus\_src1 => thor:/vol0/arrow  
**Source Admin Username:** Administrator  
**Destination Admin Username:** Administrator

	Total Moved	Throughput (/sec)
Bytes	3KB	0B
Files	10	0
Directories	2	0

**Node:** node1  
**Operational State:** complete  
**Created:** 1/30/14 16:06:09  
**Started:** 1/30/14 16:07:49  
**Completed:** 1/30/14 16:08:44  
**File Logging:** enabled  
**Logging File:** mirror\_job.txt  
**Moving:** /difftest1.txt  
**Note:** CIFS migration

**Status Message:** N/A  
**Overwrite Mode:** always

At the bottom, there are buttons for 'Create', 'Start', 'Dismiss', and an 'Actions' dropdown menu with options: 'Abort', 'Pause', 'Stop', 'Reverse', and 'Transition'.

➤ To create a FlashMove or FlashMirror job:

1. From the **Data Management** tab, click on the **Create** button at the top. The **Add New Data Management Job** wizard starts.

The screenshot shows a window titled "Add New Data Management Job" with a close button (X) in the top right corner. The window contains a "Configuration" section with the following fields and options:

- Initial state:** A dropdown menu set to "Stopped".
- Job type:** A dropdown menu set to "Mirror".
- Overwrite mode:** A dropdown menu set to "Filehandle, size or modification time changed".
- Enable file logging:** A checked checkbox.
- Log filename (optional):** A text input field containing "mirror\_job.txt".
- Preserve CIFS ACLs:** A checked checkbox.
- Overwrite destination root ACL:** An unchecked checkbox.
- CIFS vserver:** A dropdown menu set to "new\_global".
- Notes:** A text input field containing "CIFS migration".

At the bottom of the window, there are three buttons: "Back", "Next", and "Add Job".

2. Enter the information about how the job will be performed:

- **Initial State** – This will be **Stopped** by default.

Select **Start Now** if you want the migration to start when you create the job, rather than using the **Start** button later.

- **Job Type** – Either **Move** or **Mirror**.
- **Overwrite Mode** – How files on the destination core filer will be treated by the source core filer.
  - **Always** – Use this option when you are using FlashMove or FlashMirror to migrate data to a location that has not been previously populated by the source core filer. This copies all data in all cases. If you are unsure about which option to use, use this one (although it might take a bit longer than the other options).
  - **Filehandle, size, or modification time changed** – Use this option if you have halted and re-created a migration. This will only migrate data that has not yet been transferred, unless it has changed since the cancellation. This can save time if you cancelled a migration when a large amount of the operation was complete.
  - **Size or modification time changed** – Use this option when you have copied files to the new core filer in another way (for example, a simple copy), but want to update anything that has changed on the source core filer. The option checks normal NFS attributes on a file, and copies files if the size or date have changed.
- Check **Enable file logging** if you want a log file to be created. If you do this, the **Log filename** field will appear, allowing you to enter a name for the file. If you do not enter a name, Avere OS will use `.avere_log_migration[ID]`. The file is created in the destination export directory, with no additional file extension.
- Check **Preserve CIFS ACLs** if you want the access control settings on CIFS shares to be the same as the original settings.



### Important

CIFS migrations cannot be used with NFSv4 ACLs. The source and destination core filers must be configured to use CIFS ACLs.

Two additional fields appear:

- **Overwrite destination root ACL** – Check this option if you want to transfer the ACL settings to the destination. This will either overwrite or transfer the security settings to the destination.



### Note

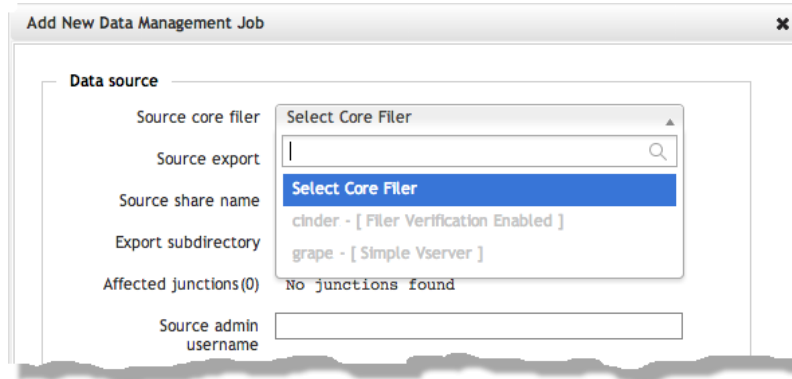
If you are migrating between different types of core filers, you should select this option, or the migration is likely to fail.

- **CIFS vservers** – Select the GNS vserver containing the CIFS share from this drop-down list. (Only GNS vservers will be included in the list.)
  - **Notes** – Enter any notes, such as a reason for data management job, in this field.
3. After you have entered the information in the first window, click **Next**.

## Enter Data Source Parameters

4. Enter the **Data Source** information. This will be the location of the dataset that will be migrated.
- **Source Core Filer** – From this drop-down list, choose the core filer from which you want to transfer information, any of the core filers you have added to the cluster.

If any of the core filers are configured so that they cannot be migrated, they will be unselectable, with the reason for their unavailability listed by the name, as shown in the following image:



## Enter Data Destination Parameters

5. Enter the **Data Destination** information. This will be the location *to where* the dataset will be copied.

**Add New Data Management Job**

**Data source**

Source core filer: cinder

Source export: /

Source share name: newg3.1share

Export subdirectory: /vol/cifshares

Affected junctions(0): No junctions found

Source admin username: Administrator

**Data destination**

Destination core filer: new\_core

Destination export: /vol0/1

Destination share name: migrated\_share

Export subdirectory: /vol/cifshares

Affected junctions(0): No junctions found

Destination admin username: Administrator

Back Next Add Job

- **Destination Core Filer** – From this drop-down list, choose the core filer to which you want to copy, any of the core filers you have added to the cluster.
- **Destination Export** – From this drop-down list, choose the export (the filesystem) you want to copy to.



### Note

You can type part of the export name to filter the list options.

- **Export Subdirectory** – Optionally, enter any subdirectory to which you want to move or mirror the data. If you do not enter a subdirectory, the data will be migrated to the root of the destination core filer. If you enter a subdirectory that does not yet exist, the Avere OS will create it.
- **Affected Junctions** (informational) – Links from one filesystem to a directory in another filesystem.



### Important

The destination core filer cannot have a junction pointing to itself; for example, Filer1:/vol1/data/ cannot be moved (or mirrored) to Filer1:/vol1/newdata.

6. Click the **Add Job** button.

The data management job (FlashMove or FlashMirror) is added to the **Data Management** tab. A window appears, informing you when the job has been added successfully. Choose **Close** to dismiss the window.

The screenshot shows the 'Data Management' tab in the LiGo Cluster interface. A job with ID 2 is selected, showing details for a move operation from 'grape:/vol/haus\_src1' to 'thor:/vol0/arrow'. The job status is 'move completed'. The details panel includes fields for Source Admin Username, Destination Admin Username, Node, Operational State, Created, Started, Completed, File Logging, Logging File, Moving, and Note. A table shows 'Total Moved' and 'Throughput (/sec)' for Bytes, Files, and Directories. The 'Actions' dropdown menu is open, showing options: Abort, Pause, Stop, Reverse, and Transition.

Total Moved		Throughput (/sec)	
Bytes	3KB		0B
Files	10		0
Directories	2		0

## Running the Data Management Job

➤ To start a FlashMove or FlashMirror job:

1. On the **Data Management** tab, click on the job row. It is highlighted in blue.



### Note

You can click on the arrow to the right of the job to view status information about the job.

2. If you chose **Start Now** as the **Initial state** in the first wizard window, the migration starts. Otherwise, click the **Start** button on the **Data Management** tab. As the job progresses, its status will appear in the job information.

You then have the following options on the **Action** drop-down list (top or bottom of the panel):

- **Abort** – Completely stops the data management job (you cannot re-start the job). You will need to choose **Dismiss** to remove the job from the list.
- **Pause** – Continues mirroring data already transferred, but stops transferring data. Choose **Action > Start** to continue a paused job from where it paused.
- **Stop** – Stops transferring *and* mirroring data. Choose **Action > Start** to restart a stopped job from the beginning.
- **Reverse** – Swaps the source and destination exports, while continuing to keep them in sync. In other words, the “authoritative” copy is switched.
- **Transition** – Finishes the migration in a current FlashMirror job, and then halts mirroring (transitions from a FlashMirror job to a FlashMove job).

In addition, you can click on the **Dismiss** button to remove a stopped, aborted, or completed job from the list.

## Chapter 9. Monitoring the Cluster (Dashboard Tab)

Use the **Dashboard** tab to monitor how well a cluster is operating, and for problems. Situations that need attention are displayed as conditions and alerts.

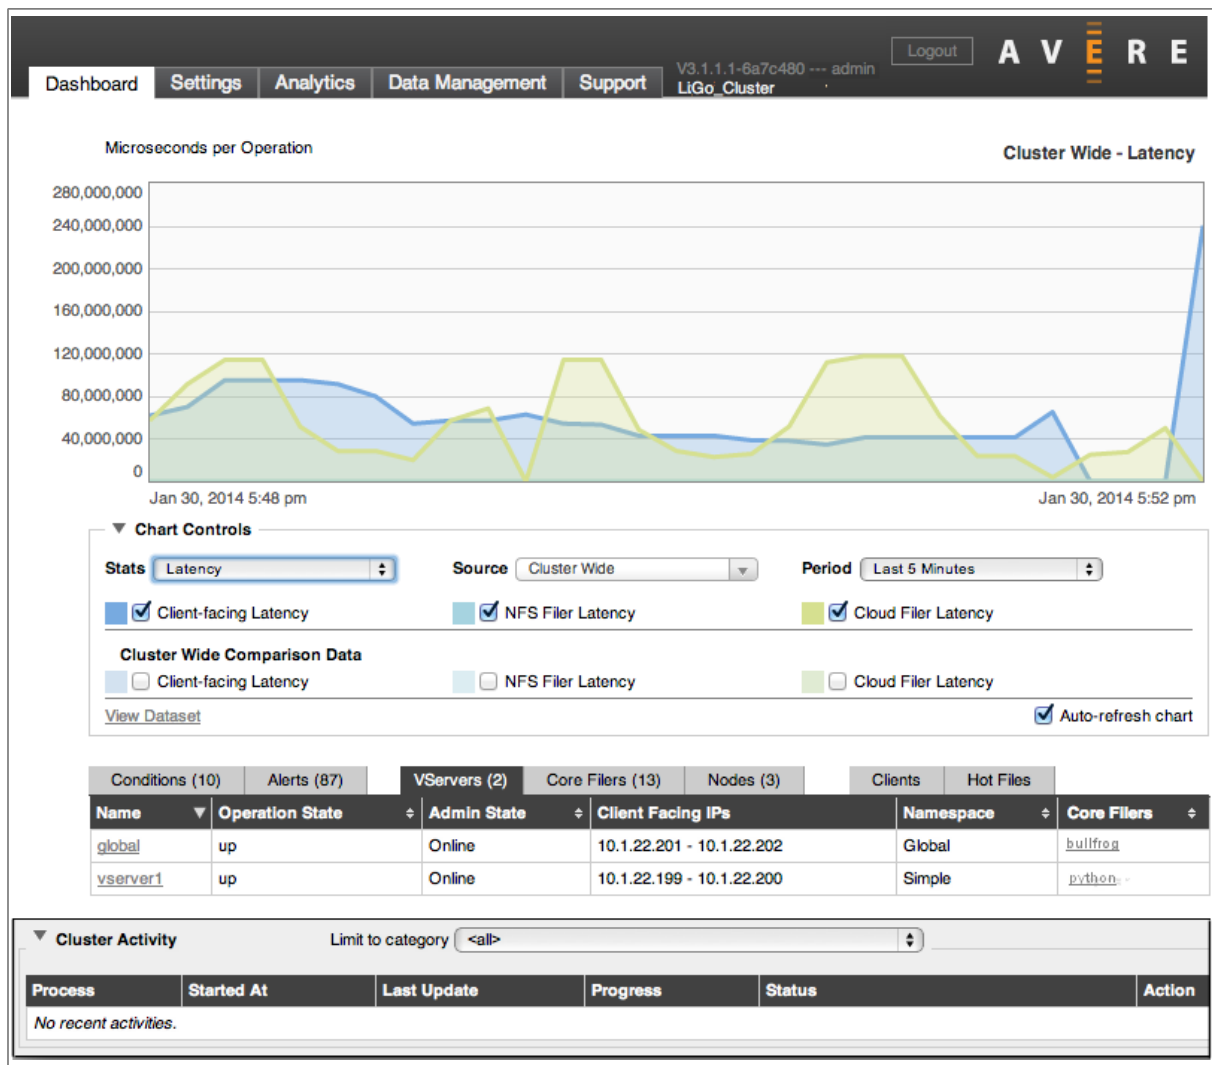


### Important

Avere recommends that you check the **Dashboard** tab at least daily and investigate any red system errors or yellow alerts. For more information, refer to Section 9.3, “Monitoring Conditions and Alerts” on page 121.

### 9.1. Overview of the Dashboard

The **Dashboard** tab is divided into the following sections:



- In the center, the performance graph displays system performance by user-selected criteria, as described in Section 9.2, “Viewing System Performance” on page 118.



## Note

The Advanced Statistics grapher has been superseded by the standard Avere Control Panel Dashboard. If you have questions about transitioning from the Advanced Statistics grapher to the Dashboard, contact Avere Global Services.

- Under the performance graph, the status bar provides easy access to information about the cluster. It includes the following tabs, described later in this chapter:
  - **Conditions** – Current system errors and alerts that are affecting the operation of the cluster.
  - **Alerts** – A list of resolved system alerts and notifications.
  - **VServers** – A list of virtual servers (vservers), including each virtual server's name, operational state, administrative state, client-facing IP addresses, type of namespace (simple or global), and core filers associated with it.
  - **Core Filers** – A list of core filers, including each core filer's internal name, operational state, network name, number and age of modified files not yet written back to it, and vservers associated with it.
  - **Nodes** – The nodes contained in the cluster, including each node's name, state, client-facing IP addresses, current CPU usage, and performance summary.
  - **Clients** – User-selectable information about clients using the cluster.
  - **Hot Files** – User-selectable information about the most active files on the cluster.
- Click the triangle to the immediate left of the **Cluster Activity** label at the bottom of the **Dashboard** tab. This will expand the bar to show the following current and recent cluster-wide activities.
  - Process name
  - Time at which the process started
  - Last update time for the process
  - Progress for the activity, if available
  - Current process status
  - Action to take, if applicable

Choose <all> from the the **Limit to category** drop-down list to display all current and recent activities, or choose <in progress> to display only current activities.

## 9.2. Viewing System Performance

You can view the performance of the FXT system by setting what is displayed on the Dashboard's graph. For more information on graphs, refer to Chapter 10, *Using Graphs (Analytics Tab)* on page 135.

- To expose the chart controls, click the triangle to the immediate left of the **Chart Controls**.
  - The **Auto-refresh chart** checkbox is enabled by default to automatically refresh the display. You can disable and re-enable it at will.
  - Optionally, select different colors for one or more displayed statistics by clicking the colored square next to the statistic's listing in the **Chart Controls** area and selecting a new color from the color grid. To keep the same color, click the statistic's original colored square again.

## 9.2.1. Choosing the Statistics to be Displayed

- To set what statistics will be shown on the performance graph:
1. From the **Stats** drop-down list, choose the *type* of statistics you want to include on the graph. Choices include:
    - **Ops/Second** for the number operations the selected source (cluster, vservers, node, or core filer) is making each second. The options for this selection are listed in Section 9.2.1.1, “Graphing Options for Ops/Sec” on page 119.
    - **Throughput** for throughput. The options for this selection are listed in Section 9.2.1.2, “Graphing Options for Throughput” on page 120.
    - **Latency** for latency. The options for this selection are listed in Section 9.2.1.3, “Graphing Options for Latency” on page 121.
  2. From the **Period** drop-down list, choose the time period over which to display statistics. Choices range from the last five minutes to the last seven days, with a **Date/Time** option that allows you to set a customized time period.
  3. From the **Source** drop-down list, choose the *source* of the statistics you want to include on the graph. Choices include:
    - **Cluster Wide** for combined statistics for all vservers, core filers, and nodes in the cluster.
    - **VServers > vservers#** for statistics from the selected virtual server.
    - **Core Filer > core filer#** for statistics from the selected core filer.
    - **Nodes > node\_name** for statistics from the selected node.
  4. Check **Auto-refresh chart** to automatically update the graph. Graphs will update every 30 seconds for users with read-only access to the Avere Control Panel, and every 10 seconds for users with read/write access.

### 9.2.1.1. Graphing Options for Ops/Sec

The screenshot shows the 'Chart Controls' panel for 'Ops / Second' statistics. It includes three main sections: 'Stats' (set to 'Ops / Second'), 'Source' (set to 'Cluster Wide'), and 'Period' (set to 'Last 5 Minutes'). Below these are three columns of checkboxes for enabling or disabling specific statistics. The first column, under 'Client Ops', has a checked box for 'Client Ops' and an unchecked box for 'Data Management Ops'. The second column, under 'Sync Core Filer Ops', has a checked box for 'Sync Core Filer Ops' and an unchecked box for 'Data Management Ops'. The third column, under 'Async Core Filer Ops', has a checked box for 'Async Core Filer Ops' and an unchecked box for 'Data Management Ops'. At the bottom, there is a 'View Dataset' link and a checked 'Auto-refresh chart' checkbox.

Chart Controls		
<b>Stats</b> Ops / Second	<b>Source</b> Cluster Wide	<b>Period</b> Last 5 Minutes
<input checked="" type="checkbox"/> Client Ops	<input checked="" type="checkbox"/> Sync Core Filer Ops	<input checked="" type="checkbox"/> Async Core Filer Ops
<input checked="" type="checkbox"/> Data Management Ops		
<b>Cluster Wide Comparison Data</b>		
<input type="checkbox"/> Client Ops	<input type="checkbox"/> Sync Core Filer Ops	<input type="checkbox"/> Async Core Filer Ops
<input type="checkbox"/> Data Management Ops		
<a href="#">View Dataset</a>		<input checked="" type="checkbox"/> Auto-refresh chart

When you select **Ops/Sec**, you can enable or disable the display of the following statistics on the graph:

- **Client Ops** for total operations exchanged between clients and the Avere cluster
- **Sync Core Filer Ops** for synchronous operations from the cluster to the core filer
- **Async Core Filer Ops** for asynchronous operations from the cluster to the core filer

If you are viewing statistics from a source other than the cluster, and want to compare the current source's statistics to cluster-wide statistics, you can enable or disable the same statistics by using the checkboxes under the **Cluster Wide Comparison Data** heading.

### 9.2.1.2. Graphing Options for Throughput

▼ Chart Controls

Stats Throughput

Source Cluster Wide

Period Last 5 Minutes

☒ Client Read
 ☒ Read-ahead
 ☒ Fill from Core Filer

☒ Node-to-node Fill
 ☒ Node-to-node Write
 ☒ Client Write

☒ Sync Write to Core Filer
 ☒ Async Write to Core Filer
 ☒ Data Management Read

☒ Data Management Write

Cluster Wide Comparison Data

☒ Client Read
 ☐ Read-ahead
 ☒ Fill from Core Filer

☒ Node-to-node Fill
 ☐ Node-to-node Write
 ☐ Client Write

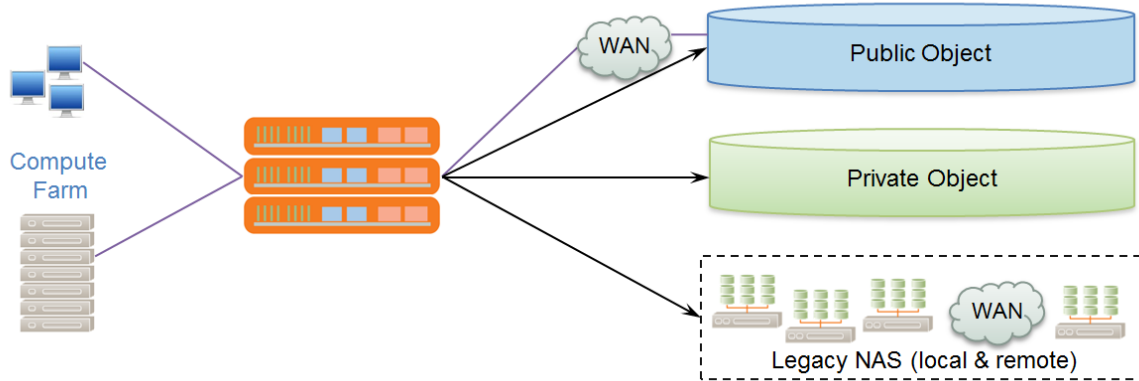
☐ Sync Write to Core Filer
 ☐ Async Write to Core Filer
 ☐ Data Management Read

☐ Data Management Write

[View Dataset](#)
☒ Auto-refresh chart

Throughput statistics tell you how well communication is flowing between clients, the cluster, and the core filers.

#### Workstations



Choose **Throughput** to graph any of the following statistics:

- **Client Read** – READ operations *from* clients *to* the cluster.
- **Read-ahead** – Read-ahead operations from clients to the cluster. These operations occur when clients request access to data that will is not currently on the cluster, and the cluster needs to then retrieve the data from the core filer. Requesting data ahead of time speeds up access.
- **Fill from Core Filer** – READ operations *from* the cluster *to* the core filer. These operations occur when clients request access to data that is not currently on the FXT Series cluster.
- **Node-to-node Fill** – Cluster READ operations among nodes in the cluster. These operations distribute file information across the cluster for load balancing and optimal performance.
- **Client Write** – WRITE operations from clients to the cluster.
- **Sync Write to Core Filer** – Synchronous WRITE operations from the cluster to the core filer. These operations occur when clients change data and the cache policy is set to **Read**, as discussed in Chapter 6, *Setting the Cache Policy* on page 84. The cluster sends write operations directly to the core filer.

- **Async Write to Core Filer** – Asynchronous WRITE operations from the cluster to the core filer. These operations occur when the cache policy is set to **Read/Write**, as discussed in Chapter 6, *Setting the Cache Policy* on page 84.

If you are viewing statistics from a source other than the cluster, and want to compare the current source's statistics to cluster-wide statistics, you can enable or disable the same statistics by using the checkboxes under the **Cluster Wide Comparison Data** heading.

### 9.2.1.3. Graphing Options for Latency

The screenshot shows the 'Chart Controls' panel. At the top, there are three dropdown menus: 'Stats' set to 'Latency', 'Source' set to 'Cluster Wide', and 'Period' set to 'Last 5 Minutes'. Below these are three checkboxes, all of which are checked: 'Client-facing Latency' (blue), 'NFS Filer Latency' (teal), and 'Cloud Filer Latency' (green). Under the 'Cluster Wide Comparison Data' heading, there are three more checkboxes, all of which are unchecked: 'Client-facing Latency' (light blue), 'NFS Filer Latency' (light teal), and 'Cloud Filer Latency' (light green). At the bottom left is a 'View Dataset' link, and at the bottom right is a checked 'Auto-refresh chart' checkbox.

Throughput statistics tell you how quickly the clients, the cluster, and the core filers are communicating.

Choose **Latency** to graph any of the following statistics:

- **Client-facing Latency** – Latency between clients and the cluster.
- **Core Filer Latency** – Latency between the cluster and the core filer.

If you are viewing statistics from a source other than the cluster, and want to compare the current source's statistics to cluster-wide statistics, you can enable or disable the same statistics by using the checkboxes under the **Cluster Wide Comparison Data** heading.

### 9.2.2. Downloading Statistics Data

You can view and download the currently displayed set of statistics in comma-separated values (CSV) format. Click **Download Dataset** at the bottom of the **Chart Controls** area, and a separate browser window or tab opens with the CSV-formatted statistics displayed in text. Use your browser's "Save As" function to save the statistics to a separate file.

## 9.3. Monitoring Conditions and Alerts

Conditions and alerts are displayed in the Avere Control Panel to give you information about unusual behavior in the cluster, including the severity of the notification, the time of the notification, and a description of the condition or alert, with suggestions for corrective action when available.

An *alert* is a transient and self-limited situation, often a one-time occurrence. An example of an alert is a software-image upgrade. An alert will remain visible until it is dismissed.

A *condition* is a situation that will probably require some action on the part of the cluster administrator. Conditions can be hidden from the display, but not dismissed (removed).


Conditions and alerts can be one of three levels of severity:

- Informational. This has no color coding.

Conditions (1)	Alerts (0)	VServers (2)	Core Filers (13)	Nodes (1)	Clients	Hot Files
<div> Show Hidden (0) Hide Hide All Occurrences Manage Hidden </div>						
<input checked="" type="checkbox"/>	2012/10/03_13:37:19	The active cluster software image is transitioning from AvereOS_V3.0.0.1-ab1ff63-cc to AvereOS_V3.0.0.1-1dea01d.				

- A *system error*, color-coded red, indicates that the cluster cannot serve some or all data. An example of a system error is loss of communication with one of the nodes in the cluster.


When a system error occurs, a image is displayed at the top of the Dashboard, *and* information about the error is displayed on either the **Conditions** or **Alerts** tab.



Conditions (1)	Alerts (28)	VServers (1)	Core Filers (13)	Nodes (1)	Clients	Hot Files
<div> Show Hidden (0) Hide Hide All Occurrences Manage Hidden </div>						
<input checked="" type="checkbox"/>	<input type="checkbox"/>	2012/10/03_13:50:44 The node vantaudi3 currently is not ready to serve data for one or more vservers. Some client interfaces might be affected. <a href="#">[more details]</a>				

- A *performance notification*, color-coded yellow, indicates that performance, data access, or both are possibly degraded. An example of an alert is the temporary suspension of virtual-server (vserver) access while the cluster automatically rebalances its virtual interfaces (vifs).

When a performance notification occurs, a image is displayed at the top of the Dashboard, *and* information about the error is displayed on either the **Conditions** or **Alerts** tab.



Conditions (1)	Alerts (28)	VServers (1)	Core Filers (13)	Nodes (1)	Clients	Hot Files
<div> Dismiss Auto Dismiss Select all Deselect all Alert History Manage Auto Dismiss </div>						
<input type="checkbox"/>	<input type="checkbox"/>	2012/10/01_15:49:24 The Filesystem Service on node vantaudi3 has been restarted. <a href="#">[more details]</a>				

### 9.3.1. Managing Conditions

When a condition occurs, the Avere Control Panel provides detailed information under the **Conditions** tab on the Dashboard's status bar. This information allows you to decide what, if anything, to do about the condition. In most cases, conditions that require action provide suggestions in the description text. If you need help to understand the message or resolve the issue, contact Avere Global Services.

Conditions (1)	Alerts (28)	VServers (1)	Core Filers (13)	Nodes (1)	Clients	Hot Files
<div> Show Hidden (0) Hide Hide All Occurrences Manage Hidden </div>						
<input checked="" type="checkbox"/>	<input type="checkbox"/>	2012/10/03_13:50:44 The node vantaudi3 currently is not ready to serve data for one or more vservers. Some client interfaces might be affected. <a href="#">[more details]</a>				

Unlike alerts, conditions cannot be simply dismissed; they must resolve themselves or be resolved by an administrator before they are removed from the list of active conditions. They can, however, be hidden, so that older persistent conditions will no longer be listed on the main tab.

### 9.3.1.1. Hiding Conditions

- To hide one or more conditions, do one of the following:
  - Click on the checkbox to the left of the condition time and choose **Hide**. Any selected conditions will no longer be displayed on the **Conditions** tab.
  - Choose **Hide All Occurrences**. All active conditions will be removed from the **Conditions** tab.

To display hidden conditions, choose **Manage Hidden**. You will be taken to the same page that allows you to change automatic dismissal of alerts, described in Section 9.3.3, “Unhiding Conditions and Canceling Auto-Dismiss” on page 125.

### 9.3.1.2. Clearing Conditions



#### Caution

In general, conditions will provide you with valuable information about your cluster’s performance. You should only clear unresolved conditions in very limited cases.

- To clear all conditions:
  1. Navigate to the **Settings > Administration > System Maintenance** page.
  2. In the **Monitoring** area, choose **Clear all conditions**.

### 9.3.1.3. Condition Examples

The following example shows a condition that automatically cleared; no intervention is needed:

2009/07/09_12:29:12	<i>Cleared alert:</i> A tokenmgr on node averel is currently initializing. It is common for a subset of the data in the cluster to be inaccessible during this time (tokenmgr=b17323f4-a3f4-4dac-acac-30c589e6814e).
2009/07/09_12:29:06	A tokenmgr on node averel is currently initializing. It is common for a subset of the data in the cluster to be inaccessible during this time (tokenmgr=b17323f4-a3f4-4dac-acac-30c589e6814e).

The following example is a condition that, if it persists, requires administrative investigation and possible intervention:

2009/07/06_17:44:43	The cluster vifs are being rebalanced. All servers will be suspended until the rebalance has completed.
---------------------	---

The following example is a system error that cannot be resolved without intervention, possibly including assistance from Avere Global Services:

2009/07/07_16:39:36	The node avere2 could not be initialized. The NVRAM header is invalid.
---------------------	--



### 9.3.3. Unhiding Conditions and Canceling Auto-Dismiss

➤ To display hidden conditions, or to cancel auto-dismissal of alerts:

1. Do one of the following:

- From the **Conditions** tab, choose **Manage Hidden**.
- From the **Alerts** tab, choose **Manage Auto Dismiss**.

You will be taken to **Hidden Alerts** on the **Settings** tab, where you can select individual conditions and alerts.

	Description	Type	Occurrences To Hide/Dismiss
<input type="checkbox"/>	HA is now fully configured.	alert	all
<input type="checkbox"/>	The local directories feature is currently enabled and there are only two nodes in the cluster. If one of the nodes is lost, while the other is not able to service local directory operations, the local directories state could become inconsistent and data could be lost. Add another node to the cluster to prevent this possible failure.	condition	2014/01/29_11:25:28
<input type="checkbox"/>	High Availability (HA) is enabled in a two node cluster. Cluster operations and data access will be disrupted if another failure occurs. Add another node to the cluster, or configure a core filer for use with HA <a href="/txt/ha.php">here</a> .	condition	2014/01/29_11:25:29

Deselect All   Select All   Unhide/Stop Auto-Dismissing

2. Select any of the conditions or alerts you want to remove from the **Description** list.

3. Choose **Unhide/Stop Auto Dismissing**. Selected conditions will now appear on the **Conditions** tab, and selected alerts will no longer be automatically dismissed.

### 9.3.4. Viewing Condition and Alert History

- To view a list of resolved and dismissed conditions and alerts, choose **Alert History** from the **Alerts** tab.

You are taken to the Alert History page.

Conditions are kept on the Alert History page until they are resolved; alerts are retained until the history buffer is filled, at which point the oldest alerts are purged to make room for new conditions and alerts. The number of conditions and alerts the history buffer can hold varies by type, length, and other factors.

Dashboard Settings Analytics Data Management Support V3.0.0.1-1dea01d --- admin		
Alert History		
	2012/10/03_15:24:19	<i>Cleared alert:</i> A Cluster Data Manager on node L1 is currently initializing. It is common for a subset of the data in the cluster to be inaccessible during this time. <a href="#">[more details]</a>
■	2012/10/03_15:24:17	A Cluster Data Manager on node L1 is currently initializing. It is common for a subset of the data in the cluster to be inaccessible during this time. <a href="#">[more details]</a>
	2012/10/03_15:24:07	<i>Cleared alert:</i> The active cluster software image is transitioning from AvereOS_V3.0.0.1-1dea01d to AvereOS_V3.0.0.1-1dea01d.
	2012/10/03_15:24:04	<i>Cleared alert:</i> The server global on node L1 is initializing. Access is suspended for this node.
■	2012/10/03_15:23:58	The server vserver1 on node L1 is initializing. Access is suspended for this node.
	2012/10/03_15:23:56	<i>Cleared alert:</i> The node L1 is not currently serving data because the log is replaying.
■	2012/10/03_15:23:50	The node L1 is not currently serving data because the log is replaying.
	2012/10/03_15:23:28	<i>Cleared alert:</i> Unable to validate cluster DNS server settings. <a href="#">[more details]</a>
■	2012/10/03_15:23:27	Unable to validate cluster DNS server settings. <a href="#">[more details]</a>
	2012/10/03_15:23:26	The cluster images have been updated. The active image is AvereOS_V3.0.0.1-1dea01d; the alternate image is AvereOS_V3.0.0.1-1dea01d.
	2012/10/03_15:20:12	The active cluster software image is transitioning from AvereOS_V3.0.0.1-1dea01d to AvereOS_V3.0.0.1-1dea01d.

If a condition or alert is listed as a *Cleared alert*, the cluster automatically resolved the condition or alert and removed it from the list of active conditions and alerts. A cleared alert often negates a previous red or yellow condition.

- Use your browser's **Back** button to return to the Dashboard.

## 9.4. VServers Tab

Conditions (0)	Alerts (2)	VServers (2)	Core Filers (3)	Nodes (1)	Clients	Hot Files
Name ▼	Operation State ↕	Admin State ↕	Client Facing IPs	Namespace ↕	Masses ↕	Hot Clients ↕
<a href="#">global</a>	up	Online	10.1.22.201 - 10.1.22.202	Global	<a href="#">thor</a>	No
<a href="#">vserver1</a>	up	Online	10.1.22.199 - 10.1.22.200	Simple	<a href="#">grape</a>	No

Information about vservers is available from the **VServers** tab in the status bar on the Dashboard. When you click the tab, the Dashboard displays the following information for each vserver:

- **Name** – The name of each vserver. You can click on the name and go to the **Vserver Details** page on the **Settings** tab, further described in Section 4.2.2, “Configuring and Modifying a Virtual Server (VServer Details Page)” on page 68.
- **Operation State** – Whether the virtual server is **up** or **down**.
- **Admin state** – The administrative state of the virtual server, any of the following:
  - **online**: The vserver is available to the cluster.
  - **removing**: The vserver is in the process of being removed from the cluster.
  - **removed**: The vserver was previously recognized by the cluster, but has been removed from the cluster configuration.
  - **suspended**: The vserver is not available to the cluster, but is still recognized by the cluster.
  - **flushing**: The vserver is transferring information from the cache to a core filer.
- **Client Facing IPs** – The IP address range available to the client machines from that vserver.
- **Namespace** – The type of namespace (simple or global). For more information on namespaces, refer to Section 4.3, “Creating and Maintaining a Global Namespace” on page 72.
- **Core Filer** – core filer associated with the vserver. You can click on the name of the core filer, and go to the **Settings > Core Filer > Name** page, further described in Section 3.3, “Maintaining Core Filers” on page 60.

Refer to Section 2.6, “High Availability” on page 33 and Section 4.1, “Virtual Servers and Namespaces” on page 64 for more information about vservers.

## 9.5. Core Filers Tab

Conditions (0)	Alerts (3)	VServers (2)	Core Filers (2)	Nodes (1)	Clients	Hot Files
Name ▼	Admin State ⇅	Cache Policy ⇅	Local Dirs ⇅	Network Name	Modified Data ⇅	VServers ⇅
<a href="#">grape</a>	Online	<a href="#">Read-Write</a> Writeback: 12 hrs	No	grape	-	<a href="#">vserver1</a>
<a href="#">thor</a>	Online	<a href="#">Read</a> Verification: 30 secs	No	thor.com	-	<a href="#">global</a>

Information about core filers is available from the **Core Filers** tab in the status bar on the Dashboard. When you click the tab, the Dashboard displays the following information for each core filer:

- **Name** – The name of each core filer. You can click on the name of the core filer to go to the core filer’s **Core Filer Details** page on the **Settings** tab, further described in Chapter 2, *Configuring the Cluster (Settings Tab / Cluster)* on page 10.
- **Admin State** – The administrative state of the core filer, any of the following:
  - **online**: The core filer is available to the cluster.
  - **removing**: The core filer is in the process of being removed from the cluster.
  - **removed**: The core filer was previously recognized by the cluster, but has been removed from the cluster configuration.
  - **invalidating**: The core filer is removing all data from the cache.
  - **suspended**: The core filer is not available to the cluster, but is still recognized by the cluster.
  - **flushing**: The cluster is transferring information from the cache to the core filer.
- **Cache Policy** – The cache policy associated with each core filer. You can click on the policy link to go to the core filer’s **Cache Policy** page, further described in Chapter 6, *Setting the Cache Policy* on page 84.
- **Local Dirs** – Whether or not the cluster can create directories on the core filer. This is defined in the **Cache Policy** page on the **Settings** tab, further described in Chapter 6, *Setting the Cache Policy* on page 84. **Local Dirs** must be enabled for FlashMove and FlashMirror operations. For more information, refer to Chapter 8, *Moving and Mirroring Data on Core Filers (Data Management Tab)* on page 110.
- **Network Name** – The host name or the IP address of the core filer.
- **Modified Files** – Number and age of files not yet written back to the core filer. These operations occur when the cache policy is set to **Read/Write**, as discussed in Chapter 6, *Setting the Cache Policy* on page 84.
- **Vservers** – Any vservers associated with the core filer (there will be a minimum of one). You can click on the name of the vserver to go to the **VServer Details** page, further described in Section 4.2.2, “Configuring and Modifying a Virtual Server (VServer Details Page)” on page 68.

## 9.6. Nodes Tab

Conditions (0)		Alerts (2)	VServers (2)		Core Filers (2)	Nodes (1)	Clients	Hot Files
Name ▾	Model ▾	State ▾	Client IPs			CPU	Performance	
<a href="#">L1</a>	FXT	up 23h 17m	vserver1:10.1.22.199 vserver1:10.1.22.200 global:10.1.22.201 global:10.1.22.202			16%	Instant: 0 ops/sec; 100% hit rate; 0 ms latency. 1m avg: 0 ops/sec; 100% hit rate; 0 ms latency.	

Currently connected to node **L1**.

When you select the **Nodes** tab, the Dashboard displays the following information for each node in the cluster:

- **Name** – You can click on the name and go directly to the detailed information page, which is the same one you would display from the **Settings** tab. For more information on managing nodes, refer to Section 2.2, “Managing FXT Nodes” on page 13.
- **Model** – The type of hardware that is running on the node, generally an FXT type node.
- **State** – Whether the node is up, and how long it has been running.
- **Client IPs** – Client IP addresses currently located on the node.
- **CPU** – Approximate CPU usage, displayed as a percentage.
- **Performance** – Performance summary, at the exact moment, and averaged over the previous minute. These numbers can provide information about whether the nodes are communicating properly within the cluster.
  - Number of operations (read, write, metadata) per second.
  - Percentage of client requests being serviced by the node and not forwarded to the core filer.
  - Latency, measured in milliseconds.

## 9.7. Clients Tab

Information about clients is available from the **Clients** tab in the status bar on the Dashboard. The **Clients** tab initially displays information about NFS clients.

Conditions (0)	Alerts (16)	VServers (1)	Core Filers (3)	Nodes (2)	Clients	Hot Files
Source	new_global	Show	Hot Clients	Collect "Hot" Information	<input checked="" type="checkbox"/> Details	Auto-refresh <input checked="" type="checkbox"/>
Client Address	FXT Node	Avere Address	Activity	Rate		
CIFS client	node7065	127.0.0.3	2 ops			
CIFS client	node7065	127.0.0.3	5 ops	Infinity ops/sec		
CIFS client	node7065	127.0.0.3	303 ops	Infinity ops/sec		

There are 3 hot clients.

By default, the client table automatically refreshes every 30 seconds. If you want to disable this behavior, deselect the **Auto-refresh** checkbox on the Clients tab of the Dashboard’s status bar. To re-enable automatic refreshing, select the checkbox again.

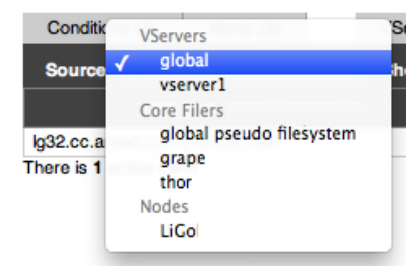


### Note

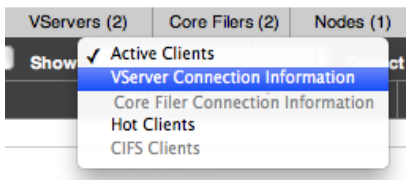
A *hot client* is a client that generates a disproportionately high amount of demand on the cluster relative to other clients. To see information about hot clients, check **Collect “Hot” Information** and then choose **Hot Clients** from the **Show** drop-down list.

You can also enable hot client collection from the **Settings > VServers > Details** page, as described in Section 4.2.2.2, “Enabling Hot Client Collection” on page 69.

From the **Source** drop-down list, you can select the vserver, core filer, or node that is connected to the client, and that you want information about.



From the **Show** drop-down list, you can select information about the source.



If you select this from the Source drop-down list:	You can choose from these options on the Show drop-down list:	And these information fields are displayed underneath:
Vserver	Active Clients, CIFS Clients	Client Address FXT Node Avere Address
	Vserver Connection Information	FXT Node Client Address Avere Address Protocol Connection Type
	Hot Clients	Client Address FXT Node Avere Address Activity Rate
Core Filer	Core Filer Connection Information	FXT Node Client Address Avere Address Protocol Connection Type
Node	Active Clients	Client Address Vserver Avere Address
	Vserver Connection Information	Vserver Client Address Avere Address Protocol Connection Type
	Core Filer Connection Information	Client Address Avere Address Protocol Connection Type

- **Activity:** Number of operations that have been done by the client in the last 30 seconds.

- **Avere Address:** Avere IP address to which the client is connected. Generally (though not always), this will also include the port.
- **CIFS Clients:** The CIFS clients accessing the selected vserver.
- **Client Address:** The client machine's IP address. Generally (though not always) this will also include the port.
- **Connection Type:** File-system connection type
- **FXT Node:** Name of the FXT Series node to which the client is connected.
- **Core Filer:** File-system connection type
- **Protocol:** IP protocol (TCP or UDP) used for the connection.
- **Rate:** Operations per second done by the client.

## 9.8. Hot Files Tab

A *hot file* is a file that has a high number of operations (read, write, metadata, or a mix) relative to other files.

Information about hot files is available from the **Hot Files** tab in the status bar on the Dashboard.

Conditions (0)	Alerts (16)	VServers (1)	Core Filers (3)	Nodes (2)	Clients	Hot Files
Select By	Ops	Cluster-wide view	Auto-refresh			
Name	Type	FXT Node	Ops/sec			
thor.com:(MASSID:2 FSID:412316926016)/	DIR	LiGo	23			
thor.com:(MASSID:2 FSID:412316926016)/deleteda	REG	LiGo	<1			
thor.com:(MASSID:2 FSID:412316926016)/deleteme	REG	LiGo	<1			

There are 3 hot files.

➤ To view information about hot files:

1. Select the **Hot Files** tab on the Dashboard's status bar.
2. From the **Select By** drop-down list, choose the measurement by which hot files are displayed:
  - **Ops:** The total operations per second performed on each file.
  - **Bytes Read/Written:** The number of bytes read or written, per second, from each file.
  - **Dir Updates:** The number of directory updates and other metadata operations performed on each file, per second.
3. Optionally, select the **Auto-refresh** checkbox to update the display automatically.
4. Optionally, select the **Cluster-wide view** checkbox to enable cluster-wide, as opposed to per-vserver, hot-file tracking.

## 9.9. Monitoring the Cluster from Outside the Avere OS

You can check or be notified of cluster alerts without logging in to the Avere Control Panel by configuring the **Settings > Cluster > Monitoring** page.

The screenshot shows the 'Monitoring' configuration page in the Avere Control Panel. The left sidebar contains a navigation menu with items: VServer, Core Filer, Cluster (expanded), General Setup, Administrative Network, Cluster Networks, FXT Nodes, High Availability, Monitoring (selected), Schedules, Directory Services, Kerberos, Login Services, Active Directory, Optimization, IPMI, Support, Licenses, Cloud Credentials, VLAN, and Administration. The main content area is titled 'Monitoring' and includes a message: 'Press the submit button to save and activate your changes.' The configuration is divided into three sections: 'Email Monitoring', 'Logs', and 'SNMP Features'. The 'Email Monitoring' section has buttons for 'Create', 'Modify', 'Remove', and 'Test'. It contains a table with one entry: 'admin@mycompany.com' for 'Cluster services,Network'. Below the table are fields for 'Mail server' (mail.mycompany.com), 'Mail-from address' (avere\_cluster1@mycompany.com), and a checked 'Include additional context' checkbox. The 'Logs' section has a 'Syslog server' field (syslog.mycompany.com) and a 'Test message' button. The 'SNMP Features' section has a checked 'Enable SNMP' checkbox and fields for 'SNMP contact' (avere\_cluster1@mycompany.com), 'SNMP location' (My Company's Data Center), 'SNMP read-only community' (public), 'SNMP trap host' (monitoring.mycompany.com), and 'SNMP trap port' (162). 'Revert' and 'Submit' buttons are at the bottom right.

Dashboard Settings Analytics Data Management Support V3.1.1.1-eabc561 --- admin LiGo\_Cluster

VServer + Monitoring

Core Filer +

Cluster -

General Setup

Administrative Network

Cluster Networks

FXT Nodes

High Availability

Monitoring

Schedules

Directory Services

Kerberos

Login Services

Active Directory

Optimization

IPMI

Support

Licenses

Cloud Credentials

VLAN

Administration +

Press the submit button to save and activate your changes.

Email Monitoring

Create Modify Remove Test

Email Addresses	Alert Categories	Actions
admin@mycompany.com	Cluster services,Network	<input type="checkbox"/>

Mail server mail.mycompany.com

Mail-from address avere\_cluster1@mycompany.com

Include additional context ☒

Logs

Syslog server syslog.mycompany.com Test message

Enable SNMP ☒

SNMP Features

SNMP contact avere\_cluster1@mycompany.com

SNMP location My Company's Data Center

SNMP read-only community public

SNMP trap host monitoring.mycompany.com

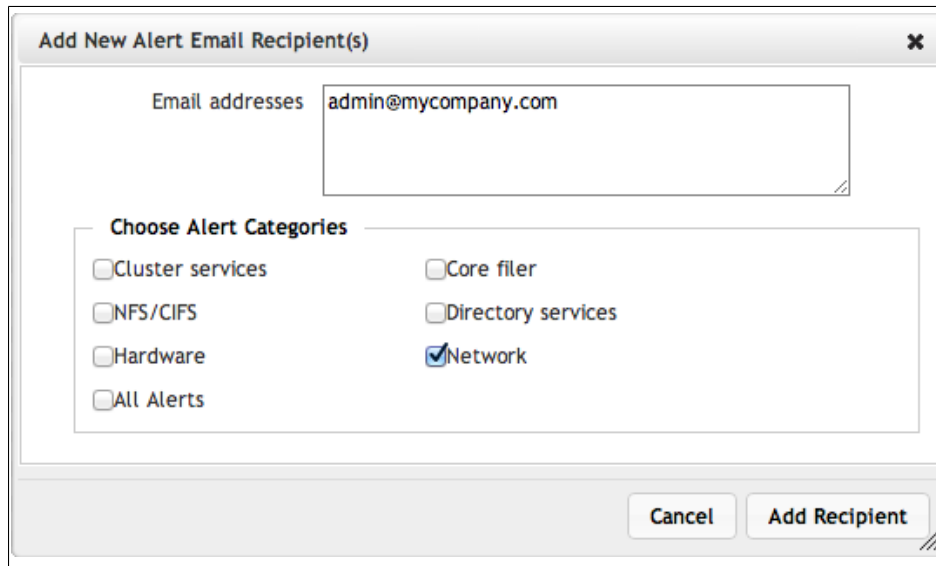
? SNMP trap port 162

Revert Submit

## 9.9.1. Sending Alerts Via Email

➤ To receive email when critical alerts occur:

1. From the **Monitoring** page, choose **Create**. The **Add New Alert Email Recipient(s)** wizard starts.



The screenshot shows a window titled "Add New Alert Email Recipient(s)" with a close button (X) in the top right corner. Inside the window, there is a text field labeled "Email addresses" containing the text "admin@mycompany.com". Below this field is a section titled "Choose Alert Categories" which contains a list of checkboxes: "Cluster services", "NFS/CIFS", "Hardware", "All Alerts", "Core filer", "Directory services", and "Network". The "Network" checkbox is checked. At the bottom right of the window are two buttons: "Cancel" and "Add Recipient".

- a. In the **Email addresses** field, enter an email address, or a comma-separated list of email addresses, to which emails are to be sent, within a given set of alert categories.
  - b. From the **Choose alert categories** area, select the types of alert that you want to have emailed, and then click on **Add Recipient**. You are returned to the **Monitoring** page.
2. In the **Mail server** field, enter the name or IP address of an SMTP server that the cluster can use to send email.
  3. In the **Mail-from address** field, enter the email address that will be included in the email's **From** field.
  4. If you want to test the email configuration, click the **Send test email** button. The cluster uses the provided information to send a test message that includes the name of the Avere cluster and the FXT node that you are logged into.
  5. Click the **Submit** button.

### 9.9.2. Specifying a Syslog Server

You can specify a syslog server to capture Avere alert messages. The syslog server must already exist.

➤ To specify a syslog server:

1. Enter the IP address or URL of an existing syslog server in the **Syslog server** field.
2. If you want to test the connectivity between the cluster and the syslog server, click the **Send test message** button. The cluster sends a test message to the syslog server.
3. Click the **Submit** button.

### 9.9.3. Configuring SNMP

You can configure the cluster to emit Simple Network Management Protocol (SNMP) messages for use by an SNMP monitor. Avere OS supports SNMPv1 and SNMPv2c. Avere's SNMP MIBs are available from the **Downloads** section of the Support portal at <https://averesystems.force.com/support/login>.

➤ To configure SNMP:

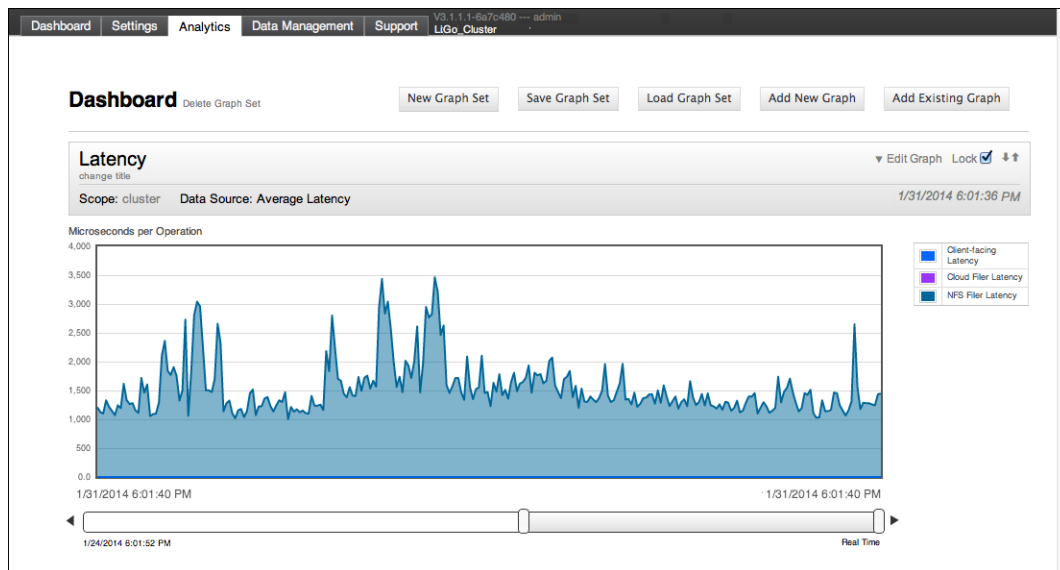
1. Select the **Enable SNMP** checkbox.
2. In the **SNMP contact** field, enter a contact name for the cluster.
3. In the **SNMP location** field, enter the location of the cluster.
4. In the **SNMP trap host** field, enter the name or IP address of the SNMP monitor host (manager) to which the cluster is to send SNMP trap notifications.
5. In the **SNMP trap port** field, enter the port number of the SNMP monitor host to which the cluster is to send SNMP trap notifications.
6. Click the **Submit** button.

## Chapter 10. Using Graphs (Analytics Tab)

The Avere Control Panel's **Analytics** tab allows you to obtain in-depth information about activity between your storage network and the FXT Series cluster. There is additional graphing information in Section 9.2, "Viewing System Performance" on page 118.

When you initially click on the **Analytics** tab, the Avere Control Panel displays the following dynamic charts by default:

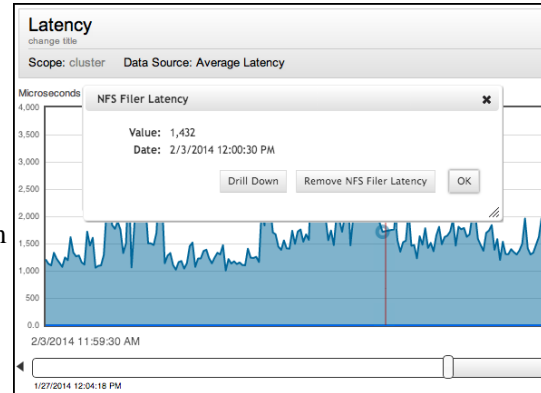
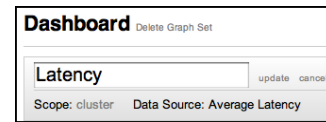
- **Latency** – Latency, in microseconds, for operations between clients and the cluster, and between the cluster and the core filer.
- **Ops/Second** – The total number of file operations per second, broken down by the number of operations between clients and the cluster, the number of synchronous write operations from the cluster to the core filer, and the number of asynchronous write operations from the cluster to the client.
- **Core Filer Throughput** – Throughput operations from the cluster, including number of read-ahead operations, number of read operations from the core filer, number of synchronous and asynchronous write operations to the core filer, and number of file operations between the cluster's constituent nodes.
- **Core Filer Ops** – Number of NFSv3 operations from the cluster to the core filer, broken down into total number of operations, number of read operations, and number of write operations.



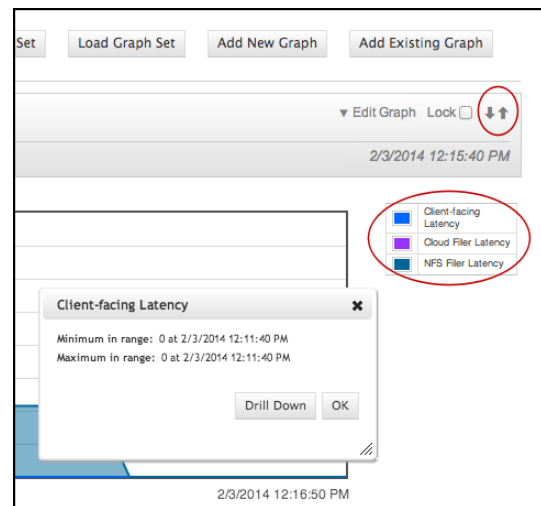
All of the charts have dynamic, real-time updating, at a standard update rate of once every five seconds.

## 10.1. Basic Graph Use

- Click the **Change title** link in the graph's title bar to change the graph's title. When you have finished, click **Update** to commit any changes.
- Click the scope listed after the **Scope** link in the graph's title bar to change the graph's scope; selections include cluster (the default), an individual node, an individual vserver, an individual mass, or an individual client.
- Click on a point on the graph's line or on a histogram square to display detailed information about that datapoint. You can also view specific data within the datapoint by clicking the **Drill Down** button and selecting the data you want to see (for example, the number of read requests from a particular NFS client).



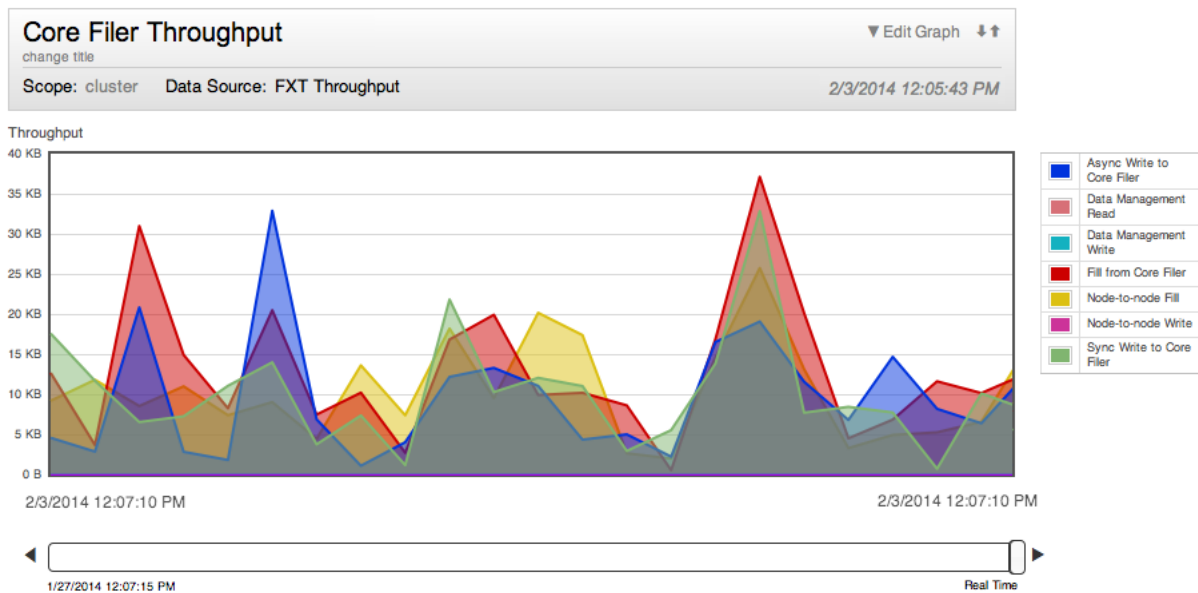
- Click on any statistic listed in the graph's key to display the minimum and maximum values recorded for the statistic and, optionally, to view more specific information.
- Click the up-arrow on the right side of the title bar to move the graph up in the currently displayed graph set, and click the down-arrow to move it down. Rearranging the graphs so that they are adjacent can allow you to more easily compare them.



- The slider at the bottom enables you to scroll backwards and forwards through the graph's history. Each graph retains a maximum of one week's worth of data. Select the **Lock** checkbox to lock any graphs to each other; that is, if you scroll one of the locked graphs, the others will scroll in tandem.
- Click the **Edit Graph** link in the graph's title bar to modify how (or whether) the graph appears. This includes the ability to customize the graph by adding and removing individual statistics, changing colors, resetting the dataset, and more. Graph editing is described in Section 10.3.1, "Editing a Graph" on page 139.

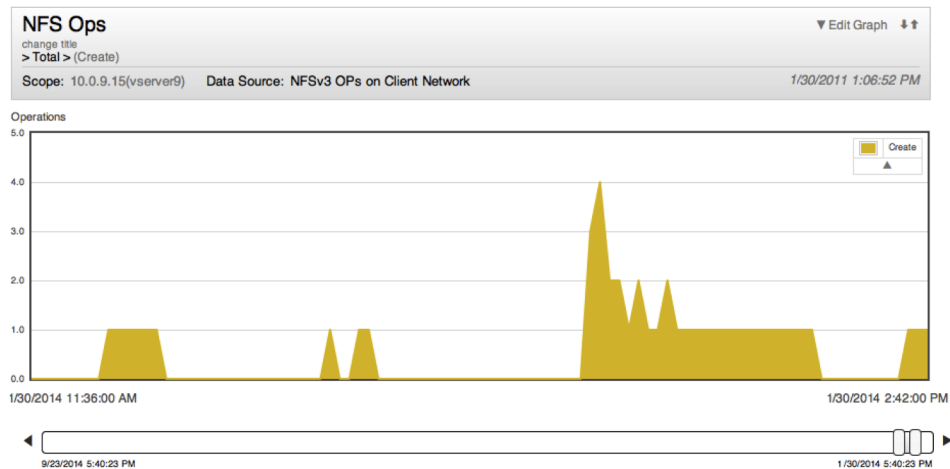
## 10.2. Sample Graphs from the Analytics Tab

### Cluster Throughput



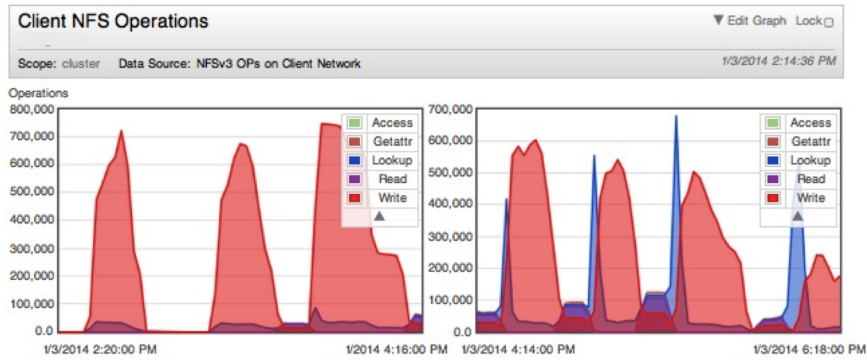
Throughput for the full FXT Series cluster

### A View of Operations Over a Short Time



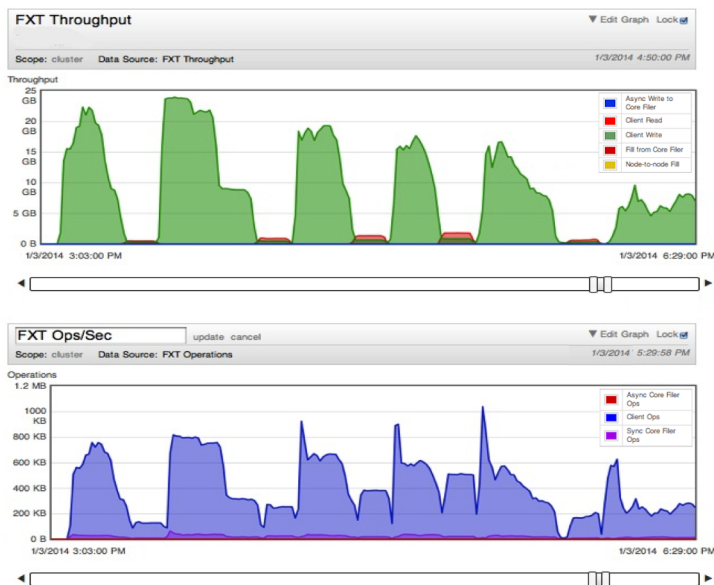
A zoomed-in graph showing NFS Create operations across a short range of time

## A Comparison of Similar Operations at Two Different Time Periods



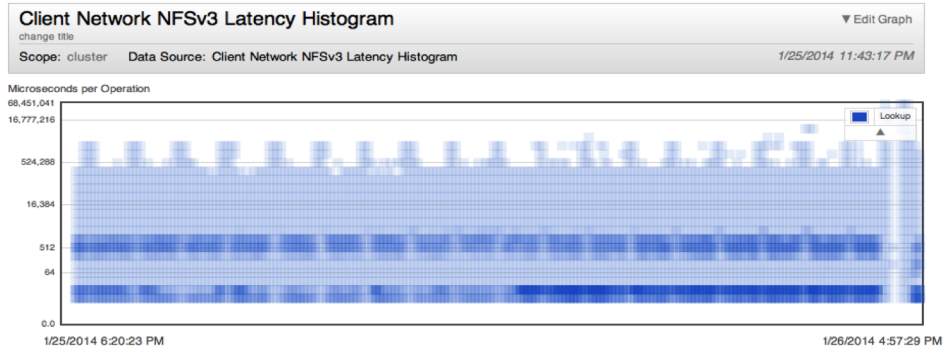
Two charts showing client NFS operations at different times

## A Comparison of Different Operations Across the Same Time Period

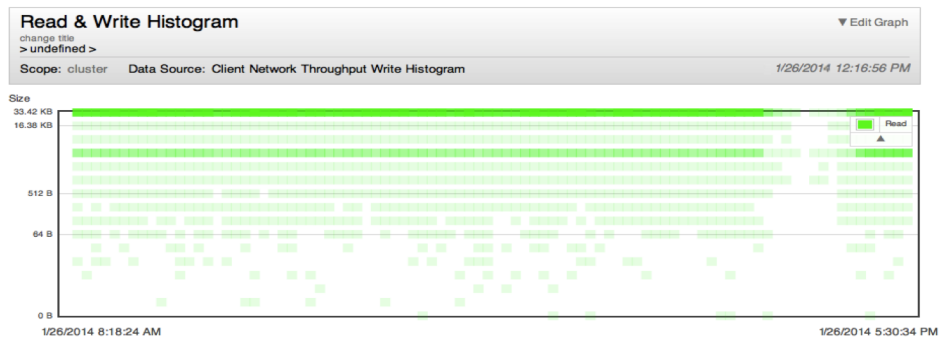


Two charts showing different types of operations (throughput and operations per second) across the same time period

## Histogram Example



## Another Histogram



A histogram of read sizes from the client network

## 10.3. Working with Graphs

This section describes the interface for working with graphs on the **Analytics** tab.

### 10.3.1. Editing a Graph

Click the **Edit Graph** link in the graph's title bar to perform any of the following operations:

- Reset the graph
- Delete the graph
- Remove the graph from the currently displayed graph set
- Stack the graph with another graph
- Modify the graph's data (that is, the statistics that feed into its display)
- Duplicate the graph
- Split the graph
- Switch from a translucent fill to solid fill or vice versa
- Edit the colors used in the graph

### 10.3.2. Adding a User-Defined Graph

- To add a user-defined graph:
  1. Click the **Add New Graph** from the top right-hand corner of the **Analytics** tab. The Avere Control Panel displays the **Add New Graph** dialog box.
  2. In the dialog box, enter a title for the new graph in the **Title** field.
  3. From the **Data Set** drop-down list, choose an initial data set to populate the graph.
  4. From the **Scope** drop-down list, choose the initial scope for the graph. The default scope is cluster. Depending on the selected data set, the scope can also include individual cluster nodes, vservers, masses, and clients.
  5. Click inside the **Stats** field and select one or more statistics that the graph is to display.
  6. In the **Date Range** area, select the date range that you want the graph to display. You can use the slider to specify a date range (time period) from any two points from one week ago to the current time. Alternatively, select a time period from the **Realtime** drop-down list to specify that the graph is to display information from the specified time (for example, 10 minutes ago) to the current time.
  7. Click the **Add Graph** button. The Avere Control Panel displays the new graph at the top of the current graph set.

### 10.3.3. Adding a Predefined Graph

Predefined graphs include the default graphs that come with the Avere OS as well as any previously created user-defined graphs.

- To add a predefined (existing) graph:
  1. Click the **Add Existing Graph** button from the top right-hand corner of the **Analytics** tab. The Avere OS displays the **Load Graph** dialog box.
  2. Choose an existing graph from the drop-down list.
  3. Click the **Load Graph** button. The Avere Control Panel displays the graph at the top of the current graph set.

### 10.3.4. Modifying an Existing Graph

- To modify an existing graph, use the **Edit Graph** link as described in Section 10.1, “Basic Graph Use” on page 136.

### 10.3.5. Working with Graph Sets

This section describes working with graph sets. A graph set is simply a grouping of graphs that appears on the Dashboard at any given time.

#### 10.3.5.1. Creating a New Graph Set

To create a new graph set, click the **New Graph Set** button at the top of the Avere Control Panel’s **Analytics** tab, then add new graphs, existing graphs, or both as described in Section 10.3.2, “Adding a User-Defined Graph” on page 140 and Section 10.3.3, “Adding a Predefined Graph” on page 140, respectively.

### 10.3.5.2. Saving a Graph Set

- To save a graph set:
  1. Click the **Save Graph Set** button at the top of the Avere Control Panel's **Analytics** tab. The Avere Control Panel displays the **Save Graph Set** dialog box.
  2. Enter a name for the new graph set in the **Graph Set Title** field.
  3. Click the **Save Graph Set** button in the **Save Graph Set** dialog box.

### 10.3.5.3. Loading a Graph Set

- To load a saved graph set:
  1. Click the **Load Graph Set** button at the top of the Avere Control Panel's **Analytics** tab. The Avere Control Panel displays the **Load Graph Set** dialog box.
  2. Choose a graph set to load from the **Load Graph Set** dialog box's drop-down list.
  3. Click the **Load Graph Set** button in the **Load Graph Set** dialog box.

# Chapter 11. Updating and System Maintenance (Settings | Administration)

This section describes how you can update the Avere OS, and how to use the **System Maintenance** page. For more information on the **Hidden Alerts** page, refer to Section 9.3.3, “Unhiding Conditions and Canceling Auto-Dismiss” on page 125.

## 11.1. Upgrading the Avere OS Software



### Note

- All nodes in a cluster must run the same software version.
- Downloading a new alternate image overwrites the previous alternate image, which cannot be recovered.

➤ To download a new alternate image:

1. Navigate to the **Settings > Administration > Software Update** page.

Dashboard

Settings

Analytics

Data Management

Support

V3.1.1.1-6a7c480 --- admin

LiGo\_Cluster

VServer

Core Filer

Cluster

Administration

System Maintenance

Software Update

Users

Hidden Alerts

Software Update

Cluster Software Update Information

Active imageAvereOS\_V3.1.1.1-03e81d0 (2014/01/28\_23:33)

Alternate imageAvereOS\_V3.1.1.1-5eb3f90

Activate in high-availability mode☒

Activate alternate imageCancel

Download New Software Image

Download URLhttp://download.averesystems.com/software/

URL usernamejuser

URL password\*\*\*\*\*

Download to alternate image locationCancel

Upload New Software Image From This Workstation

File to downloadChoose FileNo file chosen

Upload to alternate image location

Update status

Activate AvereOS\_V3.1.1.1-03e81d0 complete

Node	Software Update Status
node1	Download of AvereOS_V3.1.1.1-5eb3f90 complete
node7065	Activate AvereOS_V3.1.1.1-03e81d0 complete

2. If your cluster is using high-availability (HA), ensure that the **Activate in high-availability mode** checkbox is selected.
3. Enter the URL of the new software image in the **Download URL** field.

4. If needed, enter a username for the download URL in the **URL username** field and the user's password in the **URL password** field.
5. Click the **Download to alternate image location** button.

The download progress is shown in the **Update status** fields. This process includes initial download, verification, and installation phases on the node on which the download is initiated, and a final download phase to all nodes in the cluster.

To run an alternate software image in a cluster, you must activate it. When you activate an alternate image, the previously active image becomes the alternate image.

➤ To activate an alternate image:

1. From the **Settings > Administration > Software Update** page, click the **Activate alternate image** button. the cluster restarts and the Avere Control Panel reports messages such as `Failed to load content: communication failure` and `Data error`. There is a short interruption in cluster services while the cluster restarts.
2. When the cluster restarts, the browser is redirected to the Login page. Log in to the cluster as usual.
3. Click the **Alerts** tab on the Avere Control Panel's **Dashboard** tab's status bar and ensure that informational alerts indicate that the cluster images were updated.
4. Navigate to the **Settings > Administration > Software Update** page. Verify the following:
  - The image that was previously listed in the **Active Image** field is now listed in the **Alternate Image** field
  - The **Update status** and **Software Update Status** fields indicate that the activation of the alternate image is complete.

If you need to revert to the previous software image, click the **Activate alternate image** button again. You cannot revert to the previous alternate image if you have already downloaded a new alternate image.

## 11.2. Using the System Maintenance Page



### Caution

- When you first access this page, a message at the top of the page shows how many files on the cluster have not yet been written (flushed) to the core filer. To avoid potential data loss, do not perform any operations on the page until the message reads “All files flushed”.
- Some of the controls on this page disrupt client access and can potentially cause data loss; use them only at the direction of Avere Global Services.

➤ To access the **Maintenance Operations** page:

- Navigate to **Settings > Administration > System Maintenance**.

Here, you can perform the following actions, all of which disrupt client access to the cluster unless noted otherwise.

Dashboard

Settings

Analytics

Data Management

Support

V3.1.1-6a7c480 --- admin

LiGo\_Cluster

VServer

Core Filer

Cluster

Administration

System Maintenance

Software Update

Users

Hidden Alerts

Maintenance Operations

2 files to be flushed

Cluster Services

Perform operation in high-availability mode to avoid client access disruption.

Restart the services on every node in the cluster; no committed data will be lost.

Reboot every node in the cluster; no committed data will be lost.

Reboot the cluster as above, but force switch to the alternate software image (AvereOS\_V3.1.1-5eb3f90).

Safely power down all nodes in the cluster. User data in the cluster will be retained.

Redistribute directory managers across the cluster.

Clear all active conditions.

Restart cluster

Reboot cluster

Reboot alternate

Power down cluster

Rebalance directory managers

Clear all conditions

Access

Stop accepting access requests from client systems. Suspension is a prerequisite to invalidating cached data or erasing the cluster.

Start allowing access from client systems after suspend.

Suspend access

Unsuspend access

Data - Warning! These operations may cause loss of data.

Discard all cached data in the cluster, potentially including changed data from clients.

Reset the FXT nodes, but retain or modify the current configuration. Erase all cached data.

Reset the FXT nodes to unconfigured factory condition. Erase all cached data and the current configuration.

Reset the FXT nodes to unconfigured factory condition, then zero out all drives. Erase all cached data and the current configuration.

Invalidate cache

After reset:  
Reconfigure with the current  
Reconfigure with modification

After reset:  
Reboot cluster  
Power down cluster

After reset and zeroing:  
Reboot cluster  
Power down cluster

Currently connected to node node7065.

### 11.2.1. Cluster Services (Restarting)

- If you have high availability enabled, you can select the **Perform operation in high-availability mode to avoid client access disruption** checkbox. For more information, refer to Section 2.6, “High Availability” on page 33.

- **Restart cluster** – Restart services on every node in the cluster.

Restarting cluster services is generally a safe operation. It temporarily disrupts client access but retains client data.

- **Reboot cluster** – Restart every node in the cluster.

Restarting the cluster is generally a safe operation. It temporarily disrupts client access but retains client data.

- **Reboot alternate** – Restart the cluster to the installed alternate image.

Restarting the cluster with the alternate image is generally a safe operation. When you restart the cluster with an alternate image, the node coordinating the restart process restarts the other nodes in the cluster in parallel before restarting itself.

- **Power down cluster** – Power down all nodes in the cluster.

Powering down the cluster is generally a safe operation. When a cluster is powered down, all client data is written to FXT Series storage media but is not necessarily written to the core filer.

### 11.2.2. Monitoring (Statistics Collection)

- **Monitor stats** – Start a statistics-collection process to upload for support analysis.

Collecting statistics is a safe operation, but it can slow overall cluster response time by consuming CPU cycles.

- **Stop stats monitoring** – Stop the statistics-collection process.

- **Clear all conditions** – Removes all conditions from the Dashboard.

### 11.2.3. Client Access

- **Suspend access** – Suspend client access to the cluster.

Suspending client access to the cluster is generally a safe operation. Suspending access triggers each virtual server to write client data rapidly to its core filer but does not guarantee that all data is written to the core filer if you subsequently perform additional operations such as invalidating the cache. Suspending access is a prerequisite to invalidating the cache, erasing the cluster, and zeroing disks.

- **Unsuspend access** – Re-enable client access to a suspended cluster.

## 11.2.4. Data (Resetting Cache and Nodes)

- **Invalidate cache** – Discard all cached data in the cluster.



### Warning

Invalidating client data is an unsafe operation unless you are certain that all data has been written to the core filers. Invalidated data cannot be recovered from the cluster after this operation is performed.

Before performing this operation, you must suspend client access to the cluster by clicking **Suspend access**.

- Delete all client data but retain or modify the cluster's configuration by clicking the **Erase and reconfigure cluster** button.

This is an unsafe operation unless you are certain that all client data has been written to the core filers.

➤ To change all or part of your cluster configuration:

1. Suspend client access to the cluster by clicking the **Suspend access** button.
2. When client access is suspended, click the **Erase and reconfigure cluster** button. The Avere Control Panel displays the Erase and Reconfigure page.
3. Make the desired changes on the **Erase and Reconfigure** page. The page provides most of the same fields and options as the **Initial Setup** page, described on Section 2.1, “General Setup” on page 11. You can specify new values for any or all of the fields on this page.
4. Click the **Erase and reconfigure** button at the bottom of the page.



### Note

This action results in all nodes except the primary node being removed from the cluster and all configuration values for high availability and CIFS being deleted.

The Avere Control Panel displays the message `Erase and reconfigure cluster: Are you sure?`

5. Click the **OK** button.
6. The Avere Control Panel displays the message `Please wait a few minutes for the cluster to reboot, then re-add nodes on the Cluster Setup page. All nodes in the cluster restart.`
7. When the cluster restarts, open the Avere Control Panel and go to the Cluster Setup page.
8. Perform the following steps to restore the cluster to its previous state:
  - a. Join the cluster's other nodes to the primary node either automatically, as described in Section 2.1.2, “Allowing Unconfigured FXT Nodes Automatically Join the Cluster” on page 11, or manually, as described in Section 2.2.3, “Modifying a Node's Name/Node Detail Page” on page 15.
  - b. If the cluster had an HA data repository, reconfigure it. See Section 2.6.1, “Specifying an HA Data Repository” on page 34 for information.
  - c. If the cluster had high availability enabled, reenable it as described in Section 2.6, “High Availability” on page 33.

- d. If the cluster had CIFS enabled, reenable it as described in Chapter 7, *Configuring CIFS Access* on page 95.
9. Resume normal cluster operations.

### 11.2.5. Delete Data and Configuration

- Delete all client data and the cluster's configuration by clicking the **Erase cluster** button.

This is an unsafe operation unless you are certain that all client data has been written to the core filers and that you want to delete the cluster configuration. Neither invalidated client data nor client configuration information can be restored after this operation is performed. After client configuration information is deleted, the FXT Series nodes that constituted the cluster are no longer in communication and cannot serve data unless they are configured into a new cluster.

Before performing this operation, you must suspend client access to the cluster by clicking the **Suspend access** button.



#### Note

If you have configured CIFS access on a cluster whose configuration data you delete, the cluster's Active Directory machine trust account (MTA) is retained on the Active Directory server. You can manually delete the cluster's MTA from the Active Directory server; see the Active Directory documentation for details.

- Delete all client data, delete the cluster's configuration, and zero out the FXT Series nodes' data disks by clicking the **Erase cluster and zero disks** button.

This is an unsafe operation as described in the previous bullet. Additionally, all data on the nodes' disks is overwritten with zeros, ensuring that no data can be retrieved from the disks.

Before performing this operation, you must suspend client access to the cluster by clicking the **Suspend access** button.



#### Note

It takes approximately three hours to zero out the disks on an FXT Series node.



#### Important

After you initiate this operation, allow it to run until the Avere Control Panel becomes unavailable. The Avere Control Panel does not currently provide a progress indicator on the disk-zeroing operation. After the operation starts, the Avere Control Panel returns to the System Maintenance page, and it appears that you can perform additional actions.

## 11.3. Adding and Modifying Users

The Avere OS has a default administrative account, named `admin`. You can change the password for this account, or add additional users.

➤ To add or change user settings:

1. Navigate to the **Settings > Administration > Users** page.

Name	Permission	Actions
admin	Full Access	<button>Remove</button> <button>Password</button> <button>Permission</button>
nuser	Read-only	<button>Remove</button> <button>Password</button> <button>Permission</button>

**Add New User**  

User Name:

Permission:

User password:

Confirm password:

Add User

2.
  - For a new user:
    - a. Enter the user's name, select **Full Access** or **Read-Only** from the **Permission** drop-down list, and enter the password twice.  

The Avere Control Panel masks the password as you enter it into the password fields.
    - b. Click **Add User** to add the user to the list.
  - For an existing user, choose one of the following buttons from the **Actions** section:
    - **Remove** – The Avere OS asks if you are sure you want to remove the user. Click **OK** to continue.
    - **Password** – Fields for the new password appear. Enter the new password twice, and then choose **Change password**.
    - **Permission** – The **Permission** drop-down list appears. Select the permission you want to use, and then choose **Change permission**.

## Chapter 12. Troubleshooting and Getting Support

This section discusses troubleshooting scenarios and obtaining support for your cluster.

You can manage support settings in two places:

- By navigating to the **Settings > Cluster > Support** page.
- From the Dashboard's **Support** tab, described in Section 12.3, "Contacting Avere Global Services (Support Tab)" on page 153.

### 12.1. Possible Troubleshooting Scenarios

Issues you can encounter include but are not limited to the following. Suggestions for addressing the issues are provided where applicable. If you cannot resolve the issue on your own, contact Avere Global Services as described in Section 12.3, "Contacting Avere Global Services (Support Tab)" on page 153.

- The Avere Control Panel's **Dashboard** tab can indicate that the cluster cannot write data to the core filer at the specified interval.

To address this problem, you can increase the value of the maximum writeback delay. Consider the criteria listed in Section 6.2.2.1, "Determining the Maximum Writeback Delay" on page 87 when adjusting this value. You can also evaluate whether the incoming load from the clients has increased or whether the performance of the core filer has decreased.

- Clients can have problems reading or writing data. To address this problem, verify the following:
  - Ensure that the network is functioning correctly and that all network connections are intact, as described in the *Quick Start Guide*.
  - Ensure that the core filer is operating and available.
  - Ensure that the clients are mounted to the correct location.
  - Ensure that NFS exports are set and implemented correctly, as described in Section 2.5, "Managing Exports (Settings Tab | VServer)" on page 28.
  - Ensure that any networking or client recommendations listed in the *Release Notes* are observed.
- Conditions and alerts can appear on the **Dashboard** tab. To address the problem, read the alert message and try to determine the cause.
  - If you are able to resolve an alert, select its listing under the **Alerts** tab on the **Dashboard** tab's status bar and click **Dismiss checked alerts**.
  - If you cannot understand the message or resolve the cause, or if the alert is a condition that does not resolve on its own, contact Avere Global Services.

## 12.2. Configuring Support Settings

By configuring settings on the Avere Control Panel's **Support** page you will be able to provide information to Avere Global Services that will allow them to provide a faster response in the event of an issue with the cluster.

- To configure support settings:
  1. Navigate to the **Settings > Cluster > Support** page.

**Support**

**Customer Info**

Customer ID:

Upload info validated:  [Revalidate upload information](#)

Statistics Monitoring: ☐

General Information Upload: ☐

Crash Information Upload: ☒

Gather Rolling Trace Data: ☐

Enable Memory Debugging: ☐

Trace Level:

Core save policy:

[Revert](#) [Submit](#)

**Secure Proactive Support (SPS)**

Enable SPS Link: ☒

SPS Upload Interval:

Enable Remote Commands: ☐

**Shipping Information**

Contact Name:

Address1:

Address2:

City:

State:

ZIP:

Country:

Telephone:

[Revert](#) [Submit](#)

**Support Package Update**

Support package URL:

Current Status:  [Update](#)

Setting	Modified	Value	Check	Actions
<b>Add a Custom Setting</b>				
Caution: do not change custom settings except as advised by Avere Global Services personnel.				
Name		<input type="text" value="cluster.xdrAuthGroupsFont"/>		
Check		<input type="text" value="QC"/>		
Value		<input type="text" value="32"/>		
Note		<input type="text" value="checking the fonts"/>		
				<a href="#">Submit</a>

### 12.2.1. Support Package Update

This area allows you to upload a new support package, if you are instructed to do so by Avere Global Services. New support packages contain the latest support tools.

➤ To upload a new support package:

1. Obtain the URL of the new support package from Avere Global Services. This will generally be e-mailed to you.
2. Navigate to the **Settings > Cluster > Support** page/
3. Enter the URL in the **Support package URL** field.
4. Click the **Update** button.

### 12.2.2. Customer Info

This area allows you to configure what information is saved about your use of the cluster, so that Avere Global Services can use the information to help troubleshoot possible later problems.

- In the **Customer ID** field, enter the customer ID provided by Avere Global Services or your Avere Systems representative.
- If the value of the **Upload info validated** field does not read **yes**, click the **Revalidate upload information** button.

If the **Upload info validated** field still does not read **yes**, contact Avere Global Services for assistance. Possible issues with upload information include firewalls, other network-security mechanisms, or a more general network problem.

- Select one or more of the following items for monitoring by Avere Global Services. You can discuss the appropriate settings with your Avere Systems representative during the installation of the product and with Avere Global Services when you are troubleshooting an issue.
  - **Statistics Monitoring** – Avere Global Services often recommends that this item be enabled by default.
  - **General Information Upload** – Avere Global Services often recommends that this item be enabled by default.
  - **Crash Information Upload** – Avere Global Services generally recommends that this item be enabled only during troubleshooting.
  - **Gather Rolling Trace Data** – Avere Global Services generally recommends that this item be enabled only during troubleshooting.
  - **Enable Memory Debugging** – Avere Global Services generally recommends that this item be enabled only during troubleshooting.
- If you are instructed to do so by Avere Global Services, enter a value in the **Trace Level** field. This field is generally used only during troubleshooting. Avere Global Services will provide you with the appropriate value to enter if tracing is required.
- Click the **Save** button to save the information entered into the **Customer Info** section.

Choose a value from the **Core save policy** drop-down list. This area determines how core files are saved. The core files can be examined by Avere Global Services for more information about problems during troubleshooting. Possible values include the following:

- **Overwrite oldest**, the default value, overwrites the oldest core file when a new core file is generated. This selection enables only one core file to exist at a time.

- **Overwrite newest** overwrites the most recent core file when a new core file is generated. This selection enables only one core file to exist at a time.
- **Use available space** retains as many core files as possible before overwriting older files with newer files.

If you believe you need a value other than the default, ask Avere Global Services or your Avere Systems representative. Values other than the default value are generally used only during troubleshooting.

### 12.2.3. Secure Proactive Support (SPS)

SPS provides features for remote troubleshooting with minimal to no customer involvement. When you enable SPS, the information entered here will be transmitted to Avere using a secure link, at each time interval, or ‘heartbeat’, that you set. SPS uses the same upload mechanism as the **Customer Info** statistics, so you do not need to update. Information such as the following is sent to Avere:

- Summary statistics – those statistics collected by the Dashboard
- Any alerts or conditions
- The state and hardware configuration of the nodes in the cluster.
- Shipping information – the location and contact information any replacement hardware needed

➤ To enable Secure Proactive Support:

1. Select the **Enable SPS Link** checkbox.
2. Select the **SPS Upload Interval** from the drop-down list, anything from one minute to one day, or a custom interval. If a “heartbeat” is not able to be sent at the interval, the information will be saved for the next interval.
3. Select **Enable Remote Commands** if you want Avere to be able to run a limited number of commands on your cluster, *if the cluster software has determined this is reasonable, and asks Avere to run the commands.*

If so, the cluster (at the site) will ask if there is a command to run. In , will be executed by support people, info will be saved in a log (who requested the action, the command executed, the result of the command, the time of the command, and the reason for the command). This information will be encrypted at Avere.

4. Enter the **Shipping Information** for where any hardware should be shipped. Any alerts and conditions that might indicate hardware failures will initiate a shipment, verified by Avere and the customer.

### 12.2.4. Custom Settings



#### Note

All information for the fields in this area must be provided by Avere Global Services.

1. Type the values for the **Name**, **Check** (checksum), and **Value** fields provided to you by Avere Global Services.
2. In the **Note** field, enter a text note for the setting. You can provide your own note or leave this field blank if Avere Global Services does not give you a specific string to enter.
3. Click the **Submit Custom Setting** button.

4. Verify that the settings appear correctly in the table below the **Submit Custom Setting** button.



### Note

You can have multiple custom settings on your cluster, but all custom settings must be determined and provided by Avere Global Services

## 12.3. Contacting Avere Global Services (Support Tab)

The Avere Control Panel's **Support** tab is designed to assist your interactions with Avere Global Services. The **Support** tab lists contact information via web, phone, and email and provides mechanisms for you to upload information to Avere Global Services.

- To troubleshoot an issue with Avere Global Services:

1. Click on the Dashboard's **Support** tab, and continue as directed by your Avere Global Services representative.

The screenshot shows the Avere Control Panel's **Support** tab. The top navigation bar includes **Dashboard**, **Settings**, **Analytics**, **Data Management**, and **Support**. The user is logged in as **admin** on the **LiGo\_Cluster**. The interface is divided into several sections:

- Generate Support Information:** Provides contact details for Avere Global Services, including a web link (<http://www.averesystems.com/support>), phone numbers (1-888-88-AVERE, Option 2 (Toll-Free) and 1-412-635-7170, Option 2), and an email address ([support@averesystems.com](mailto:support@averesystems.com)). It also offers links to download **Release Notes**, **Daily Operations Guide**, and **cluster configuration**.
- Details:** Displays cluster information: **Cluster name** is `liz_newg3_1_Cluster`, **Software version** is `V3.1.1.1-6a7c480`, **Choose node** is set to **Cluster Wide**, and **Upload information validated** is `no`.
- Support Status:** A table showing support operations for nodes `node7065` and `node1`, both indicating "No support operations currently running".
- General information upload:** Includes a **Choose gather mode** dropdown set to **Normal support information**, a **Comment** text area, and an **Upload information** button.
- Advanced information gathering:** Includes a **Choose gather mode** dropdown set to **Full packet capture (50GB Ring buffer)**, a **Gather duration** dropdown set to **10 minutes**, a **Capture filter** text input, a **Comment** text area, and **Start collecting** and **Stop collecting** buttons.

2. In the **Details** section, choose the scope of troubleshooting information from the **Choose node** drop-down list. Possible values include **Cluster Wide** (the default) or any individual node in the cluster.

3. In the **General information upload** section, choose one of the following from the **Choose gather mode** drop-down list:
  - **Normal support information** (the default)
  - **Current statistic information**
  - **Historical statistic information**
  - **Minimal log collection**
  - **Data dump information**
  - **Rolling trace information**
  - **Current trace information**
4. Click the **Upload information** button.
5. In the **Advanced information gathering** section, perform the following steps as directed by Avere Global Services:
  - a. From the **Choose gather mode** drop-down list, choose one of the following options:
    - **Directory Trace**
    - **Writeback Trace**
    - **Full packet capture** (the default)
    - **Full packet capture (100MB Ring buffer)**
    - **Performance information**
    - **Memory Debugging**
    - **Read/Write Trace**
    - **Stats gathering**
    - **Timing Trace**
    - **Trace/partial packet**
    - **Trace**
  - b. From the **Gather Duration** drop-down list, choose one of the following options:
    - **Until manually stopped** (the default)
    - Any of the options ranging from **1 minute** to **30 minutes**.
  - c. Click the **Start collecting** button to start collecting the specified information
6. If you selected **Until manually stopped** from the **Gather Duration** drop-down list, click the **Stop collecting** button after an interval determined by you and Avere Global Services.

## 12.4. When to Contact Avere Global Services

Contact Avere Global Services if any of the following events occur:

- The cluster is unresponsive, unreachable, and/or not serving data.
- The **Dashboard** tab shows an alert or condition, either yellow or red, whose message you do not understand or whose cause you cannot determine or remedy.

## 12.5. Saving Cluster Configuration Data

You can save a copy of the cluster's basic configuration as an XML file. This is useful in the event that the node fails catastrophically and must be reconstructed.

➤ To save a copy of the configuration file:

1. Navigate to the Avere Control Panel's **Support** tab.
2. Click the **Download current cluster configuration** link.

Your browser either downloads the file, named `armconfig.xml`, to its standard download location or prompts you for a location on your local machine to which to save the file.

3. Save or copy the file to a well-known location.
4. Optionally copy the file to a USB memory stick.

## Appendix A. Core Filer-Specific Configuration Notes

This section provides information about specific configuration recommendations for several supported core filers when used with the Avere system.

### A.1. Configuring Export FSIDs for GNU Linux NFS Servers

Linux NFS servers automatically assign new filesystem IDs (FSIDs) to some or all of its NFS exports upon events such as reboots, the addition of volumes, and general changes to the server topography. Changing the FSID on an export whose data is served to clients through an FXT cluster can disrupt client access. To work around this issue, specify a value for the FSID for each export in the server's `/etc/exports` file, as shown in **bold** in the following example:

```
/stable *(rw,no_root_squash,insecure,fsid=8)  
/export/1 *(rw,no_root_squash,insecure,fsid=1)  
/export/2 *(rw,no_root_squash,insecure,fsid=2)  
/export/3 *(rw,no_root_squash,insecure,fsid=3)  
/ *(rw,no_root_squash,insecure,fsid=0)
```

Before editing the `/etc/exports` file, unmount all clients, include the FXT cluster, from any volume whose FSID you plan to specify. Client access can be affected if the `/etc/exports` file is edited while clients are mounted to the NFS server.

After editing the `/etc/exports` file, restart the **nfsd** and **mountd** processes on the NFS server to put the changes into effect.

For more information about the `/etc/exports` file's format and options, run the **man 5 exports** command on your Linux server.



#### Note

- If LUNs were added or changed on the core filer, FSIDs change regardless of previous edits to the `/etc/exports` file.
- Older Linux clients can have problems accessing GNS vservers. Contact Avere Global Services for a list of clients known to have this issue.

## A.2. NetApp Data ONTAP 7G and Data ONTAP 8.0 7-Mode

This section lists notes for use with Data ONTAP systems.

### A.2.1. Using Native CIFS Volumes for CIFS ACLs

Note the following when configuring CIFS shares with a Data ONTAP 7 core filer using native CIFS volumes:

- The qtree security style must be set to `ntfs`.
- Check the NetApp documentation for details about configuring NIS or LDAP for username resolution and DNS for hostname resolution.
- Check the NetApp documentation for the **exportfs** command for details about export configuration. Exports must be configured as follows:
  - The Avere cluster must be able to access the exports for which ACLs are enabled by using the `sys` security flavor.
  - The NFS export rules must be configured to provide the Avere cluster with root access. You can do this on `ntfs` volumes by using the following command:

```
options cifs.nfs_root_ignore_acl on
```

### A.2.2. Using NFSv4 for CIFS ACLs

Note the following when configuring CIFS shares with a Data ONTAP 7 core filer using NFSv4:

- The qtree security style must be set to `unix`.
- Check the NetApp documentation for details about configuring NIS or LDAP for username resolution and DNS for hostname resolution.
- Both NFSv4 and NFSv4 ACLs must be enabled by using the following commands:

```
options nfs.v4.enable on
options nfs.v4.acl.enable on
```

- Check the NetApp documentation regarding the following setting:

```
options nfs.v4.id.domain
```

Note that the value of this setting must match the value of the **NFS Domain** field on the **Directory Services** page, as discussed in Section 2.8.5, “Specifying the Source for Usernames” on page 46 and Section 7.8, “Creating CIFS Shares” on page 107.

- Check the NetApp documentation for the **exportfs** command for details about export configuration. The Avere cluster must be able to access the exports for which ACLs are enabled by using the `sys` security flavor.
- For NFSv4, Data ONTAP requires OWNER@, GROUP@, and EVERYONE@ ACEs in the default ACL for proper operation.

## A.3. ZFS on OpenSolaris and Oracle Solaris

Note the following when configuring an FXT Series cluster with a ZFS (OpenSolaris or Solaris) core filer:

- For consistent connectivity, you must combine all of your FXT Series cluster's cluster IP addresses into a single round-robin DNS (RRDNS) hostname, and each IP address in the RRDNS hostname must reverse-resolve to the hostname.

After you configure the single RRDNS hostname, use the following command on the core filer:

```
% zfs set sharenfs="rw=@FXT_RRDNS_name,root=cluster_RRDNS_name" epool/export
```

- For consistent connectivity, you must combine all of your FXT Series cluster's client-facing (vserver) IP addresses into a single round-robin DNS (RRDNS) hostname, and each IP address in the RRDNS hostname must reverse-resolve to the hostname.

After you configure the single RRDNS hostname, use the following command on the core filer:

```
% zfs set sharenfs="rw=@FXT_client-facing_RRDNS_name" epool/export
```

- If individual IP addresses are used in NFS export rules on your FXT Series cluster, they must also reverse-resolve to a valid hostname that forward-resolves back to the specific IP addresses.

After you configure each individual IP address, use the following command on the core filer:

```
% zfs set sharenfs="@individual_IP_address" epool/export
```

- If you cannot set up DNS resolution, you can use the workaround of “per-network-type” rules by specifying the full 32 bitmask per IP address, as in the following example:

```
% zfs set sharenfs="@individual_IP_address/32" epool/export
```

Because this method leads to long and confusing export policies on the ZFS core filer, it is not recommended unless no other options are available.

Note the following when configuring CIFS shares with a ZFS (OpenSolaris or Solaris) core filer using NFSv4:

- The following attributes must be on set on each ZFS volume that is to be accessed:

```
zfs set aclmode=passthrough epool/volume_name
zfs set aclinherit=passthrough epool/volume_name
```

- Add the following line to the core filer's /etc/default/nfs file:

```
NFSMAPID_DOMAIN=nfsv4_domain_name
```

where *nfsv4\_domain\_name* is the value of the **NFS Domain** field on the **Directory Services** page, as discussed in Section 2.8.5, “Specifying the Source for Usernames” on page 46 and Section 7.8, “Creating CIFS Shares” on page 107.

## A.4. EMC Isilon OneFS



### Important

The following information applies to OneFS 5.x and 6.0.x. It does not apply to OneFS 6.5.x, which uses authentication and authorization services that differ from previous versions. Versions of OneFS prior to 5.x and after 6.0.x have not been tested as of this writing.

Note the following when using LDAP and Active Directory together in a CIFS configuration with a OneFS core filer:

- Edit the `/etc/mcp/override/lsassd.xml` file on the OneFS core filer to include the following lines inside the file's `<isi-data>` tags:

```
<modify-text id="nss-priority">50</modify-text>
<modify-text id="ad-priority">75</modify-text>
```

## Appendix B. Password and Group File Formats

This section discusses the formatting requirements for external password and group files as described in Section 2.8.2, “Configuring Netgroups” on page 42 and Section 2.8.5, “Specifying the Source for Usernames” on page 46.

### B.1. General Formatting Rules

In general, password and group files need to conform to the standards of the UNIX `/etc/passwd` and `/etc/group` files, respectively. Note the following rules for how the Avere system parses both password and group files:

- The parser is as flexible as possible when parsing entries. Any line that includes the absolute minimum number of fields for the entry type is treated as valid. For example, a valid password entry can be as simple as `name::uid::gid`; a valid group entry can be no more than `name::gid`.
- The parser does minimal validation of password (user) or group names. Names that are considered invalid include the following:
  - Empty strings
  - Strings that contain non-printable ASCII characters
  - Strings that begin with “+” or “-”

If `passwd` or `group` entries conflict with system `passwd` or `group` entries, then conditions will be displayed on the Dashboard, if you have checked **Enable Dashboard Conditions** as described in Section 2.8.2, “Configuring Netgroups” on page 42.

- The parser does not cross-validate entries. If a file contains an entry that declares user `dsmith` belongs to group ID `2099`, the parser does not verify that a group with the ID `2099` exists. Similarly, if an entry for a group named `finance` lists a member with the username `alice`, the parser does not verify that a user with the name `alice` exists.
- If a numeric field (UID or GID) cannot be converted from its text representation to an integer, the record is considered to be invalid and is ignored.
- Files can contain comments. Any line that begins with zero or more whitespace characters followed by the `#` symbol is considered to be a comment and is ignored.

### B.2. Searching Rules

The Avere system searches for password and group entries in the following way:

- The system searches entries in the order in which they appear in the file.
- Searches by name are case insensitive. For example, a search for `alice` matches `alice`, `Alice`, or `ALICE`.
- Searches return successfully after finding the first matching record. Looking at the previous item, this means that a search for `alice` returns whichever capitalization variation of the name occurs first in the file and does not continue searching for other variations.

## B.3. Password-Entry Parsing

The Avere system parses password entries as follows:

- The parser ignores any fields after the `gid` field, if present.
- The parser returns the username exactly as it is listed in the password file, including any leading, trailing, or embedded spaces or tabs.
- The parser always returns the string `*` for the `password` field.
- The parser always returns an empty string for the `gecos` field.
- The parser always returns the string `/nonexistent` for the `homeDir` field.
- The parser always returns the string `/usr/sbin/nologin` for the `login` field.

### “Passwd” File Format Example

The expected format is:

```
name:password:uid:gid:login_class:gecos:home_dir:shell
```

Avere OS only uses the `name`, `uid`, and `gid` parameters; everything after `gid` is ignored.

## B.4. Group-Entry Parsing

The Avere system parses group entries as follows:

- The parser ignores any fields after the `members` field if present.
- The parser returns the group name exactly as it is listed in the group file, including any leading, trailing, or embedded spaces or tabs.
- The parser always returns the string `*` for the `password` field.
- The parser returns an empty member list if the file does not specify any members for the group or if there is no `members` field.
- The `members` field is a comma-separated list of usernames. Unlike its behavior with password and group name entries, the parser trims any leading or trailing whitespace from the entries in the `members` field.

### “Group” File Format Example

The expected format is:

```
name:password:gid:member_list
```

Avere OS only uses the `name`, `gid`, and `member_list` (if any) parameters.

## B.5. The Format of the Netgroup File

If you configure the cluster to read netgroup information from an external file, the file must be in `/etc/netgroup` format. This format differs among operating systems.

The format used by Avere OS is designed to be tolerant of other operating systems' idiosyncracies, and is as follows:

- Each line in the file starts with the name of a netgroup, in all upper-case letters, followed by the netgroup's membership definition. The netgroup's membership definition can take one of two forms:
  - The name or names of other netgroups
  - An ordered list of the form (*host, user, domain*)
    - Use a blank space or asterisk (\*) to specify a wildcard value for any member of the list.
    - Use a dash (-) to specify an "ignore" value for any member of the list.
    - The value of *host* is an IP address; a fully qualified domain name; or a partially qualified domain name that can be resolved with the DNS domain, the DNS search paths, or both.
    - The value of *user* specifies a user. In Avere OS as well as traditional NFS-export implementations, this value is not used; user-specific authentication is provided by a directory service instead of the netgroup. Specify the \* (asterisk) or - (dash) character as the value of *user*.
    - The value of *domain* indicates the domain to which the netgroup applies. It is typically specified as a wildcard.



### Note

If the netgroup is used with NIS, if the domain value does not match the domain specified on the Directory Services page, the value is ignored.

If the netgroup is used with LDAP, the domain value is ignored in all cases.

- Lines that begin with the pound (#) character are treated as comments.
- You can use a backslash (\) at the end of long lines to continue them to the next line.
- You can nest netgroup definitions by expressing multiple tuples or multiple netgroup names. The nested netgroup names must be defined before the nesting line.

### Netgroup File Example

```
# Role-based example (that is, multiple hosts/clients per group):
BUILDBOSTS (build0,,) (build1,,) (build2,,)
#
# Machine-based example (that is, one netgroup per host/client):
CYCL0 (cycl0,,)
CYCL1 (cycl1,,)
CYCL2 (cycl2,,)
#
# Definitions can be nested:
GROUPA (alpha,,) (bravo,,)
GROUPB (charlie,,) (delta,,)
GROUPC (echo,,) (foxtrot,,)
GROUPAB GROUPA GROUPB
ALL GROUPA GROUPB GROUPC
```

# Index

## A

- access control lists (ACLs): [103](#)
- access to cluster, modifying: [145](#)
- activating
  - alternate images: [143](#)
  - licenses: [6](#)
- Active Directory
  - configuring: [97](#)
  - prerequisites: [100](#)
  - username mapping: [97](#)
- adding
  - Avere OS users: [148](#)
  - core filers to cluster: [50](#)
  - graphs and graph sets: [140](#)
  - IP addresses to a vserver: [67](#)
  - junctions to a GNS: [74](#)
- advanced networking
  - about: [77](#)
  - adding IP address ranges: [67](#)
  - enabling: [78](#)
- AES-256 encryption: [56](#)
- alerts
  - dismissing: [124](#)
  - email: [133](#)
  - history: [126](#)
  - viewing on Dashboard: [121](#)
- alternate images
  - activating: [143](#)
  - downloading: [142](#)
  - rebooting into: [145](#)
- Amazon services
  - configuring: [56](#)
  - supported regions: [56](#)
- Amplidata services, configuring: [56](#)
- Analytics tab: [135](#)
- authentication
  - Kerberos: [30, 43](#)
  - prerequisites: [39](#)
  - UNIX/SYS: [30](#)
- Avere Control Panel, about: [7](#)
- Avere Global Services
  - and custom settings: [152](#)
  - contacting: [153](#)
- Avere OS
  - configuring: [4](#)
  - documentation available: [9](#)
  - reverting to previous: [143](#)
  - updating: [142](#)
  - users, adding: [148](#)

- Avere OS users
  - modifying: [148](#)

## B

- bandwidth, modifying for writes: [91](#)
- Bonjour (DNS-SD), supported: [5](#)
- bookmark for logging in: [7](#)
- browseable CIFS shares: [108](#)
- browsers, supported: [5](#)
- buckets: [54](#)
  - (see also vaults)
- Amazon services: [56](#)
- Amplidata services: [56](#)
- Cleversafe services: [56](#)
- compression: [55](#)
- configuring: [56](#)
- key files for: [57](#)
- selecting HTTP or HTTPS: [55](#)

## C

- cache policies
  - about: [84](#)
  - advanced features, cautions: [86](#)
  - and CIFS: [84](#)
  - and polling agents: [89](#)
  - cache-utilization controls: [92](#)
  - for high availability: [34](#)
  - per-core filer: [84](#)
  - preventing monopolization: [92](#)
  - read/write mode: [84, 87](#)
  - read mode: [84, 86](#)
  - specifying: [85](#)
  - writeback delay, custom: [90](#)
- caches
  - invalidating data: [146](#)
  - schedules: [36](#)
  - selecting a mode: [52, 61](#)
- cache-utilization controls
  - precautions: [93](#)
  - specifying: [92](#)
- cautions, definition: [1](#)
- CIFS
  - about: [95](#)
  - access control lists (ACLs): [103](#)
  - access to a junction: [75](#)
  - configuring: [99, 102](#)
  - configuring namespaces: [96](#)
  - constrained delegation: [104](#)
  - core filer prerequisites: [95](#)
  - enabling: [101](#)
  - limitations: [96](#)
  - native volumes: [157](#)

- recommendations: 95
- security modes: 109
- setting cache policies for: 84
- shares (see CIFS shares)
- username mapping: 97
- CIFS shares
  - accessing: 109
  - advanced properties: 108
  - browseable: 108
  - creating: 107
  - deleting: 109
  - inherited permissions: 108
  - modifying: 109
  - oplocks: 108
  - read-only: 108, 108
  - security masks and modes: 108
  - strict locking: 108
- clauses, defining: 37
- clearing conditions: 123
- Cleversafe services, configuring: 56
- client role VLANs: 82
- clients
  - modifying access to a cluster: 145
  - problems reading or writing data: 149
  - viewing on Dashboard: 129
- cloud core filers
  - buckets, about: 54
  - compression: 55
  - credentials: 63
  - details, modifying: 61
  - ports, service provider: 55
  - service name: 54
  - vaults, about: 54
  - “creating” (adding): 53
- cluster role VLANs: 82
- clusters
  - adding core filers to: 50
  - configuring: 10
  - configuring for Active Directory: 98
  - invalidating data cache: 146
  - logging into: 7
  - logging out of: 9
  - modifying client access: 145
  - modifying configuration: 146
  - monitoring from outside Avere OS: 132
  - multiple VLANs on the network: 26
  - powering down: 145
  - prerequisites: 5
  - problems writing data at given intervals: 149
  - restarting services: 145
  - saving configuration settings: 155
  - username mapping: 99
- collecting
  - hot client statistics: 69
  - statistics: 145
- Common Internet File System (see CIFS)
- compression, enabling for cloud core filers: 55
- conditions
  - clearing: 123
  - hiding: 123
  - history: 126
  - unhiding: 125
  - viewing on Dashboard: 121
- configuring
  - Active Directory: 97
  - Amazon services: 56
  - Amplidata services: 56
  - Avere OS: 4
  - buckets: 56
  - CIFS: 99, 102
  - Cleversafe services: 56
  - clusters: 10
  - custom settings: 152
  - encryption: 57
  - export FSIDs for GNU Linux: 156
  - FlashMove or FlashMirror job: 112
  - IPMI cards: 49
  - netgroups: 42
  - NIS: 42
  - saved customer information: 151
  - static routes: 80
  - support settings: 150
  - ZFS core filers: 158
- constrained delegation, using with CIFS: 104
- Control Panel overview: 8
- core filers: 50
  - (see also cloud core filers, NFS core filers)
  - adding to a cluster: 50
  - cache policies: 84
  - changing cloud filer details: 61
  - changing NFS details: 60
  - CIFS prerequisites: 95
  - cloud: 53
  - invalidating: 60
  - maximum number: 64
  - multiple VLANs on the network: 2, 26
  - NFS: 51
  - prerequisite information: 50
  - removing: 60
  - verification time, custom: 90
  - verification time, enabling: 91
  - viewing on Dashboard: 128
  - write bandwidth: 91
  - ZFS, configuring: 158

- creating
  - CIFS shares: [107](#)
  - export policies: [31](#)
  - GNS (global namespaces): [73](#)
  - vservers, requirements: [66](#)
- credentials
  - cloud core filer: [63](#)
  - modifying: [63](#)
- customer information, saving: [151](#)
- customizing graphs: [136](#)
- custom settings
  - Avere Global Services only: [152](#)
  - configuring: [152](#)

## D

- Dashboard
  - about: [8](#)
  - downloading statistics from: [121](#)
  - performance graph: [117](#)
  - statistics: [119](#)
  - status bar tabs: [118](#)
- Dashboard tab: [117](#)
- data management
  - configuring a FlashMove or FlashMirror job: [112](#)
  - destination parameters: [115](#)
  - local directories: [111](#)
  - prerequisites: [111](#)
  - source parameters: [114](#)
  - starting a FlashMove or FlashMirror job: [116](#)
- Data Management tab: [110](#)
- directory services
  - about: [39](#)
  - LDAP: [41](#)
  - NIS: [42](#)
  - selecting: [39](#)
- dismissing alerts: [124](#)
- DNS-SD (DNS Service Discovery) (see Bonjour)
- documentation, available: [9](#)
- downloading alternate images: [142](#)
- drilling down into graphs: [136](#)

## E

- early writeback, about: [91](#)
- email alerts: [133](#)
- enabling
  - advanced networking: [78](#)
  - CIFS on a vservers: [101](#)
  - core filer verification time: [91](#)
  - encryption: [56](#)
  - high availability: [34](#)
  - hot client collection: [69](#)
  - oplocks: [108](#)

- Secure Proactive Support: [152](#)
- encryption
  - configuring: [57](#)
  - enabling: [56](#)
  - passphrases: [58](#)
- enhancements, summary: [1](#)
- export policies
  - creating: [31](#)
  - deleting: [32](#)
  - deleting rules: [32](#)
  - managing: [31](#)
  - qtrees: [32](#)
- export rules
  - about: [28](#)
  - deleting: [32](#)
- exports
  - and CIFS shares: [108](#)
  - configuring FSIDs: [156](#)
  - for high availability: [34](#)
  - unmounting: [32](#)

## F

- filers (see core filers, cloud core filers, NFS core filers)
- FlashCloud
  - about: [53](#)
  - licenses: [6](#)
- FlashMove and FlashMirror
  - about: [110](#)
  - (see also data management)
  - licenses: [6](#)
- FSIDs, configuring for GNU Linux: [156](#)

## G

- GID and UID entries, parsing: [160](#)
- GNS (global namespaces)
  - about: [64](#)
  - adding a junction: [74](#)
  - configuring for CIFS: [96](#)
  - creating: [72](#)
  - designing: [72](#)
  - maintaining: [73](#)
  - specifying: [67](#)
- graphs
  - adding predefined: [140](#)
  - adding user-defined: [140](#)
  - customizing: [136](#)
  - Dashboard: [117](#)
  - default: [135](#)
  - drilling down: [136](#)
  - dynamic: [135](#)
  - editing: [139](#)
  - graph sets, about: [140](#)

- history: [136](#)
- Latency: [121](#)
- minimum and maximum statistics: [136](#)
- Ops/Second: [119](#)
- Throughput: [120](#)
- graph sets, adding: [140](#)
- group and password files, parsing: [160](#)

## H

- hiding conditions: [123](#)
- high availability
  - about: [33](#)
  - and 2-node clusters: [35](#)
  - disabling: [36](#)
  - enabling: [34](#)
  - exports: [34](#)
  - repositories, modifying: [36](#)
  - repositories, specifying: [34](#)
- history
  - alerts: [126](#)
  - conditions: [126](#)
  - graphs: [136](#)
- home nodes, setting: [71](#)
- home shares: [107](#)
- hot clients, enabling collection: [69](#)
- hot files, viewing on Dashboard: [131](#)
- HTTP and HTTPS, using in cloud services: [55](#)

## I

- images, alternate: [142](#), [145](#)
- important
  - attempted export access causes cluster downtime: [32](#)
  - CIFS and NTP: [23](#)
  - CIFS and time source: [23](#)
  - definition: [1](#)
  - export required for HA: [34](#)
  - limit VMware optimization: [48](#)
  - modifying MTU: [20](#)
  - OpenLDAP settings required: [41](#)
  - set cache policy before HA: [34](#)
  - using NTP: [23](#)
- inherited permissions: [108](#)
- invalidating
  - cluster data cache: [146](#)
  - core filers: [60](#)
- IP addresses
  - adding to a vserver: [67](#)
  - client-facing, about: [67](#)
  - modifying: [70](#)
- IPMI cards, configuring: [49](#)

## J

- Javascript, required: [5](#)
- junctions
  - adding to a GNS: [74](#)
  - CIFS access: [75](#)
  - deleting: [76](#)
  - modifying: [76](#)

## K

- Kerberos
  - enabling: [43](#)
  - keytab files: [44](#)
  - realms: [43](#)
- Kerberos authentication, selecting: [30](#)
- key files
  - creating: [57](#), [59](#)
  - passphrase: [58](#)
  - recovery: [57](#)
- keytab files, Kerberos: [44](#)

## L

- Latency, viewing: [121](#)
- Latency graph: [135](#)
- LDAP
  - configuring: [41](#), [45](#)
  - encryption: [46](#)
- licenses, activating: [6](#)
- limiting cache use: [92](#)
- local directories
  - data management: [111](#)
  - enabling: [94](#)
- LZ4 and LZ4HC compression: [55](#)

## M

- management role VLANs: [81](#)
- mapping, usernames for Active Directory: [97](#)
- masks and modes, security: [108](#)
- master key files (see key files)
- monitoring clusters from outside Avere OS: [132](#)

## N

- namespaces: [64](#)
  - (see also GNS (global namespaces))
  - global, specifying: [67](#)
  - simple: [66](#)
- netgroup files, format: [162](#)
- netgroups, configuring: [42](#)
- networking, advanced (see advanced networking)
- NFS
  - core filers, modifying details: [60](#)
  - core filers, “creating” (adding): [51](#)
  - exports and CIFS shares: [108](#)

- NFSv3 client requirements: [5](#)
- NFSv4 access control lists (ACLs): [103](#)
- NIS, configuring: [42](#)
- nodes
  - setting home: [71](#)
  - viewing on Dashboard: [129](#)
- notes, definition: [1](#)

## O

- online help, about: [9](#)
- oplocks, enabling: [108](#)
- Ops/Second, viewing: [119](#)
- Ops/Second graph: [135](#)
- Ops graph: [135](#)
- optimization, read-only: [108](#)

## P

- passphrases, for encryption key files: [58](#)
- password and group files, parsing: [160](#)
- periods, read mode: [87](#)
- permissions, inherited: [108](#)
- polling agents, implementing: [89](#)
- ports, used with cloud filers: [55](#)
- predefined graphs: [140](#)
- prerequisites
  - Active Directory: [100](#)
  - authentication: [39](#)
  - browser: [5](#)
  - core filers for CIFS: [95](#)
  - JavaScript: [5](#)
- problems reading or writing data: [149](#)
- protocols, supported client: [5](#)

## Q

- qtrees
  - and native CIFS volumes: [157](#)
  - applying: [32](#)

## R

- ranges, IP address: [67](#)
- read/write mode cache policies: [84](#), [87](#)
- read mode cache policies: [86](#)
- read mode periods and schedules: [84](#)
- read-only CIFS shares: [108](#)
- read-only optimization: [108](#)
- realms, Kerberos: [43](#)
- rebooting into alternate images: [145](#)
- regions, supported Amazon: [56](#)
- removing
  - core filers: [60](#)
  - vservers: [65](#)
- repositories, high availability

- modifying: [36](#)
- specifying: [34](#)
- restarting cluster services: [145](#)
- reverse name resolution: [43](#)
- reverting to previous Avere OS: [143](#)

## S

- schedules
  - clauses: [37](#)
  - defining: [36](#)
  - deleting: [38](#)
  - for read mode periods: [88](#)
- Secure Proactive Support
  - enabling: [152](#)
  - information sent: [152](#)
- services, restarting cluster: [145](#)
- setting
  - Dashboard statistics: [119](#)
  - home nodes: [71](#)
- settings
  - cluster: [10](#)
  - configuration, saving: [155](#)
  - custom: [152](#)
  - support: [149](#), [150](#)
- Settings tab: [8](#)
- shares
  - home: [107](#)
  - regular: [107](#)
- shares, CIFS
  - and NFS exports: [108](#)
  - creating: [107](#)
- simple namespaces, selecting: [66](#)
- SMB (Server Message Block) (see CIFS)
- SNMP, configuring: [134](#)
- specifying high availability repositories: [34](#)
- SPS (see Secure Proactive Support)
- static routes, configuring: [80](#)
- statistics
  - collecting: [145](#)
  - downloading: [121](#)
  - graph displays: [136](#)
  - viewing on Dashboard: [119](#)
- strict locking in CIFS shares: [108](#)
- support
  - customer information: [151](#)
  - package updates: [151](#)
  - settings, configuring: [150](#)
- supported
  - Amazon regions: [56](#)
  - browsers: [5](#)
  - client protocols: [5](#)
- Support tab: [149](#)

- suspending vservers: [65](#)
- syslog server, specifying: [134](#)
- System Maintenance page: [144](#)
- system performance, viewing: [118](#)

## T

- tabs, Dashboard status bar: [118](#)
- Throughput, viewing: [120](#)
- Throughput graph: [135](#)
- tooltips, viewing: [9](#)

## U

- UID and GID entries, parsing: [160](#)
- unhiding conditions: [125](#)
- UNIX/SYS authentication, selecting: [30](#)
- unmounting exports: [32](#)
- updating
  - Avere OS: [142](#)
  - support packages: [151](#)
- user-defined graphs: [140](#)
- username mapping: [97](#)

## V

- vault settings: [55](#)
- verification time, enabling and disabling: [91](#)
- virtual servers (see vservers)
- VLANs
  - about: [77](#)
  - adding IP address ranges: [67](#)
  - client roles: [82](#)
  - cluster roles: [82](#)
  - creating: [78](#)
  - management roles: [81](#)
  - multiple on core filer network: [2](#), [26](#)
  - roles: [79](#)
  - static routes: [80](#)
- VMware optimization: [48](#)
- vservers
  - about: [64](#)
  - adding IP address ranges: [67](#)
  - configuring: [68](#)
  - creating: [66](#)
  - enabling CIFS: [101](#)
  - namespace type: [66](#)
  - no QoS parameters: [64](#)
  - removing: [65](#)
  - renaming: [68](#)
  - suspending: [65](#)
  - viewing on Dashboard: [127](#)

## W

- writeback, early: [91](#)

- writeback delays, custom: [90](#)
- write bandwidth, modifying: [91](#)

## Z

- ZFS core filers, configuring: [158](#)