VDURA Data Platform User Guide

Monitoring and managing the VDURA Data Platform

Documentation for PanFS 12.0.0

© 2002–2026 VDURA Inc. All rights reserved. VDURA, the VDURA logo, VDURA Data Platform, VeLO, VPOD, DirectFlow, PanView, PanMove, and Multi-Level Erasure Coding (MLEC) are trademarks or registered trademarks of VDURA, Inc. in the U.S. and/or other countries. All other trademarks, registered trademarks, trade names, company names, and service marks are the respective properties of their holders.

Parts of this product are derived from Berkeley Database software, licensed from Sleepycat Software. Parts of this product are derived from Berkeley BSD systems, licensed from the University of California. The BDB software is also BSD licensed, and carries the following copyrights:

- Copyright © 1990–2009 Oracle. All rights reserved.
- Copyright © 1990, 1993, 1994, 1995 The Regents of the University of California. All rights reserved.
- Copyright © 1995, 1996 The President and Fellows of Harvard University. All rights reserved.
- Copyright © 2000–2005 INRIA, France Telecom

Parts of this product are derived from gd Version 1.8.4, licensed by Boutell.Com, Inc.

Copyright © 1982, 1986, 1991, 1993 The Regents of the University of California. All rights reserved.

© UNIX System Laboratories, Inc.

All or some portions of this file are derived from material licensed to the University of California by American Telephone and Telegraph Co. or Unix System Laboratories, Inc. and are reproduced herein with the permission of UNIX System Laboratories, Inc.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Disclaimer

VDURA has carefully checked the information in this document and believes it to be accurate. However, VDURA assumes no responsibility for any inaccuracies that this document may contain. In no event will VDURA be liable for direct, indirect, special, exemplary, incidental, or consequential damages resulting from any defect or omission in this document, even if advised of the possibility of such damages.

Any alterations or inconsistent replications of this document cannot be considered accurate and are void. VDURA may not be able to support issues related to altered documents that are not approved by subject matter experts.

Document version 12.0.0-2026-10-19-18:18:38

1. Introduction 1
1.1 About This Guide 1
1.1.1 Organization of This Guide 1
1.1.2 Document Conventions 2
1.1.3 Navigation Instructions 2
1.2 Additional Resources 4
1.2.1 Online Help 4
1.2.2 Product Documentation 4
1.2.3 Release Notes 4
1.3 Contact VDURA 4
2. Overview of the VDURA System 6
2.1 VDURA System Architecture 6
2.1.1 Single Namespace 6
2.1.2 Object Storage 6
2.2 VDURA V5000 Certified Platform Hardware 8
2.2.1 V5000 Director Nodes 8
2.2.2 V5000 Flash Storage Nodes 8
2.2.3 V5000 Capacity Expansion Enclosures 10
2.3 VDURA Software 10
2.3.1 DirectFlow Client 10
3. PanActive Manager 14
3.1 Connect to PanActive Manager 14
3.2 Navigate PanActive Manager 15
3.2.1 Important Links 16
3.2.2 Tabs, Menus, and Pages 17
3.2.3 Additional Links and Tools 18
4. Network Access 22
4.1 Configure System Settings 22
4.1.1 Administrator Access 23
4.1.2 System Name, IP Addresses, and Network Information 26
4.1.3 IP Address Ranges for Nodes 32
4.1.4 IP Aliases 34
4.1.5 Default Routers 36
4.1.6 Static Routes 37

4.1.7 Subnets ... 38
4.1.8 NIC Configuration ... 39
4.1.9 NTP and Time Zone Settings ... 39
4.1.10 DNS Settings ... 42
4.1.11 DNS Delegation ... 43
4.1.12 Lock Management ... 51
4.1.13 Configure Notification Services ... 52
4.1.14 Configure SNMP ... 61

4.2 Configure VDURA System Access ... 63
4.2.1 DirectFlow Client Support ... 63
4.2.2 Configure Netgroups ... 78
4.2.3 S3 Connectivity ... 80

4.3 Permissions ... 80
4.3.1 VDURA ACL-Based Permissions ... 80

4.4 UNIX Name Resolution Services ... 91
4.4.1 Configure Name Service Preferences ... 91
4.4.2 Configure NIS ... 93
4.4.3 Configure LDAP ... 94

4.5 SSH Key Management ... 101
4.5.1 Generate a New System Public SSH Key ... 101
4.5.2 Configure Admin and Guest User SSH Keys ... 102

5. Configure Storage Data ... 104
5.1 StorageSets and Volumes ... 104
5.1.1 About StorageSets ... 104
5.1.2 About Volumes ... 105
5.1.3 Understanding Capacity in StorageSets and Volumes ... 106

5.2 Configure and Manage StorageSets ... 109
5.2.1 List StorageSets ... 110
5.2.2 Create StorageSets ... 113
5.2.3 Manage StorageSets ... 114
5.2.4 Reconstruct Failed Storage Nodes in a StorageSet ... 121

5.3 Configure and Manage Volumes ... 121
5.3.1 List Volume Details ... 122
5.3.2 Create Volumes ... 126
5.3.3 Manage Volumes ... 136

5.3.4 Configure Automatic Volume Services Failover 144

5.4 Volume Data Compression 150
5.5 Configure Erasure Coding with a RAID Layout 151
5.5.1 Extended File System Availability (EFSA) 152
5.5.2 RAID 6+ 153
5.5.3 RAID 5+ 153
5.5.4 RAID 10+ 154

5.6 Horizontal Parity 154
5.7 Multi-Level Erasure Coding 154
5.8 User and Group Quotas 155
5.8.1 Configure User and Group Quotas 156
5.8.2 Common User Quota Limit Configurations 164
5.8.3 Monitor User and Group Quotas 167

5.9 Enclosure Migration 169
5.9.1 Before You Begin 169

5.10 Realm Management 175
5.10.1 Add a Director to the Repset 177

6. Security Enhanced Linux in VDURA Systems 178
6.1 Terminology 178
6.2 SELinux Overview 179
6.3 PanFS and MAC 180
6.4 Configuration and Use 180
6.5 Important Considerations 181

7. Manage Storage Hardware 182
7.1 Monitor and Manage the VDURA System 182
7.1.1 Manage System-Level Hardware 191

7.2 Monitor and Manage Enclosures 192
7.3 Monitor and Manage Nodes 197
7.3.1 Manage Node-Level Hardware 199
7.3.2 Node Management Tasks 209
7.3.3 Remove a Lost Node 211
7.3.4 Disk Firmware 211
7.3.5 Node Temperature 212
7.3.6 Replace a Node 212

7.4 Network and NIC Status 212

8. Monitor Performance 215
8.1 Display Performance Status 215
8.1.1 System-At-A-Glance 215
8.1.2 Storage Statistics 219
8.1.3 Drive Statistics 222
8.1.4 Director Statistics 224
8.1.5 DirectFlow Clients 225
8.2 Create Reports 226
8.2.1 Report Intervals 226
8.2.2 Performance Report 227
8.2.3 System Statistics Report 227
8.2.4 Capacity Report 229
8.2.5 Asset Report 230
8.2.6 Volume Report 231
9. Upgrade VDURA Software 233
9.1 Upgrade System Software 233
9.1.1 Before You Begin 233
9.2 Upgrade and Uninstall DirectFlow Client Software 237
9.2.1 Upgrade DirectFlow Client Software 238
10. End-to-End Encryption 239
10.1 Introduction and Overview 239
10.1.1 What is Encrypted 240
10.2 Key Management 240
10.2.1 Role-Based Authentication 240
10.2.2 Encryption Domain 240
10.2.3 KMS Integration 241
10.3 DirectFlow Client Behavior 241
10.4 Gateway Protocols 241
10.5 Configuration Considerations 241
10.5.1 External KMS Configuration 241
10.5.2 Key Management 242
10.5.3 Encrypted Volume Configuration 242
10.5.4 Client Configuration 242
10.5.5 Secure Status 242
10.5.6 Mounting PanFS 242

10.6 Failure Modes 242
10.6.1 KMS Network Connection Issues 243
10.6.2 Lost or Destroyed Keys 243
10.6.3 Compromised Keys 243
10.6.4 Reading Data without a Key 243
10.6.5 Reading Data with the Wrong Key 243

10.7 Performance Considerations 243
10.8 Upgrades and Downgrades 243
10.9 Encryption Configuration Procedures 244
10.9.1 KMIP Server Configuration 244
10.9.2 Realm-Wide Encryption Configuration 245
10.9.3 Volume-Level Encryption Configuration 245
10.9.4 Client Interface Configuration 246

11. Kubernetes CSI Driver for PanFS 248
11.1 The PanFS CSI Driver 248
11.2 Installation, Configuration, and Diagnostics 248

12. Remote Direct Memory Access in VDURA Systems 250
12.1 Introduction 250
12.2 Prerequisites 250
12.3 Enable or Disable RDMA on the Realm 251
12.4 Mount the RDMA Client 252
12.4.1 Mount Options 252
12.4.2 Error Messages 254
12.4.3 Known Limitations 254

13. Troubleshoot Your System 256
13.1 System Registration 256
13.2 View Events 256
13.2.1 Monitor Login Activity with the Historical Event Log 261

13.3 Perform Network Diagnostic Tests 261
13.4 Perform System Diagnostics 263
13.4.1 Select a Diagnostic Test 263
13.4.2 Download a Diagnostic Test 265

13.5 Troubleshooting 265
13.5.1 Collect State Information 266
13.5.2 Enable Diagnostic Tracing 267
13.5.3 Collect DirectFlow Client Traces 267

14. PanMove Sync 268
14.1 PanMove Sync Basics 268
14.1.1 Operation 269
14.1.2 Compatibility 271
14.1.3 Operation Requirements 271
14.2 Configuration and Use 272
14.2.1 Before Using PanMove Sync Manager 272
14.2.2 Start PanMove Sync Manager 273
14.2.3 Verify Client Setup 281
14.2.4 Create a Client-List File 281
14.2.5 Specify the Number of Processes to Start on Each Client 283
14.2.6 Replicate or Exclude Specific Files 283
14.2.7 Resume an Interrupted Replication 284
14.2.8 Retry Replication of Failed Files 284
14.2.9 Terminate a Replication Process 285
14.2.10 Manually Start a Worker 285
14.2.11 Specify a Log File Location 286
14.2.12 Example Scenarios 287
14.3 Logging and Error Handling 287
14.4 Limitations 288
15. System Limits and Requirements 290
16. VDURA System Error Messages 292

1. Introduction

Welcome to the *VDURA Data Platform User Guide*. This introduction includes the following sections:

- About This Guide
- Additional Resources
- Contact VDURA

1.1 About This Guide

The *VDURA Data Platform User Guide* describes how to use PanActive Manager™ to monitor and manage your VDURA® system. This guide is written for system and network administrators who are familiar with Network Attached Storage (NAS) systems and Storage Area Networks (SANs). Users should also be familiar with administering and managing file systems. A working knowledge of network protocols, network administration and management, and the Linux®, UNIX®, and/or Windows® operating systems is also required.

1.1.1 Organization of This Guide

The *VDURA Data Platform User Guide* includes the following chapters:

- Chapter 1, "Introduction" is an introduction to the rest of the *VDURA Data Platform User Guide*.
- Chapter 2, "Overview of the VDURA System" provides an introduction to the VDURA Data Platform™.
- Chapter 3, "PanActive Manager" provides an introduction to PanActive Manager, which enables you to monitor and manage performance of your VDURA system.
- Chapter 4, "Network Access" describes how to configure network administration and management settings and how to set up DirectFlow® clients to access the system.
- Chapter 5, "Configure Storage Data" describes how to create and monitor StorageSets and volumes, configure RAID, set user and group quotas, perform disk migration, and configure realm management.
- Chapter 6, "Security Enhanced Linux in VDURA Systems" explains the VDURA implementation of Security Enhanced Linux (SELinux).
- Chapter 7, "Manage Storage Hardware" describes how to manage nodes.
- Chapter 8, "Monitor Performance" describes how to use PanActive Manager to display status summaries, event listings, and reports to monitor performance of the VDURA system.
- Chapter 9, "Upgrade VDURA Software" describes how to maintain and upgrade VDURA system software.

- Chapter 10, “End-to-End Encryption” provides information about and instructions for managing VDURA’s end-to-end encryption implementation.
- Chapter 11, “Kubernetes CSI Driver for PanFS” explains how to configure and use CSI with PanFS.
- Chapter 12, “Remote Direct Memory Access in VDURA Systems” explains the VDURA implementation of Remote Direct Memory Access (RDMA).
- Chapter 13, “Troubleshoot Your System” describes how to run diagnostics and troubleshoot problems you might experience with your system.
- Chapter 14, “PanMove Sync” describes how to use PanMove Sync, a Linux-based client-side parallel data replication tool that allows customers to replicate and migrate data across a LAN.
- Chapter 15, “System Limits and Requirements” describes system limits and requirements.
- Chapter 16, “VDURA System Error Messages” provides details about VDURA system error messages.

1.1.2 Document Conventions

This manual uses the following standard set of typographical conventions to introduce new terms, illustrate screen displays, and describe command syntax.

Convention	Meaning
boldface	Identifiers (names of classes, objects, constants, events, functions, program variables, forms, labels, and reports), environment variables, filenames, table names, column names, icons, menu items, command names, and other similar terms appear in boldface.
`courier`	Information that the product displays and information that you are instructed to enter appear in a courier typeface.

1.1.3 Navigation Instructions

VDURA user documentation uses a simple, consistent format to tell you how to get to the correct PanActive Manager page for the task described. Navigation instructions adhere to the following format:

Click a tab > click a left menu item > click a specific page.

For example, to generate a capacity report, this guide will use **Status & Reports > Reports > Capacity** to navigate you to the correct page.

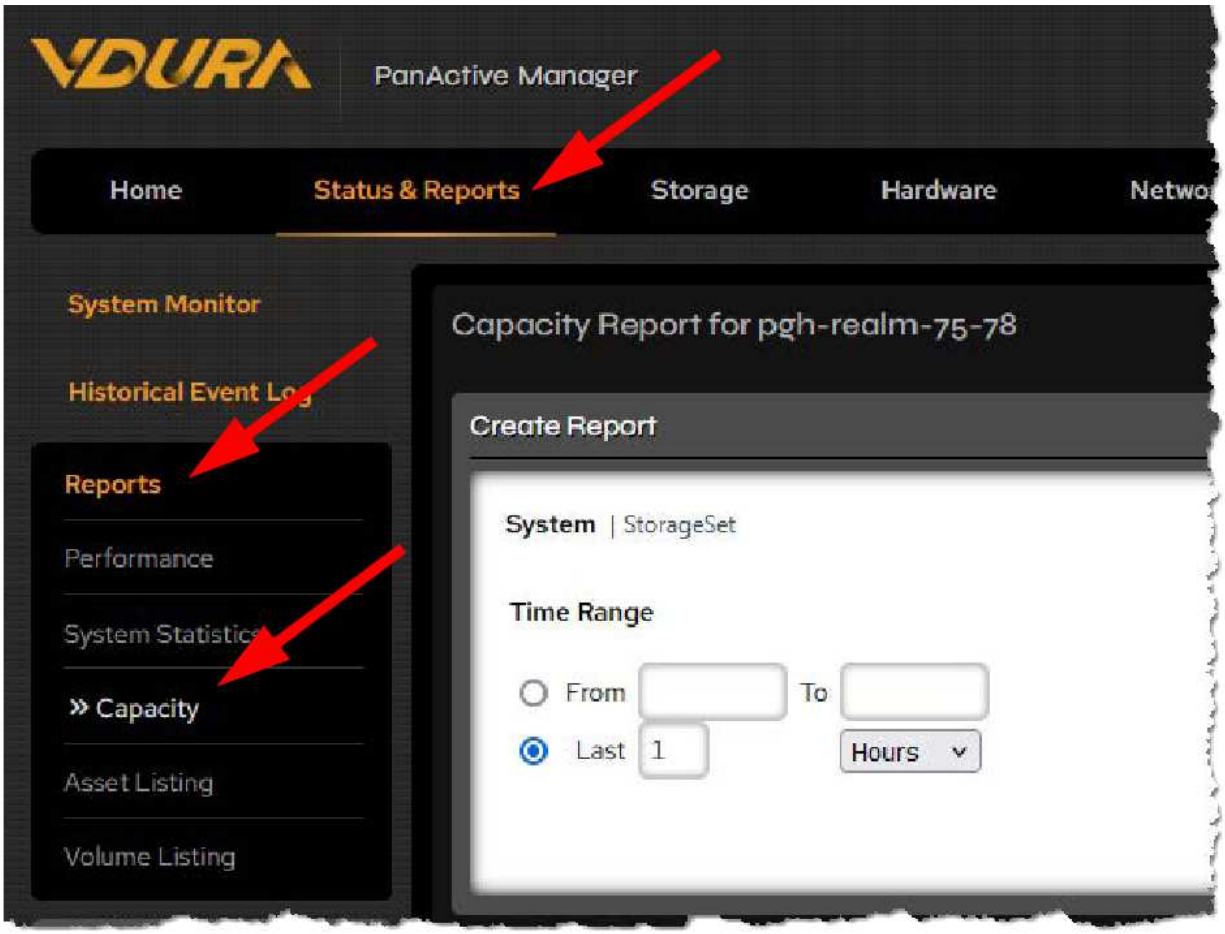


Figure 1-1. Example navigation.

An exception to the navigation rules occurs when a left menu item does not have pages listed under it. Some left menu items have only one page that appears by navigating to **tab > left menu item**. For example, **Status & Reports > Historical Event Log**.

> **NOTE:** Clicking a left menu item will by default open the first page in that category. For example, **Storage > Volumes > Listing** takes you to the same page as **Storage > Volumes**.

1.2 Additional Resources

1.2.1 Online Help

Online help is available in PanActive Manager in two ways:

- **Entire online help contents**: From the **Welcome** page, click the **Help** link to view the entire online help contents.
- **Page-level help**: From a particular page in PanActive Manager, click the **Help** link to view help for that page.

1.2.2 Product Documentation

The full customer documentation set for VDURA products is available at vdura.com/resources/documentation.

1.2.3 Release Notes

Release notes supplement the information in this manual by describing new features, identifying fixed issues, and explaining known problems and their workarounds. Release notes are available at vdura.com/resources/documentation. Be sure to look for the correct version(s) by noting the version number in the release notes file names.

Be sure to examine the release notes before you begin the installation and setup process, as they contain important information.

1.3 Contact VDURA

There are several ways to contact VDURA, depending on your needs:

Internet	The **support.VDURA.com** website provides access to submit a service request and track orders online. **Support.VDURA.com** also offers access to custom software downloads.
Technical Support	Get technical support and file service requests using any of these methods: • **Telephone—North America**: 1-888-726-2727 • **Telephone—International**: 1.412.323.6439 • **Email**: technical-support@vdura.com

File a Service Request	**support.VDURA.com**

2. Overview of the VDURA System

This chapter provides an introduction to VDURA products, architecture, and operating environment. It includes the following topics:

- VDURA System Architecture
- VDURA V5000 Certified Platform Hardware
- VDURA Software

2.1 VDURA System Architecture

The VDURA system architecture provides a unified, scalable storage system that:

- Virtualizes the system into a single namespace so that it is easy to manage and administer.
- Clusters across all storage elements to scale capacity and performance.
- Optimizes the cluster design to ensure dynamic load balancing.
- Balances resources to optimize performance.
- Scales capacity symmetrically with processing, caching, and network bandwidth.
- Integrates new storage elements into the storage fabric automatically.
- Delivers higher bandwidth within existing Ethernet networks.

VDURA system architecture supports native DirectFlow file access while virtualizing data across all available Virtualized Protected Object Devices (VPODs).

2.1.1 Single Namespace

VDURA provides native DirectFlow file access within a single virtualized namespace, avoiding the costly management issues associated with managing multiple servers with multiple namespaces.

2.1.2 Object Storage

The **object** is the fundamental unit of data storage in the VDURA system. Unlike files or blocks, which are used as the basic components in conventional storage systems, an object is a combination of file data plus a set of attributes (metadata) that define various aspects of the data. These attributes can define, on a per-file basis, the RAID levels, data layouts, and quality of service. Unlike conventional block storage where the storage system must track all of the attributes for each block in the system, the object maintains its own attributes to communicate to the storage system how to manage a particular piece of data. This simplifies the task of the storage system and increases its flexibility by distributing management of the data with the data itself.

Each VPOD efficiently transfers data between the Ethernet or InfiniBand network and block storage devices. Aggressive data and metadata caches optimize storage device access. SSDs allow **write-behind** caching, and object-aware **read prefetching** optimizes seek overhead.

The VPOD implementation also minimizes disk space fragmentation problems. When writing to a disk, VPOD software determines the best possible location on the disk for the size, type, and location of data object being created. Regardless of file system protocol, the VPOD has knowledge of the state of the drive and is able to optimize every data object layout decision. When fragmentation occurs, the VPOD knows what data comprises complete data objects. It can dynamically redistribute the data layout of an object to minimize fragmentation. This design improves performance and reliability.

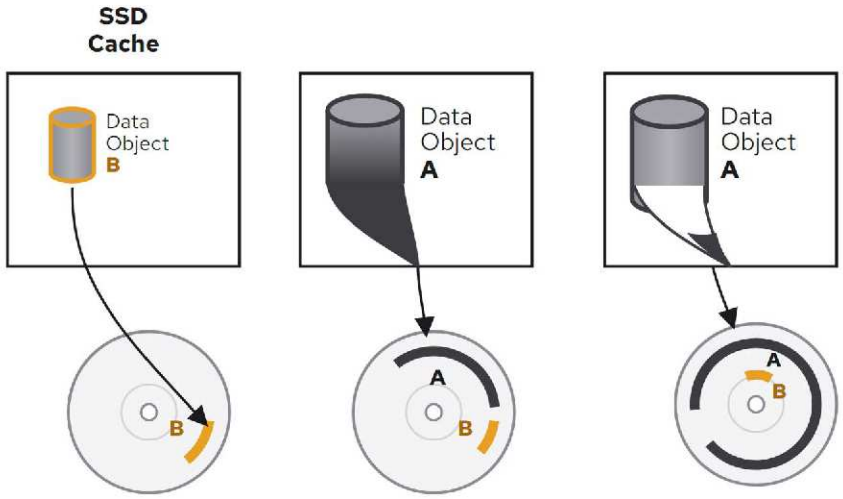


Figure 2-1. Data layout, caching, and prefetching algorithms.

When a client reads part of a data object, the VPOD uses its knowledge of the layout of the entire data object to prefetch the entire object, even if that object is placed on discontiguous disk sectors. This prefetching ability allows subsequent client read requests to be serviced at cache memory speeds. Write performance is also improved as writes are cached in the VPOD's SSD, allowing applications to quickly return to processing while the data is written to disk in the background.

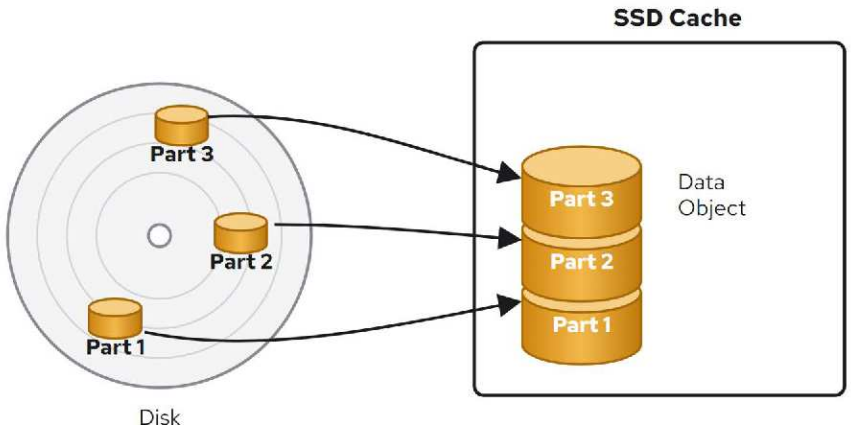


Figure 2-2. Data prefetching.

2.2 VDURA V5000 Certified Platform Hardware

VDURA V5000 Certified Platform™ hardware is an all-NVMe flash appliance engineered for AI and HPC pipelines that demand relentless GPU feed rates. Built on industry-standard servers, V5000 pairs flash performance with optional HDD capacity expansion, giving organizations a cost-balanced path from pilot to petabytes.

2.2.1 V5000 Director Nodes

V5000 Director Nodes anchor the VDURA control plane, running the flash-optimized VeLO Metadata Engine™ that coordinates billions of file operations and keeps AI pipelines saturated. Each compact server combines high-core CPUs, plentiful DRAM, and dual 200 Gb/s links for ultra-low-latency orchestration.

Instant secure-erase and optional self-encrypting drives safeguard data, while toolless service and redundant power maximize uptime for demanding AI and HPC deployments.

2.2.2 V5000 Flash Storage Nodes

Flash Storage Nodes form the performance-driven foundation of the data plane, dedicated exclusively to storing user data. Available in configurations of either all-NVMe flash for extreme performance or NVMe flash with HDD capacity expansion for high-performance, economical bulk storage, flash Storage Nodes deliver versatile and optimized infrastructure.

VPODs

Each VDURA Storage Node hosts 12 Virtualized Object Storage Devices™ (VPODs) rather than treating the entire server as a single failure domain.

VPODs are discrete, virtualized, and protected storage units that form the foundation of data storage. VPOD architecture ensures linear scalability and consistent parallel performance, accommodating thousands of nodes seamlessly within a single cluster.

V5000 Storage Nodes contain 12 NVMe SSD drives and can optionally be connected to a JBOD containing either 78 or 108 HDDs. Each VPOD in the Storage Node is associated with an SSD. If the Storage Node is a capacity expansion node, each VPOD will also be associated with several HDDs in the JBOD.

Each VPOD operates independently, but shares processor, memory, and network resources with the rest of the VPODs in the node.

In both PanActive Manager and PanCLI, a Storage Node appears as a combination of 12 VPODs, each with own IP and data serial node, plus one special item labeled "VPOD Container", also with its own IP. The VPOD Container represents the hypervisor role and main OS of the Storage Node. It runs PanFS software that deploys, configures, and monitors VPODs and also publishes hardware monitoring data (JBOD, sensors, fans, etc.) to the realm.

Since each VDURA Storage Node hosts multiple VPODs, this architecture introduces a finer unit of failure isolation.

- The unit of failure is now one VPOD, not the physical server.
- More VPODs per node increases operational granularity and cluster flexibility.
- Device-level failures are isolated to a single VPOD, eliminating the risk of full node failure.
- Storage reconstruction only affects the failed VPOD's component objects, not the entire node.
- Large files are striped across component objects in multiple VPODs using N+2 erasure coding (large files are striped across a set of VPODs plus additional VPODs into that stripe that store the P and Q data protection values of an N+2 erasure coding scheme).
- Small POSIX files are triple-replicated across VPODs.

If a VDURA Storage Node fails, the system must reconstruct only those VPODs that were on the failed node.

> **NOTE: VPOD** (Virtualized Object Storage Device) is VDURA's replacement for **OSD** (Object Storage Device). In this guide, you will notice that the terms OSD and VPOD are used interchangeably throughout PanActive Manager. For the purposes of the UI, they are interchangeable.

2.2.3 V5000 Capacity Expansion Enclosures

The VDURA V5000 system has two capacity expansion enclosure options: high-density enclosures containing 78 HDDs or ultra-dense enclosures with 108 HDDs. Both are 4U, 19" rackmount enclosures that come fully populated with HDDs. The 4U78 storage enclosure is ideal for one-meter deep racks, while the 4U108 fits 1.2-meter deep racks. V5000 storage enclosures support a wide range of HDDs. Both enclosure types also include two Titanium-level redundant power supplies.

2.3 VDURA Software

2.3.1 DirectFlow Client

Similar to massively parallel supercomputer architectures, VDURA DirectFlow software supports fully-parallel reads and writes from DirectFlow clients to OSDs without encountering bottlenecks. The DirectFlow software module is a loadable kernel module that transparently adds support for the VDURA file system protocols to Linux compute nodes. DirectFlow provides high-performance, high-fidelity POSIX semantics for large-scale distributed systems based on a cache consistency protocol with the VDURA metadata managers. In addition, there are specialized interfaces that optimize the large shared file workloads found in MPI-IO applications. For more information about the DirectFlow Client, see DirectFlow Client Support.

PanActive Manager

PanActive Manager provides an easy-to-use browser-based GUI to manage your VDURA system. The interface makes initial configuration, ongoing monitoring, report generation, and troubleshooting a smooth and understandable set of tasks. The SNMP interface allows integration of the storage system into enterprise data center management frameworks.

PanCLI

PanCLI is the VDURA Data Platform command-line interface. Administrators use it over ssh to monitor and manage the system from the command line instead of the PanActive Manager GUI.

PanCTL

PanCTL is the new command-line client for administering VDURA storage systems. It manages tenants, volumes, nodes, storage sets, and more from a terminal or script, using the same PanAPI backend as Vantage.

PanFS

PanFS is an object-based, parallel file system. Files are divided into chunks that are stored in different objects and accessed in parallel for performance and redundancy purposes. Data objects are containers for data and attributes that describe how the data fits into the overall distributed file system.

Files are striped across objects, and each file has its own map that defines the striping pattern. This allows different objects to have different striping patterns and provide different performance and redundancy characteristics. Because maps are just a small array of object identifiers and associated access capabilities, they are efficient to manage.

For example, small files and directory objects are mirrored so that they are cheap to create and cheap to update with small I/Os often associated with small files and directories. Large files are striped more widely using RAID on a per-file basis to provide redundancy with less capacity overhead than mirroring. Very large files use a two-level striping pattern that optimizes performance in large storage systems.

The storage map for a file can change dynamically to accommodate changes in the available storage resources. For example:

- If a storage node fails, component objects will automatically be reconstructed on alternate VPODs.
- The addition of new storage resources will result in automatic rebalancing of capacity so that all VPODs are loaded uniformly.

All data management is automatic and performed in the background while the system is online.

PanFS is built as a true parallel file system, engineered to handle the intense I/O demands of modern AI and HPC workloads. Each stored file is individually striped across many Storage Nodes, allowing each component piece of a file to be read and written in parallel, increasing the performance of accessing every file.

VDURA's parallel architecture dramatically accelerates data access, significantly boosting performance and throughput. Unlike other enterprise systems which route data through limited head nodes, causing potential bottlenecks and requiring additional backend network infrastructure, VDURA's DirectFlow Client communicates directly with all relevant Storage Nodes. Each compute server directly accesses the nodes holding the data, bypassing intermediary bottlenecks. Director Nodes manage metadata and coordinate system activity out-of-band, ensuring efficient data flow without interference or congestion.

This direct and parallel design eliminates traditional NAS hotspots, ensures predictable and scalable performance, and simplifies infrastructure by removing the need for a separate, costly backend network. VDURA architecture delivers seamless scalability, consistently high performance, and exceptional efficiency across every stage of the AI pipeline, from ingest and training to inference and long-term data retention.

PanFS delivers true linear scalability across both metadata and data services without compromise or complexity. AI workloads evolve fast, from early experimentation to scaled production across global clusters. Add Director Nodes to boost throughput for metadata-heavy tasks like model versioning and checkpoint tracking. Add Storage Nodes to scale bandwidth and capacity to support more training data, inference logs, or multi-tenant pipelines.

VDURA enables linear scalability, and growth is seamless and predictable. A 50 percent increase in Storage Nodes delivers 50 percent more throughput and capacity with no bottlenecks or architectural redesigns.

This flexibility is powered by VDURA's fully virtualized, distributed software stack:

- VeLO Metadata Engine instances run in-memory across Director Nodes, scaling to billions of parallel metadata operations, perfect for high-frequency file creation, access pattern analysis, and rapid AI job cycles.
- VPODs manage user data in independently scalable units, each with its own erasure-coded stripe and logic, ideal for bursty checkpoint writes, long-term data lake retention, or active model training sets.

Together, these software-defined services form an intelligent, self-balancing system that grows effortlessly with the pipeline. Whether scaling GPU clusters, expanding training data, or retaining more checkpoints, VDURA adapts instantly.

To ensure that problems are detected before they cause a failure, PanActive Manager notifies administrators of impending failures, allowing them to identify and prevent problems before they affect the system. Systems can be monitored for disk and fan failures, as well as temperature and other environmental threats.

Realm Management

VDURA realm management software logically clusters Directors so that all Directors are fully engaged in workload processing.

To enable scalability, realm management is designed to allow hundreds of VDURA Directors to cluster together while evenly and effectively distributing the workload. Realm management provides the following advantages:

- Initial distribution and ongoing management of files is optimized across all hardware elements while minimizing interdependence.
- System operation is dynamically improved by adapting to changing client usage patterns and changes in file sizes.
- A highly available file system automatically adapts to Director failures.

For more information on realm management, see Realm Management.

Sentinel

VDURA Sentinel is a cloud-based fleet-monitoring console for VDURA storage systems. It gives customers a single place to see the health of every site and realm they run. From there they can review alerts, open support tickets, and capture and send diagnostics. An on-site agent reads each realm's PanActive and Observability APIs read-only, and it sends a summarized snapshot to the cloud over outbound HTTPS only. Sharing is optional, and nothing connects back into the customer's network. On V5000 systems the agent is built into PanFS from 12.x, and on legacy hardware it's a standalone container.

Vantage

Vantage is VDURA's newest system management UI and will become a full replacement for PanActive Manager in further releases.

VeLO

Velocity Layered Operations™ (VeLO) is a key value store used within the platform's Director layer for handling high-performance metadata operations. Optimized for flash, VeLO supports the diverse range of workloads found in AI by supporting an infinitely-scalable number of VeLO instances in the same global namespace. Data is triplicated across each instance for superior durability.

3. PanActive Manager

This chapter introduces PanActive Manager, the browser-based GUI that enables you to monitor and manage the performance of your VDURA system. It contains the following topics:

- Connect to PanActive Manager
- Navigate PanActive Manager

> **NOTE:** You can also use a command-line interface to perform monitoring and management procedures. For information on the VDURA command-line interface (PanCLI), see the online help or the *VDURA PanCLI Reference Guide*, available at vdura.com/resources/documentation.

> **NOTE: VPOD** (Virtualized Object Storage Device) is VDURA's replacement for **OSD** (Object Storage Device). In this guide, you will notice that the terms OSD and VPOD are used interchangeably throughout PanActive Manager. For the purposes of the UI, they are interchangeable.

3.1 Connect to PanActive Manager

Network access is configured when the VDURA system is initially installed, as described in the *VDURA V5000 Certified Platform Installation and Service Guide*. To change these settings with PanActive Manager, see Configure System Settings.

> **NOTE:** PanActive Manager supports HTTPS protocols and NIS restrictions.

You can connect to PanActive Manager using any current major browser. Configure browser preferences as follows:

- Enable cookies.
- Enable JavaScript.
- Disable caching. Because information changes frequently on some PanActive Manager status pages, your browser's cache may cause out-of-date data to be displayed. If you choose to enable caching, configure your browser to retrieve the fresh page from the HTTP server whenever you want current data.

- Download and import the CA certificate. To ensure a secure connection to VDURA realms, download the CA certificate and import the certificate file into your browser to avoid having to manually set an exception in your browser. At the PanCLI command line, do one of the following:
  - Enter **cacert** to display the CA certificate information. Next, copy and paste the information into a text editor and save it as **cacert.pem** on the local machine. Finally, refer to your browser's help instructions to import the file.
  - Enter **cacert ?[user]@?[host]:[file]** to download the CA certificate file to a remote machine using **scp**. Refer to your browser's help instructions to import the file. When saving your file, be aware that the system does not support Unicode.

Enter the IP address for PanActive Manager in your browser. This is the IP address you assigned the VDURA system during installation. PanActive Manager opens, displaying the **Home** page.

In the **Sign In** box, enter **admin** for the user name and the password you set when you initially configured your VDURA system, then click **Sign In**.

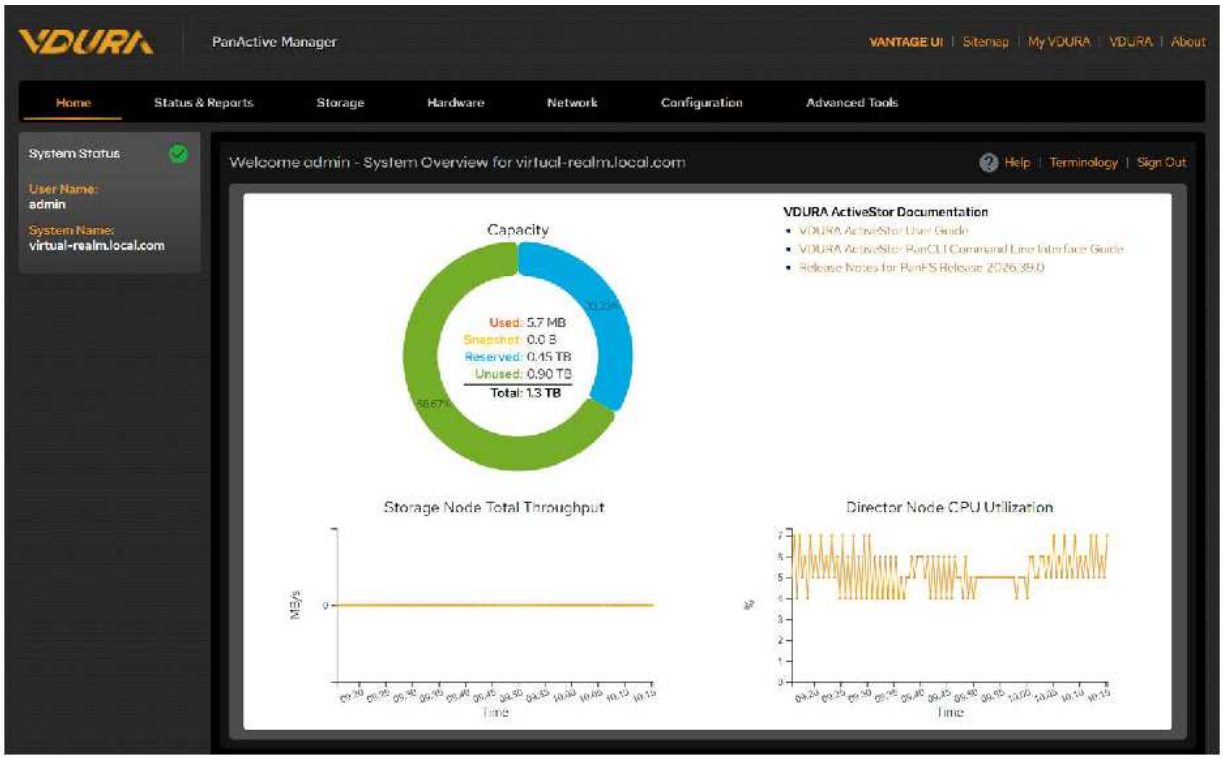


Figure 3-1. PanActive Manager Home page.

3.2 Navigate PanActive Manager

The **Home** page summarizes the current status of your system, including capacity and charts of storage node throughput and Director CPU utilization, and provides links to VDURA documentation.

3.2.1 Important Links

PanActive Manager contains several sets of links that provide a variety of information.

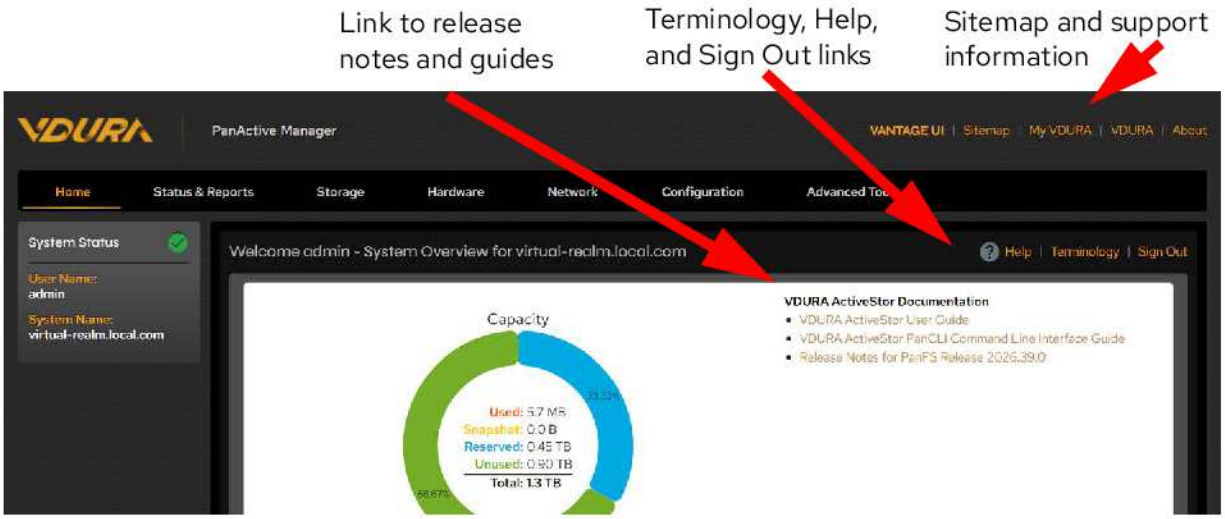


Figure 3-2. Helpful links.

Sitemap and Support Information

Links at the top right of most pages take you to the Vantage UI, a site map of PanActive Manager, support pages for VDURA products, and software version information.

- **VANTAGE UI**: Switch to the Vantage UI for this system.
- **Sitemap**: Open a clickable map of the main pages in PanActive Manager.
- **My VDURA**: Access product documentation, submit a service request, and download software releases.
- **VDURA**: Go to the VDURA corporate website.
- **About**: Display version information and your Customer ID.

Terminology and Sign Out Links

The following links appear in the right corner of each page below the menu tabs:

- **Help**: Click to display help topics relevant to the current page you are viewing.
- **Terminology**: Click to display a glossary of commonly used system terms.
- **Sign Out**: Click to end your session.

Documentation

The full customer documentation set for VDURA products is available at vdura.com/resources/documentation.

3.2.2 Tabs, Menus, and Pages

The tabs across the top of every page allow you to navigate PanActive Manager. Each tab displays left menu options for the tasks you can perform from that tab. Many of the left menu items contain additional subsets of links to particular screens.

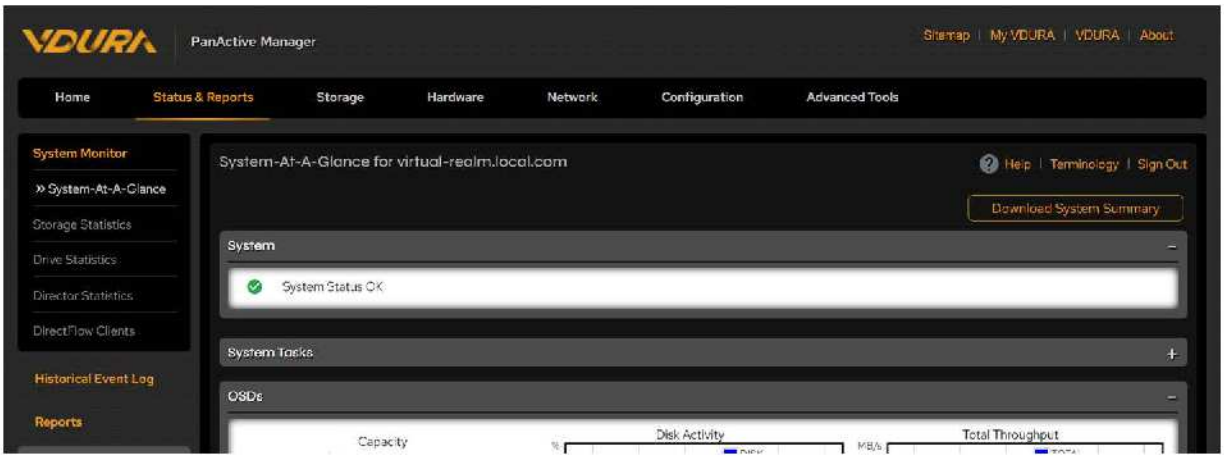


Figure 3-3. Tabs, menus, and pages.

The following table describes the navigation tabs that appear on each page.

Tab	Function
Home	Display the **Welcome** page containing System Status, System Overview, and VDURA documentation links, as well as total throughput and CPU utilization for the system.
Status and Reports	Display status summaries, event listings, and reports regarding your VDURA system. See Monitor Performance.
Storage	Configure StorageSets, volumes, user and group quotas, DirectFlow exports, and Netgroups. See Configure Storage Data and Network Access.

Tab	Function
Hardware	Display system-, enclosure-, and node-level reports and controls for VDURA hardware; configure realm management; and display hardware diagnostics. See Manage Storage Hardware and Configure Storage Data.
Network	Configure IP ranges, static routes, aliases, routers, and subnets; monitor NIC status and perform network diagnostics; enable storage network failover; configure enclosure- and node-level network settings; set NTP servers and timezones, DNS servers, and services; and configure name services, NIS, and LDAP. See Network Access.
Configuration	Configure network and administrator security settings, SSH keys, notification recipients, SNMP, and system registration. See Network Access.
Advanced Tools	Check version history, upgrade software, view file lock management, and perform diagnostic tracing. See Upgrade VDURA Software, Lock Management, and Troubleshoot Your System.

3.2.3 Additional Links and Tools

A **System Status** box appears on every page in the left pane. Click the status icon to view system status information.

If the system is locked, a link will appear to unlock the system.

On some pages, a System Tasks heading appears with a plus icon. Click the plus icon to display any tasks

currently in progress.

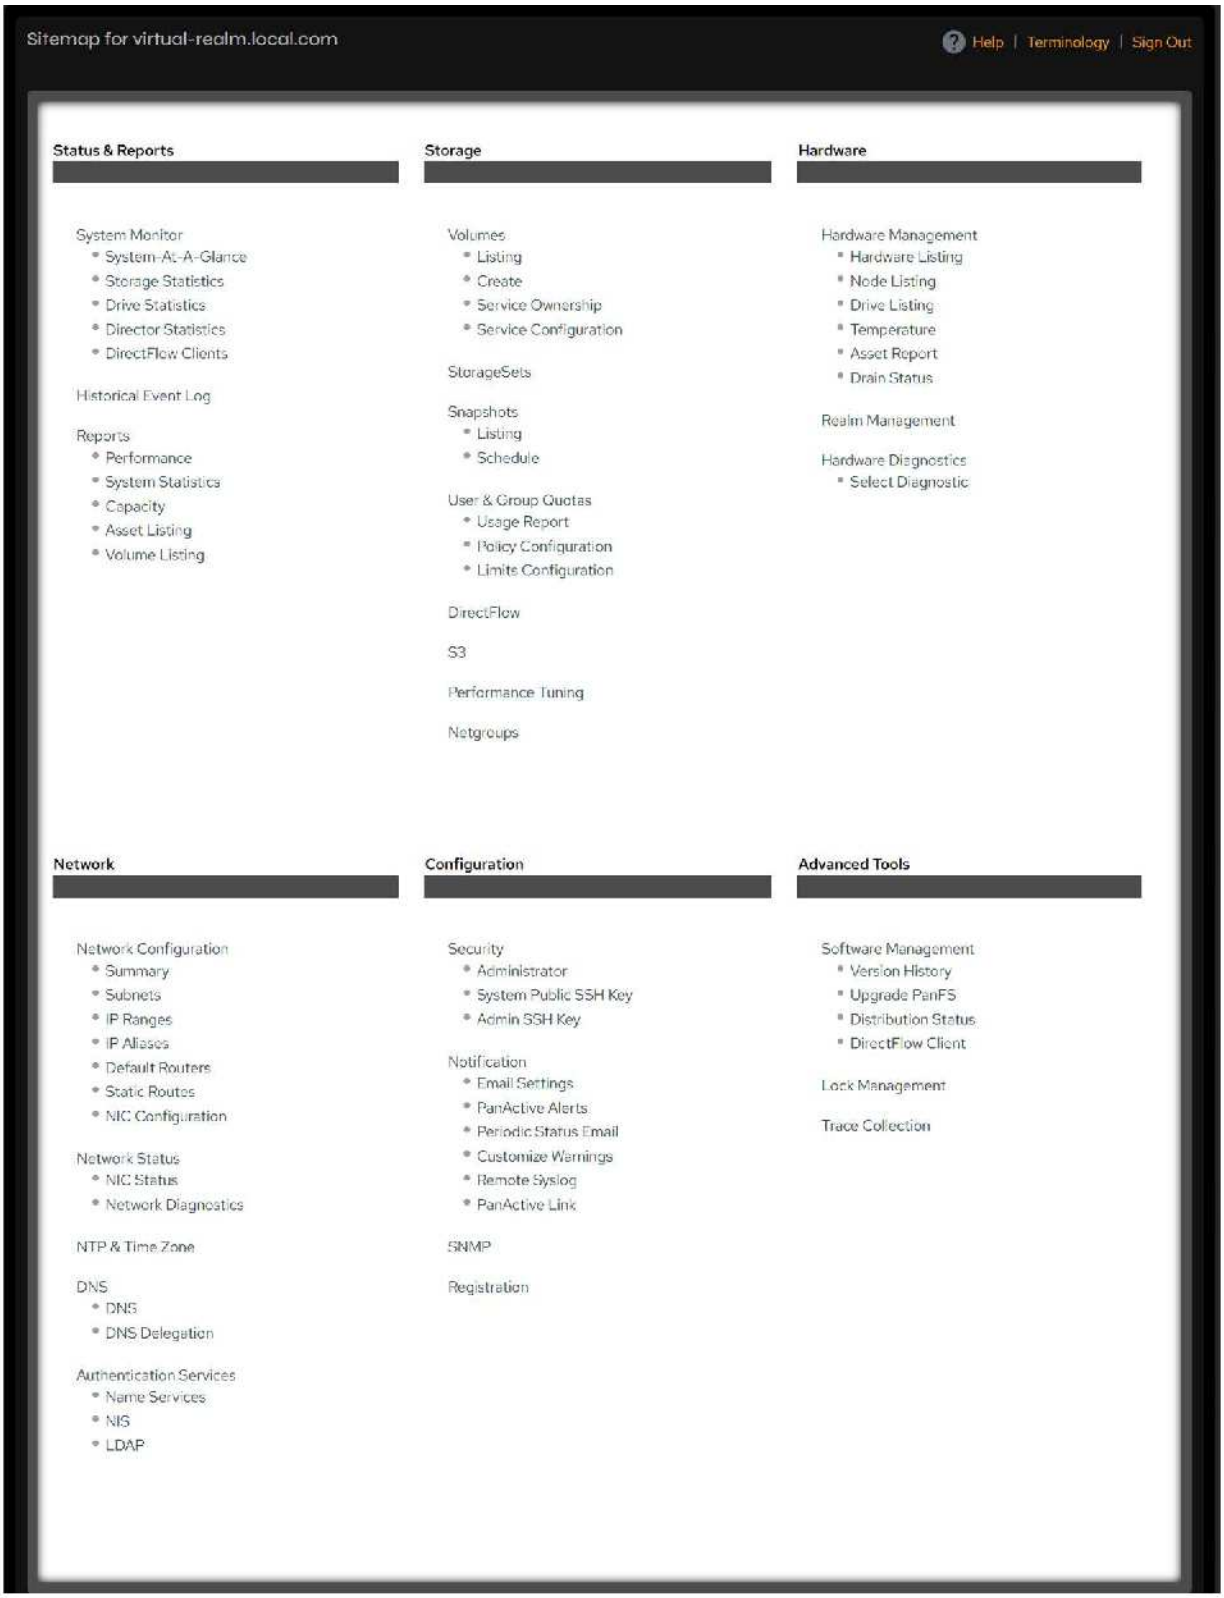


Figure 3-4. Sitemap page.

4. Network Access

This chapter describes how to use PanActive Manager to configure and manage network access to your system. It contains the following sections:

- Configure System Settings
- Configure VDURA System Access
- Permissions
- UNIX Name Resolution Services
- SSH Key Management

4.1 Configure System Settings

Network and administrator settings for your VDURA system are initially configured during installation using the Install Wizard. To modify the initial settings, use PanActive Manager or PanCLI. You can reconfigure the following network and administrator settings:

- Administrator Access
- System Name, IP Addresses, and Network Information
- IP Address Ranges for Nodes
- IP Aliases
- Default Routers
- Static Routes
- Subnets
- NIC Configuration
- NTP and Time Zone Settings
- DNS Settings
- DNS Delegation
- Lock Management
- Configure Notification Services
- Configure SNMP

4.1.1 Administrator Access

Click **Configuration > Security > Administrator**.

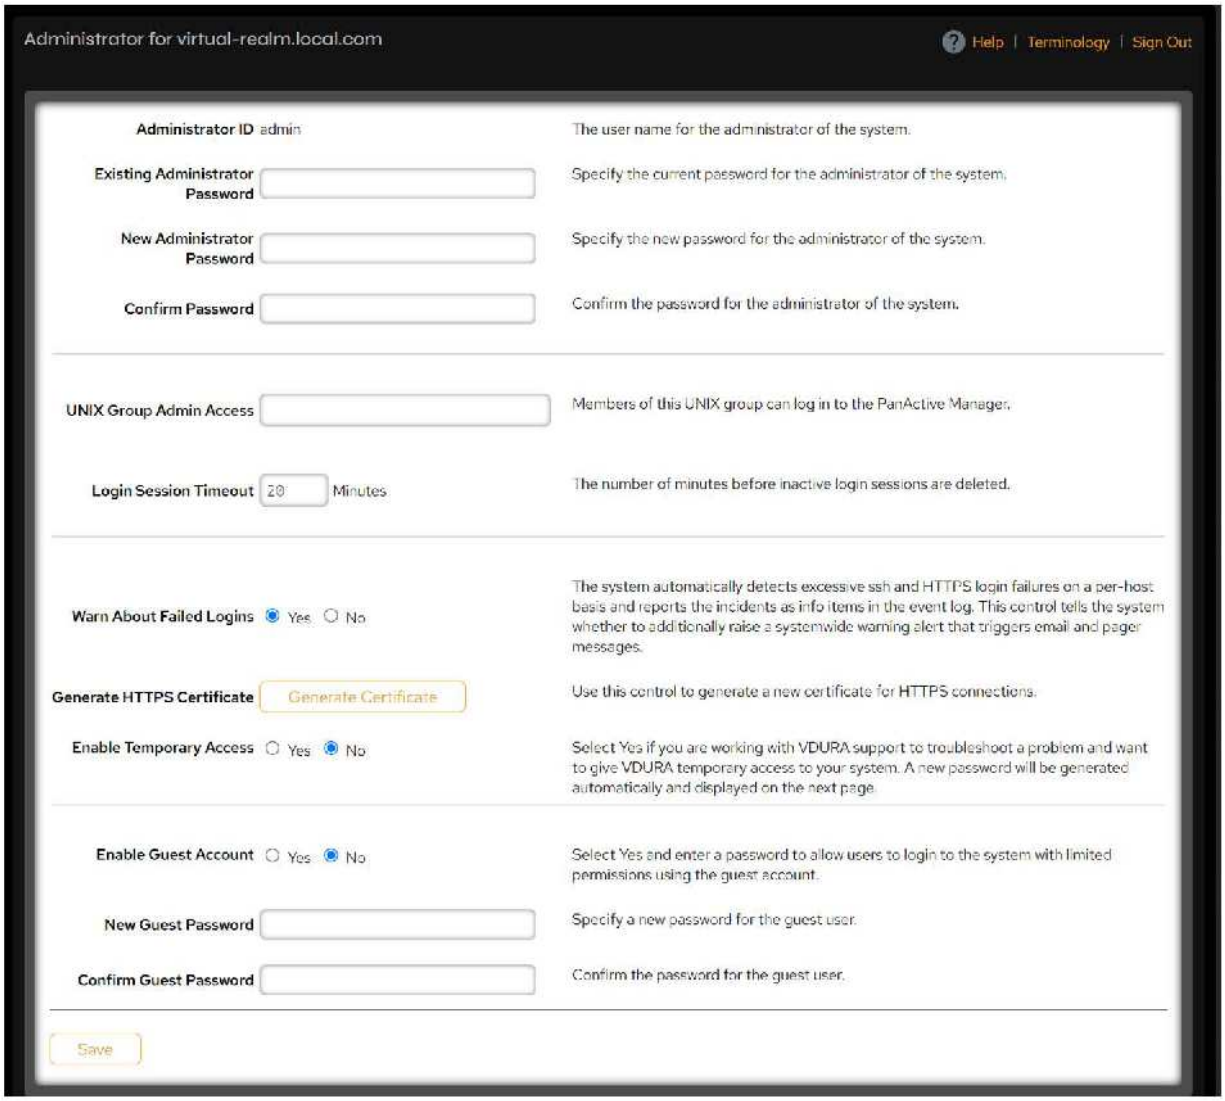


Figure 4-1. Administrator page.

Use the following controls to change administrator settings.

Control	Description
Administrator ID	The user name for the administrator of the VDURA system. The **Administrator ID** is hard-coded as “admin”. You cannot change this ID.
Existing Administrator Password	Enter the existing administrator password.

Control	Description
New Administrator Password	Enter the new password. VDURA recommends the administrator password contain at least six characters, with at least one letter and at least one number or symbol. The special characters ?, |, <, >, & are not allowed, nor are empty passwords allowed.
Confirm Password	Enter the new administrator password again for confirmation.
UNIX Group Admin Access	Members of this UNIX group can log in to PanActive Manager using their UNIX user name and password. If this field is blank, you can only log in using the default administrator ID.
Login Session Timeout	The number of minutes before inactive login sessions are deleted.
Warn About Failed Logins	The system automatically detects excessive ssh and HTTPS login failures on a per-host basis and reports the incidents as items in the Historical Event Log. When set to **Yes**, this control tells the system whether to additionally raise a system-wide warning alert that triggers email and pager messages.
Generate HTTPS Certificate	Click **Generate Certificate** to generate a new security certificate for HTTPS access. For more information, see HTTPS Security Certificate.
Enable Temporary Access	Select **Yes** if you are working with VDURA Technical Support to troubleshoot a problem and want to give VDURA temporary access to your system. The temporary user name and password are displayed on the screen when you save your changes. This account is automatically disabled 48 hours after you click **Yes**.

Control	Description
Enable Guest Account	Select Yes, then enter and confirm a password in the New Guest Password and Confirm Guest Password fields, to allow users to log in to the system with limited permissions. For more information, see Guest Accounts.

Click **Save** to commit your changes when finished.

> **NOTE:** For information on changing administrator settings with the command-line interface, see the PanCLI **admin** command in the *VDURA PanCLI Reference Guide*.

> **NOTE:** For information regarding login failure warnings and security logging, see View Events and Monitor Login Activity with the Historical Event Log.

HTTPS Security Certificate

With the **Generate Certificate** button, PanActive Manager can generate a security certificate that clients must accept whenever connecting to PanActive Manager via HTTPS. The first time clients connect to PanActive Manager by HTTPS, they are prompted to accept a certificate. The prompt indicates that the certificate was not signed by a trusted issuer and may not be authentic. Users can ignore the warning and safely accept the certificate.

Guest Accounts

Using PanActive Manager, system administrators can provide limited access for specific users by creating a guest account. This account permits read-only access to the system through PanActive Manager and PanCLI. Users assigned a guest account can view data and system configurations, but are not allowed to modify or change system configurations.

4.1.2 System Name, IP Addresses, and Network Information

Click **Network > Network Configuration > Summary**.

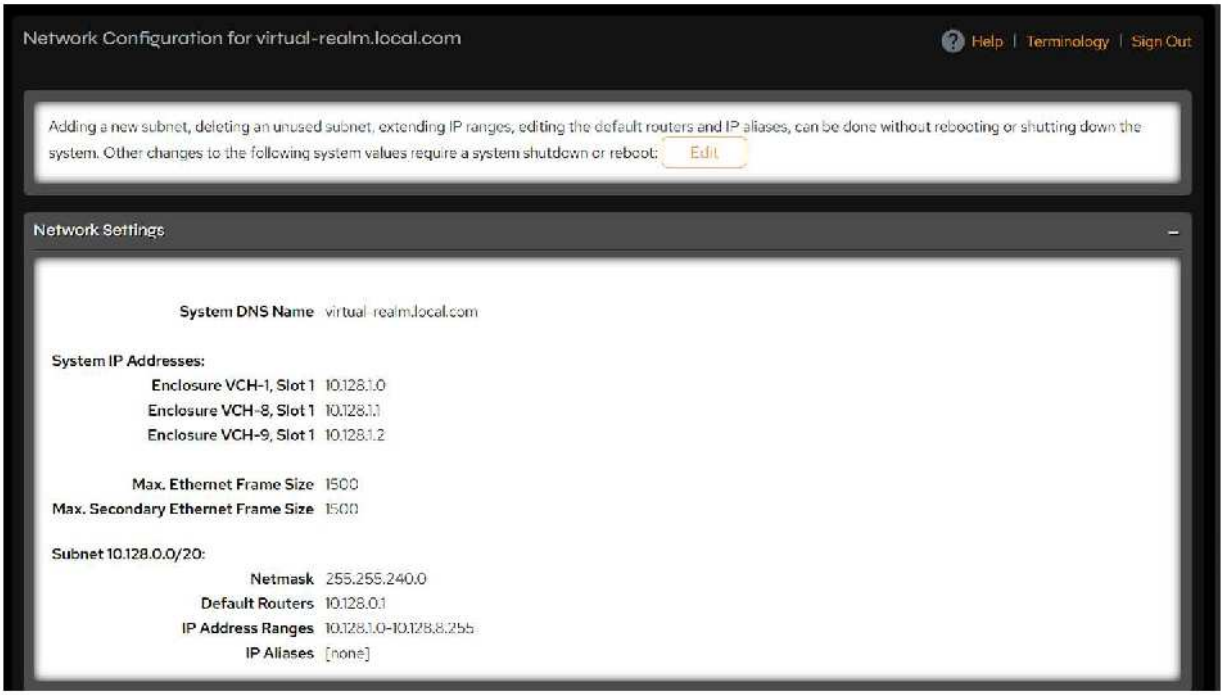


Figure 4-2. Network Configuration page.

Click the **Edit** button to open the **Edit Networking** page and revisit the settings that were initially configured using the Install Wizard, as well as other settings.

Changes made using the Edit button require a system shutdown or reboot. Changes to individual configurations are more easily accomplished by going directly to the relevant page. Use the pages under **Network Configuration** in the left menu instead of the **Edit** button to make changes directly to those areas without needing to reboot or shut down the system.

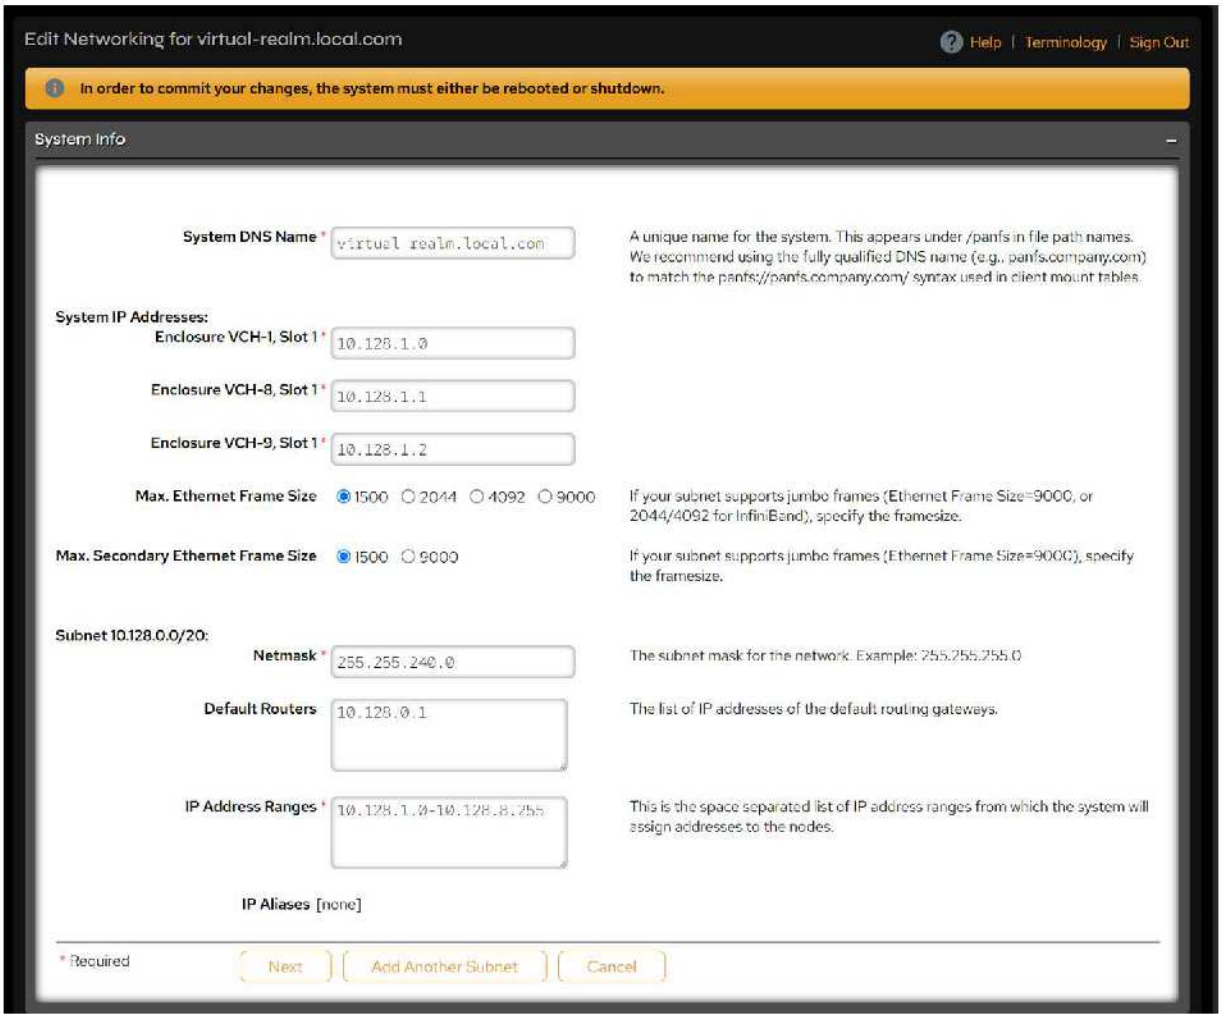


Figure 4-3. Edit Networking page.

Use the following controls to change settings.

Control	Description
System DNS Name	Specify a unique name for the system. This appears under **/panfs** in file path names. VDURA recommends using the fully qualified DNS name (e.g., panfs.company.com) to match the **panfs://panfs.company.com/** syntax used in client mount tables.
System IP Addresses Enclosure VCH-1, Slot 1	Specify the IP addresses for the system. These are the IP addresses you configured for the first director installed on your system. If you have enabled realm management, a separate **System IP Address** field is displayed for each director in the system. You must specify an IP address for each director in the realm. See Realm Management for more information.
Max. Ethernet Frame Size	If your subnet supports jumbo frames, select a size (1500, 2044, 4092, or 9000). Extending the frame size reduces server overhead and increases data throughput. To enable jumbo frames, you must restart the realm in PanActive Manager. **CAUTION**: If you enable jumbo frames and there are clients that access the VDURA system across links that do not support jumbo frames, the system will be inaccessible to those clients.

Control	Description
Max. Secondary Ethernet Frame Size	This option applies to the secondary interface. If your subnet supports jumbo frames, select a size (1500 or 9000). Extending the frame size reduces server overhead and increases data throughput. To enable jumbo frames, you must restart the realm in PanActive Manager. **CAUTION**: If you enable jumbo frames and there are clients that access the VDURA system across links that do not support jumbo frames, the system will be inaccessible to those clients.
Netmask	Specify the subnet mask for the network. Example: `255.255.255.0`
Default Routers	Specify the IP address(es) of the default routing gateway(s). Example: `10.10.10.1`
IP Address Ranges	This is the list of IP address ranges from which the system will assign addresses to the nodes. Addresses must be separated by spaces. Example: `10.1.45.11-10.1.45.114` **NOTE**: PanActive Manager will not allow an IP address that is in use to be removed. Any changes that force a node to get a new IP address will require a system reboot.
IP Aliases	If IP aliases have been assigned, they will appear here. Otherwise, the field will display **[none]**. Radio buttons beneath the IP aliases allow you to keep or clear them.

> **NOTE:** For information on configuring this information with the command-line interface, see the **netconfig** PanCLI command in the *VDURA PanCLI Reference Guide*.

Add or Remove Subnets

To add another subnet, click the **Add Another Subnet** button, then complete the information for **Netmask**, **Default Routers**, and **IP Address Ranges** for the subnet marked "**Subnet ???**".

Click the **Remove a Subnet** button to remove a subnet. This button only appears when there is more than one subnet.

You can also add a new subnet from the **Network > Network Configuration > Subnets** page by completing the information for **Netmask**, **Default Routers**, and **IP Address Ranges**, and then clicking **Add**. If you add a subnet from this page, a system reboot is not required.

DNS, NIS, HTTP, SMTP, and NTP Server Information

Click the **Next** button to view pending changes and to make additional changes using the following controls.

Control	Description
DNS Domain Names	The text box provides a search list of DNS domain names. Modify as needed. Although a domain name is not required, its absence will prevent use of DNS delegation.
DNS Server IP Addresses	Enter IP addresses for DNS servers that the system uses for host name lookups. Although the DNS server is not required, its absence will prevent the use of hostnames for NTP, SMTP, NIS, and secure web proxy servers, as well as DNS delegation.
Enable DNS Delegation	Delegation gives control over a subset of the DNS namespace to a DNS server running on the VDURA system. The VDURA system DNS server resolves DNS names for system entities and provides intelligent load balancing. For more information, see DNS Delegation.

Control	Description
Zone Name	Specify the root of the DNS namespace being delegated.
NIS Domain Name	Specify the Network Information Service (NIS) domain name.
NIS Server Name	Specify the NIS server name or IP address.
NIS Hostname Resolution	Select **Enable** or **Disable**.
Secure Web Proxy Server	Specify the name or IP address of a secure web proxy server to be used to make outbound HTTPS requests to VDURA servers if you have chosen to enable PanActive Link. For information about PanActive Link, see Configure PanActive Link.
Secure Web Proxy Port	Specify the secure web proxy port number.
SMTP Server	Specify the name or IP address of the mail server.
Administrator Email	Specify a single email address for the administrator.
NTP Servers	Specify a space-separated list of Network Time Protocol (NTP) server names and/or IP addresses. Two or three servers are recommended for redundancy. A list of NTP servers is available at https://tf.nist.gov/tf-cgi/servers.cgi. For a randomly chosen public server, specify "pool.ntp.org". If no NTP servers are specified, local timing takes over. One Director is automatically chosen to operate as a stratum 0 time source for the system.

To implement the changes, click the **Commit Changes** button. In order to commit your changes, the system must either be rebooted or shutdown.

4.1.3 IP Address Ranges for Nodes

The VDURA system uses DHCP to assign IP addresses to nodes within the system. As your system grows or network infrastructure changes, you can use PanActive Manager to add or modify the address ranges assigned to nodes.

> **NOTE:** The VDURA system's DHCP server listens on a non-standard port number and assigns IP addresses only to nodes on the system. It will not conflict with any other DHCP server that might be running on your network.

Click **Network > Network Configuration > IP Ranges**.

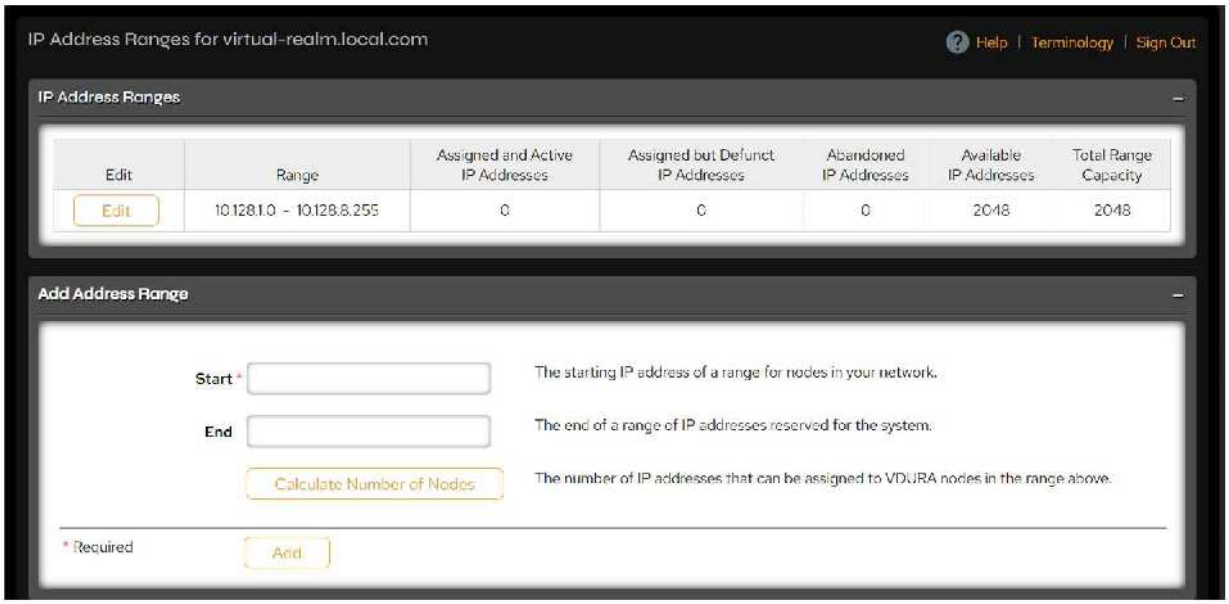


Figure 4-4. IP Address Ranges page.

To add a range of IP addresses, complete the **Add Address Range** fields and click **Add**.

Field	Description
Start	Specify the start of a range of IP addresses for the VDURA system. VDURA recommends that you assign the VDURA system its own IP subnet with at least as many available addresses as the maximum predicted number of nodes in the storage system. By default, the **Start** address is the address that you entered when you configured the first director.
End	Specify the end of a range of IP addresses for the VDURA system. The **End** address indicates the last address in a block of IP addresses that you preassign for the VDURA system. Make sure that you plan for growth in the system. For example, if your system will include 100 nodes, your start IP address would be 10.10.10.2 and the end IP address would be 10.10.10.102. If the end address is not specified, only one IP address is added to the range.
Calculate Number of Nodes	Click the **Calculate Number of Nodes** button to have PanActive Manager count the number of IP addresses in the new address range before you add it to your system configuration.

To edit IP ranges, click the **Edit** button and either edit, split, or delete an existing IP address range. All nodes using addresses within the selected range are displayed, along with MAC addresses, assigned IP addresses, and expiration dates. Click **Submit** to implement the changes.

Be sure not to overlap your existing IP range, as this will produce an error.

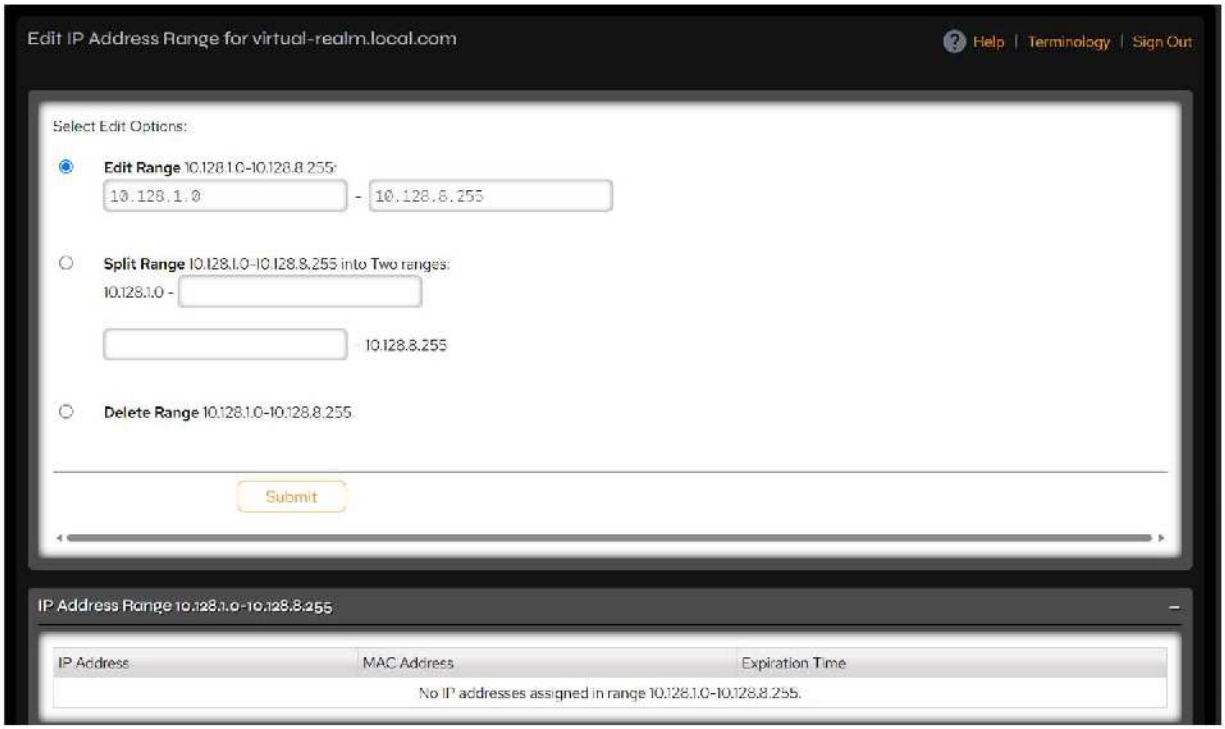


Figure 4-5. Edit IP Address Range page.

> **NOTE:** For information on changing these settings with the command-line interface, see the **netconfig** PanCLI command in the *VDURA PanCLI Reference Guide*.

4.1.4 IP Aliases

Use IP aliasing to assign multiple IP addresses on different subnets to the same network interfaces on the Directors. Remote clients on different subnets can then access the Director. This address range can be configured on the primary or secondary network. See IP Alias Failover.

> **NOTE:** Name resolution for ACLs assigned to files and directories that are accessed from the secondary network will fail if a subset of the Directors that manage volumes are not directly connected to the secondary network.

Please check the *VDURA V5000 Certified Platform Installation and Service Guide*, available at vdura.com/resources/documentation, for more details on how to configure IP aliases.

Add an IP Alias

Click **Network > Network Configuration > IP Aliases**.

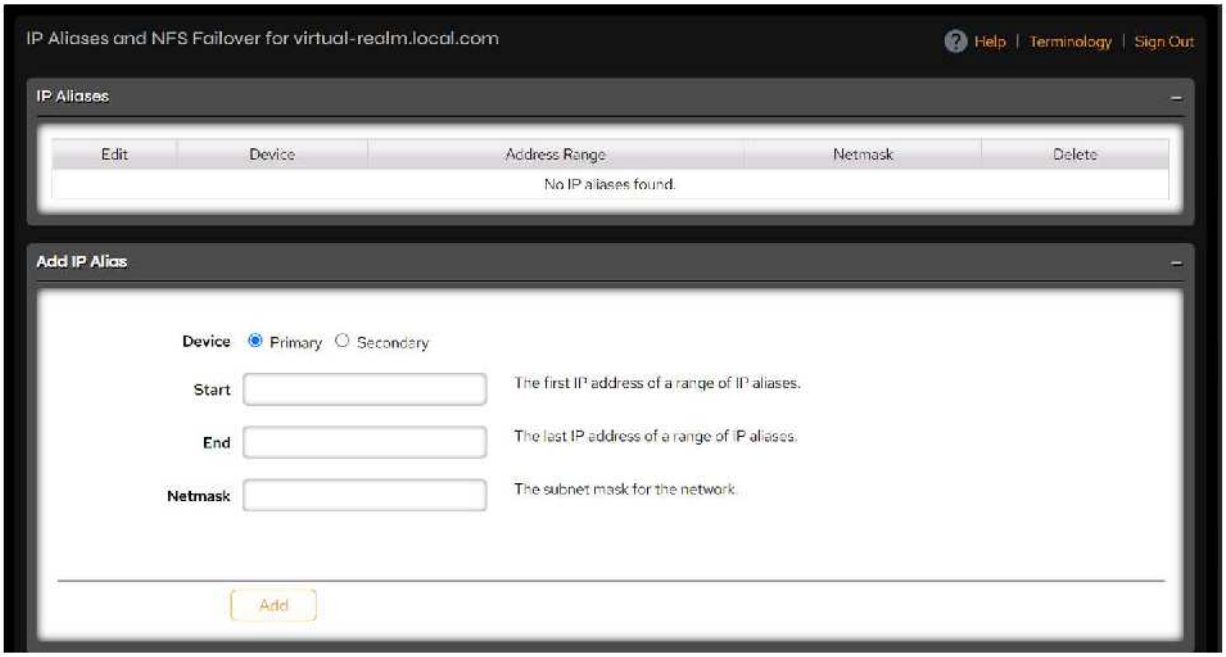


Figure 4-6. IP Aliases and NFS Failover page.

> **NOTE:** The title of this page refers to NFS failover. CIFS and NFS have been deprecated in PanFS as of 12.0.0.

To add primary or secondary IP aliases, select the Device (Primary or Secondary) and use the fields provided to specify the start and end of the IP range. If the **End** field is left blank, the system will assume that there is only one IP address in the range.

Enter the netmask and click **Add** to save the changes. A message will appear at the top of the page confirming that the IP alias was successfully added.

Edit an IP Alias

If you wish to modify existing IP aliases, click the **Edit** button next to the range to modify the configuration. Click **Submit** to save the new configuration.

Delete an IP Alias

To delete IP aliases, click the **Delete** button. A message will appear at the top of the page confirming that the IP alias was successfully deleted.

IP Alias Failover

IP alias failover can prevent clients from hanging indefinitely when the IP address they have mounted becomes unavailable due to the failure of a Director. If IP alias failover is configured, detection of a Director failure will automatically cause the system to reassign the IP aliases previously assigned to that Director to other Directors in the system. After the IP address is reassigned, clients can reestablish their connections to the new Director handling that address.

To use IP alias failover, the administrator must configure an additional range of IP addresses described in IP Aliases. The additional range can also be configured using the **netconfig alias** command in PanCLI.

Although access is still available through addresses in the PanDHCP-managed primary address range(s) typically used for DirectFlow and management access, these PanDHCP-managed IP addresses are not failed over in the event of a Director failure. Clients that have mounted a PanDHCP-managed primary address will hang if the Director to which that address is assigned crashes or fails.

4.1.5 Default Routers

Click **Network > Network Configuration > Default Routers** to change or add IP addresses to the default routing gateways.



Figure 4-7. Edit Default Routers page.

Add or change IP address(es) in the **Default Routers** text box for the appropriate subnet, then click **Submit** to save your changes.

4.1.6 Static Routes

Static routes allow manual manipulation of the routing table. You can use this option to add or delete a route or to connect to remote networks.

Click **Network > Network Configuration > Static Routes**.

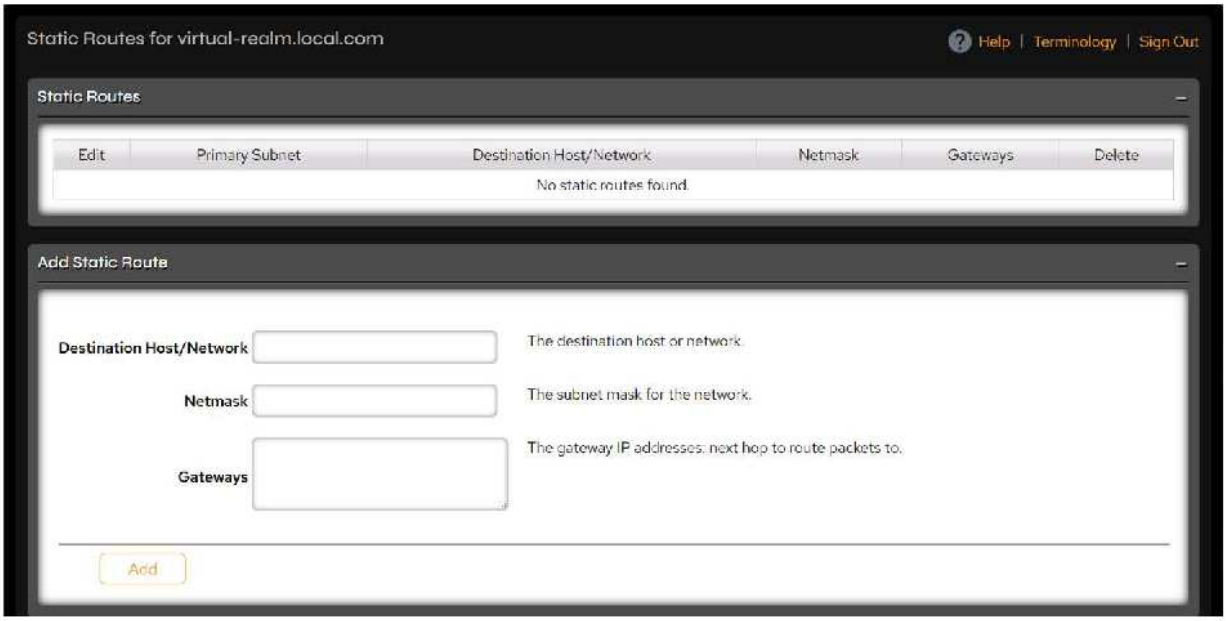


Figure 4-8. Static Routes page.

To add a new static route, complete the **Destination Host/Network**, **Netmask**, and **Gateways** fields and click **Add**. To edit an existing static route, click the **Edit** button next to the static route you want to change, make your changes, and click **Submit**.

> **NOTE:** For information on configuring this information with the command-line interface, see the **netconfig staticroute** PanCLI command in the *VDURA PanCLI Reference Guide*.

4.1.7 Subnets

Click **Network > Network Configuration > Subnets**.

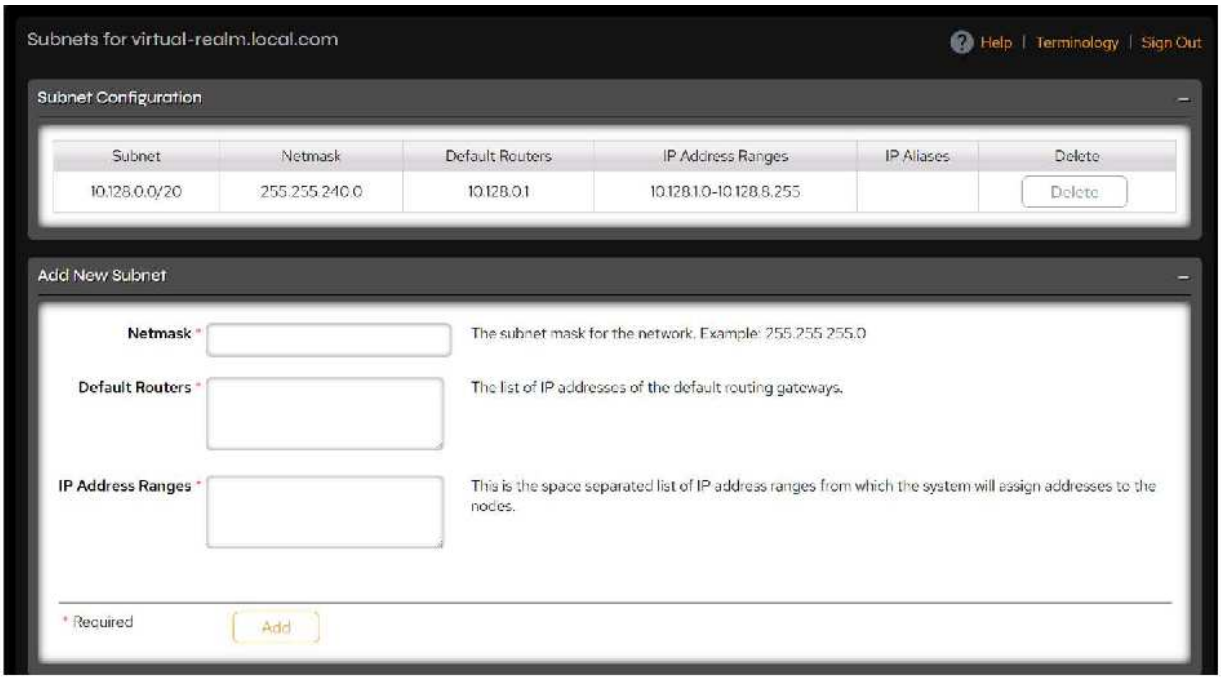


Figure 4-9. Subnets page.

Add a Subnet

To add a subnet, complete the information required in the following fields.

Control	Description
Netmask	Specify the subnet mask for the network. Example: `255.255.255.0`
Default Routers	Specify the IP address(es) of the default routing gateway(s). Example: `10.10.10.1`

Control	Description
IP Address Ranges	Specify a space separated list of IP address ranges from which the system will assign addresses to the nodes. Example: `10.1.45.11-10.1.45.114`

When finished, click **Add**. The realm does not require a reboot when adding a new subnet from this page.

Delete a Subnet

To delete a subnet, click the **Delete** button next to the subnet to be removed in the **Subnet Configuration** table.

4.1.8 NIC Configuration

Use the NIC Configuration page to configure the network interfaces of an enclosure. Select the enclosure, make your changes, and click **Save Configuration** or click **Cancel** to discard them.

Click **Network > Network Configuration > NIC Configuration**.



Figure 4-10. NIC Configuration page.

4.1.9 NTP and Time Zone Settings

The VDURA system can access Network Time Protocol (NTP) servers to synchronize its time setting with other systems. Because system operation requires an NTP server, in the absence of an external NTP server, one Director will be automatically chosen to operate as a stratum 0 time source for the system. To intentionally implement “local timing”, leave the NTP server field blank to allow a Director to automatically be chosen. Do not enter the IP address of a specific Director.

> **NOTE:** When configuring NTP, be aware that NTP server time on a client and the VDURA system must be no more than ten minutes apart, and that closer is better.

> **NOTE:** For information on changing this setting with the command-line interface, see the **ntp** PanCLI command in the *VDURA PanCLI Reference Guide*.

Click **Network > NTP & Time Zone**.

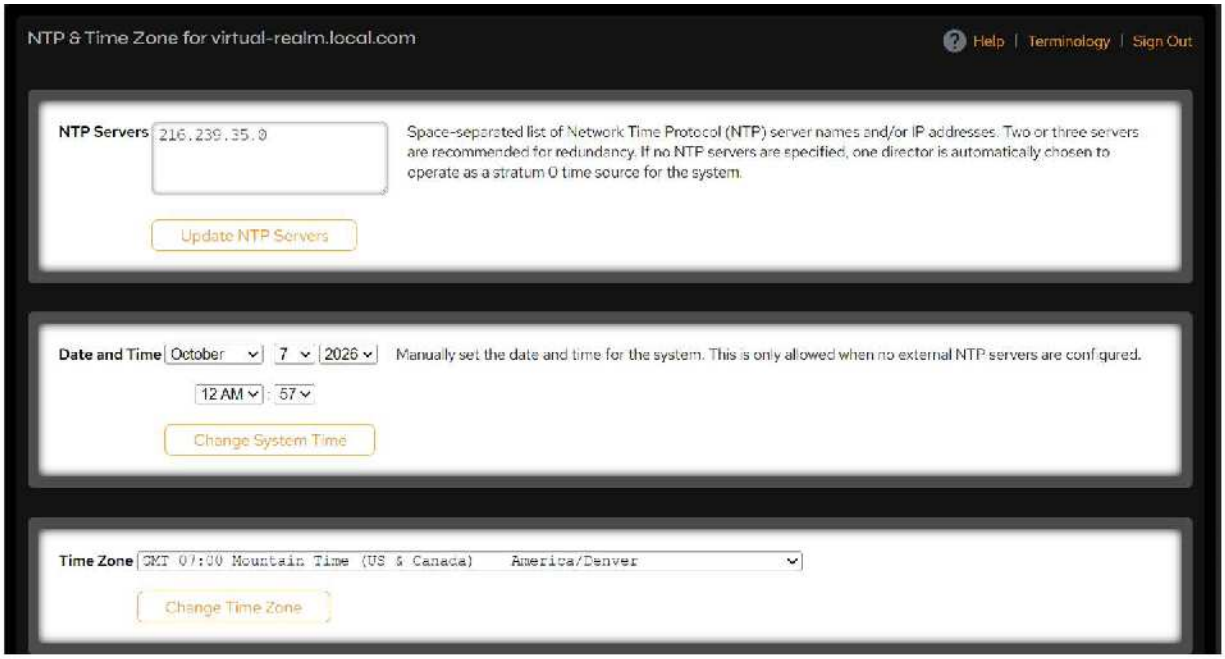


Figure 4-11. NTP & Time Zone page.

Use the following controls to change time settings.

Control	Description
NTP Servers	Specify a space-separated list of Network Time Protocol (NTP) server names and/or IP addresses. Two or three servers are recommended for redundancy. A list of NTP servers is available at https://tf.nist.gov/tf-cgi/servers.cgi. For a randomly chosen public server, specify **pool.ntp.org**. If no NTP servers are specified, local timing takes over. One Director is automatically chosen to operate as a stratum 0 time source for the system. To implement local timing, leave the NTP server field blank and a director will automatically be chosen. Do not enter the IP address of a specific Director. Click the **Update NTP Servers** button to save the server list.
Date and Time	Manually set the date and time using the drop-down menus provided. Use the manual date and time setting only if there is no external NTP server available. Click the **Change System Time** button to save your changes.
Time Zone	Specify the time zone for your network clock. Click the **Change Time Zone** button to save your changes.

4.1.10 DNS Settings

The VDURA system uses the standard Domain Name Server (DNS) protocol to map names to IP addresses and vice versa. Although you are not required to configure DNS server information for your VDURA system, not doing so prevents the use of hostnames for NTP, SMTP, NIS, and HTTP proxy servers, as well as preventing the use of DNS delegation.

A maximum of three DNS server IP addresses is supported.

> **NOTE:** For information on changing these settings with the command-line interface, see the **dns** PanCLI command in the *VDURA PanCLI Reference Guide.*

Valid System and DNS Names

Valid DNS names are currently defined as:

- containing a maximum of 255 characters.
- composed of one or more nonempty DNS components separated by "." characters. A DNS component is defined as:
  - having a maximum of 63 characters.
  - using ASCII characters (rule imposed by RM).
  - not using the following characters: `/ @ , . " < > \t \n \r`

Change DNS Settings

Click **Network > DNS**.

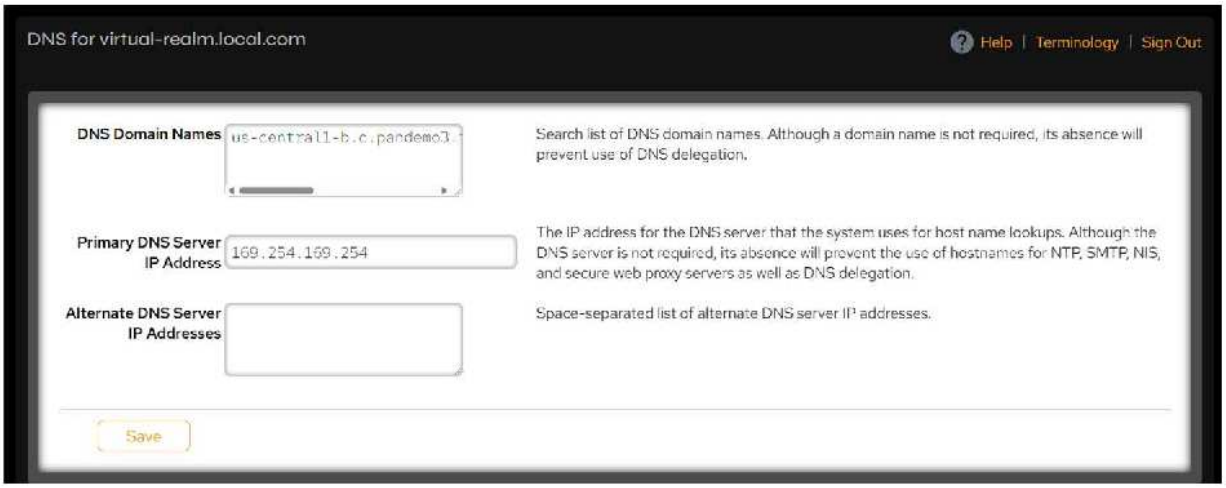


Figure 4-12. DNS page.

Use the following controls to modify DNS settings, then click **Save** when finished.

Control	Description
DNS Domain Names	Specify a space-separated list of domain names to search. Although a domain name is not required, its absence will prevent use of DNS delegation. Example: `mycompany.com example.com`
Primary DNS Server IP Address	Specify the IP address for the DNS server that the system uses for host name lookups. Although the DNS server is not required, its absence will prevent the use of hostnames for NTP, SMTP, NIS, and secure web proxy servers, as well as DNS delegation. Example: `10.10.10.1`
Alternate DNS Server IP Addresses	Specify a space-separated list of alternate IP addresses for DNS servers that the system can use for host name lookups. The order you specify sets the order in which servers are consulted for lookups. You can specify up to two IP addresses. Example: `10.10.10.2 10.10.10.3`

4.1.11 DNS Delegation

DNS delegation allows the VDURA system to serve as a DNS server, resolving DNS names for devices and services in your system. Enabling DNS delegation provides several benefits:

- When IP aliases and dynamic DNS are configured, DDNS distributes mounts over all IP aliases. Those addresses transfer transparently from a failed Director to another Director. When IP aliases are not configured, DDNS distributes mounts over all of the service addresses. Please check the *VDURA V5000 Certified Platform Installation and Service Guide*, available at vdura.com/resources/documentation, for more details on how to configure IP aliases and Configure the DNS Server with DNS Delegation Enabled, for information on configuring DNS to allow DDNS updates.

- Changes to your system (for example, adding more Directors) require no changes to your site's DNS servers or clients that access the system, unless you change the address of one or more repset members.

Some sites have policies restricting DNS delegation. If DNS delegation is not available at your site, you can disable the DNS delegation feature. However, doing so then requires you to manually reconfigure your DNS server when you add, remove, or reconfigure directors in your system.

The DNS delegation option you select determines the proper configuration of your site's DNS server(s) and the clients that access the system. Remember that changing the DNS delegation settings on the VDURA system also requires changes to your site's DNS server configuration and potentially the clients that access the system. Improper DNS server or client configuration can significantly decrease the performance of your VDURA system.

> **NOTE:** For information on changing DNS delegation settings with the command-line interface, see the **dns** PanCLI command in the *VDURA PanCLI Reference Guide*.

Enable DNS Delegation

Enabling DNS delegation on the VDURA system makes the system responsible for a sub-zone of your domain. Typically, the sub-zone should be your VDURA system name followed by the local domain name (for example, **system.mycompany.com**). Properly configuring the environment involves the following steps:

1. Enable DNS delegation on the VDURA system and specify the name of the sub-zone for which it is responsible.
2. Configure the authoritative DNS server for the zone containing the system to delegate responsibility of the sub-zone to the system.
3. Configure clients of the VDURA system to access the system appropriately.

Enable DNS Delegation on the VDURA System

Prior to this step, follow the instructions in DNS Settings to properly configure your DNS settings. Be sure that the first entry in **DNS Domain Names** is the name of the zone in which the system is installed (for example, **mycompany.com**) and that the primary DNS server IP address is the IP address of the authoritative DNS server for the zone

Click **Network > DNS > DNS Delegation**.

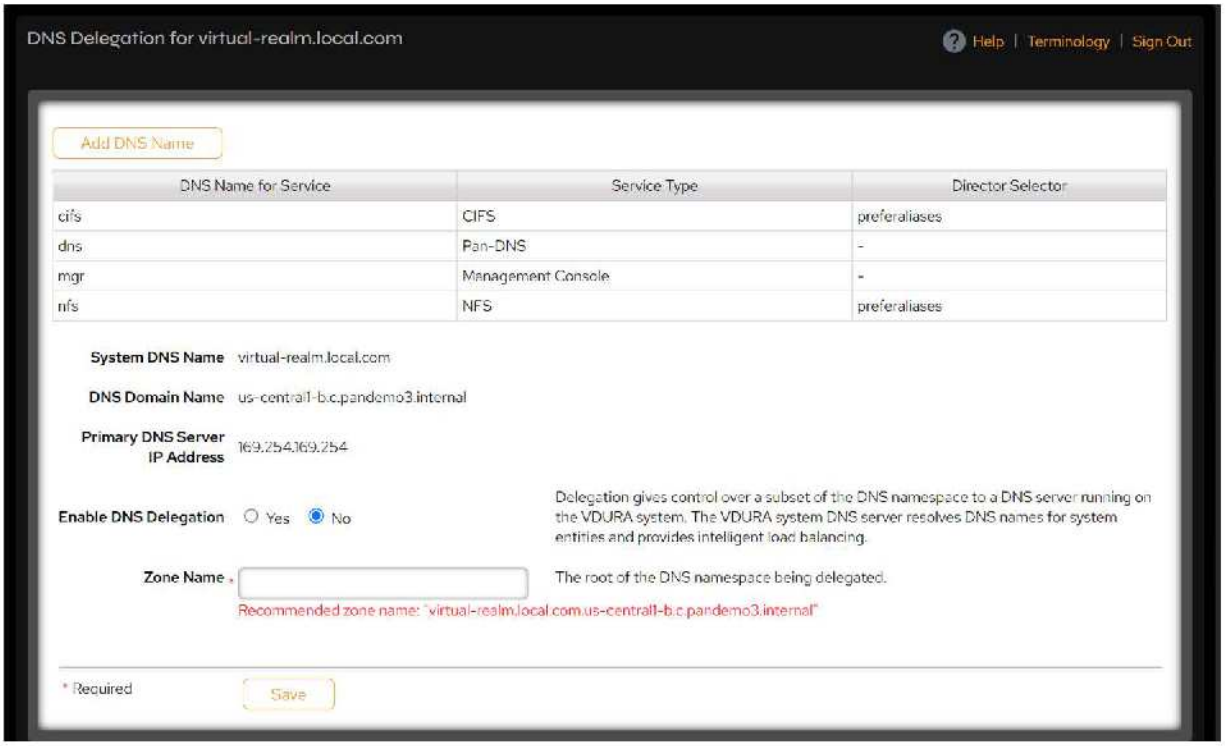


Figure 4-13. DNS Delegation page.

> **NOTE:** The DNS service table on this page lists default cifs and nfs entries. You do not need to use or modify these entries. CIFS and NFS have been deprecated in PanFS as of 12.0.0.

In the **Zone Name** field, enter the name of the DNS zone you want to delegate to the system. Typically, this should be your VDURA system name followed by the local domain name (for example, **system.mycompany.com**). The system will typically suggest a zone name. When finished, click **Save**.

Next, add a DNS service. Click the **Add DNS Name** button at the top of the page to open the **Add DNS Service** page.

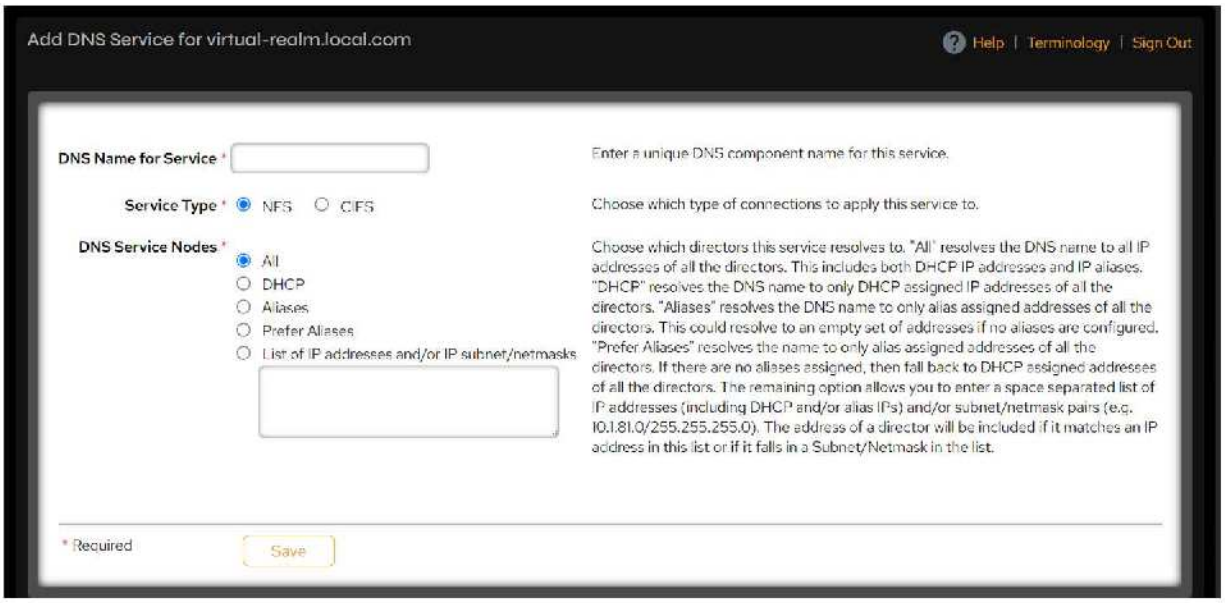


Figure 4-14. Add DNS Service page.

> **NOTE:** The Service Type options on this page include NFS and CIFS. CIFS and NFS have been deprecated in PanFS as of 12.0.0.

Use the following controls to add a DNS service, then click **Save** when finished.

Control	Description
DNS Name for Service	Enter a unique DNS component name for this service.
Service Type	Choose which type of connections to apply this service to: NFS or CIFS. CIFS and NFS have been deprecated in PanFS as of 12.0.0.

Control	Description
DNS Service Nodes	Choose which directors this service resolves to: • **All**: Resolves the DNS name to all IP addresses of all Directors. This includes both DHCP IP addresses and IP aliases. • **DHCP**: Resolves the DNS name only to DHCP-assigned IP addresses of all Directors. • **Aliases**: Resolves the DNS name only to alias-assigned addresses of all Directors. This could resolve to an empty set of addresses if no aliases are configured. • **Prefer Aliases**: Resolves the name only to alias-assigned addresses of all Directors. If no aliases assigned, then fall back to DHCP-assigned addresses of all Directors. • The remaining option allows you to enter a space separated list of IP addresses (including DHCP and/or alias IPs) and/or subnet/netmask pairs (e.g., 10.1.81.0/255.255.255.0). The address of a Director will be included if it matches an IP address in this list or if it falls in a subnet/netmask in the list.

Configure the DNS Server with DNS Delegation Enabled

The next step is to configure the DNS server to delegate responsibility for the specified sub-zone. The examples in this section assume that you are using the **BIND** DNS server.

- In the DNS server's configuration file, add an **NS** record to delegate the zone, and an **A** record to specify the address of the VDURA system.
- You must precede the system name with **ns.** when configuring the **NS** and **A** records for the system. The **ns.** service name identifies the single Director in the system running the DNS service. The system name itself can resolve to multiple Directors if you have enabled realm management, providing transparent load balancing of system services.
- VDURA suggests a value of 86,400 seconds (1 day) as the **Time to Live** (TTL) value, but you can adjust the value in accordance with your local policies.

Example DNS configuration lines may look like the following:

```
system    86400 IN NS ns.system.mycompany.com
ns.system 86400 IN A  10.1.1.11
```

With realm management configured, system DNS resolution services can transfer transparently from a failed Director to another Director in the system. (See Realm Management for more information.) However, this feature also requires you to configure your DNS server to allow Dynamic DNS Updates (DDNS) from the VDURA system. You can do so by either allowing updates from the system's entire subnet or by allowing updates explicitly from all Directors in the system. For example, to allow updates from all Directors in the 10.1.1.0 subnet:

```
zone "mycompany.com" in {
    type master;
    file "db.mycompany";
    allow-update {
        10.1.1/24;
    };
};
```

If you decide not to enable DDNS, and the Director hosting the DNS delegation service fails, you must reconfigure your DNS server to delegate your system's sub-zone to the IP address of the Director now hosting the DNS delegation service. Until you reconfigure your DNS server, existing client mounts may continue to work, but new clients will not be able to mount system volumes.

When finished, reload your DNS server with the new configuration information. Use **kill -HUP**, **rndc reload**, or whatever is appropriate for your DNS server.

Configure Clients to Access the System with DNS Delegation Enabled

DirectFlow clients should use the system name in their mount point specifications, not the system's IP address. For example, the mount entry in **/etc/fstab** should be of the form:

```
panfs://system.mycompany.com:global /panfs panfs ...
```

instead of the form:

```
panfs://10.1.1.11:global /panfs panfs ...
```

Delete a DNS Name

On the **Network > DNS > DNS Delegation** page, click the **Delete** button in the row of the DNS name you want to delete.

You can only delete a DNS name that was previously manually added. The default named DNS services are not deletable and will never display a delete button.

Disable DNS Delegation

Disabling DNS delegation on the VDURA system means either that your site's DNS server must take responsibility for resolving names when accessing the system, or that you must hard-code system IP addresses in your client configurations. Properly configuring the environment requires the following steps:

1. Disable DNS delegation on the VDURA system.
2. Configure the authoritative DNS server for the zone containing the system to resolve system names for client access.
3. Configure clients of your VDURA system to access the system appropriately.

Detailed steps are provided below.

Disable DNS Delegation on the VDURA System

On the **Network > DNS > DNS Delegation** page, select the **No** radio button next to **Enable DNS Delegation**, then click **Save**.

Configure the DNS Server with DNS Delegation Disabled

Configure your DNS server to resolve system names for client access. The examples in this section assume that you are using the **BIND** DNS server.

In the DNS server's configuration file, add an **A** record to specify the address of the VDURA system. If you have enabled realm management, you can add an **A** record for each Director in the system. With this configuration, you can still access system services if a system Director fails. VDURA suggests a value of 86,400 seconds (1 day) as the **Time to Live** (TTL) value, but you can adjust the value in accordance with your local policies. For example:

```
system    86400 IN A  10.1.1.11
system    86400 IN A  10.1.1.12
system    86400 IN A  10.1.1.13
```

If you configure your DNS server with multiple **A** records for realm Directors, you must manually update the DNS server configuration when you add or remove Directors from the realm or if the IP address of a realm Director changes.

When finished, reload your DNS server with the new configuration information. Use **kill -HUP**, **rndc reload**, or whatever is appropriate for your DNS server.

Configure Clients to Access the System with DNS Delegation Disabled

How you access system volumes depends on the type of client and how you configured your site's DNS server.

DirectFlow clients should do one of the following:

- Mount the system using its DNS name. For example, the mount entry in **/etc/fstab** should be of the form:

  ```
  panfs://system.mycompany.com:global /panfs panfs ...
  ```

- If you are not using DNS, you must use the system's IP address. For example:

  ```
  panfs://10.1.1.11:global /panfs panfs ...
  ```

- If you are not using DNS and the VDURA system has realm management enabled, you can use the **rmlist** mount option to specify the IP addresses of all Directors in the realm. For example:

  ```
  panfs://10.1.1.11:global /panfs panfs rw,noauto,
  rmlist=(10.1.1.11;10.1.1.12;10.1.1.13)  0  0
  ```

4.1.12 Lock Management

Shared files can be read by multiple users, but are locked for writes to only one user at a time. Normally, when a client machine shuts down gracefully, all file locks are released. However, if the client panics or shuts down ungracefully, the locks may not release properly. To break a file lock using PanActive Manager, thus making it available for write by another user, follow the process below.

Click **Advanced Tools > Lock Management**. Any locked files will appear in the **Lock Management** table.



Figure 4-15. Lock Management page.

This page provides information about the system's locked files.

Column	Description
Volume/path	Volume and path (where known) to the locked file.
Lock Type	Either **byte_lock read** or **byte_lock write**.
Handle	Name of the locked file.
Offset	Part of the file that is locked.
Length	Number of bytes between the offset and the end of the file.

Put a check in the boxes next to the files you want to unlock and click the **Break Locks** button.

4.1.13 Configure Notification Services

The VDURA system creates alerts with two levels of priority and provides several ways to view them. **Critical** alerts generate if system availability or data integrity is at imminent risk, while **Warning** alerts generate when there is a developing problem that might lead to impairment of system availability or data integrity. You can configure PanActive Manager to send email and/or pager notifications to users for these alerts, as well as notify VDURA Technical Support. Periodic status emails can also be configured that provide text data for parsing. Aside from email, alert data can be sent to a remote syslogging facility (see Set Remote Syslog Collectors and Facility). The system keeps an event history that can be reviewed and exported by the system administrator (See View Events).

> **IMPORTANT:** All email notifications require access to an SMTP mail server.

> **NOTE:** For information on changing these settings with the command-line interface, see the **notify** PanCLI command in the *VDURA PanCLI Reference Guide*.

Configure Email (SMTP Server) Settings

Click **Configuration > Notification > Email Settings**.

The Recent Email Events table at the top of the page lists recent email notification events, with columns for event category, code, time, user name, component, type, node serial, enclosure, hardware serial, IP address, StorageSet, volume service, and volume.

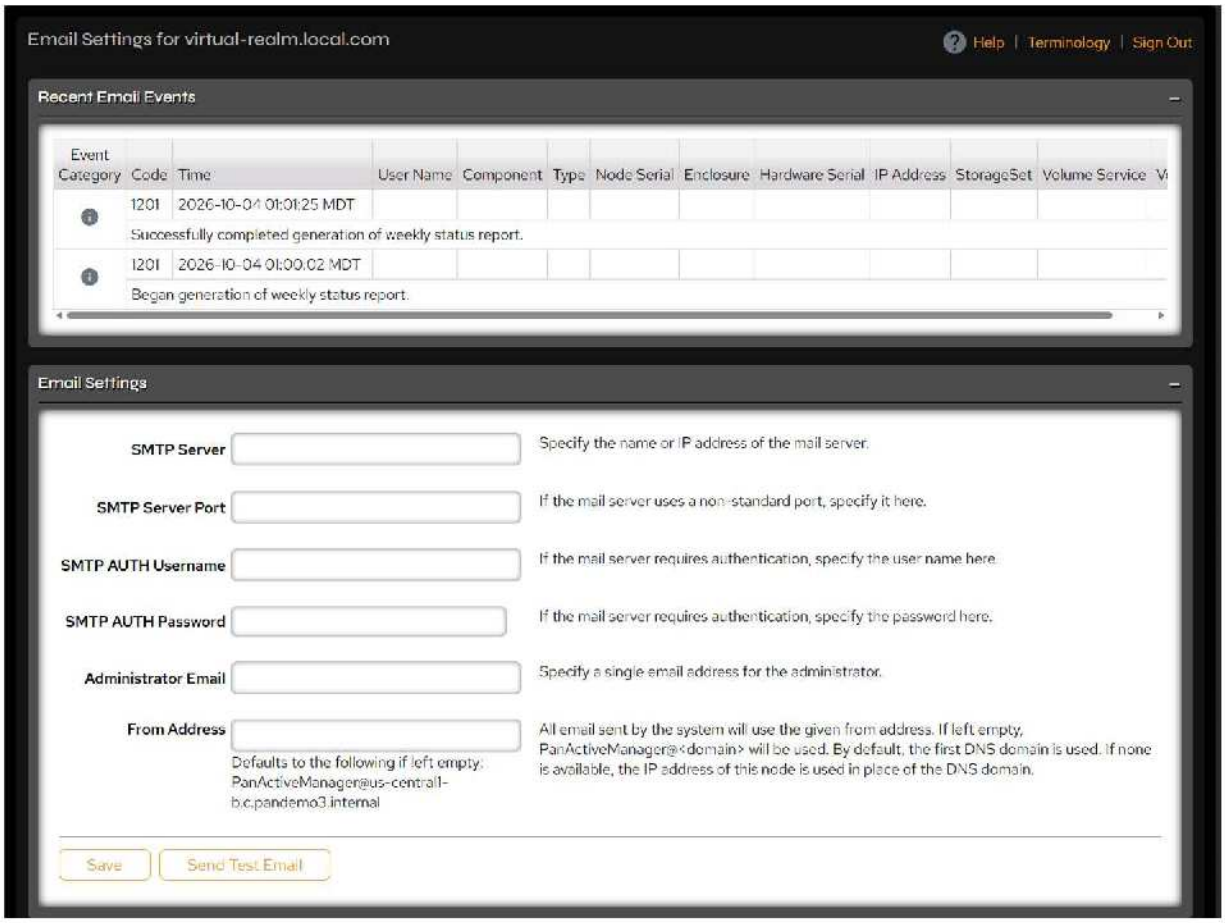


Figure 4-16. Email Settings page.

Use the following controls to configure the SMTP server and administrator email address:

Control	Description
SMTP Server	Specify the name or IP address of the mail server. The system uses the Simple Mail Transport Protocol (SMTP) to send mail about problems it encounters. For example: `mail.yournetwork.com`. **Note:** The VDURA system does not use an internal SMTP server. Depending on your SMTP server, you may need to update your SMTP server to accept email requests coming from repset IP addresses.
SMTP Server Port	If the mail server uses a non-standard port, specify it here.
SMTP AUTH Username	If the mail server requires authentication, specify the user name here.
SMTP AUTH Password	If the mail server requires authentication, specify the password here.
Administrator Email	Specify a single system administrator's email address, including the full domain name. For example: `joe@yournetwork.com`. If left blank, your **Administrator ID** forms the first part of the email address. For example: `admin@yournetwork.com`.
From Address	All email sent by the system will use the given "From" address. If left empty, **PanActiveManager@<domain>** will be used. By default, the first DNS domain is used. If none is available, the IP address of this node is used in place of the DNS domain.

When finished, click **Save**, then click **Send Test Email** to confirm the settings.

> **NOTE:** Whenever you add or remove a node from the replication set, make sure the SMTP server configuration is updated to allow email requests from the new repset IP address list.

Set Alert Recipients

You can configure PanActive Manager to email event notices to network administrators.

Click **Configuration > Notification > PanActive Alerts**.

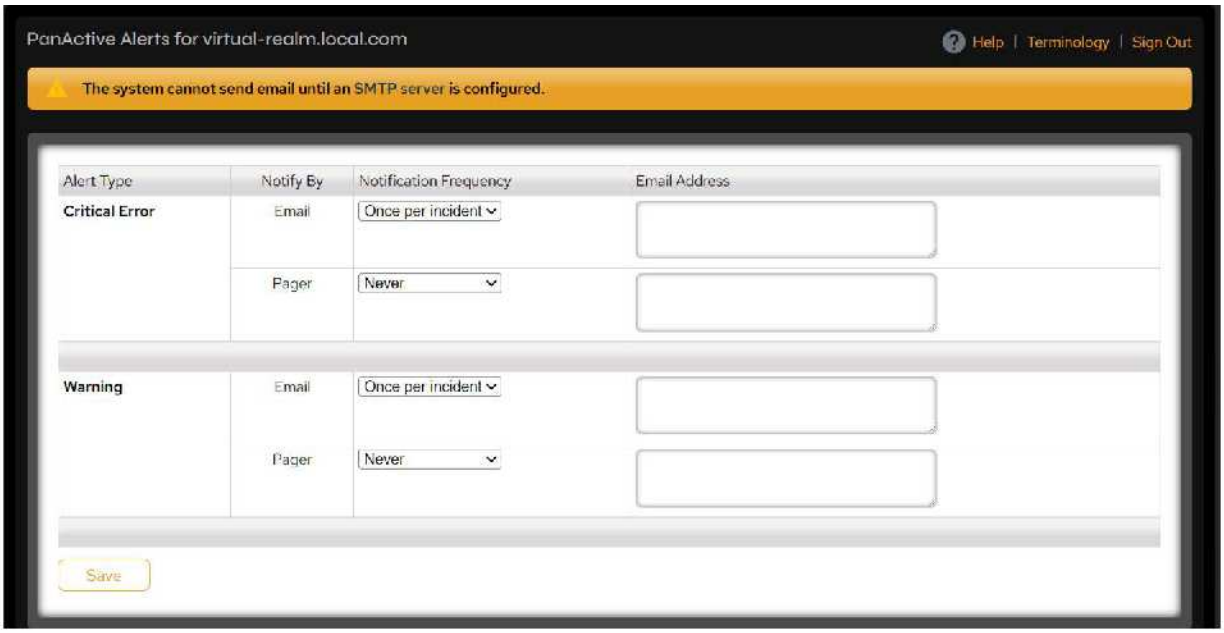


Figure 4-17. PanActive Alerts page.

Use the controls to add alerts and alert recipients to the notification list.

Control	Description
Alert Type	• **Critical Error**: Errors are critical if system availability or data integrity is at imminent risk. • **Warning**: Warnings indicate a developing problem that might lead to impairment of system availability or data integrity.
Notify By	Select **Email** or **Pager** (an email address configured for paging).

Control	Description
Notification Frequency	Select Once per incident, Once a day, Once an hour, Every 5 minutes, or Never.
Email Address	A comma-separated list of recipients. For example: `admin@vdura.com`, `syslist@vdura.com`.

When finished, click **Save**.

Set Status Recipients

You can configure PanActive Manager to email periodic status reports with information about StorageSets, volumes, nodes, and gateway performance to network administrators.

Click **Configuration > Notification > Periodic Status Email**.



Figure 4-18. Periodic Status Email page.

Use the following controls to configure the frequency of status emails and add recipients to the notification list.

Control	Description
Frequency	Select **Once a week**, **Once a day**, **Once an hour**, or **Never**.

Control	Description
Send To	A comma-separated list of recipients email addresses. For example: `admin@yourcompany.com`, `syslist@yourcompany.com`.
Send Status Email	Click **Send Status Email** to send a status email immediately to all listed recipients.

When finished, click **Save**.

Warn About AWOL Clients

PanActive Manager automatically detects the failure of clients to respond and reports those failures in the **Historical Event Log**. However, you can also customize the notifications to additionally raise a system-wide alert that triggers email and pager messages when clients do not respond.

Click **Configuration > Notification > Customize Warnings** and select the **Yes** radio button next to **Warn about AWOL Clients**.

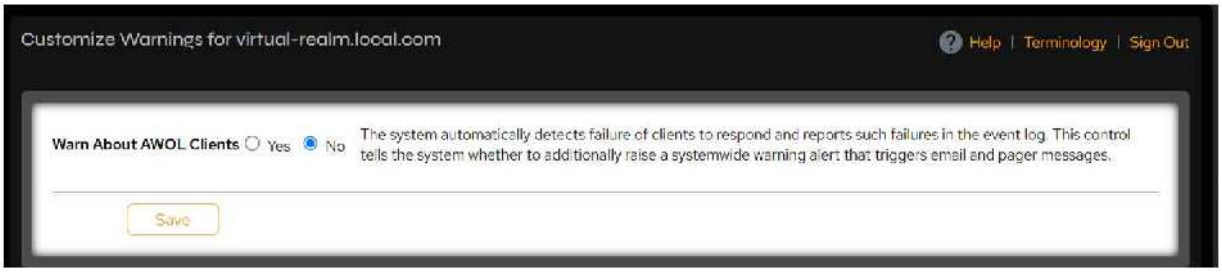


Figure 4-19. Customize Warnings page.

When finished, click **Save**.

Set Remote Syslog Collectors and Facility

Alerts in the **Historical Event Log** can be sent to remote syslog collectors.

Click **Configuration > Notification > Remote Syslog**.

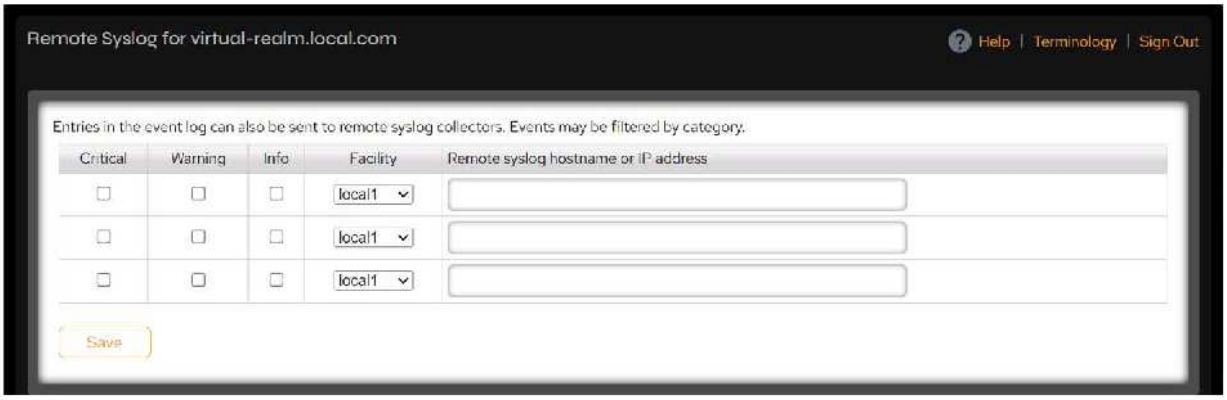


Figure 4-20. Remote Syslog page.

Use the following controls to configure which entries in the Historical Event Log to send to the remote syslog collectors:

Control	Description
Alert Criteria	Place a check in the boxes beneath the type of alert data you want to be sent: • **Critical**: Errors are critical if system availability or data integrity is at imminent risk. • **Warning**: Warnings indicate a developing problem that might lead to impairment of system availability or data integrity. • **Info**: Informational messages about the system. If no boxes are checked, the remote syslog collector will be unfiltered, and all events in the **Historical Event Log** will be sent to that host.
Facility	Select the facility from the list (auth, authpriv, console, cron, daemon, ftp, kern, lpr, mail, news, ntp, security, syslog, user, uucp, or local0 through local7). The default facility is local1.

Control	Description
Remote Syslog Hostname or IP Address	Specify the hostname or IP address of the remote syslog collector. If the host is already listed, the facility and event categories for this host will be updated. If the host is not already listed, it is added, along with the specified facility and event categories.

When finished, click **Save**.

If you are using either **rsyslogd(8)** on a Linux system or **syslogd(8)** on a Linux or other Unix-like system, you should include the following configuration in your **/etc/rsyslog.conf** (or **/etc/syslog.conf,** as appropriate) file:

```
local1.*        /var/log/panasas.log
```

If you do not specify a facility, the system will use the default **local1**. You may choose a different file name for logging for your environment, if needed. For backwards compatibility with **syslog(8)**, ensure that the specified log file exists and is writable by the user under which the **rsyslogd/syslogd** daemon runs.

If you are using some other logging solution that understands how to consume and process syslog data, you must specify some valid disposition for syslog data sent for the **local1** facility.

Configure PanActive Link

When enabled, PanActive Link allows the system to contact VDURA Technical Support engineers to assist in monitoring and maintaining your VDURA system. Enabling PanActive Link also facilitates software updates, provides up-to-date system information and error reports, and ensures optimal system performance, up-time, and reliability.

> **NOTE:** If the system is behind a firewall, you may need to specify a web proxy host and port.

PanActive Link can be configured for email with or without HTTPS functionality. Depending upon how it is configured, PanActive Link will contact VDURA for the following reasons:

- By email:
  - Send alert (error/warning) emails in batches every five minutes.
  - Send stack trace email when a crash is detected.
  - Send weekly status emails.
- By HTTPS:
  - Fetch the latest available upgrade information (currently not available).
  - Download software images (currently not available).

> **NOTE:** These HTTPS features are currently not available: the system cannot download from VDURA. Download software from **support.VDURA.com** instead.

Click **Configuration > Notification > PanActive Link**.

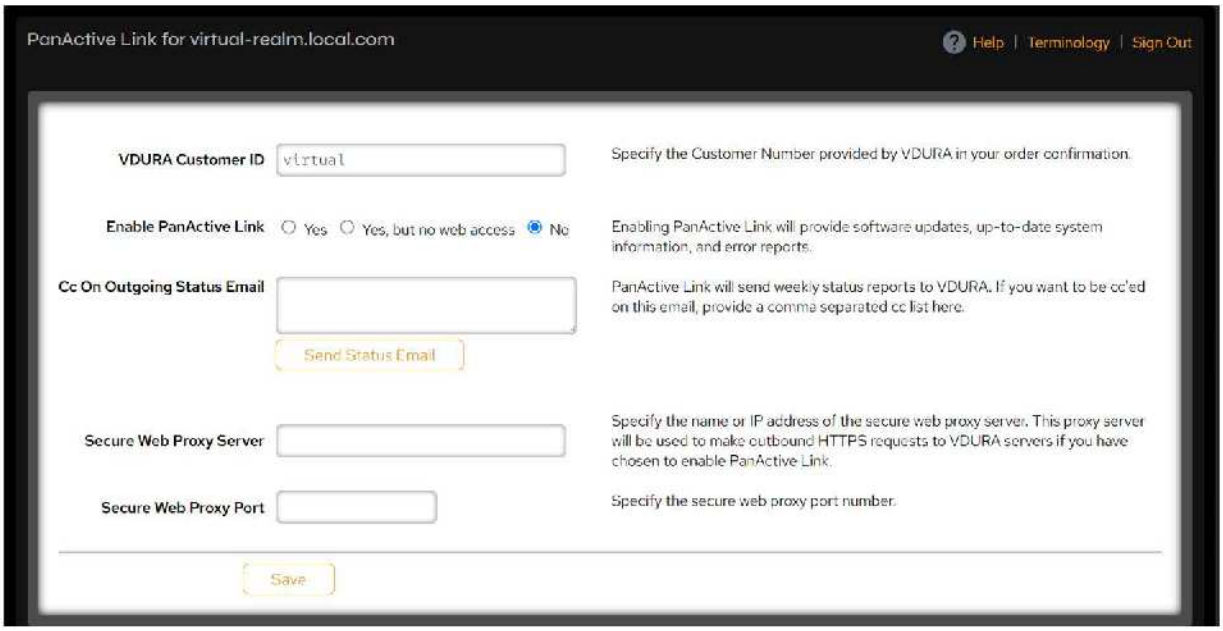


Figure 4-21. PanActive Link page.

Use the following controls to configure PanActive Link.

Control	Description
VDURA Customer ID	Specify the Customer ID number provided by VDURA in your order confirmation. You can find your customer ID by clicking the **About** link in the upper right corner of each PanActive Manager page.
Enable PanActive Link	• **Yes**: Enable PanActive Link support (both email and web). • **Yes, but no web access**: Enable PanActive Link support (email only). • **No:** Do not enable PanActive Link.
Cc On Outgoing Status Email	PanActive Link will send weekly status reports to VDURA. To be cc'ed on this email, provide a comma separated cc list here.
Secure Web Proxy Server	Specify the name or IP address of the secure web proxy server. This proxy server is used to make outbound HTTPS requests to VDURA servers if you have chosen to enable PanActive Link with web access.
Secure Web Proxy Port	Specify the secure web proxy port number.

When finished, click **Save**.

4.1.14 Configure SNMP

Simple Network Management Protocol (SNMP) is a UDP-based network protocol management system. VDURA currently supports SNMP v1. In the VDURA system, this protocol provides a means to monitor and control network devices and to manage configurations, statistics collection, performance, and security. Administrators can use PanActive Manager to change SNMP settings or define new traps, which are notifications that report events to the SNMP manager with a specified message.

The target for SNMP queries should be the **mconsole** node. The **mconsole** node can be identified by issuing the PanCLI command **syscontrol realmmgr** and looking for the **mconsole** service running on one of the replication set members.

To enable SNMP on the system, click **Configuration > SNMP**.

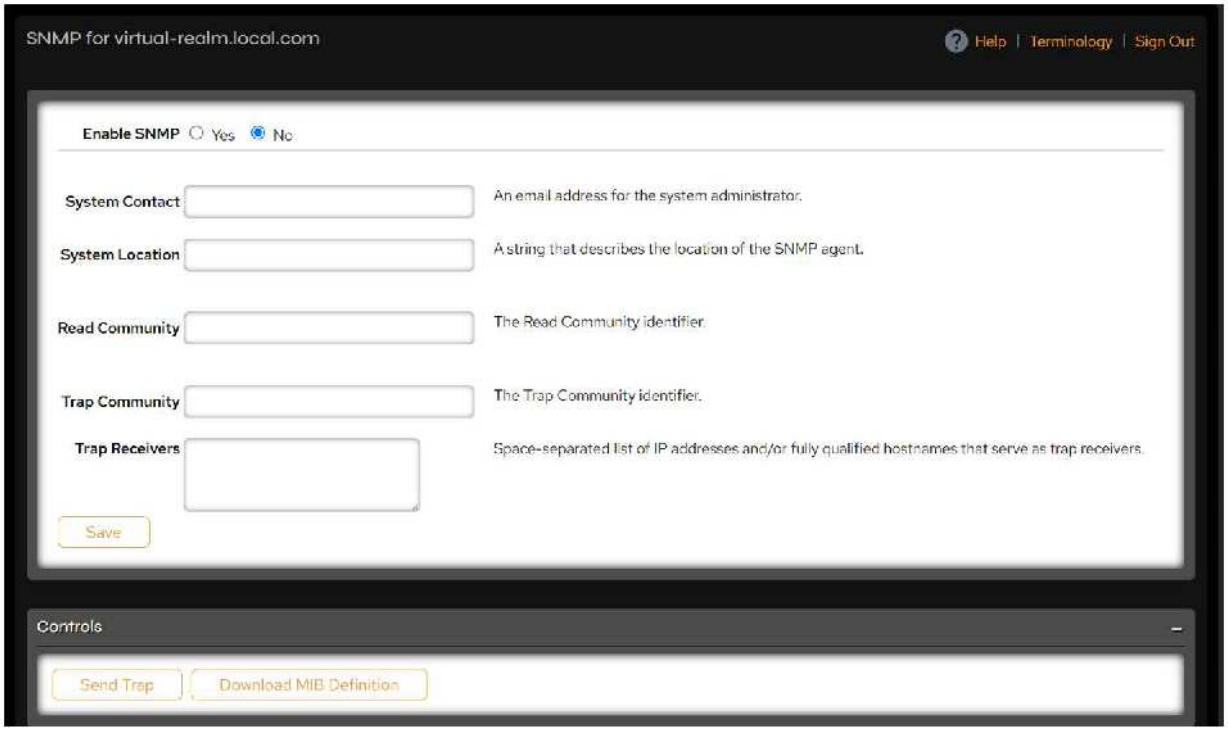


Figure 4-22. SNMP page.

Click the **Yes** radio button for **Enable SNMP** at the top of the page, then use the following controls to enter required SNMP configuration values:

Control	Description
System Contact	Enter an email address for the system administrator.
System Location	Enter a string that describes the location of the SNMP agent.
Read Community	Enter the Read Community (RC) identifier.
Trap Community	Enter the Trap Community identifier (SNMP TRC).
Trap Receivers	Enter a space-separated list of IP addresses and/or fully qualified host names that serve as trap receivers; these are intelligent network devices used to validate SNMP requests.

When finished, click **Save**.

Send Traps

Once the SNMP values have been saved, you can send SNMP traps.

From the **Controls** field at the bottom of the page, click the **Send Trap** button. A message confirming that your trap was sent will appear at the top of the page.

SNMP MIBs

SNMP uses an extensible design where the available information is defined by management information bases (MIBs). MIBs describe the structure of the management data of a device subsystem. There are currently ten VDURA MIB definitions along with the **PANASAS-MIBS-README** file, which defines the order in which the MIBs are to be loaded.

> **IMPORTANT:** Read the PANASAS-MIBS-README file before loading MIBs into your SNMP system!

To download the MIB definitions, click **Configuration > SNMP**, then click **Download MIB Definition** in the **Controls** section. A dialog box opens asking whether you want to save or open the **PANASAS-MIBS.tgz** file.

4.2 Configure VDURA System Access

This section contains information about the following topics for accessing the VDURA system:

- DirectFlow Client Support
- Configure Netgroups
- S3 Connectivity

4.2.1 DirectFlow Client Support

DirectFlow for Linux enables direct access to the VDURA system from supported hosts, improving performance and availability by providing direct and highly parallel access to the Director and Storage Nodes in the VDURA system. DirectFlow for Linux provides unmatched performance for Linux clusters of all sizes and is ideal for the high-performance Linux clusters deployed in technical computing environments.

DirectFlow software supports fully parallel reads and writes from DirectFlow clients to Storage Nodes. It is a loadable kernel module that plugs into the standard kernel infrastructure, allowing the user to mount a VDURA file system. Therefore, you interact with a mounted VDURA file system just like any other kernel file system.

To use the DirectFlow Client, ensure that the following ports are open and accessible:

Client	Direction	Server
*	<–>	3095
*	<–>	10622
*	<–>	3260
3095	<–>	*
ICMP	<–>	ICMP

> **NOTE:** The **seclabel** mount option is needed for DirectFlow to support SELinux in enforcing mode and is automatically added by SELinux. For example: `$ mount | grep panfs`
> `panfs://192.168.91.1:global on /panfs type panfs(rw,relatime,seclabel,callback-network-allow=192.168.100.0/24,rmlist (192.168.91.1;192.168.91.2;192.168.91.3))`

DirectFlow Client Ports

DirectFlow Client ports are built to match Linux distribution, kernel, and processor types. For distributions and kernels that do not support weak-module loading, the DirectFlow port kernel, processor type, and distribution must match exactly. Distributions and kernels that support weak-module loading (such as RHEL 7 and SLES 12) may be able to use a single port on a variety of kernel versions. Upon installation of the software, the package automatically checks that the software is compatible with the kernel.

Supported Distributions

DirectFlow for Linux supports the following distributions:

- RHEL 8.10 and higher, including the following derivatives:
  - Rocky Linux
  - Alma Linux
- SLES 15 SP7 and higher

- Ubuntu 22.04 and higher
- Debian 13 and higher on x86_64 architecture

OS support is for both x86_64- and ARM-based systems.

All distributions except Ubuntu and Debian support RPM for package management and have known software development tool versions. Ubuntu and Debian support DEB for package management. In addition, weak module loading is supported on RHEL 8.10 and above and SLES 15 and above.

In addition, please be aware of the following compatibilities for users wishing to implement VDURA's E2EE feature:

- RHEL on x86 architecture: Supported
- Ubuntu on x86 architecture: Supported
- Any distributions on ARM architecture: Not Supported

Download the DirectFlow Client

This feature has been deprecated in PanActive Manager as of PanFS 12.0.0 but has not yet been removed from the UI. Contact VDURA Technical Support to obtain upgrade images.

You can use PanActive Manager to download DirectFlow Client software from **support.VDURA.com** to your local computer. PanActive Link must be enabled and have web access in order for downloads and version information to be available via PanActive Manager or PanCLI. In addition, if the system is behind a firewall, you may need to specify a web proxy host and port.

> **NOTE:** For information on downloading software images with the command-line interface, see the **directflow** PanCLI command in the *VDURA PanCLI Reference Guide*.

Click **Advanced Tools > Software Management > DirectFlow Client** to view the list of versions available for download.



Figure 4-23. DirectFlow Client page.

All released versions compatible with your system software will be displayed as follows:

- If PanActive Link is enabled, the table will be populated with available DirectFlow Client versions. Pay particular attention to the columns that indicate compatibility with your system. Click the **Download** button to save DirectFlow Client software to your local computer.
- If PanActive Link is disabled, either enable it from **Configuration > Notification > PanActive Link** or go to the **Downloads** page of **support.VDURA.com**, click **DirectFlow Client for Linux** in the left pane, and select a compatible version to download.

If necessary, copy the software to the machine you want to install it on.

Install the DirectFlow Client

Unmount the realm from the client machine using **umount /panfs**.

For RPM installations with a previous version of DirectFlow installed, uninstall the old module before installing the new one with **rpm -e [panfs-*]**, where **[panfs-*]** is the currently installed module.

> **NOTE:** The **rpm -U** command is not supported; **rpm -e** will delete all VDURA mount lines from the **/etc/fstab** file, but will backup the file.

Install the DirectFlow Client software:

- For RPM files, where **[panfs*.rpm]** is the name of the RPM file you downloaded, use **rpm -hiv [panfs*.rpm]**.
- For DEB files, where **[path/to/deb]** is the name and location of the DEB file you downloaded, use **apt install /[path/to/deb]**.

> **NOTE:** When installing DirectFlow on SLES 11, the kernel will not allow the DirectFlow module to load. To allow loading, go to **/etc/modprobe.d/unsupported-modules** and set **allow_unsupported_modules** to 1.

Configure Secure Boot Client (Optional)

If secure boot is enabled on the client, the kernel prompts you to enroll the public PanFS key to the shim security keyring. Secure boot is provided by a BIOS and kernel supporting UEFI 2.3.1 Errata C or higher specification. When the client boots, the shim bootloader must authenticate the operating system for the

OS to load. The kernel then only loads kernel modules that it can authenticate using the shim security database. The Machine Owner Key utility allows the root user to manage the public keys for installed kernel modules.

Once DirectFlow is installed, register the DirectFlow key with the following command:

```
# mokutil --import /etc/pki/panfs/panfs-mok.der
```

You will be prompted to enter the client root user password twice. The command will not run without entering the root password. If you get a prompt to run **mokutil** and cannot run it, DirectFlow will not load. To avoid this issue, do not run **config_panfs** until you have the client root password.

Once the command is run, the output indicates that the key was successfully enrolled.

Subsequent DirectFlow installations on this client should not require a public key enrollment, as the public key for any DirectFlow port is the same.

Configure the DirectFlow Client

Before configuring the DirectFlow Client, make sure the client's system is synchronized with the same time server as the VDURA system (e.g., via NTP).

Once the software is successfully installed, configure it:

```
/usr/local/sbin/config_panfs
```

You will be prompted for the DNS name or IP address of the VDURA system. Typically, you should provide the DNS name unless your site's DNS server is not configured to resolve the system name. If you have configured your VDURA system for realm management with multiple Directors and you decide to provide the IP address, enter the IP address of the president Director.

The configuration script loads the DirectFlow kernel module and creates a VDURA mount entry in **/etc/fstab.** For example:

```
# cat /etc/fstab | grep panfs

panfs://service-45:global   /panfs  panfs   rw,noauto 0 0
```

The **panfs** specification in **/etc/fstab** is set to **noauto** by default. If you want this file system to automatically mount at boot, change the **noauto** switch to **noauto,panauto** for all installations except Ubuntu and Debian. To configure automount for an Ubuntu or Debian installation, use **_netdev** rather than **panauto** (e.g., **rw,_netdev** in the example above).

Once mounted, you can access the VDURA file system on your system through the **/panfs** mount point. In the path, follow **/panfs** with the DNS name or IP address of your system to access the root volume. If your site has multiple systems, you can access them through the same global **/panfs** mount point by specifying the appropriate DNS name or IP address.

Customize the Number of Asynchronous Writepages

The DirectFlow Client's writepages function is used by the Linux kernel to write dirty pages to storage. Writepages are allowed to begin writing pages to storage and return to the calling function before writes have completed. An application writing a file can be copying data into the page cache and writing dirty pages to storage simultaneously.

Single-client and multi-client workloads tend to need different optimizations. Single-client workloads require more asynchronous work to improve performance. Multi-client workloads usually need throttling mechanisms to prevent the system from becoming overloaded. As a result, a single-client optimization is not necessarily a good multi-client optimization.

The default number of asynchronous writepages is set to zero. To customize the system to your particular requirements, add the following mount option text to **config_panfs,** where **n** is the number of writepages you want to allow (see Configure the DirectFlow Client):

```
max-async-writepages=n
```

VDURA recommends that the number of asynchronous writepages be set to **2**, with a maximum of **32**. Higher values may cause the DirectFlow Client to hang.

Global Versus Local Mounts

The default mount point **/panfs** created by the configuration script is a global mount. It mounts the root volume of a system, providing access to all volumes on a system. You can mount multiple systems through the global **/panfs** mount point.

You can also create local mounts. You can mount a specific volume from a system, and you can mount the volume at any path on the client.

For example:

```
panfs://10.1.3.11/home  /myhome  panfs  rw  0  0
```

> **NOTE:** For more information on mount options, refer to the **mount.panfs** man page, which is installed in **/usr/local/man/man8.**

IP-Based Access Control for DirectFlow

IP-based access control allows administrators to grant or deny access to hosts on a per-volume basis (including the root volume) based on their IP addresses when accessing the file system via DirectFlow. Clients accessing storage from an authorized IP address are granted read-write or read-only access, or they are denied access depending on the rules specified by the system administrator. Different IP addresses and subnets can be granted different levels of access to the same volume based on their configuration.

IP-based access for DirectFlow can be configured using either PanActive Manager or PanCLI commands. Configuration of per-volume rules up to the maximum number of volumes (1200) is supported. Each volume supports 128 rules matching client IP addresses or subnets. The limit on rules per volume applies to all types of rules for that volume in aggregate (sum of read-write, read-only, and deny rules). Similarly, 128 volume wildcard rules (which apply to all volumes without specific rules) are supported.

Important Rules for Determining Client Access

The allowed level of access for a particular client to a volume is determined as follows:

- If no rules for a volume are found for a particular client or for a realm default wildcard client, but there are any rules at all for the volume in question, then the volume wildcard rules are never used.
- If no rules for a volume are found for a particular client or for a realm default wildcard client, and there are no rules at all for the volume in question, then the volume wildcard rules are used.
- In both of the above cases, the first matching rule determines whether access is granted read-only, granted read-write, or denied.
- See "**Using the Client Wildcard"** and "**Using the Volume Wildcard"** for important instructions on configuring and using wildcards.
- If no volume-specific or realm default rules are found for a client, access is denied.

> **NOTE:** For more information about export file rules, see DirectFlow Exports File Realm Default Rules.

If an administrator attempts to load a new exports file with rules that do not match any existing volumes, the system warns the administrator about these rules, but it does not require additional confirmation nor does it prevent the installation of the new exports. This allows an administrator to configure export rules for a volume before creating that volume and preventing it from being undesirably exposed to realm-wide defaults at any time during its lifecycle.

When DirectFlow exports are installed, they are immediately propagated through the system. The administrator does not need to perform any server-side restarts, nor does the administrator need to remount any clients. The changes, however, will not take effect until any outstanding capabilities time out (for example, in moving from read-write or read-only to deny access, the change will not take effect until the current capabilities expire).

Configure IP-Based Access

Click **Storage > DirectFlow**.

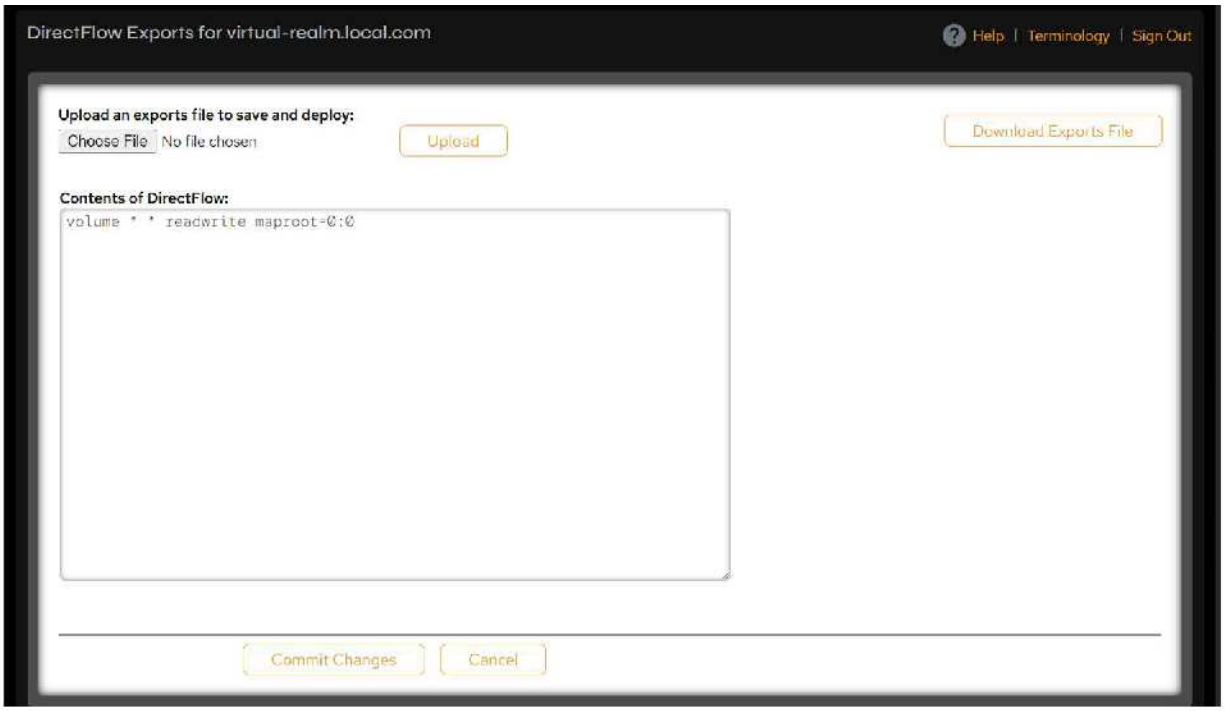


Figure 4-24. DirectFlow Exports page.

You have three options for creating or modifying an exports file:

- Edit the exports definition in the **Contents of DirectFlow: text box**, then click **Commit Changes** to save the new definition.

- Edit the exports file offline by clicking the **Download Exports File** button. This will download the current exports file as a TXT file. Make your changes, then upload the updated file by browsing to it with the **Browse** button and clicking **Upload**.
- Upload a new exports file by clicking the **Browse** button and navigating to it, then clicking the **Upload** button.

To configure the exports file with PanCLI, use the following command:

```
realm exports upload [file]
```

The exports file is in ASCII format, consistent with many UNIX configuration files. The sections below provide detailed examples of various supported file configurations.

DirectFlow Exports File Realm Default Rules

The default exports file contains lines using the following format:

```
volume * * readwrite maproot=0:0
```

*	Volume wildcard that grants access to all volumes without explicit rules.
*	Client wildcard that grants access to all clients.
readwrite	Specifies the type of access. (Other options include **readonly** and **deny**.)
maproot=0:0	Specifies the UID and GID mappings for incoming requests for UID 0. The general form of this directive is **maproot=<uid>:<gid>**. The **maproot** directive may be replaced with the **mapall** directive, which specifies a UID:GID mapping for all incoming requests matching this rule.

A simple example configuration file using this format is:

```
volume "/home" 10.65.0.0/16 readwrite maproot=0:0 # lets the work cluster read and write as root

volume "/home" 10.67.121.3 readwrite maproot=65534:65534 # specifies that Bob needs to write from his desktop

volume "/home" * readonly  maproot=65534:65534 # gives everyone else readonly access with root squashed

volume * * readonly maproot=0:0 # everything else is read-only
```

DirectFlow Exports File Formatting Rules

General formatting rules for the DirectFlow exports file include:

- Each line (as separated by a new line and/or carriage return) is parsed independently.
- Empty lines or lines that contain only white space are ignored.
- Comments are denoted by **#** characters. When a **#** is encountered outside a quoted string, that character and everything after it is treated as a comment, meaning that it is completely ignored by the parser.
- Rules begin with a keyword denoting the target of the rule. At this time, the only supported target is **volume**, meaning the rule applies to DirectFlow client access.
- The format of a volume rule is:

  ```
  volume "<volume_name>" <subnet> <permission> [<id_option>]
  ```

The literal string **volume** denotes that this is a volume rule.

- The volume name is contained in quotes (volume names may not themselves contain quotes), so there is no need for escaping within the quoted volume name. Using an asterisk in place of a quoted volume name denotes a realm-wide default (see "**Using the Volume Wildcard"**). Note that the volume name is either a realm-wide default or an exact match; no regular expressions or wildcards are supported.

Read access to the root volume (/) must always be configured in order to mount subvolumes (**/vol1**). Because **mount.panfs** does lookups and reads on directories and links in the root volume in order to locate the root object for **/vol1**, if it has no access to the root volume, then it cannot navigate to **/vol1**.

Per Volume (including root volume)

In the following example, all groups of clients have read-only access to the root volume. In addition, each of the compute clusters owns its own volume, and none of the other client groups can read the volume of another group.

```
volume "/" 192.168.0.0/16 readonly # root volume read-only for all clients in 192.168 class B
subnet

volume "/vol1" 192.168.1.0/24 readwrite

volume "/a/vol2" 192.168.2.0/24 readwrite

volume "/a/vol3" 192.168.3.0/24 readwrite
```

Because the absence of a matching rule means access is implicitly denied, clients in the subnets 192.168.0.0/16 but outside of groups 1, 2, and 3 have read-only access to files and directories of the root volume (/) but no access to **/vol1**, **/a/vol2**, or **/a/vol3**.

Suppose that clients in groups 192.168.2.0/24 and 192.168.3.0/24 need to be able to read the other groups' files. This is accomplished by listing the volume multiple times, adding access rules for a different set of clients.

```
volume "/" 192.168.0.0/16 readonly

volume "/vol1" 192.168.1.0/24 readwrite

volume "/a/vol2" 192.168.2.0/24 readwrite

volume "/a/vol2" 192.168.3.0/24 readonly # share read-only with 192.168.3.0/24

volume "/a/vol3" 192.168.3.0/24 readwrite

volume "/a/vol3" 192.168.2.0/24 readonly # share read-only with 192.168.2.0/24
```

When determining the allowed level of access for a particular client to a volume, the set of rules pertaining to the volume are evaluated as follows:

- The rules pertaining to the volume are evaluated in order. The first matching rule determines whether access is granted read-only, granted read-write, or denied.
- If no rules for the volume are found for this client, then the realm default rules are evaluated in order. The first matching rule determines whether access is granted read-only, granted read-write, or denied.
- If no volume-specific or realm default rules are found for this client, access is denied.

An advantage of the first matching rule semantic is that it allows carving out a subset within a subnet and granting the subset read-write access, while the outer subnet gets a lesser level of access, as shown in the following example:

```
volume "/" 192.168.0.10 readwrite # admin workstation .10

volume "/" 192.168.0.11 readwrite # admin workstation .11

volume "/" 192.168.0.0/16 readonly
```

Using the Client Wildcard

In the following example, each group of clients has read-write access to its own volume, but every other group is given read-only access to the volume. This is done by using a wildcard (*) client IP specification that implies “all clients”.

```
volume "/" 192.168.1.0/24 readonly
volume "/" 192.168.2.0/24 readonly
volume "/" 192.168.3.0/24 readonly
volume "/vol1" 192.168.1.0/24 readwrite
volume "/vol1" * readonly # all other clients
volume "/a/vol2" 192.168.2.0/24 readwrite
volume "/a/vol2" * readonly # all other clients
volume "/a/vol3" 192.168.3.0/24 readwrite
volume "/a/vol3" * readonly # all other clients
```

Because the client wildcard matches any IP address, subsequent rules for the volume will never be evaluated. For example:

```
volume "/a/vol2" 192.168.2.0/24 readwrite
volume "/a/vol2" * readonly # all other clients
volume "/a/vol2" 192.168.0.10 readwrite # not evaluated
```

Administrators are allowed to create these situations because they may find it desirable to temporarily insert such rules at the beginning or in the middle of a file and later remove them. This allows them to preserve the rules within their realm’s configuration.

Using the Volume Wildcard

The volume wildcard (*) can be used to establish a default access policy for all volumes without enumerating them explicitly. Perhaps the administrator anticipates adding volumes and wants them to be covered by some access control rules without manual updates to the DirectFlow exports configuration. In the following example, the group of clients are granted read-write access to all volumes:

```
volume * 192.168.1.0/24 readwrite
volume * 192.168.2.0/24 readwrite
volume * 192.168.3.0/24 readwrite
```

Volume wildcard (*) rules apply only to volumes without any explicit rules. One can get the effect of a read-only root volume with all other volumes shared read-write with:

```
volume * 192.168.1.0/24 readwrite
volume * 192.168.2.0/24 readwrite
volume * 192.168.3.0/24 readwrite
volume "/" 192.168.1.0/24 readonly
volume "/" 192.168.2.0/24 readonly
volume "/" 192.168.3.0/24 readonly
```

Every volume will either have some volume-specific rules or not. If there are one or more volume-specific rules for the volume, then those rules govern access. If there are none, access is governed by the volume wildcard (*) rules. There is no automatic fallback to wildcard rules when a volume specific rule exists.

> **NOTE:** If there is a rule for a specific volume, that volume will not be evaluated in any rule that uses * as the volume specification.

Negative Rules

System administrators may wish to carve out exceptions, denying access to one or more clients that would match a subsequent rule.

```
volume "/" 192.168.1.0/24 readonly
volume "/" 192.168.2.0/24 readonly
volume "/" 192.168.3.0/24 readonly
volume "/vol1" 192.168.1.0/24 readwrite
volume "/vol1" 192.168.2.1 deny
volume "/vol1" 192.168.3.1 deny
volume "/vol1" * readonly # all other clients
volume "/a/vol2" 192.168.2.0/24 readwrite
volume "/a/vol2" 192.168.1.1 deny
volume "/a/vol2" 192.168.3.1 deny
volume "/a/vol2" * readonly # all other clients
volume "/a/vol3" 192.168.3.0/24 readwrite
volume "/a/vol3" 192.168.1.1 deny
volume "/a/vol3" 192.168.2.1 deny
volume "/a/vol3" * readonly # all other clients
```

Here, three specific clients (192.168.1.1, 192.168.2.1, and 192.168.3.1) which would otherwise fall in a subnet allowed read-only access by the client wildcard (*) are denied access. Specifically, 192.168.1.1 is denied access to **/a/vol2** and **/a/vol3**, 192.168.2.1 is denied access to **/vol1** and **/a/vol3**, and 192.168.3.1 is denied access to **/vol1** and **/a/vol2**.

Identity Mapping Options

In the situation where each group of clients has read-write access to its own volume, we can grant every other group limited read-write access by mapping all remote users to a **<user>:<group>** of **65534:65534**:

```
volume "/" 192.168.0.0/16 readonly

volume "/vol1" 192.168.1.0/24 readwrite

volume "/vol1" 192.168.0.0/16 readwrite mapall=65534:65534

volume "/a/vol2" 192.168.2.0/24 readwrite

volume "/a/vol2" 192.168.0.0/16 readwrite mapall=65534:65534

volume "/a/vol3" 192.168.3.0/24 readwrite

volume "/a/vol3" 192.168.0.0/16 readwrite mapall=65534:65534
```

> **NOTE:** When mapping identity options, the numerical form of the UID/GID is supported, but the user/group name is not (e.g., **mapall=65534:65534** is supported; **mapall=nobody:nobody** is not).

maproot=65534:65534 is the default behavior for all DirectFlow exports, assuming no other **maproot** or **mapall** directive was specified. To allow each group of clients read-write access (as root) to its own volume:

```
volume "/" 192.168.0.0/16 readonly

volume "/vol1" 192.168.1.0/24 readwrite maproot=0:0 # root access 1

volume "/vol1" 192.168.0.0/16 readwrite mapall=65534:65534

volume "/a/vol2" 192.168.2.0/24 readwrite maproot=0:0 # root access 2

volume "/a/vol2" 192.168.0.0/16 readwrite mapall=65534:65534

volume "/a/vol3" 192.168.3.0/24 readwrite maproot=0:0 # root access 3

volume "/a/vol3" 192.168.0.0/16 readwrite mapall=65534:65534
```

If an existing export entry contains a host IP address, it is flagged as a conflict for a new export containing a different volume name, but the same host IP address. Use the following export entry:

```
/<original_volume> <options> /<new_volume> <options> <IP
address>
```

The DirectFlow exports configuration language does not support rules that only determine access permission or identity mapping. DirectFlow exports rules have effect on both, even if that effect is the default effect, read-write access or root-squashing to **65534:65534**.

When handling a request from a client to access a volume, the previously described matching process is performed once, yielding both an access permission for the client (read-write, read-only, deny) and an identity mapping policy for the access (different accesses from the same client will be mapped differently when **maproot** is in play).

4.2.2 Configure Netgroups

Use **netgroups** to group a number of hosts under a netgroup name.

Click **Storage > Netgroups**.

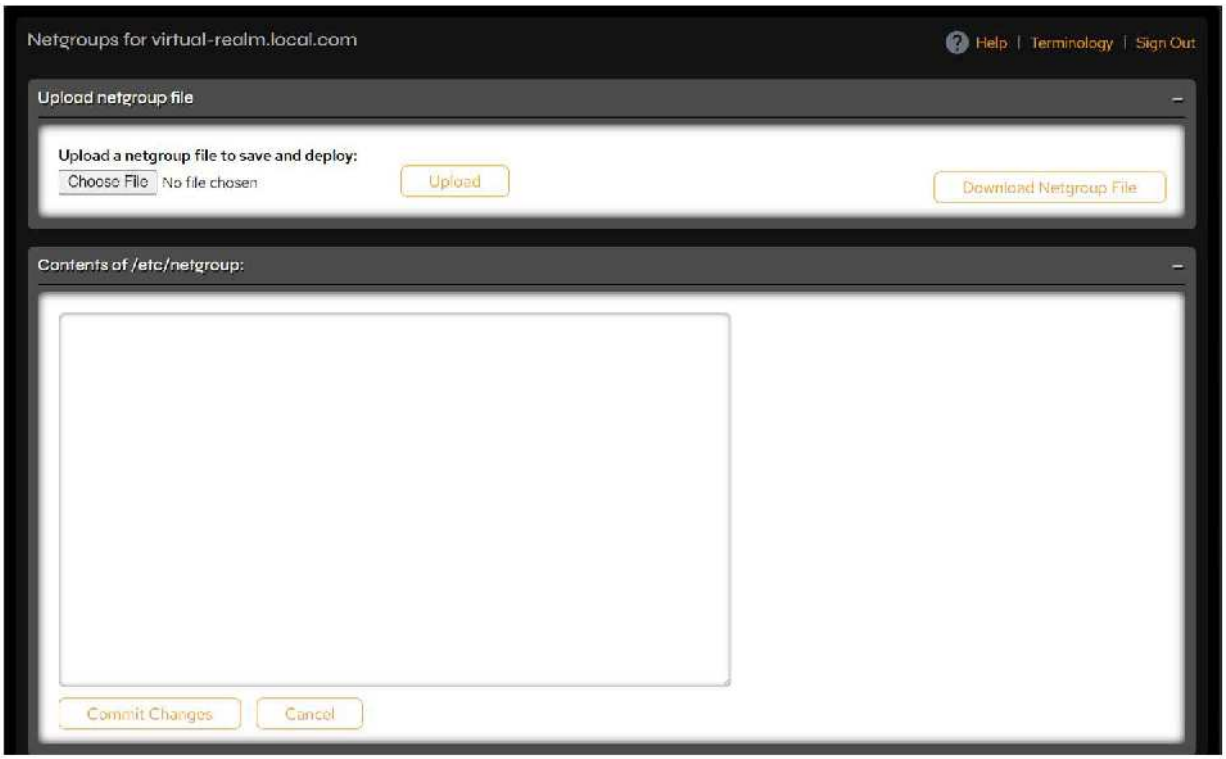


Figure 4-25. Netgroups page.

Create or modify the **/etc/netgroup** file to specify **netgroups** (sets of host, user, and domain) that are to be given similar network access.

- Each line in the file should consist of a **netgroup** name, followed by a list of the members of the **netgroup**:

  ```
  myNetGroup1 (hostname1,user1,my.company.com)
  ```

- Each member can be either the name of another **netgroup** or a specification such as **(host**, **user**, **domain)** where the host, user, and domain are character string names for the corresponding component:

  ```
  myNetGroup2 myNetGroup1 (hostname3,user3,my.company.com)
  ```

- Any of the comma-separated fields may be empty to specify a wildcard value or may be "-" to specify no valid value:

```
myNetGroup1 (hostname1,,)
```

- The members of the list may be separated by whitespace or commas, and the "\" character can be used at the end of a line to indicate line continuation.

You have three options for creating/modifying an **/etc/netgroup** file:

- Edit the netgroup definition in the **Contents of /etc/netgroup: text box**, then click **Commit Changes** to save the new definition.
- Edit the **netgroup** file offline by clicking the **Download Netgroup File** button. This will save the current netgroup file as a TXT file. Make your changes, then upload the updated file by using the **Browse** button to navigate to it and clicking the **Upload** button.
- Upload a new **netgroup** file by clicking the **Browse** button and navigating to it, then clicking the **Upload** button.

If you do not specify host or network information in the **/etc/netgroup** file, PanActive Manager assumes that you want no access restrictions on client mounting and configures the system accordingly. If you specify an invalid host or network in the **/etc/netgroup** file, PanActive Manager rejects your update to the file and informs you of the error.

> **WARNING:** Netgroup files must be 53999 bytes or smaller to successfully deploy. File sizes of 54000 bytes or larger will result in an error and fail to deploy.

4.2.3 S3 Connectivity

As of PanFS 12.0.0, S3 configuration is done through PanCTL. The PanCTL **s3** command group enables or disables the S3 gateway for the realm, sets the gateway port range and path prefix, and configures S3 per tenant. For details, see the **s3** command group in the *PanCTL User Guide*.

4.3 Permissions

4.3.1 VDURA ACL-Based Permissions

VDURA uses a rich Access Control List-based (ACL-based) permissions model, intended to provide both flexible access control and compatibility with other systems. Interoperating well with other permission models is a primary goal. Additionally, the system is intended to be flexible enough to allow users to directly specify what they want without being overly complex or cumbersome.

VDURA allows ACLs on both directories and file objects. Every object has an ACL that contains an ordered list of Access Control Entries (ACEs), each specifying a type of permission, optional inheritance flags, an identity to grant permissions to, and a bitmask of permissions to grant.

ACEs with inheritance control flags (any flag but the Inherited flag) are called inheritable ACEs. ACEs that do not have any inheritance flags are called explicit ACEs. Inherited ACEs are inherited from the parent directory and have the Inherited flag set. Effective ACEs affect access to the object and do not have the Inherit Only flag set. PanFS defines canonical mappings between PanFS ACLs and the permissions models of client systems.

> **NOTE:** VDURA recommends inheritable ACEs be limited to 125 within a volume. This allows for the additional ACEs added automatically during other processes to not exceed the VDURA limit of 128 ACEs.

VDURA ACL Structure

A VDURA ACL is an ordered list of items called ACEs: **<type, inheritance flags, identity, permissions>**. Each ACE represents the set of rights of a given type that a VDURA identity holds, and does not necessarily encapsulate that identity's overall rights on the object. The following are the main characteristics of ACLs:

- **Uniqueness**: The ACEs in an ACL shall be unique. No ACL may contain more than one ACE for a given **<type, inheritance flag, identity>** combination.
- **Ordering**: ACEs are ordered as follows:

  1. Explicit deny ACEs
  2. Explicit allow ACEs
  3. Inherited deny ACEs
  4. Inherited allow ACEs

- **Permissions**: ACLs should not contain conflicting permissions for a given identity (e.g., both positive and negative read permissions), but in practice this is hard to enforce, so ACL parsing rules handle conflicting permissions in a consistent manner according to the required ordering. While some permissions apply only to certain types of objects, any ACE may contain any of the permissions, since it is legal to specify a directory-only permission in a file ACL and vice versa. Permissions in an ACE that do not apply to the object's type are ignored. The VDURA permission set supports the full permission set as that of Windows.

- **Object creation**: All file system objects in VDURA have an ACL. If an ACL is not specified at creation time, objects receive a copy of the ACEs that have inheritance flags set from the parent directory. The ACEs are ordered at the time of creating the object. It is also possible to have objects that have an ACL but no ACEs; this is called an empty ACL. In this case, the owner of the object shall have implicit permissions to modify the ACL.

VDURA ACE Types

VDURA defines the following four ACE types:

- **Positive (+)**: rights that are specifically given to an entity.
- **Negative(~)**: rights that are specifically denied an entity.
- **Audit(@)**: provided for Windows compatibility and intentionally not interpreted or given any specific meaning by VDURA.
- **Alarm(!)**: also provided for Windows compatibility and intentionally not interpreted or given any specific meaning by VDURA.

VDURA ACL Inheritance Types

The ACL inheritance feature is based on Windows semantics. When objects are created in any directory, the ACEs that have inheritance flags set are applied to the object being created. The various inheritance flags are defined below. Each one has different semantics based on the object type that is being created.

Inheritance Flag	Meaning
Object Inherit (OI)	Non-container child objects inherit the ACE as an effective ACE. For child objects that are containers, the ACE is inherited as an inherit-only ACE unless the No Propagate Inherit flag is also set.
Container Inherit (CI)	Child objects that are containers, such as directories, inherit the ACE as an effective ACE. The inherited ACE is inheritable unless the No Propagate Inherit flag is also set.

Inheritance Flag	Meaning
No Propagate Inherit (NP)	If the ACE is inherited by a child object, the system clears the Object Inherit and Container Inherit flags in the inherited ACE. This prevents the ACE from being inherited by subsequent generations of objects.
Inherit Only (IO)	Indicates an inherit-only ACE that does not control access to the object to which it is attached. If this flag is not set, the ACE is an effective ACE which controls access to the object to which it is attached. Both effective and inherit-only ACEs can be inherited depending on the state of the other inheritance flags.
Inherited (ID)	Indicates that the ACE was inherited. The system sets this flag when it propagates an inherited ACE to a child object.

VDURA ACE Identities

The identity component of an ACE is a UNIX identity (UID or GID), a Windows identity (SID), or a special VDURA identity.

Some identities are treated specially by the VDURA access control model in order to allow general ACL settings such as "give read permission to any user" or "only allow read for administrative users". The only defined special identity is the group "everybody", which corresponds to all users in the realm.

> **NOTE:** All Directors require access to the NIS and/or LDAP identity mapping servers (see UNIX Name Resolution Services).

The table below describes the string format for the types of identities supported by VDURA.

String Format	Description
uid:*uid*	UNIX user with ID **uid**
gid:*gid*	UNIX group with ID **gid**

String Format	Description
user_sid:*sid*	Windows user with SID **sid**
group_sid:*sid*	Windows group with SID **sid**
pan_gid:1	Everyone or Other group

VDURA Permissions

VDURA permissions are a bitmask specifying the access given to a specific identity. Since VDURA ACLs may be applied to both directory and file objects, the permission set includes permissions for both directory content manipulation/access and object manipulation/access.

Other systems overload permission bits to have meaning in both contexts, but the VDURA ACL inheritance model, wherein objects by default inherit a copy of their parent directory's ACL, prevents this from being a suitable solution. The intent of the VDURA permissions model is that a permission bit's meaning should be unambiguous: either it only applies in a specific context, or it has exactly the same meaning in every context.

Directory-only permissions may be specified on a non-directory object and will have no effect. Due to the inheritance model a file could easily inherit a parent ACL with directory permissions, so there is no reason to prohibit this on file ACLs.

VDURA defines the following permissions:

Permission	Abv.	Effect on Object	Effect on Directory
read/list directory	r	read data of object	list directory
write	w	modify data of object	n/a
execute	x	execute object	list directory
append	a	append data to object	n/a
create	c	n/a	create files in the directory
delete	d	delete files	delete directory
lock	k	lock	n/a

Permission	Abv.	Effect on Object	Effect on Directory
change ACL (permissions)	p	modify ACL	modify ACL
take ownership (owner)	o	change owner	change owner
sandbox (see below)	s	n/a	see below
create dir	C	n/a	create subdirectories in directory
read named attributes	n	read named attributes of an object	n/a
write named attributes	N	modify the named attributes of an object	n/a
delete child	D	n/a	delete file or subdirectory within a directory
read attributes	R	read attributes of an object	read attributes of an object
write attributes	W	write attributes of an object	modify attributes of an object
read ACL	P	read the ACL of an object	read the ACL of an object
synchronize	S	n/a	n/a
write retention	h	n/a	n/a
write retention hold	H	n/a	n/a

Sandbox permission is a special case intended to provide an equivalent to the UNIX sticky bit on directories. It allows users to delete files and subdirectories that they own from a directory, but not those owned by other users. Essentially, it acts as a more limited version of the delete permission more suitable

for shared directories. If a user has both sandbox and delete permissions on a directory, the delete permission takes precedence; sandbox should be interpreted as a permission grant (the permission to delete files you own) and not a restriction (a restriction from deleting files owned by other users).

The string format of the ACE permission field is the concatenation of the ACE's permission abbreviations. The strings **:all:** and **:none:** represent all the permissions and none of the permissions, respectively.

Modify ACLs with PanCLI

An administrator can view and modify ACLs using the PanCLI **objutil listacl** and **objutil setacl** commands respectively. The **objutil setacl** command may add, remove, or modify an ACE depending on whether the ACL contains an ACE with a matching type, identity, and inheritance flags. If no matching ACE is found, a new ACE will be added to the ACL. If a match is found, the permissions will be updated unless **:none:** is specified, as **:none:** will delete the ACE from the ACL.

> **NOTE:** For more information and an example of modifying a VDURA ACL using PanCLI, see the **objutil** command in the *VDURA PanCLI Reference Guide*.

Modify VDURA ACLs from Clients

VDURA ACLs may be modified from certain clients.

DirectFlow Clients

The ACL can be viewed and modified using the **getfattr** and **setfattr** utilities. The ACL is stored in the **user.panfs.acl** attribute. The format of the attribute is a space-separated list of ACEs. The format of the ACE is:

[type]identity,permissions[,inheritance flags][*]. The ACE type is assumed to be positive if no type is specified. **Setfattr** will replace the current ACL with the provided ACL. The asterisk (*) indicates the ACE is for the owner or primary group and should only be used for one of the owner and primary group ACEs.

The example below demonstrates removing the **everyone** ACE and adding an inheritable ACE for UID 640 with all permissions.

```
[root@pa-quad-15d ~]# getfattr -d /panfs/REALM114/home
getfattr: Removing leading '/' from absolute path names
# file: panfs/REALM114/home
user.panfs.acl="v2: +pan_gid:1,rxnRP,I:OICI
+gid:0,rxnRP,I:OICI* +uid:0,:all:,I:OICI* +group_sid:S-1-5-
21-1214440339-1390067357-682003330-512,:all:,I:OICI"
```

```
user.panfs.agg_layout_info="RAID-1 replicas=2"
user.panfs.agg_layout_info_files="NULL map"
user.panfs.flags="volume_root"
[root@pa-quad-15d ~]# setfattr -n user.panfs.acl -v "v2:
+gid:0,rxnRP,I:OICI* +uid:0,:all:,I:OICI* +group_sid:S-1-5-
21-1214440339-1390067357-682003330-512,:all:,I:OICI
+uid:640,:all:,I:OICI" /panfs/REALM114/home
[root@pa-quad-15d ~]# getfattr -d /panfs/REALM114/home
getfattr: Removing leading '/' from absolute path names
# file: panfs/REALM114/home
user.panfs.acl="v2: +gid:0,rxnRP,I:OICI*
+uid:640,:all:,I:OICI +uid:0,:all:,I:OICI* +group_sid:S-1-5-
21-1214440339-1390067357-682003330-512,:all:,I:OICI"
user.panfs.agg_layout_info="RAID-1 replicas=2"
user.panfs.agg_layout_info_files="NULL map"
user.panfs.flags="volume_root"
```

Merge ACEs

When adding an explicit ACE to an ACL where the type and identity match and the inheritance flags and permissions are a subset of another ACE, the two ACEs may be merged into a single ACE. However, because some versions of Windows do not support the merging of inherited ACEs when an object is created, VDURA will not merge inherited ACEs.

Inheritable ACEs and Chmod

To support mixed Windows and POSIX environments where objects have inheritable ACEs, VDURA adds the **Inherit Only** flag to any inheritable ACE for the owner, primary group, and everyone. This gives the POSIX user the desired **chmod** behavior while allowing inheritable ACEs to be inherited by child objects.

Default ACL for a New Volume

When a volume is created, the ACL contains inheritable ACEs for owner, primary group, and everyone.

For example, the ACL list for **/home** looks like this:

```
[pancli_service] objutil listacl /home

ACL for /home is:

1  +uid:0,:all:

2  +gid:0,rxnRP

3  +pan_gid:1,rxnRP
```

> **NOTE:** For more information and an example of modifying a VDURA ACL using PanCLI, see the **objutil** PanCLI command in the *VDURA PanCLI Reference Guide*.

Access Permissions in a VDURA ACL

The following table shows the VDURA permissions and the characters by which they are represented when an ACL is displayed.

Windows Permissions	VDURA Permissions	VDURA Display Char	VDURA Value	Object Type
FILE_LIST_ DIRECTORY	PAN_ACL_PERM_LIST_DIR	r	0x00000001	Directory
FILE_READ_DATA	PAN_ACL_PERM_READ	r	0x00000001	File

Windows Permissions	VDURA Permissions	VDURA Display Char	VDURA Value	Object Type
FILE_EXECUTE	PAN_ACL_PERM_EXECUTE	x	0x00000002	File, Directory
FILE_WRITE_DATA	PAN_ACL_PERM_WRITE	w	0x00000004	File
FILE_APPEND_DATA	PAN_ACL_PERM_APPEND	a	0x00000008	File
FILE_ADD_FILE	PAN_ACL_PERM_CREATE	c	0x00000010	Directory
DELETE	PAN_ACL_PERM_DELETE	d	0x00000020	File, Directory
–	PAN_ACL_PERM_LOCK	k	0x00000040	–
WRITE_DAC	PAN_ACL_PERM_CHANGE_ACL	p	0x00000080	File, Directory
WRITE_OWNER	PAN_ACL_PERM_TAKE_OWNER	o	0x00000100	File, Directory
–	PAN_ACL_PERM_SANDBOX	s	0x00000200	–
FILE_ADD_SUBDIRECTORY	PAN_ACL_PERM_CREATE_DIR	C	0x00001000	Directory
FILE_READ_EA	PAN_ACL_PERM_READ_NAMED_ATTRS	n	0x00002000	File, Directory
FILE_WRITE_EA	PAN_ACL_PERM_WRITE_NAMED_ATTRS	N	0x00004000	File, Directory
FILE_DELETE_CHILD	PAN_ACL_PERM_DELETE_CHILD	D	0x00008000	Directory
FILE_READ_ATTRIBUTES	PAN_ACL_PERM_READ_ATTRS	R	0x00010000	File, Directory
FILE_WRITE_ATTRIBUTES	PAN_ACL_PERM_WRITE_ATTRS	W	0x00020000	File, Directory

Windows Permissions	VDURA Permissions	VDURA Display Char	VDURA Value	Object Type
READ_CONTROL	PAN_ACL_PERM_READ_ACL	P	0x00040000	File, Directory
SYNCHRONIZE	PAN_ACL_PERM_ SYNCHRONIZE	S	0x00080000	File
–	PAN_ACL_PERM_WRITE_ RETENTION	h	0x00100000	–
–	PAN_ACL_PERM_WRITE_ RETENTION_HOLD	H	0x00200000	–

For example, the command

```
panfs_getattr /panfs/REALM#/home/foo.txt
```

displays the ACL permissions as:

```
ACL identity: +user_sid:S-1-5-21-1214440339-1390067357-
682003330-1104,rwxadponNDRWP

ACL identity: +group_sid:S-1-5-21-1214440339-1390067357-
682003330-512,:all:,I:ID

ACL identity: +uid:0,:all:,I:ID

ACL identity: +gid:0,rxnRP,I:ID

ACL identity: +pan_gid:1,rxnRP,I:ID
```

4.4 UNIX Name Resolution Services

The system uses information obtained from the NIS and/or LDAP servers for the following authentication purposes:

- **User information**: Used for specifying the initial owner for a new volume.
- **Group information**: Used to determine group membership for file access checks, authorize PanActive Manager logins, and specify the initial group for new volumes.
- **Netgroups**: Used to group hosts under a netgroup name.

NIS and LDAP can both be used alone, or a resolution order can be configured to search multiple services in a specified order. Setting a resolution order can aid in migration where a secondary source is still in use for users not yet migrated.

LDAP services may be hosted either on UNIX systems or on Windows systems configured to support Unix users.

4.4.1 Configure Name Service Preferences

Name resolution service preferences set the order of priority in which your authentication services should be used by the system. A name resolution service must be enabled to be used by the system for authentication.

Click **Network > Authentication Services > Name Services**.

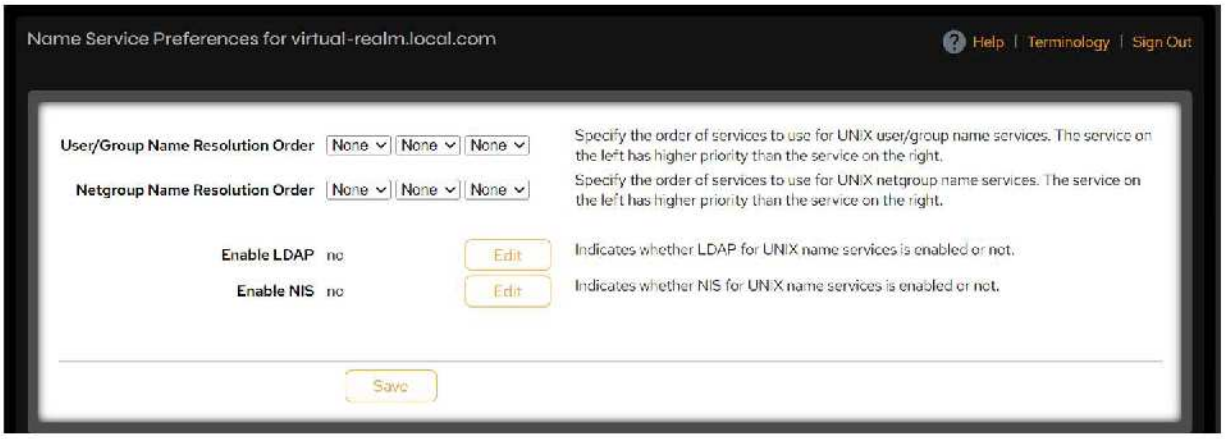


Figure 4-26. Name Service Preferences page.

Use the following controls to set the **Name Service Preferences** values.

Control	Description
User/Group Name Resolution Order	Specify the order of services to use for UNIX user/group name resolution services. The service on the left has higher priority than the service on the right.
Netgroup Name Resolution Order	Specify the order of services to use for UNIX netgroup name resolution services. The service on the left has higher priority than the service on the right. Unless a name resolution service has been enabled, it will not be used for authentication even though it is in the **Netgroup Name Resolution Order** list.
Enable LDAP	Indicates whether LDAP is enabled or not. Use the **Edit** button to modify.
Enable NIS	Indicates whether NIS is enabled for UNIX user/group/netgroup name resolution services. Use the **Edit** button to modify.

When finished, click **Save**.

4.4.2 Configure NIS

Click **Network > Authentication Services > NIS**.

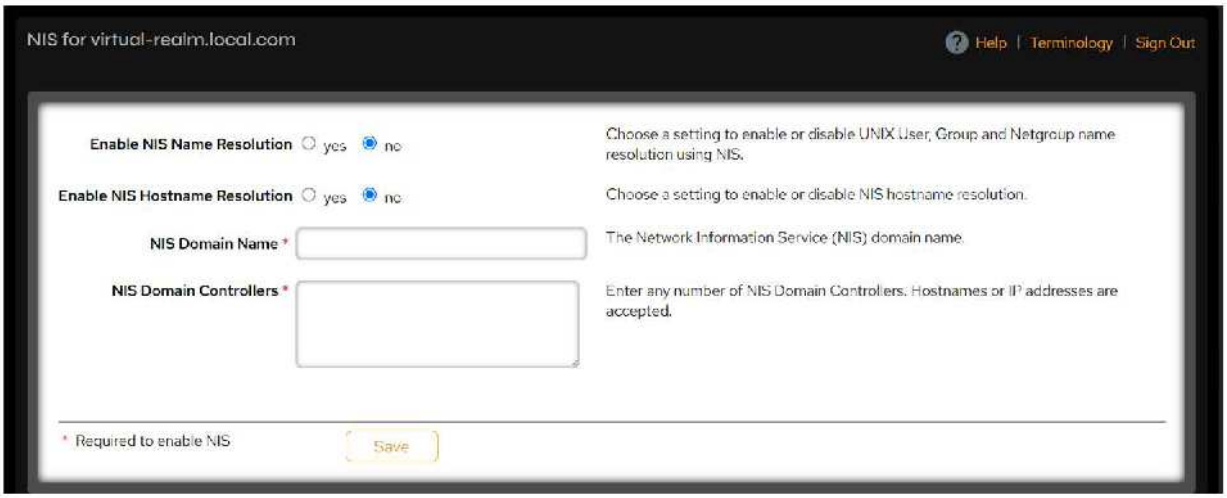


Figure 4-27. NIS page.

Use the following controls to configure NIS settings.

Control	Description
Enable NIS Name Resolution	Choose to enable or disable UNIX User, Group, and Netgroup name resolution services using NIS.
Enable NIS Hostname Resolution	Choose to enable or disable NIS hostname resolution.
NIS Domain Name	Specify the Network Information Service (NIS) domain name. Example: `mycompany.com`
NIS Domain Controllers	Specify any number of NIS server names or IP addresses separated by spaces. Example: `172.10.102.17 nis-ca.vdura.com nis-mycompany.com`

When finished, click **Save**. The system will confirm your configuration with a message at the top of the page.

4.4.3 Configure LDAP

VDURA supports servers that comply with the Lightweight Directory Access Protocol (LDAP) v3 standard (RFC 4510 and related). The LDAP schema defaults to the schema defined in RFC 2307/2307bis. VDURA does not support LDAP v2.

If you have restricted the set of IP addresses that can access the LDAP server, you must allow all addresses in the system's DHCP range, since each director will make its own queries to your site's LDAP server.

Click **Network > Authentication Services > LDAP**.

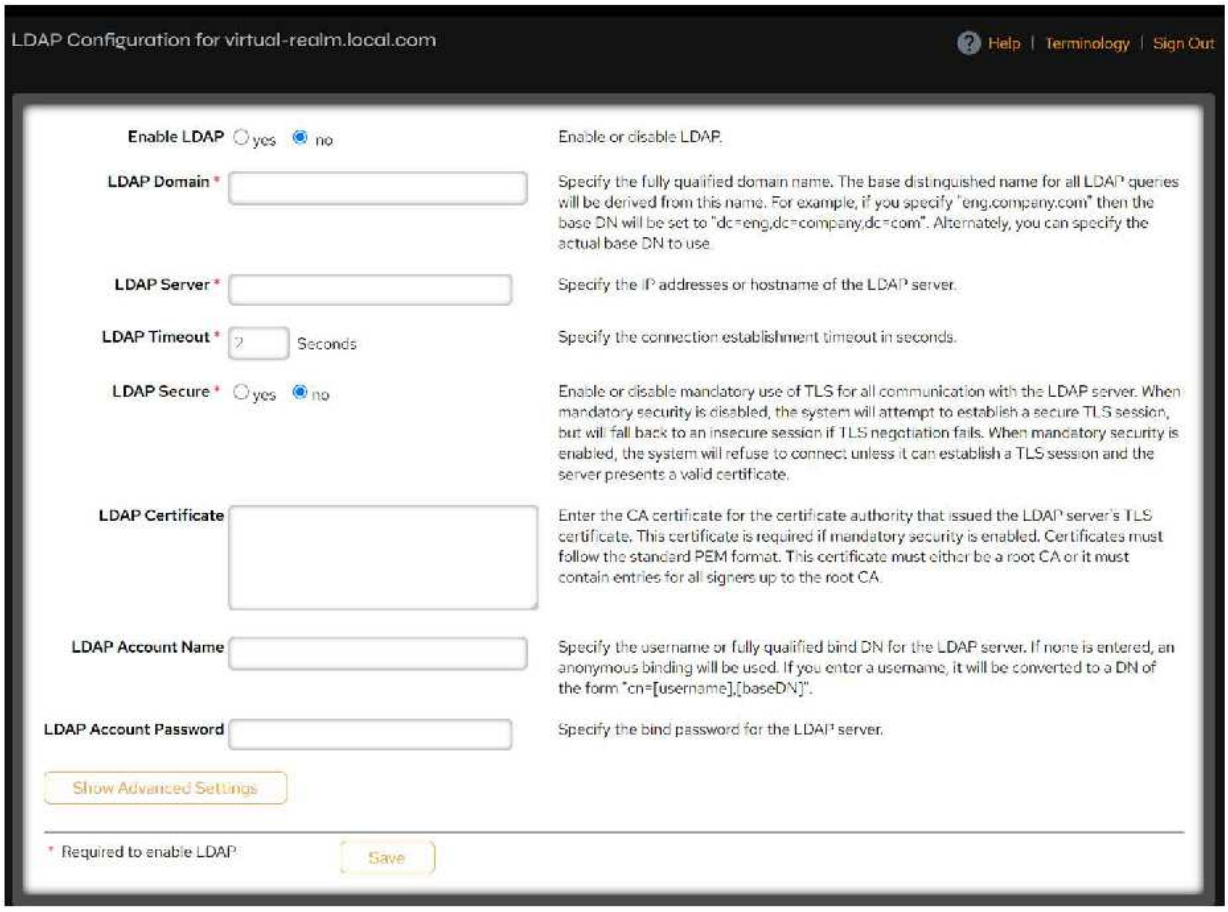


Figure 4-28. LDAP Configuration page.

Configure the following controls for LDAP settings:

Control	Description
Enable LDAP	Select **Yes** or **No** to enable or disable LDAP.
LDAP Domain	Specify the fully qualified domain name. The base distinguished name (DN) for all LDAP queries will be derived from this name. For example, if you specify **eng.company.com,** then the base DN will be set to **dc=eng,dc=company,dc=com**. Alternately, you can specify the actual base DN to use. VDURA uses the same base DN for all searches (users, groups, and netgroups). The search scope is always set to **subtree**.
LDAP Server	Specify the IP address or hostname of the LDAP server. The name specified must be resolvable in DNS but does not have to be the fully qualified name of the server. For secure LDAP setups, this should be the name as specified in the LDAP servers certificate.
LDAP Timeout	Specify the connection establishment timeout in seconds.
LDAP Secure	Enable or disable mandatory use of TLS for all communication with the LDAP server. When mandatory security is disabled, the system will attempt to establish a secure TLS session, but will fall back to an insecure session if TLS negotiation fails. When mandatory security is enabled, the system will refuse to connect unless it can establish a TLS session and the server presents a valid certificate.

Control	Description
LDAP Certificate	Enter the CA certificate for the certificate authority that issued the LDAP server's TLS certificate. This certificate is required if mandatory security is enabled. Certificates must follow the standard PEM format. **Important**: This certificate must either be a root CA or it must contain entries for all signers up to the root CA.
LDAP Account Name	Specify the username or fully qualified bind DN for the LDAP server. If none is entered, an anonymous binding will be used. If you enter a username, it will be converted to a DN of the form **cn=[username],[baseDN]**.
LDAP Account Password	Specify the bind password for the LDAP server. Since the SIMPLE bind method sends the bind password in the clear over the network, VDURA recommends using TLS for non-anonymous connections.

To configure query/search attributes, click the **Show Advanced Settings** button and proceed to the next section. Otherwise, click **Save** to commit your changes. The system will confirm your configuration with a message at the top of the page.

The default LDAP parameters are suitable for a server that uses the RFC2307 schema for user accounts, groups, and netgroups. If your LDAP server uses a different or nonstandard schema, you can configure the object class and attribute names for each type of record. These settings are presented when you click the **Show Advanced Settings** button on the **LDAP Configuration** page and are preconfigured by the

system.

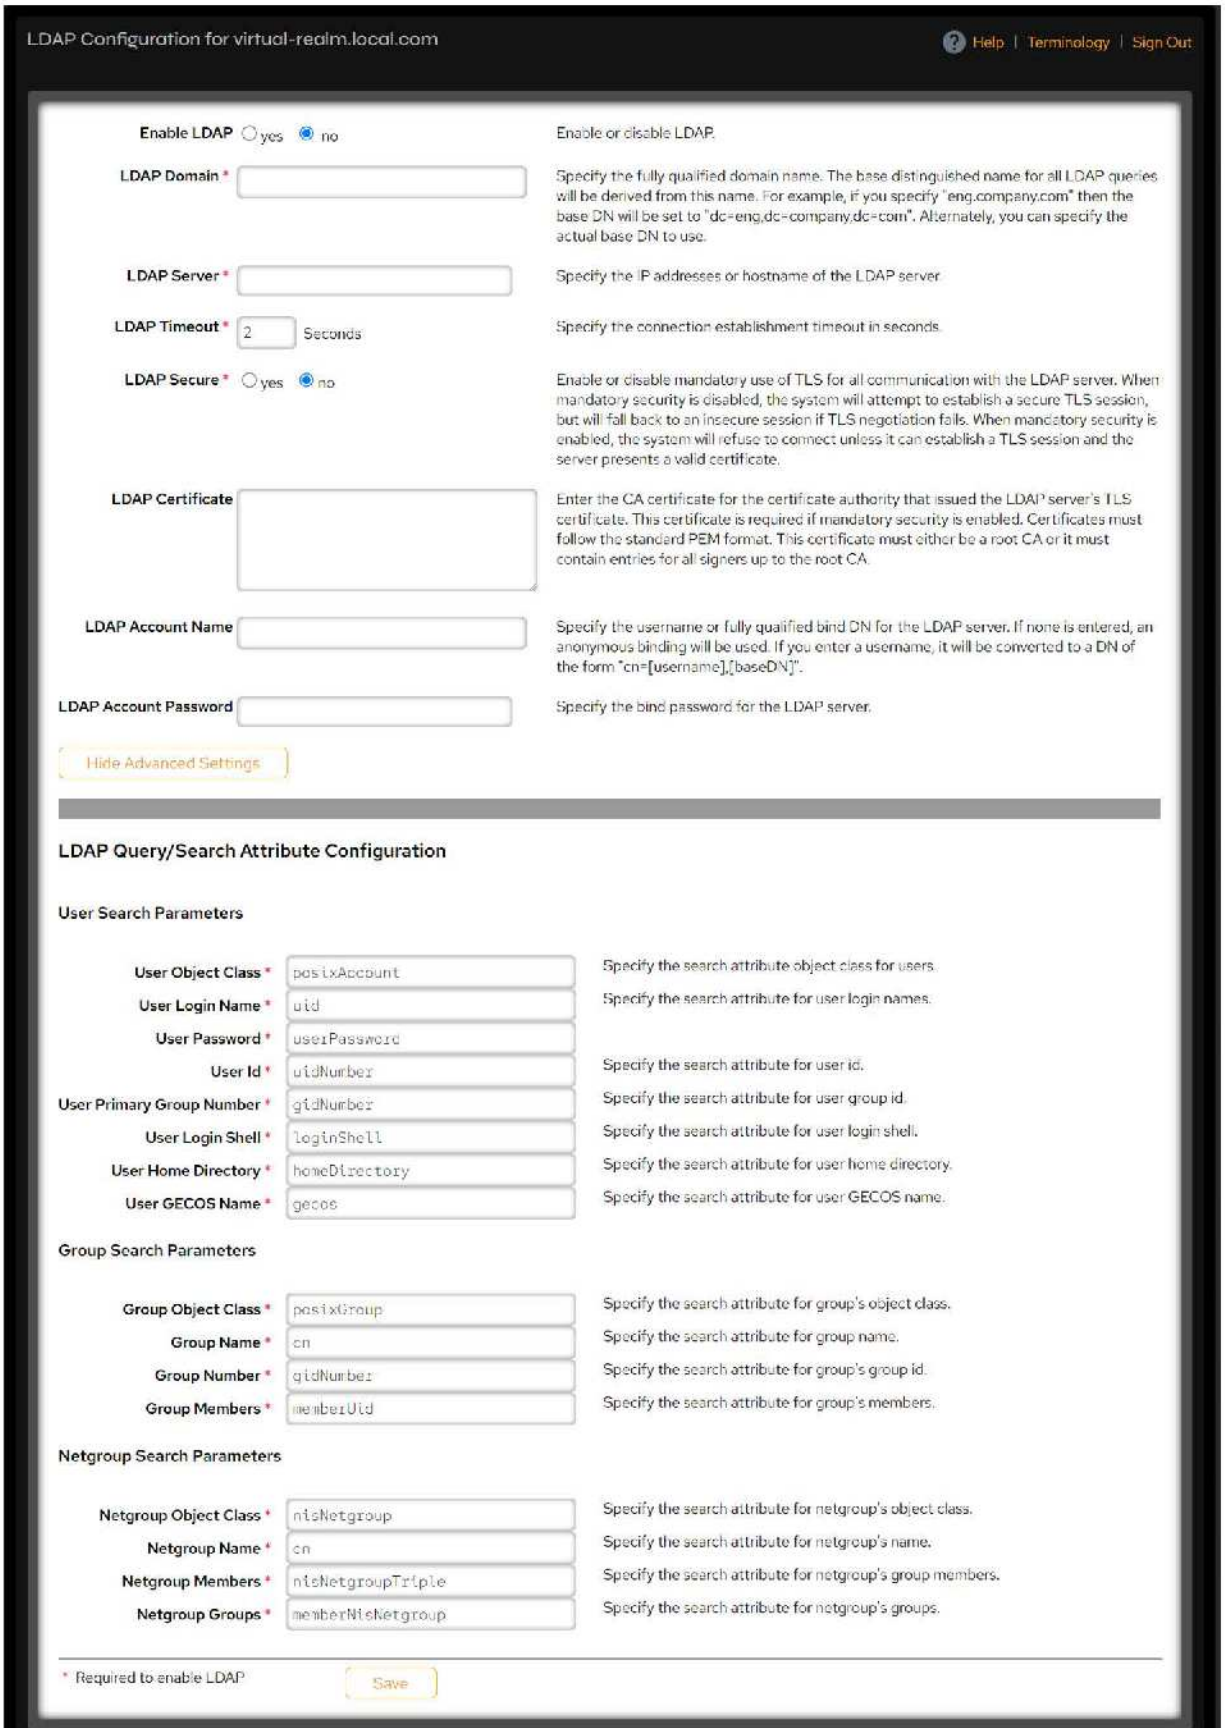

LDAP Configuration for virtual-realm.local.com
Help | Terminology | Sign Out
Enable LDAP yes no
Enable or disable LDAP.
LDAP Domain *
Specify the fully qualified domain name. The base distinguished name for all LDAP queries will be derived from this name. For example, if you specify "eng.company.com" then the base DN will be set to "dc=eng,dc=company,dc=com". Alternately, you can specify the actual base DN to use.
LDAP Server *
Specify the IP addresses or hostname of the LDAP server.
LDAP Timeout * 2 Seconds
Specify the connection establishment timeout in seconds.
LDAP Secure * yes no
Enable or disable mandatory use of TLS for all communication with the LDAP server. When mandatory security is disabled, the system will attempt to establish a secure TLS session, but will fall back to an insecure session if TLS negotiation fails. When mandatory security is enabled, the system will refuse to connect unless it can establish a TLS session and the server presents a valid certificate.
LDAP Certificate
Enter the CA certificate for the certificate authority that issued the LDAP server's TLS certificate. This certificate is required if mandatory security is enabled. Certificates must follow the standard PEM format. This certificate must either be a root CA or it must contain entries for all signers up to the root CA.
LDAP Account Name
Specify the username or fully qualified bind DN for the LDAP server. If none is entered, an anonymous binding will be used. If you enter a username, it will be converted to a DN of the form "cn=[username],[baseDN]".
LDAP Account Password
Specify the bind password for the LDAP server.
Hide Advanced Settings
LDAP Query/Search Attribute Configuration
User Search Parameters
User Object Class * posixAccount
Specify the search attribute object class for users.
User Login Name * uid
Specify the search attribute for user login names.
User Password * userPassword
User Id * uidNumber
Specify the search attribute for user id.
User Primary Group Number * gidNumber
Specify the search attribute for user group id.
User Login Shell * loginShell
Specify the search attribute for user login shell.
User Home Directory * homeDirectory
Specify the search attribute for user home directory.
User GECOS Name * gecos
Specify the search attribute for user GECOS name.
Group Search Parameters
Group Object Class * posixGroup
Specify the search attribute for group's object class.
Group Name * cn
Specify the search attribute for group name.
Group Number * gidNumber
Specify the search attribute for group's group id.
Group Members * memberUid
Specify the search attribute for group's members.
Netgroup Search Parameters
Netgroup Object Class * nisNetgroup
Specify the search attribute for netgroup's object class.
Netgroup Name * cn
Specify the search attribute for netgroup's name.
Netgroup Members * nisNetgroupTriple
Specify the search attribute for netgroup's group members.
Netgroup Groups * memberNisNetgroup
Specify the search attribute for netgroup's groups.
* Required to enable LDAP
Save

Figure 4-29. Advanced LDAP Configuration fields.

User Search Parameters

The default search attributes for user accounts are:

Parameter	Description
User Object Class	posixAccount
User Login Name	uid
User Password	userPassword
User ID	uidNumber
User Primary Group Number	gidNumber
User Login Shell	loginShell
User Home Directory	homeDirectory
User GECOS Name	gecos

Group Search Parameters

The default search attributes for groups are:

Parameter	Description
Group Object Class	posixGroup
Group Name	cn
Group Number	gidNumber
Group Members	memberUID

Netgroup Search Parameters

The default search attributes for netgroups are:

Parameter	Description
Netgroup Object Class	nisNetgroup
Netgroup Name	cn
Netgroup Members	nisNetgroupTriple
Netgroup Groups	memberNisNetgroup

Normally, these default settings are not changed. However, you can make modifications and click **Save** to commit your changes.

4.5 SSH Key Management

> **IMPORTANT:** SSH v2 is required. SSH v1 is not supported.

The VDURA system automatically generates a permanent public/private SSH key pair so that nodes can access external hosts safely and securely via **ssh** and **scp** without having to provide a password. In order for nodes to have password-free access to the root account on an external host, the administrator must copy the system public key to the file on that external host:

```
/root/.ssh/authorized_keys
```

SSH keys may be used to log into all directors and OSDs in the system. This feature does not compromise the security of the VDURA system, as it only allows outbound access to the system and never inbound access.

4.5.1 Generate a New System Public SSH Key

The administrator can generate a new public SSH key for the system at any time by navigating to **Configuration > Security > System Public SSH Key** and clicking the **Generate New System Public SSH Key** button. When a new key is generated, the old one is automatically no longer valid.

To view and generate the system public SSH key using PanCLI, enter the following command:

```
netconfig sshkey publickey
```

4.5.2 Configure Admin and Guest User SSH Keys

Employing administrator and guest user SSH keys is useful for administrators who wish to automate the running of periodic PanCLI scripts to monitor their system remotely without supplying passwords.

In order to use password-free **ssh**/**scp** access to PanCLI, the user's public SSH key must be present in the VDURA system's authorized keys list, and their SSH client must be configured to use the corresponding private key when authenticating to the VDURA system. There are separate authorized keys for the admin and guest accounts.

To configure admin or guest user access, click **Configuration > Security > Admin SSH Key**.

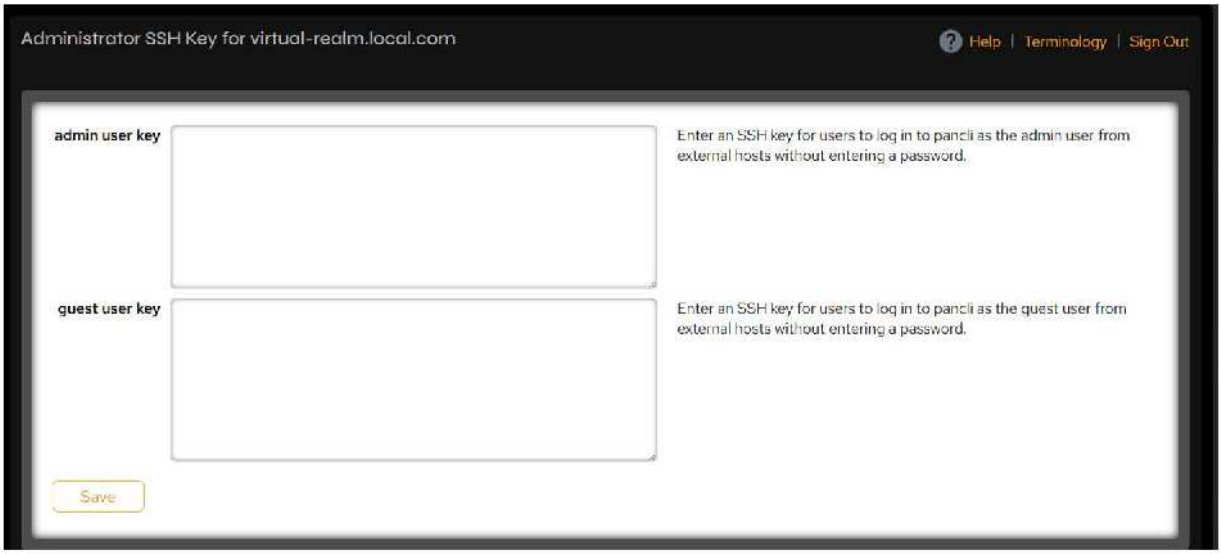


Figure 4-30. Administrator SSH Key page.

You can enable or disable access for admin and guest users. To enable access for an admin or guest user, add their SSH public key to the corresponding authorized key list. To disable access for an admin or guest user, remove their SSH public key from the corresponding authorized key list. When finished, click **Save**. Removing a user's public key does not terminate any active login sessions by that user.

> **NOTE:** The PanCLI command to control admin/guest SSH keys is **netconfig sshkey**.

The OpenSSH tool **ssh-keygen** or other commonly available tools can be used to generate a key pair for a user for whom password-free access to PanCLI is desired. Once the public/private key pair is generated, the system can be configured to allow access from users using the private key that matches the given public key. In order to maintain the security of PanCLI, the generated private key should never be placed on shared storage or in other locations that are not secure. The private key should be protected in the same way as a file containing your password.

SSH also provides the ability to encrypt the private key file with a passphrase and to require the user to provide this passphrase before using the key. VDURA recommends administrators require the use of this capability to enhance the security of their system.

After entering admin/guest SSH keys, if the SSH login continues to request a password, try logging in with **ssh -v** to see debugging information.

5. Configure Storage Data

This chapter describes how to use PanActive Manager to configure and manage storage data. It includes the following sections:

- StorageSets and Volumes
- Configure and Manage StorageSets
- Configure and Manage Volumes
- Volume Data Compression
- Configure Erasure Coding with a RAID Layout
- Horizontal Parity
- Multi-Level Erasure Coding
- User and Group Quotas
- Enclosure Migration
- Realm Management

5.1 StorageSets and Volumes

VDURA provides two mechanisms to manage capacity and the namespace: StorageSets and volumes.

5.1.1 About StorageSets

The StorageSet is a physical mechanism that groups Storage Nodes into a uniform storage pool. It is a collection of one or more enclosures grouped together to store data.

You can grow a StorageSet by adding more hardware, and you can migrate data within a StorageSet from older hardware to newer hardware with our data migration tool (see Enclosure Migration).

The VDURA system automatically creates the first StorageSet during initial installation. When you add more storage to your system, it remains unassigned and unused until you assign it to a StorageSet.

Two major factors must be considered when choosing the size of a StorageSet:

- **Performance**: The size of the StorageSet affects the maximum stripe width of data objects. A larger stripe width allows higher per-file bandwidth and faster reconstruction. Thus, larger StorageSets tend to provide better performance.

- **Data Security**: A StorageSet may contain one or more volumes. If one or more Storage Nodes fail in the StorageSet, the effects of the failure are evaluated for each volume in the StorageSet. RAID 6+ volumes can recover from up to two overlapping OSD failures in the StorageSet, while RAID 5+ volumes can recover from only a single failure in the StorageSet. If a RAID+ volume has exceeded its fault tolerance by one, EFSA keeps the volume online, but some of the objects may not be available. If a RAID+ volume exceeds its fault tolerance by more than one, the volume becomes unavailable. Thus, dividing up your storage into smaller StorageSets may decrease the likelihood of data loss or unavailability.

5.1.2 About Volumes

A volume is a logical mechanism, a sub-tree of the overall system directory structure. A read-only top-level root volume (**/**), under which all the other volumes are mounted, is created during setup. A **/home** volume is also created during setup. All other volumes are created by the user on a particular StorageSet. A volume cannot span multiple StorageSets, nor can it be moved from one StorageSet to another. The exception to this rule is the root volume, which can be moved to another StorageSet. For more information about moving the root volume, see Move the "/" Root Volume. Note that a realm can contain up to 1200 volumes, however additional volumes can be added with low-level system commands.

A volume does not have a fixed space allocation, but instead has an optional quota which can, in effect, set a maximum size on the volume. You can think of a volume as a flexible container for files and directories, and the StorageSet as a fixed size container for multiple volumes.

Because volumes cannot span multiple StorageSets, the size of a StorageSet determines the maximum amount of data that can be stored in the volumes it contains. However, you can increase the size of a StorageSet by adding new enclosures.

Planning Volume Creation

Several considerations should be taken into account when planning volume creation:

- **Use volumes to manage capacity**: Create volumes based on projects, users, or groups. You can then manage capacity by applying quotas to limit usage.
- **Create volumes to complement your backup strategy**: Use multiple volumes to facilitate multiple simultaneous backup streams, reducing overall backup time. If recovering from tape, you can recover critical volumes first.
- **Enhance performance by assigning volumes to different Directors**: Each Director runs a metadata service that manages one or more volumes; the responsibility for a volume is not spread across multiple Directors. If your system has several Directors, you should define a like number of volumes to each Director to ensure that all Directors can be assigned some of the metadata workload.

5.1.3 Understanding Capacity in StorageSets and Volumes

Available capacity in a StorageSet is shared by all volumes in it. Used capacity, however, is individual for each volume. As a result, it is possible to oversubscribe the available capacity on the StorageSet if the sum of the quotas for all volumes in the StorageSet is more than its available capacity. Care should be taken to manage the hard quotas per volume and monitor the total used capacity. Setting hard quotas on each volume can prevent this condition by ensuring that the sum of the available capacity of all the volumes is not larger than the available capacity on the StorageSet. For more information on setting hard quotas, see Create Volumes.

Capacity also includes the overhead associated with RAID 5+ and RAID 6+ parity data. RAID 6+ files offer greater protection from failures, but have a higher capacity overhead than RAID 5+ files. This should be kept in mind when calculating capacity usage and quotas in your system.

Two utilities are available for examining capacity information. The Unix **df** utility is commonly used by end users from DirectFlow file system client machines, but is restrictive in terms of what type of information can be reported. VDURA also supplies the **pan_df** utility as part of each DirectFlow package, which returns more information than the **df** utility. When using **pan_df**, be aware that the information will differ depending upon whether it is a global (overall system) mount or a volume-level mount.

Both utilities report the same three values: total capacity, capacity used, and capacity available. However, they differ on what information they use to compute those three values. The **df** utility uses the information returned by **statfs()**, which is total capacity and available capacity, and then calculates used capacity based on those two values. While the values returned through the **df** utility are intended to enable programs and scripts that run on DirectFlow clients to operate correctly, they cannot capture the full picture. On the other hand, the **pan_df** utility interacts with DirectFlow and gets data about the global mount or the volume level, depending upon what the customer is using. Hence, the values used by **pan_df** for total capacity and available capacity are different from what **df** would use.

For example, using the **pan_df** utility, the reported value for total capacity is calculated as the sum of the used and free capacity data. For a given volume (or for the overall system), the used capacity is unambiguous.

However, free capacity is more complex because there is more than one type of free capacity. There is the per-volume free capacity that relates only to the given volume. There is also the overall system free capacity that is the sum over all volumes. The sum of the reported free capacity over all volumes will typically exceed the true free capacity in the overall system. This occurs because some free capacity will be counted twice, since free capacity is associated with each StorageSet, rather than each volume.

Capacity information reported by PanActive Manager can also be obtained via the command-line interface using the **pan_pm_report_space** command. This command can produce information for the overall system or for a particular volume and StorageSet(s).

Capacity information is reported in bytes or using suffixes such as KB, MB, etc. In the latter case, the command allows for the suffixes to be defined as KB = 1000 (base 10) or KB = 1024 (base 2).

Example of pan_pm_report_space for Overall System

```
[pancli_service] diag_exec "pan_pm_report_space"
Mon Apr 21 23:12:46 GMT 2014: Executing
"pan_pm_report_space"...
Mon Apr 21 23:12:46 GMT 2014: Executed "pan_pm_report_space"
successfully:
PM 0x0500000000000002(PM) reports space (all) over 30 blades
Total Capacity: 230730033070080 bytes
Total Available Capacity: 230605356498944 bytes
Total Used Capacity: 124676571136 bytes
Usable Capacity (bucket 0):
Total: 207657029468160 bytes
Available: 207532352897024 bytes
Used: 124676571136 bytes
Reserved Capacity (bucket 1):
Total: 23073003601920 bytes
Available: 23073003601920 bytes
Used: 0 bytes
```

Example of pan_pm_report_space for a Particular Volume

```
[pancli_service] diag_exec "pan_pm_report_space -v 3"
Mon Apr 21 23:15:05 GMT 2014: Executing "pan_pm_report_space
-v 3"...
Mon Apr 21 23:15:05 GMT 2014: Executed "pan_pm_report_space -
v 3" successfully:
PM 0x0500000000000002(PM) reports space (volume) over 30
blades
Total Capacity: 230730033070080 bytes
Total Available Capacity: 230605356498944 bytes
Total Used Capacity: 124676571136 bytes
Usable Capacity (bucket 0):
Total: 207657029468160 bytes
Available: 207532352897024 bytes
Used: 124676571136 bytes
Reserved Capacity (bucket 1):
Total: 23073003601920 bytes
Available: 23073003601920 bytes
Used: 0 bytes
storageset: 1; snapshot count: 0; vol state: 3; vol flags:
0x0
live: 393216 bytes
snapshot: 0 bytes
parity: 0 bytes
hard quota: 0 bytes
soft_quota: 0 bytes
Usable Capacity (bucket 0)
free min: 6917744951296 bytes
free max: 6917745131520 bytes
percent imbalanced: 0
amount imbalanced: 1111969792 bytes
max amount imbalanced: 74166272 bytes
```

The **live**, **snapshot**, **parity**, **hard_quota**, and **soft_quota** lines are the actual volume-level capacities. **Free min** refers to the free usable (bucket 0) capacity on the most-full individual Storage Node in the StorageSet, while **free max** is correspondingly the most-empty individual Storage Node.

5.2 Configure and Manage StorageSets

Before configuring StorageSets, be sure to read StorageSets and Volumes. Always make sure the system is online via the **Hardware Management** page before completing any of the following procedures.

This section includes the following topics:

- List StorageSets
- Create StorageSets
- Manage StorageSets
  - Add Storage to a StorageSet
  - Edit a StorageSet
  - Merge StorageSets
  - Balance StorageSet Capacity
  - Delete StorageSets
- Reconstruct Failed Storage Nodes in a StorageSet

5.2.1 List StorageSets

Click **Storage > StorageSets**.

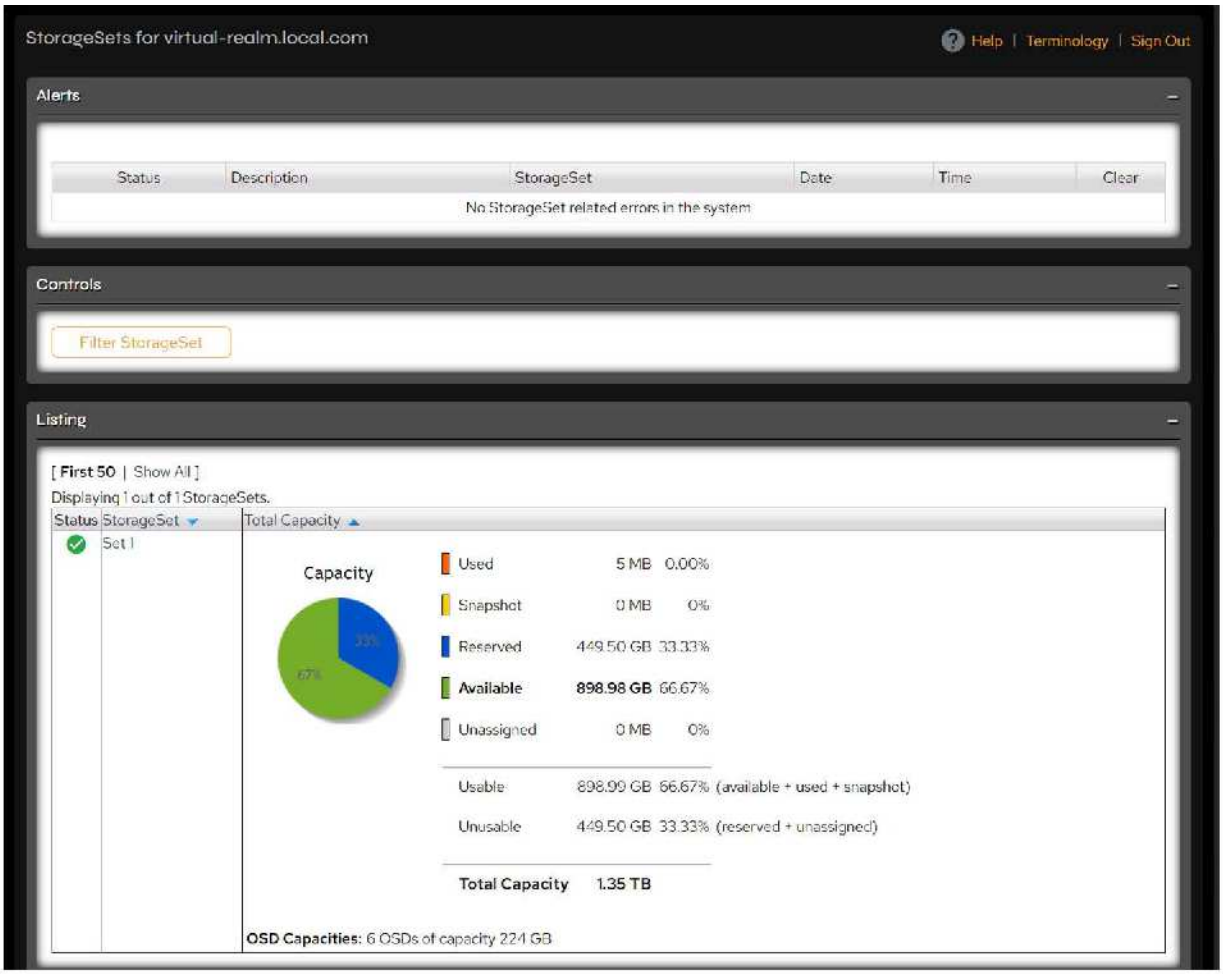


Figure 5-1. StorageSets page.

> **NOTE:** Capacity information in the GUI is reported using base 10. For example, a number optionally followed by KB, MB, GB, TB, PB or EB where KB = 1000, MB = 1000 * KB, etc.

The StorageSets page includes the following statistics and controls.

Alerts

Column	Description
Status	Shows an icon depicting error level. Click the icon to see a full description of the error.

Column	Description
Description	Brief textual description of the error state.
StorageSet	StorageSet containing error(s). Click the name for details.
Date	Date of error.
Time	Time of error.
Clear	Some non-critical errors and information messages can be removed by checking the box and clicking the **Clear** button.

Controls

Control	Description
Create StorageSet	Click to open the **Create StorageSet** page. For information on creating StorageSets, see Create StorageSets. This control will only be available if there is free hardware from which a new StorageSet can be created.
Filter StorageSet	Click to display filter fields where you can specify patterns and display/sort options to find a particular StorageSet. The following wildcards can be used in the StorageSet Glob Pattern field: * – Match any number of any set of characters. ? – Match exactly one character. [chars] – Match any character enclosed in brackets. \ – Match a literal *, ?, or bracket (used to escape wildcards).

Listing

Column	Description
Status	Displays an icon depicting error level.
StorageSet	Click the StorageSet name to view details about it and to perform functions such as editing or adding storage. See Manage StorageSets for more information.
Total Capacity information:	
Used	Utilized capacity of a StorageSet.
Snapshot	Snapshots are not currently supported in this release. They will be reimplemented in an upcoming release. Capacity consumed by snapshots of StorageSet data.
Reserved	Capacity reserved for use by reconstruction.
Available	Capacity ready to be consumed by growth in either **Used** or **Snapshot** capacity for any volume.
Unassigned	Total capacity minus space assigned to quotas on volumes or reserved.
Usable	**Available** plus **Used** plus **Snapshot**.
Unusable	**Reserved** plus **Unassigned**.
Total Capacity	Total amount of space in a StorageSet.
OSD Capacities	The number of Storage Nodes in the StorageSet and their capacities.

5.2.2 Create StorageSets

Use PanActive Manager to create a new StorageSet when new storage devices are added to the system. This is the only time you can increase the number of spares allotted to the system.

> **NOTE:** For information on configuring and managing StorageSets using the command-line interface, see the **storageset** command in the *VDURA PanCLI Reference Guide*.

Click **Storage > StorageSets**. PanActive Manager automatically detects free enclosures. If any are available, the **Create StorageSet** button appears under **Controls**. Click **Create StorageSet** to display the **Create StorageSet** page.

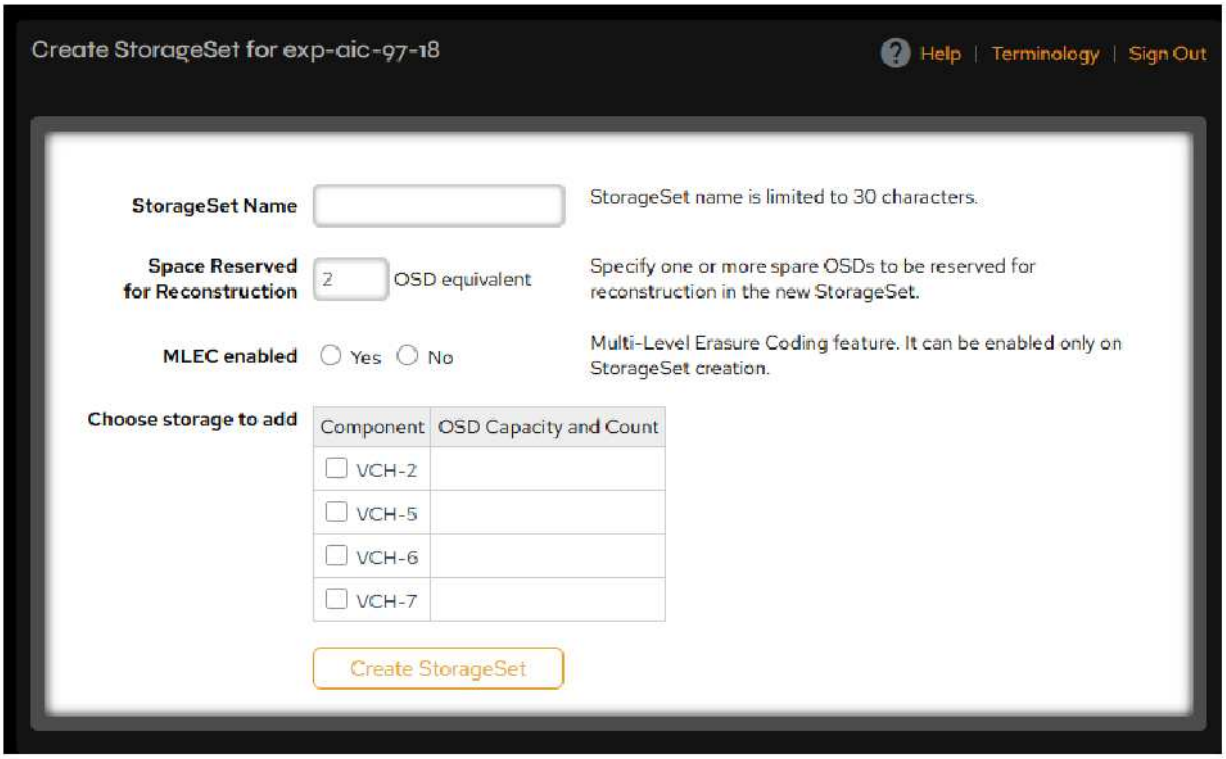


Figure 5-2. Create StorageSet page.

Enter a descriptive name up to 30 characters in length for the StorageSet in the **StorageSet Name** text box. All characters must be ASCII characters.

Enter the amount of space in OSDs (VPODs) to be reserved for reconstruction. The amount of space you reserve for reconstruction determines the number of consecutive VPOD failures the StorageSet can withstand while still providing data redundancy. Be sure to reserve sufficient space when you create the StorageSet, because additional space can only be configured when you add additional storage. For more information, see Edit a StorageSet.

If desired, enable Multi-Level Erasure Coding (MLEC). This feature provides a second layer of erasure coding and reconstruction inside VPODs. This setting cannot be changed after the StorageSet is created. For more information, see Multi-Level Erasure Coding.

Select available storage from the **Component** list and click **Create StorageSet** when finished.

5.2.3 Manage StorageSets

In the **Listing** section on the **StorageSets** page, click a name in the **StorageSet** column to open the **StorageSet Detail** page. This page contains additional information and the following controls:

Control	Description
Add Storage	Click to open the **Add Storage** page where you can add storage to the StorageSet. For more information, see Add Storage to a StorageSet. This button only appears when there is unused storage available.
Create Volume	Click to open the **Create Volume** page where you can create a new volume on the StorageSet. For more information, see Create Volumes.
Volume Recovery Priority	Click to open the **Volumes Listing** page where you can set the recovery priority for a volume in the **Listing** table. For more information, see Set Volume Recovery Priority.
Edit StorageSet	Click to open the **Edit StorageSet** page, where you can update specific attributes of the StorageSet. For more information, see Edit a StorageSet.
Merge StorageSets	Click to merge a StorageSet into the selected StorageSet. For more information, see Merge StorageSets. This button only appears when there are StorageSets that can be merged.

Control	Description
Turn Off/On Capacity Balancing	By default, capacity balancing is enabled on all StorageSets, but you can turn off capacity balancing on individual StorageSets as needed. If you have capacity balancing turned off, PanActive Manager generates a warning when a StorageSet's capacity is unbalanced. For more information, see Balance StorageSet Capacity.
Delete StorageSet	Click to delete the StorageSet. For more information, see Delete StorageSets. This button only appears when there is a StorageSet that can safely be deleted (i.e., a StorageSet with no volumes).
Reconstruct	Click to reconstruct data on reserved OSD (VPOD) space after a failure. For more information about reconstructing a StorageSet, see Reconstruct Failed Storage Nodes in a StorageSet. **Note:** Reconstruction runs on the live file system and snapshots.

Additional information field in the **Detail** area:

Additional Field	Description
Status Icon	Displays an icon depicting error level.
Set Name	Displays StorageSet name.
Status	Displays number of OSDs (VPODs) online out of total available.
View Volumes in StorageSet	Click to view the **Volumes** page for the StorageSet.
Historical Event Log	Click to go to the **Historical Event Log** page.

Additional Field	Description
Spare OSDs for Reconstruction	Displays amount of requested and available OSDs (VPODs) for reconstruction.
Capacity Imbalance	Current Imbalance: Amount of imbalance. Balancing rate: Rate at which system is correcting imbalance, if in progress.
MLEC	Enabled or disabled.
Enclosure with OSDs in this StorageSet:	Displays names of enclosure names with OSDs (VPODs) in the current StorageSet.

Avoid filling any StorageSet to more than 90 percent of capacity.

PanActive Manager issues warnings if a Storage Node or StorageSet exceeds 90 percent of capacity. Should the data on a VPOD or StorageSet approach 90 percent of capacity, consider adding storage capacity to your StorageSet or removing unneeded data.

Add Storage to a StorageSet

You can increase the capacity of an existing StorageSet by adding new, unassigned components to it. If there are unassigned VPODs available in the system, PanActive Manager displays an **Add Storage** button on the **Edit StorageSet** page. Click this button to open the **Add Storage** window.

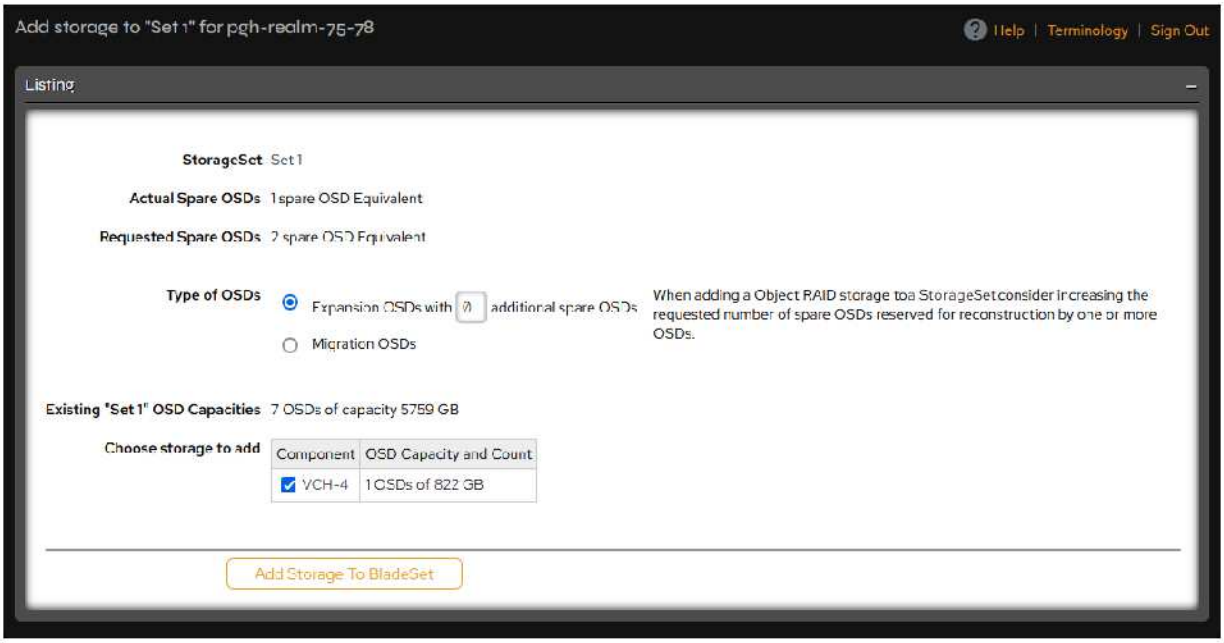


Figure 5-3. Add Storage page.

Select whether the storage will be used for expansion or migrating data from other nodes. If you will use the storage for expansion, consider assigning additional OSDs as spares reserved for reconstruction. If you are migrating data, follow the instructions in Enclosure Migration.

Unassigned storage discovered by PanActive Manager is shown in the **Choose Storage to Add** table. Select the storage you want to add, then click the **Add Storage to StorageSet** button to add it.

Edit a StorageSet

You can use PanActive Manager to edit characteristics of an existing StorageSet.

Click **Storage > StorageSets**, then click the name of the StorageSet you want to edit. The corresponding **StorageSet Detail** page appears. Click the **Edit StorageSet** button to display the **Edit StorageSet** page. This page allows you to change the StorageSet name and other characteristics described below. When finished, click **Update StorageSet**.

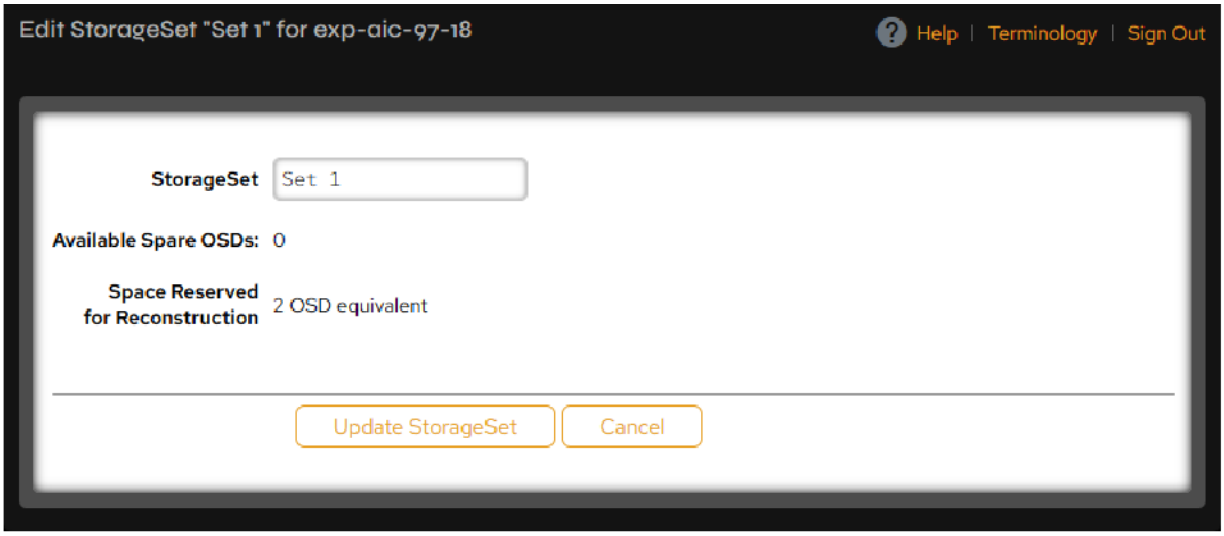


Figure 5-4. Edit StorageSet page.

The **Available Spare OSDs** field displays the current amount of space, in units of equivalent OSDs (VPODs), reserved for reconstruction. The reserved space is distributed over all OSDs in the StorageSet, rather than being restricted to individual OSDs. This is called **distributed sparing**.

The system will reserve the space equivalent of two OSDs for reconstruction per StorageSet by default. It is recommended that you increase the total reserved space as your system grows, using the guidelines and factors outlined below as you add hardware to the system.

The number of OSDs reserved determines the number of OSD failures that the StorageSet can automatically reconstruct from before you must physically replace an OSD. For example, if you reserve two OSDs of space, the StorageSet will be able to automatically reconstruct from up to two OSD failures before you must physically replace an OSD. If you have larger StorageSet, statistically you are more likely to lose more than one OSD at a time, so more reserved space than just the default is recommended.

The number of OSDs reserved also independently affects the ability of each volume in the StorageSet to reconstruct. If OSD failures overlap or are simultaneous, the availability of each volume in the StorageSet is evaluated independently for data reconstruction. Object RAID 6+ volumes can reconstruct from up to two overlapping OSD failures, so at least two OSDs worth of reserved space is recommended to allow the system to automatically reconstruct an Object RAID 6+ volume from such a failure.

The EFSA feature (see Extended File System Availability (EFSA)) available with RAID+ volume layouts can also be affected by the number of OSDs that you reserve. During EFSA, reconstruction is stopped, but the system must reconstruct any file that you attempt to modify. At least one additional spare may be necessary in order to perform this on-demand reconstruction. Therefore, if you have an Object RAID 6+ volume, it is recommended that you have at least three reserved OSDs to satisfy the EFSA feature.

Note that you may also decrease the number of OSDs reserved for reconstruction. When changed, this takes effect immediately. However, attempts to increase the amount of space reserved will not be implemented until you add additional storage to the system.

To change the amount of space reserved for reconstruction, enter the new amount of space, in units of OSDs, in the **Space Reserved for Reconstruction** field. Again, the recommended minimum number is two for StorageSets with RAID 6+ volumes and one for StorageSets with only RAID 5+ or 10+ volumes.

Merge StorageSets

This operation results in a single StorageSet that combines the components from the merged StorageSets and contains the volumes from both. You may want to merge StorageSets that share a common purpose, ownership, or management approach.

> **WARNING:** You cannot undo the merge operation.

On the **StorageSet Detail page**, click the **Merge StorageSets** button. This button only appears if there are StorageSets that can be merged. The **Merge a StorageSet into "Current StorageSet"** page appears; do the following:

- Modify the **Requested Spare OSD Count**, if desired. When merging StorageSets, consider increasing the requested number of spare nodes reserved for reconstruction to the sum of requested spares of the two StorageSets.
- Select the StorageSet you want to merge.
- Click the **Merge StorageSet** button.

Balance StorageSet Capacity

In normal operation, the system stripes data evenly across OSDs in a StorageSet. However, operations such as the ones listed below can result in unbalanced capacity across OSDs:

- Adding a Storage Node to a StorageSet's enclosure
- Adding a new enclosure to a StorageSet
- Merging StorageSets

- Reconstructing volumes

The system cannot provide maximum performance if a StorageSet is significantly unbalanced. To remedy this, the system can perform capacity balancing, redistributing data evenly on the OSDs in a StorageSet. However, capacity balancing consumes system processing resources and network bandwidth as data is moved between OSDs.

Capacity balancing attempts to evenly distribute the amount of free space over all OSDs in the system. If you have Storage Nodes with different capacities in your realm, the ones with larger capacities will be filled with data to even the amount of free space available to those with lesser capacity.

By default, capacity balancing is enabled on all StorageSets, but you can turn off capacity balancing on individual StorageSets as needed. If you have capacity balancing turned off, PanActive Manager generates a warning when a StorageSet's capacity is unbalanced. In response, you can turn on capacity balancing for that StorageSet.

Click **Storage > StorageSets**, then click the name of a StorageSet to display its **StorageSet Detail** page.

Click the **Turn Off/On Capacity Balancing** button to toggle capacity balancing off or on for the StorageSet.

Delete StorageSets

You might want to delete a StorageSet if you created it in error, if it is not used any longer, or if you want to increase space in other StorageSets. When you delete a StorageSet, the OSDs in that StorageSet become unavailable.

You can delete a StorageSet when the following two conditions have been met:

- There are no volumes in the StorageSet. Before trying to delete a StorageSet, make sure any volumes have first been deleted. Since the root volume cannot be deleted, you cannot delete a StorageSet that hosts the root volume. You must first move the root volume to another StorageSet. For more information, see Move the "/" Root Volume.
- This is not the only StorageSet in the system. If all enclosures in the system are part of this one StorageSet, the system must be reinitialized.

Click **Storage > StorageSets**, then click the name of a StorageSet to display its **StorageSet Detail** page.

Click the **Delete StorageSet** button to delete the StorageSet. The **Delete StorageSet** button will only appear when there is a StorageSet that can safely be deleted.

5.2.4 Reconstruct Failed Storage Nodes in a StorageSet

When a Storage Node within a StorageSet fails, the data on it is reconstructed using parity information and data from the other parts of the file to recreate it. The recreated data is then stored on another OSD. This process causes the volumes to transition through various states and can impact the performance and responsiveness of the system. An available spare is required for reconstruction to proceed. The availability of each volume in the StorageSet is evaluated independently for reconstruction. Volumes eligible for reconstruction are reconstructed in order of recovery priority.

In general, access to files not impacted by the failure (i.e., not on the failed OSD) is still normal. Once reconstruction starts, access to files impacted but already reconstructed is also normal. In addition, reads to files impacted but not yet reconstructed should progress, but at a slower rate. Writes to impacted but unreconstructed files will stall.

If there is an available spare, reconstruction of a failed Storage Node usually happens automatically. However, if manual intervention is required, a **Reconstruct** button will appear on the **StorageSet Detail** page. Reconstruction can be monitored in PanActive Manager under **Storage > Volumes > Listing, in the Tasks in progress column**.

5.3 Configure and Manage Volumes

Before configuring volumes, be sure to read StorageSets and Volumes.

This section includes the following topics:

- List Volume Details
- Create Volumes
- Manage Volumes
  - Edit Volume Settings
  - Delete a Volume
  - Take a Volume Offline/Online
  - Download a Volume Rescue Report
  - Move the "/" Root Volume
  - Recover Volumes
  - Handle Damaged Files
- Configure Automatic Volume Services Failover
  - Metadata Transaction Log Size
  - Move a Volume Service Manually

> **NOTE:** For information on creating and managing volumes with the command-line interface, see the PanCLI **volume** command in the *VDURA PanCLI Reference Guide*.

5.3.1 List Volume Details

The **Storage > Volumes > Listing** page displays a list of all volumes in your system.

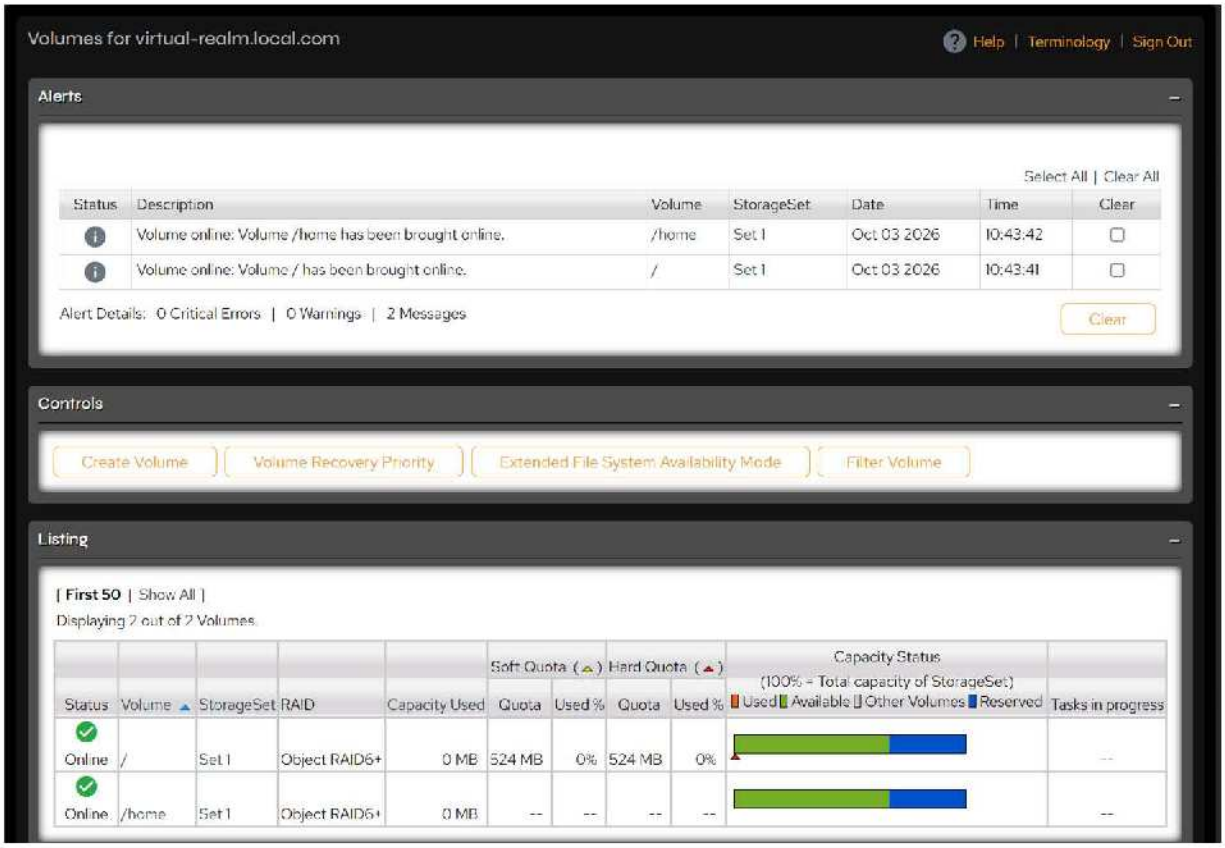


Figure 5-5. Volumes page.

The **Volumes** page includes the following information about volumes in each StorageSet.

Alerts

Column	Description
Status	Shows an icon depicting the error level. Click the icon to see a full description of the error.
Description	Brief textual description of the error state.

Column	Description
Volume	Name of volume with errors. Click the name to get details about the volume.
StorageSet	StorageSet containing error volume. Click the name to get details about the StorageSet.
Date	Date of error.
Time	Time of error.
Clear	Remove some non-critical errors and information messages by checking the box and clicking the **Clear** button.

Controls

Control	Description
Create Volume	Click to open the **Create Volume** page. For information on creating volumes, see Create Volumes.
Volume Recovery Priority	Click to display the **Recovery Priority** column in the **Listing** table where you can set the recovery priority for a volume. For more information, see Set Volume Recovery Priority.
Extended File System Availability Mode	Click this button to view or modify the **Extended File System Availability** setting. For more information, see Set the Extended File System Availability Mode.

Control	Description
Filter Volume	Click to display filter fields where you can specify patterns and display/sort options to find a particular volume. The following wildcards can be used in the Volume Glob Pattern and StorageSet Glob Pattern fields: * – Match any number of any set of characters. ? – Match exactly one character. [chars] – Match any character enclosed in brackets. \ – Match a literal *, ?, or bracket (used to escape wildcards).

Listing

Column	Description
Status	• **Online**: The system is operating normally. • **Warning**: There are system warnings that indicate a developing problem that might impair system availability or data integrity. • **Error**: The system is experiencing difficulty that might imminently impair system availability or data integrity.
Volume	Volume name.
StorageSet	Name of the StorageSet containing the volume.
RAID	RAID level for the volume (Object RAID 6+, RAID 5+, or RAID 10+).
Capacity Used	Amount of space used.

Column	Description
Soft Quota	The soft quota for the volume, if set. If the amount of data stored on the volume exceeds the soft quota, PanActive Manager sends an email notification to recipients identified when you set the quota but continues to allow data to be written to the volume. **Note:** Soft quotas are set when you create a volume or changed later on the **Edit Volume** page.
Hard Quota	The hard quota for the volume, if set. If the amount of data stored on the volume exceeds the hard quota, PanActive Manager sends an email notification to recipients identified when you set the quota and prevents any additional data from being written to the volume. **Note:** Hard quotas are set when you create a volume or changed later on the **Edit Volume** page.
Capacity Status	A band chart shows volume capacity in relation to the total capacity of the StorageSet. The bands on the chart indicate the space used by this volume, the space used by other volumes in the StorageSet, the space available, and the space reserved for reconstruction and other system use.
Tasks in Progress	Describes various tasks being performed, such as FSRC or reconstruction. User quota calculations are also listed in this column.

5.3.2 Create Volumes

Two volumes are automatically created during setup:

- The **/** root volume, under which all other volumes are mounted. The root volume is designed only to be used as an attachment point for the rest of the volumes and should never be used for storing files.
- The **/home** volume for file storage.

To create additional volumes, click **Storage > Volumes > Create**.

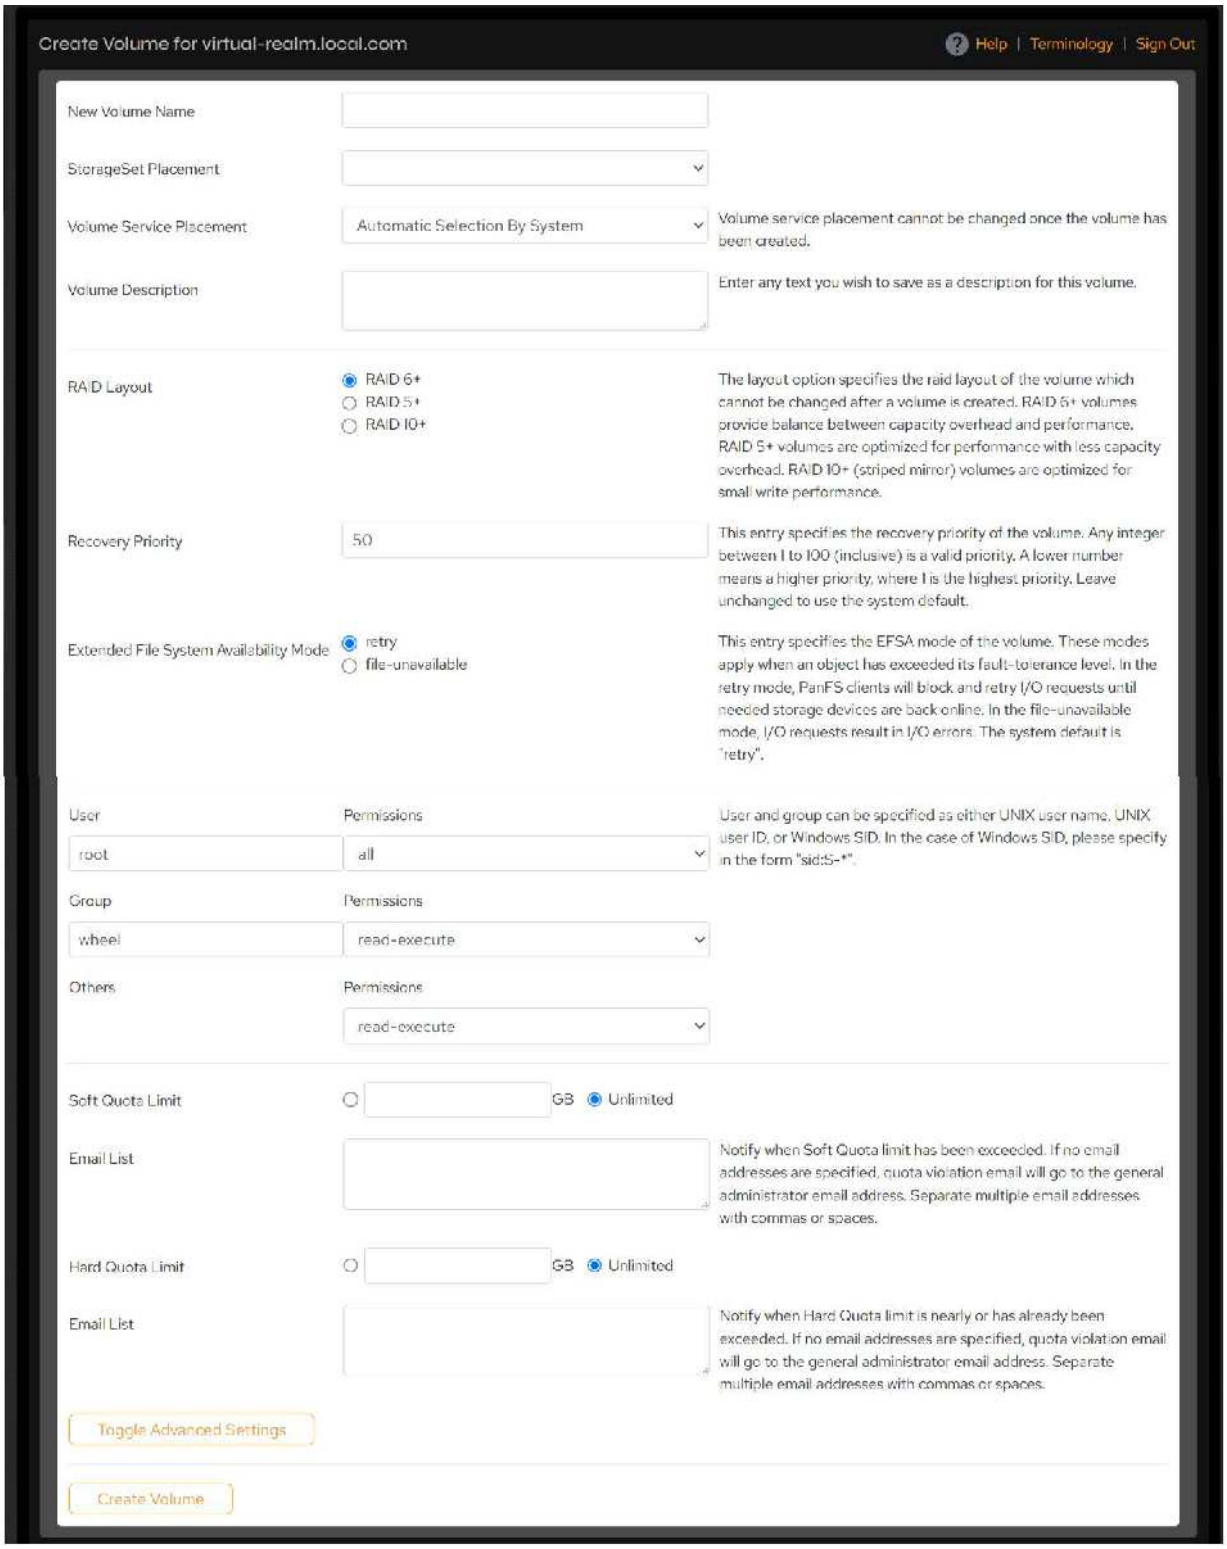


Figure 5-6. Create Volume page.

Use the controls below to set volume properties. Click **Create Volume** when finished.

Control	Description
New Volume Name	Enter a descriptive name for the volume. The following restrictions apply to volume names: • Volume name cannot be empty. • Neither **.** nor **..** can be volume name. • Volume name cannot start with **^**$**$$or$$**$.^. • Volume name cannot contain **** or **?** or more than one semicolon. • Volume names have a maximum of 255 characters.
StorageSet Placement	Select placement from the drop-down menu options. If there is only one StorageSet, it will automatically be listed. StorageSet placement cannot be changed once the volume has been created.
Volume Service Placement	This option provides some control over volume management load balancing across Directors. Each Director runs a metadata service that manages one or more volumes. The responsibility for a volume is not spread across multiple Directors. If your system has several directors, you should define a like number of volumes to ensure that all Directors can be assigned some of the metadata workload in order to enhance performance. Select **Automatic Selection by System** or select the placement manually from the drop-down list. Volume service placement cannot be changed once the volume has been created.

Control	Description
RAID Layout	The layout option specifies the RAID layout of the volume. This option cannot be changed after the volume is created. • **RAID 6+**: Volumes provide balance between capacity overhead and performance. • **RAID 5+**: Volumes are optimized for performance with less capacity overhead. • **RAID 10+** (striped mirror): Volumes are optimized for small write performance. For more information on RAID configuration, see Configure Erasure Coding with a RAID Layout.
Recovery Priority	Specify the recovery priority of the volume. Any integer from 1 to 100 is a valid priority. A lower number means a higher priority, where 1 is the highest priority. Leave unchanged to use the system default of 50. For more information, see Set Volume Recovery Priority.
Extended File System Availability Mode	Select the path to pursue when an object has exceeded its fault-tolerance level: • **retry**: VDURA clients will block and retry I/O requests until needed storage devices are back online. • **file-unavailable**: I/O requests result in I/O errors. The system default is **retry.** For more information, see Set the Extended File System Availability Mode.

Control	Description
User and Group	Set the initial ownership for the volume's root user. After creation, any client with proper permissions can mount the volume and change these settings with standard commands, such as **chown** and **chmod**. User and group can be specified as either Unix user name, Unix user ID, or Windows SID. In the case of Windows SID, please specify in the form **sid:S-***.
User/Group/Others Permissions	Permission settings include: • **None**: Specified user cannot access this volume. • **Read-only**: Specified user can read files in the volume. • **Write-only:** Specified user can create, link, remove, unlink, and rename files in the directory. • **Execute-only**: Specified user can run executables in the directory. • **Read-write**: Specified user can read, create, link, remove, unlink, and rename files in the directory. • **Read-execute**: Specified user can read files and run executables. • **Write-execute**: Specified user can create, link, remove, unlink, rename files in the directory, and run executables. • **All**: Specified user has full access to this volume.

Control	Description
Soft Quota Limit*	Enter a soft quota limit in gigabytes or click the **Unlimited** checkbox to place no limits on the volume. When the amount of data on the volume exceeds the soft quota limit, PanActive Manager sends a notification email to the people listed in the **Email List** field, but continues to allow data to be written to the volume.
Hard Quota Limit*	Enter a hard quota limit in gigabytes or click the **Unlimited** checkbox to place no limits on the volume. When the amount of data on the volume exceeds the hard quota limit, PanActive Manager sends a notification email to the people listed in the **Email List** field and prevents any additional data from being written to the volume. **Note**: The aggregate quotas of all volumes on a StorageSet can exceed the size of the StorageSet. For example, five volumes can each have a 2 TB hard quota on a 5 TB StorageSet.
*These quota limits apply to volume quotas only. For information on assigning user quotas, see Configure User and Group Quotas.	

Click the **Toggle Advanced Settings** button to customize the RAID striping, if desired.

Control	Description
Maximum Width	Specifies the maximum number of OSDs to stripe over. Any integer 1 or greater is a valid width. Leave empty to allow striping across the entire StorageSet.

Control	Description
Stripe Unit	Specifies the number of bytes written to an OSD before advancing to the next OSD in the stripe. Enter a numeric value and select KB or MB, such that the stripe unit is divisible by 16KB, with a minimum of 16KB and a maximum of 4MB. Leave empty to use the system default. For RAID 5+/RAID 6+, default is 64K.
RAID Group Width	Specifies the number of OSDs to stripe over in a single RAID group. Any integer between 3 and 20 (inclusive) is a valid width. Leave empty to use the system default. The default value is between 8 and 11 OSDs for RAID 5+ volumes and between 10 and 11 OSDs for RAID 6+ volumes. This option is only available for volumes with a RAID 5+ or RAID 6+ layout.
RAID Group Depth	This entry specifies the number of stripes written to a RAID parity group before advancing to the next parity group. Any integer 1 or greater is a valid depth. Leave empty to use the system default (16384 stripes). This option is only available for volumes with a RAID 5+ or RAID 6+ layout.

Set Volume Recovery Priority

When creating a new volume from the **Create Volume** page, you can set the recovery priority for that volume from 1 to 100, with 1 being the highest priority and 50 as the default. At any time, you can specify or change the priority with which volumes should be recovered. However, the root volume is hard-coded with a recovery priority of 1 and cannot be changed.

If two or more volumes have the same recovery priority, ties are broken using these rules in the following order:

1. The volume closest to exceeding its fault tolerance will be recovered first (only valid for reconstruction).
2. The volume with the lowest volume ID number will be recovered first.

You can also change the recovery priority from the **Storage > StorageSets > StorageSet Detail** page or from the **Storage > Volumes > Listing** page by clicking the **Volume Recovery Priority** button. A **Recovery Priority** column will appear in these pages' **Listing** tables, where you can either enter a number or click the **Default** button to return the volume to the default setting. The checkbox next to each number will automatically be marked to indicate that you have changed a value and need to save for it to be applied. Once you are finished setting recovery priorities, click **Save Changes** at the bottom of the listing to implement the settings.

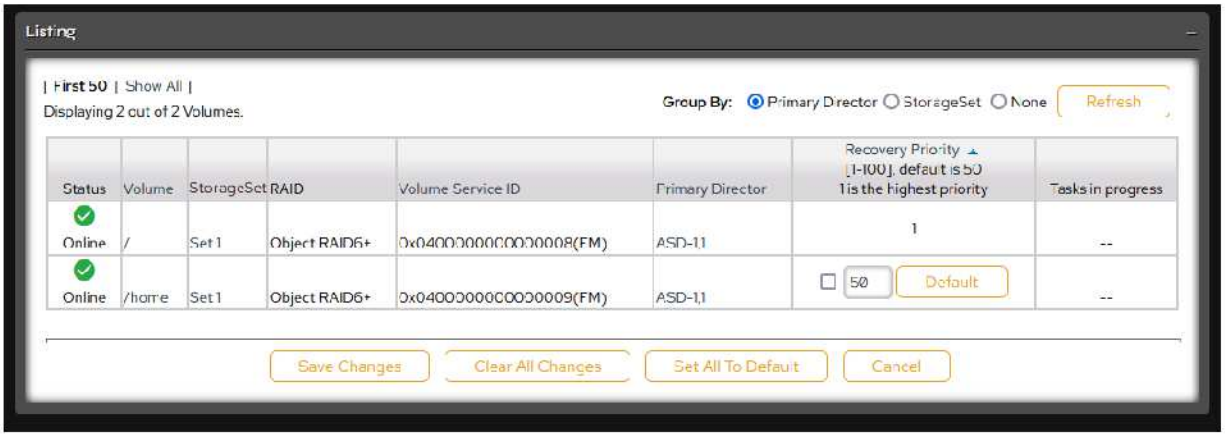


Figure 5-7. Listing page with Volume Recovery priority column.

Set the Extended File System Availability Mode

Extended File System Availability (EFSA) is defined by the ability of a RAID+ volume to remain available for reading and writing even when the fault tolerance of the volume has been exceeded by one.

- RAID 6+ volumes remain available even after three failures.
- RAID 5+ and RAID 10+ volumes remain available even after two failures.

> **WARNING:** The **ls** command may hang during an EFSA event. This command behaves differently on different platforms during EFSA. The **ls command** is often aliased in a way that makes it automatically **stat** files. You can run **unalias** to prevent this aliasing. Any time the client attempts to **stat** a file that has suffered one too many failures, it is subject to the EFSA mode for the volume (**retry** or **file-unavailable**). Therefore, any unintended **stat**-ing of a file as described here could trigger the designated behavior.

When a volume has exceeded its fault tolerance by one, you can expect the following behaviors from the file system due to EFSA:

- Volume reconstruction will stop.
- File creates and operations on newly created files will continue.
- Attempts to read or write to files whose fault tolerance have not been exceeded will also continue. However, files that have suffered one or more faults will be reconstructed on-demand before write access can be granted.
- Attempts to read or write to files whose fault tolerance have been exceeded will return an error. The error can be configured via the CLI or GUI as described in the following procedure.

Two EFSA modes can be set for a volume: **retry** or **file-unavailable**. These modes apply when an object has exceeded its fault-tolerance level.

To set or change the EFSA mode, click the **Extended File System Availability Mode** button in **Storage > Volumes**. The **Extended File System Availability Mode** column appears in the **Listing** table.

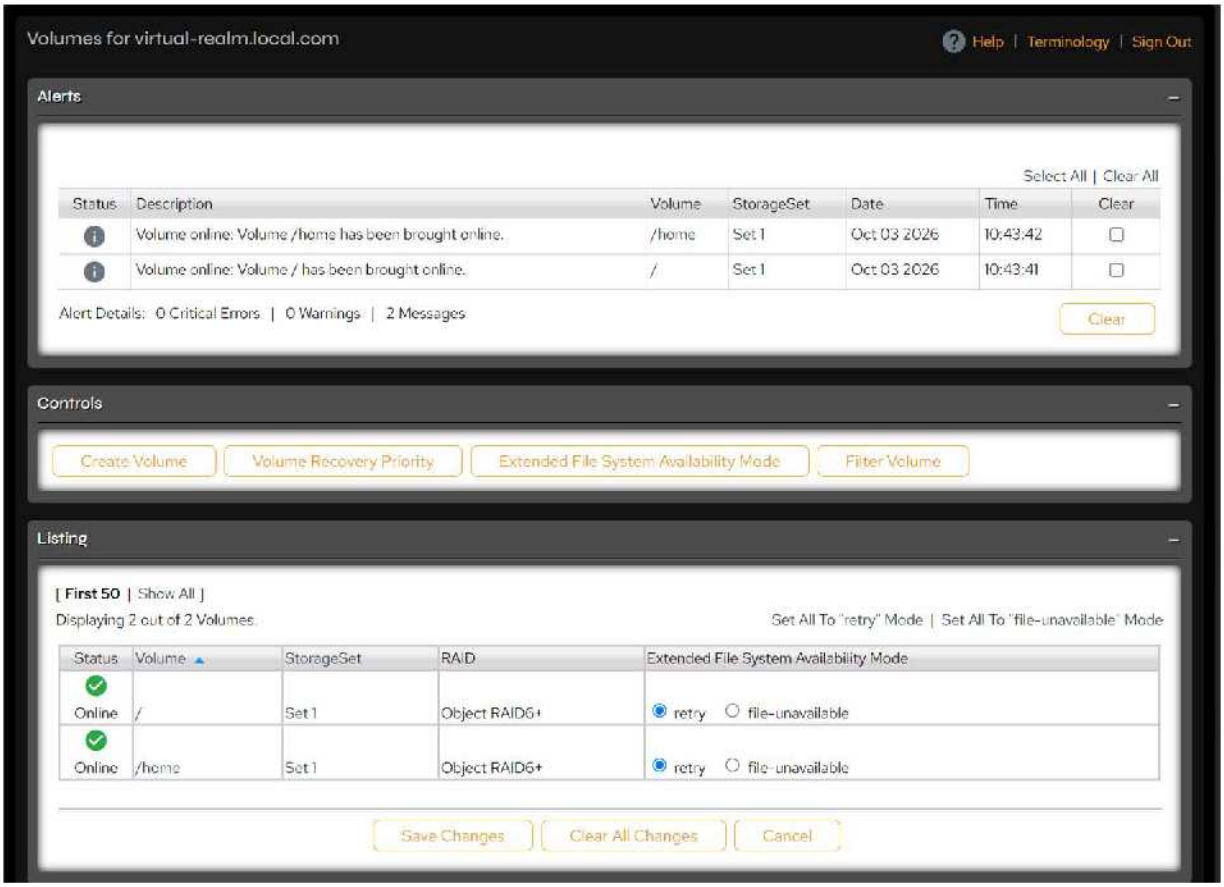


Figure 5-8. Volumes page with Extended File System Availability Mode column.

Each volume has two options in the **Extended File System Availability Mode** column:

- **retry**: VDURA clients will block and retry I/O requests until needed storage devices are back online. This is the default setting. Volumes configured for **retry** mode cause clients to get a **retry-able** error upon accessing a file that has exceeded its fault tolerance while in EFSA. When configured in this mode, access to a file that has exceeded its fault-tolerance level will appear to hang until the issue is resolved. With the **retry** setting, a DirectFlow client will continue to retry as long as it continues to get delay-and-retry from the software.
- **file-unavailable**: I/O requests result in I/O errors. Volumes configured with the **file-unavailable** error mode will cause client access to fail with a permanent error when attempting to access a file that has exceeded its fault tolerance during EFSA. When configured in this mode, access to a file that has exceeded its fault-tolerance level will result in an immediate failure, such as an I/O error, until the issue is resolved.

Select each volume as appropriate. To set all volumes to **retry** or **file-unavailable**, click the desired link on the right above the table. When finished, click **Save Changes**.

> **NOTE:** The EFSA mode can also be set while creating a volume in **Storage > Volumes > Create Volume** or from the CLI using the **volume create** command.

5.3.3 Manage Volumes

Use PanActive Manager to perform a number of management tasks on a selected volume.

Edit Volume Settings

Click **Storage > Volumes**, then click the name of the volume you want to manage in the **Listing** table **Volume** column to display its **Edit Volume** page.

Use this page to modify a number of volume settings, including volume name, access permission settings, and volume quota settings. You can also navigate to the **User and Group Quotas Enforcement Policy page** from this page to modify the user quota settings, as well as go back to the **Volume** page to modify

the recovery priority.

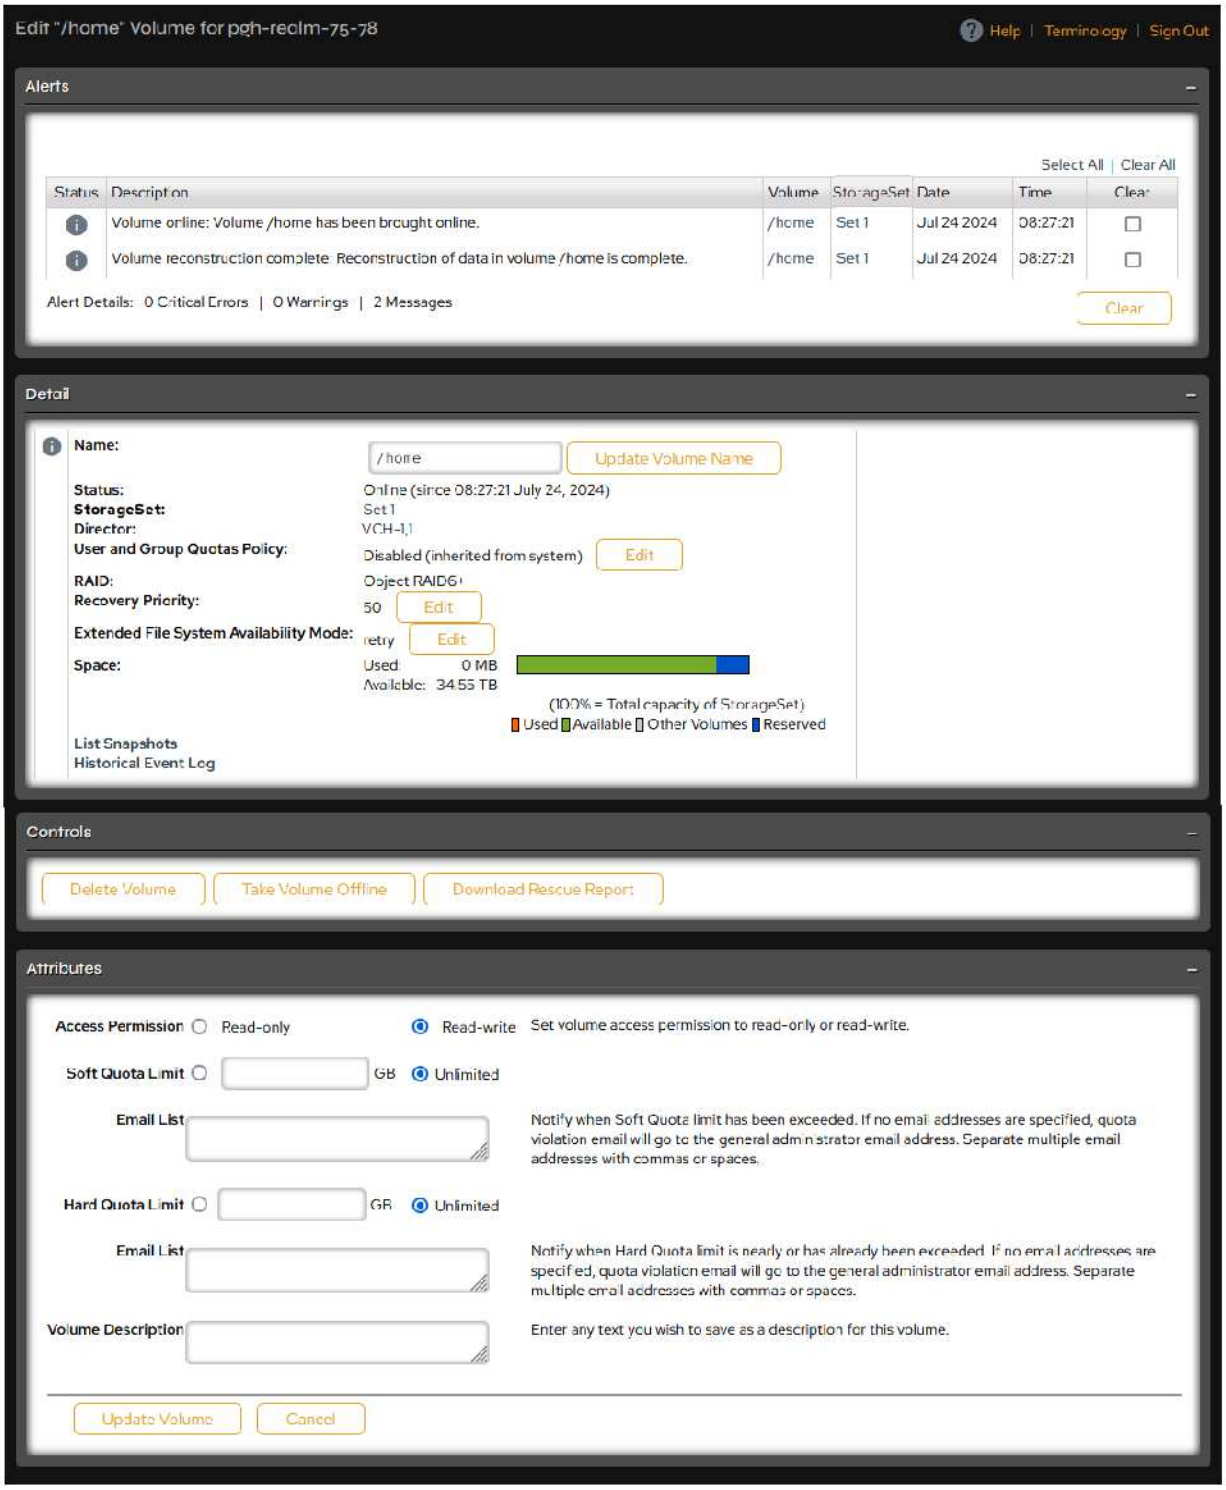


Figure 5-9. Edit Volume page.

Action	Procedure
Change volume name	In the **Name** field of the **Detail** table, enter the new name of the volume, then click **Update Volume Name**. A message at the top of the page will confirm that the volume name has been changed successfully.
Change user and group quotas policy	Click the **Edit** button next to the **User and Group Quotas Policy** field in the **Detail** table. On the **User and Group Quotas Enforcement Policy page** that opens, modify settings as desired, then click **Save**. For more information on user and group quotas, see Configure User and Group Quotas.
Change volume recovery priority	Click the **Edit** button next to **Recovery Priority** in the **Detail** table. On the **Volumes** page that opens, scroll to the **Listing** table and modify the recovery priority for the volume you want to change, then click **Save Changes**.
Change Extended File System Availability mode	Click the **Edit** button next to the current Extended File System Availability mode in the **Detail** table. On the **Volumes** page that opens, change the EFSA mode for the appropriate volume in the **Listing** table, then click **Save Changes**.
Change access permission	In the **Attributes** table, click the desired radio button next to the **Access Permission** field to change the permissions to that state. Click **Update Volume** when finished.

Action	Procedure
Change volume quota limit settings and email recipients of quota violations	Soft and hard quota limits are not dependent on each other. You can set a soft quota limit without setting a hard quota limit, and vice versa. In the **Attributes** table: • Enter a soft quota limit in gigabytes or click the **Unlimited** checkbox to exempt the volume from soft quota limits. When the amount of data on the volume exceeds the soft quota, PanActive Manager sends a notification email to the people listed in the **Email List** field, but continues to allow data to be written to the volume. • Enter a hard quota limit in gigabytes or click the **Unlimited** checkbox to exempt the volume from hard quota limits. When the amount of data on the volume exceeds the hard quota, PanActive Manager sends a notification email to the people listed in the **Email List** field and prevents any additional data from being written to the volume.

Click **Update Volume** when finished.

Delete a Volume

When you delete a volume, all data is lost. Make sure you have backed up the data or that you absolutely no longer need it. This process may take a significant amount of time to complete.

From the **Edit Volume** page, click the **Delete Volume** button. In the dialog box that opens, type **YES** and click **OK** in order to delete the volume and all of its data.

Take a Volume Offline/Online

From the **Controls** table of the **Edit Volume** page, click the **Take Volume Offline/Take Volume Online** toggle button to take the volume offline or bring it back online. Click **OK** in the dialog box to confirm your action.

Download a Volume Rescue Report

After a volume has exceeded its fault-tolerance level, the **volume rescue** feature can be used to clean up the volume and provide a damage report. This feature is destructive, as files that have exceeded their fault tolerance are deleted and the namespace is modified. Therefore, it is only available as a CLI service mode command. However, the report can be downloaded from PanActive Manager.

The **volume rescue** feature functions in two phases:

- First, a rescue reconstruction is performed on the volume. This special mode of reconstruction reconstructs all files and directories that have not exceeded their fault tolerance, and deletes all objects that have exceeded their fault tolerance. The rescue reconstruction is an offline process.
- Second, a mandatory rescue FSRC repairs the namespace of the volume and creates a log file containing the full path and file name of every object that was lost. In some cases, the full path and file name are not attainable, i.e. if a directory was lost. In these cases, the object ID is added to the log file.

Volume rescue may only be performed on a volume that has exceeded its fault-tolerance level. The system enforces this behavior.

If you are rescuing a RAID+ volume and have only exceeded the volume's fault tolerance by one, then no directories will be lost and the system will be able to give you the full path and file name of every file that was lost.

To download a volume rescue report, click the **Download Rescue Report** button on the **Edit Volume** page.

Move the "/" Root Volume

The root (**/**) volume is the exception to the rule that volumes cannot be moved to a new StorageSet. There are times when it is desirable to move the root volume, such as removing an old StorageSet that contains the root volume. Note that this is not truly moving a volume, but rather deleting the root volume and creating a new root volume on a different StorageSet. Therefore, keep in mind that moving the root volume automatically deletes all files, symbolic links, and directories that are not parts of paths to other volumes' mountpoints. Other volumes will not be affected.

Click **Storage > Volumes**, then click the root volume name **/** in the **Listing** table **Volume** column to display the **Edit "/" Volume** page.



Figure 5-10. Edit "/" Volume page.

Click the **Move Volume** button to open the **Move "/" Volume** page.

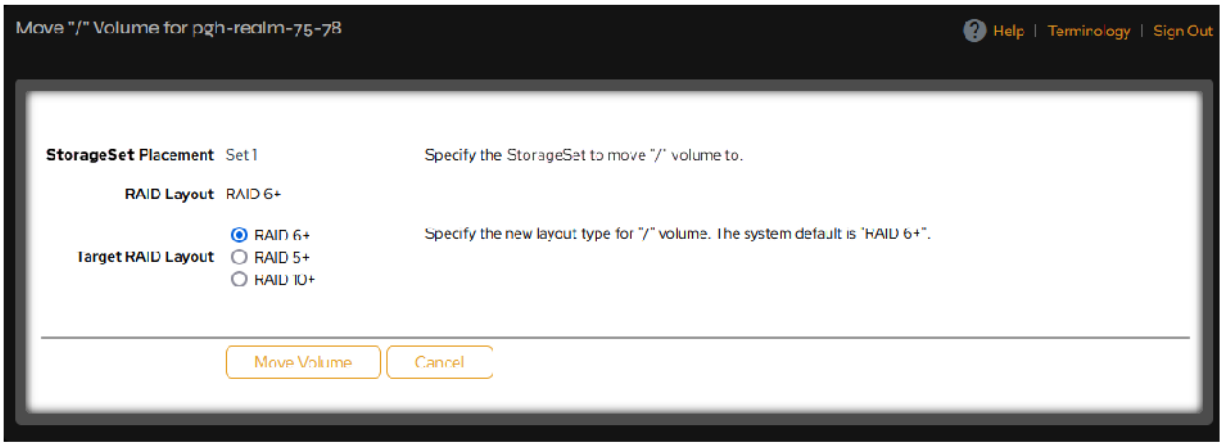


Figure 5-11. Move "/" Volume page.

Specify the new StorageSet (if appropriate) and the RAID layout for the volume, then click **Move Volume**. A warning will appear to remind you that this operation automatically deletes all files, symbolic links, and directories in the root volume that are not parts of paths to other volume mountpoints. To continue, type **YES** and click the **OK** button.

> **NOTE:** This function may also be performed using the **volume moveroot** CLI command.

If you have volumes named **/directory/vol0** and **/directory/vol1**, and the root volume is moved via the volume **moveroot** command or within the application, the permission on **/directory** will likely change after using the **moveroot** command, and you may not be able to access the volume. After the **moveroot**, modify the permissions on directories of this type to allow access to the volume.

Recover Volumes

You can use PanActive Manager to start a file system recovery (FSRC) on a volume to perform an integrity check on it. However, a file system recovery should never be performed unless requested by the system or a VDURA service representative.

Recovered files and directories are linked to a subdirectory of the volume root directory named **Lost+Found.n**, where **n** is a number that increases incrementally each time the volume is recovered.

> **WARNING:** The volume is taken offline to recover data. Data services will not be available on this volume for some time during FSRC.

Click **Storage > Volumes** then click the name of the volume you want to recover in the Listing table Volume column to display its Edit Volume page. The **Recover Volume** button will only be displayed if a volume needs recovery. Click **Recover Volume** to recover the volume.

Handle Damaged Files

The **Edit Volume** page also allows you to take corrective action should your system detect any damaged files associated with the volume during FSRC. Files can be damaged in rare circumstances, typically when a client application exits abnormally while writing or modifying a file stored on the volume. If the system detects a file with damaged data or attributes, it locks the file to prevent clients from accessing it and logs a warning in the error table on the volume.

When a volume has damaged files associated with it, the **Edit Volume** page includes a **Damaged Files in Volume** link. Click the link to display a list of all damaged files associated with the volume.

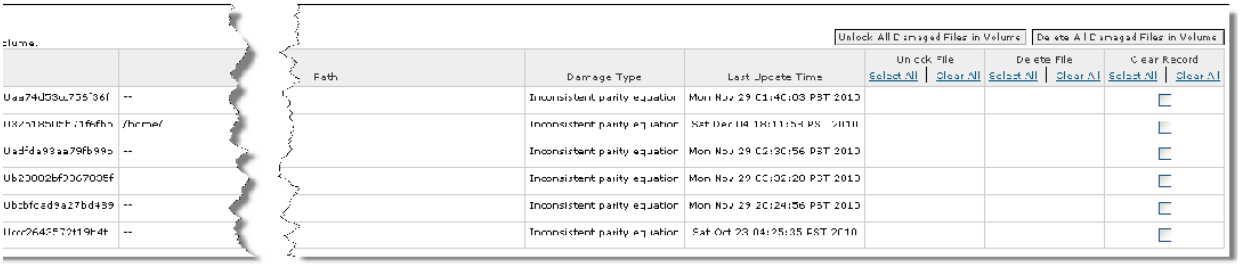


Figure 5-12. Damaged files in volume.

If some of the data might be retrievable, an **Unlock File** checkbox appears next to the file. Select this control to allow client applications to access the file again. You can then inspect the file and determine whether or not you can salvage information from the damaged file.

A **Delete File** checkbox appears next to all damaged files. If no data from the file is retrievable, it is the only control that appears next to the file name. Select this control to delete the damaged file.

After you have salvaged or deleted the damaged file(s), select the **Clear Record** box to clear the record of the damaged file.

Once you have taken corrective action on damaged files, the **Edit Volume** page includes a **History of Damaged Files** link. Click it to display a log of corrective actions taken for all damaged files that have been associated with the volume. This control does not appear if no corrective actions have been taken on this volume.

5.3.4 Configure Automatic Volume Services Failover

Each data service has responsibility for a set of volumes, and each Director can host one or more data services. Data services can be moved from both failed and running Directors to another Director. This moves the metadata management responsibility for all volumes controlled by the data service; there is no data movement.

Click **Storage > Volumes > Service Configuration**. Select the **Enabled** radio button, then click **Change Volume Service Failover**. The following message will appear: "Successfully enabled volume services failover. Backup services should come online shortly."

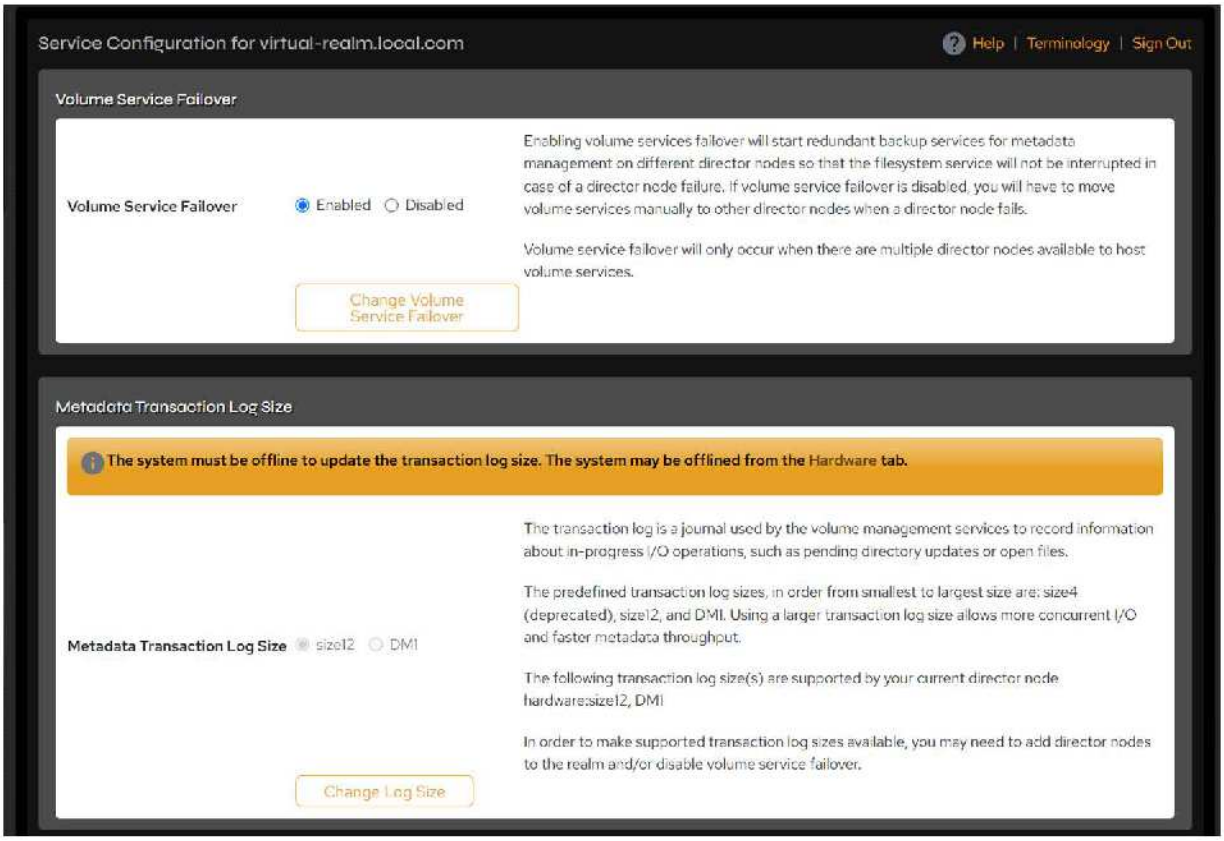


Figure 5-13. Service Configuration page.

In the **Volumes** menu, select **Service Ownership** to open **Volume Services Ownership** page. Under the **Volume Services Listing,** the Backup Director Node and Backup Director Node IP columns will populate as backups are created for the volume services.

PanActive Manager will automatically create the most efficient load balancing scenario for your system. However, you can also manually move volume services and/or backups using the fields on this page.

At least two Directors must be eligible to host volume services in order to support the volume services failover feature.

Metadata Transaction Log Size

The transaction log is a journal used by the volume management services to record information about in-progress I/O operations, such as pending directory updates or open files. This setting affects all volume services in the system.

The available journal sizes are size4 (deprecated), size12, and DM1; the sizes supported depend on the Director hardware. Using DM1 allows high concurrent I/O and faster metadata throughput. The system must be offline to change this setting. Select a size and click **Change Log Size**.

Move a Volume Service Manually

For users who have disabled automatic Volume Services Failover, volume services may be moved manually. Click **Storage > Volumes > Service Ownership** and scroll to **Move A Volume Service**

Manually at the bottom of the page.

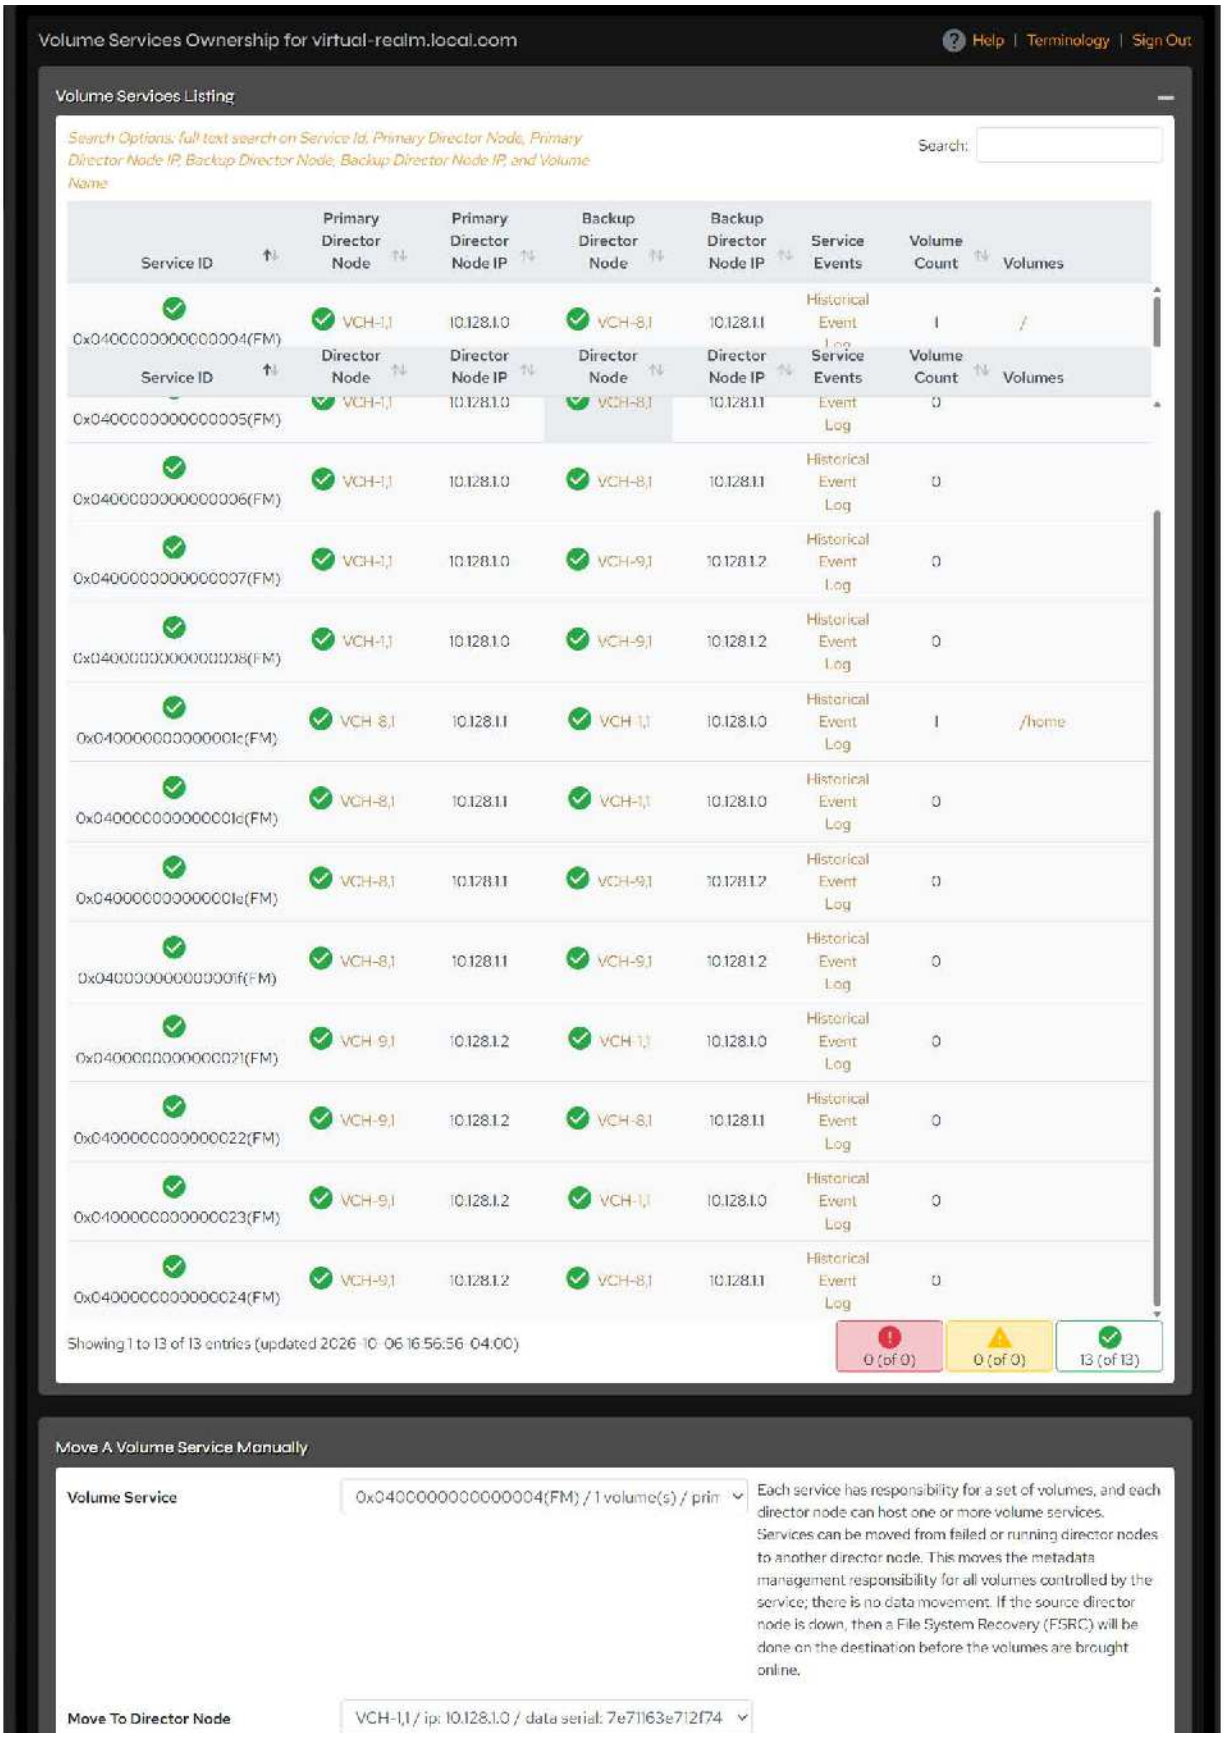

Volume Services Ownership for virtual-realm.local.com
Help | Terminology | Sign Out
Volume Services Listing
Search Options: full text search on Service Id, Primary Director Node, Primary Director Node IP, Backup Director Node, Backup Director Node IP, and Volume Name
Search:
Service ID
Primary Director Node
Primary Director Node IP
Backup Director Node
Backup Director Node IP
Service Events
Volume Count
Volumes
0x0400000000000004(FM) VCH-1,1 10.128.1.0 VCH-8,1 10.128.1.1 Historical Event Log 1 /
Service ID
Director Node
Director Node IP
Director Node
Director Node IP
Service Events
Volume Count
Volumes
0x0400000000000005(FM) VCH-1,1 10.128.1.0 VCH-8,1 10.128.1.1 Event Log 0
0x0400000000000006(FM) VCH-1,1 10.128.1.0 VCH-8,1 10.128.1.1 Historical Event Log 0
0x0400000000000007(FM) VCH-1,1 10.128.1.0 VCH-9,1 10.128.1.2 Historical Event Log 0
0x0400000000000008(FM) VCH-1,1 10.128.1.0 VCH-9,1 10.128.1.2 Historical Event Log 0
0x040000000000001c(FM) VCH-8,1 10.128.1.1 VCH-1,1 10.128.1.0 Historical Event Log 1 /home
0x040000000000001d(FM) VCH-8,1 10.128.1.1 VCH-1,1 10.128.1.0 Historical Event Log 0
0x040000000000001e(FM) VCH-8,1 10.128.1.1 VCH-9,1 10.128.1.2 Historical Event Log 0
0x040000000000001f(FM) VCH-8,1 10.128.1.1 VCH-9,1 10.128.1.2 Historical Event Log 0
0x0400000000000021(FM) VCH-9,1 10.128.1.2 VCH-1,1 10.128.1.0 Historical Event Log 0
0x0400000000000022(FM) VCH-9,1 10.128.1.2 VCH-8,1 10.128.1.1 Historical Event Log 0
0x0400000000000023(FM) VCH-9,1 10.128.1.2 VCH-1,1 10.128.1.0 Historical Event Log 0
0x0400000000000024(FM) VCH-9,1 10.128.1.2 VCH-8,1 10.128.1.1 Historical Event Log 0
Showing 1 to 13 of 13 entries (updated 2026-10-06 16:56:56 -04:00)
0 (of 0)
0 (of 0)
13 (of 13)
Move A Volume Service Manually
Volume Service
0x0400000000000004(FM) / 1 volume(s) / prin
Each service has responsibility for a set of volumes, and each director node can host one or more volume services. Services can be moved from failed or running director nodes to another director node. This moves the metadata management responsibility for all volumes controlled by the service; there is no data movement. If the source director node is down, then a File System Recovery (FSRC) will be done on the destination before the volumes are brought online.
Move To Director Node
VCH-1,1 / ip: 10.128.1.0 / data serial: 7e71163e712f74



Figure 5-14. Volume Services Ownership page.

Select the service you want to move from the **Volume Service** field. In the **Move To Director Node** field, select the new Director location for the service. When ready, click the **Move Service** button to transfer management of the volume service to the selected Director. If **Automatic Volume Service Failover** is configured, you can also move the backup that was created by clicking the **Move Backup** button.

If all volume services have been removed from a node in the past, the system disables volume service creation for the source Director. A list of those Directors for which volume service creation has been disabled is shown on **Storage > Volumes > Service Ownership** page in the **Enable Automatic Volume Service Creation** table. To re-enable a Director for automatic volume service creation, mark the **Enable** checkbox, then click **Enable Automatic Volume Service Creation for Selected Node(s)**.

> **WARNING:** If a volume service is moved from a failed Director, the volume will run FSRC before coming online. This may take considerable time depending on the number of files in the volume. If the source Director is online, no FSRC is required.

Enable Automatic Volume Service Creation

A list of Directors off of which all volume services have been moved appears at the bottom of the **Storage > Volumes > Service Ownership** page in the **Enable Automatic Volume Service Creation** table. To re-enable a Director for automatic volume service creation, mark the **Enable** checkbox, click **Enable Automatic Volume Service Creation for Selected Node(s)**, and click **Save Changes**.

5.4 Volume Data Compression

> **NOTE:** This feature is currently only available as a PanCLI command. The ability to manage volume compression in PanActive Manager will be included in a forthcoming release.

VDURA performs data compression on the Storage Nodes, ensuring zero impact on client-side CPU or memory resources. Unlike architectures that shift compression tasks to the client, consuming valuable compute and memory, VDURA handles all compression operations within the storage layer itself. This design keeps GPU and application nodes fully dedicated to user workloads, maximizing performance and system efficiency. The data compression feature can be toggled on or off at any time via PanCLI.

Use the **volume compression** PanCLI command to display whether compression is currently enabled or disabled on a volume and to enable or disable data compression on a volume.

View compression status for a volume:

```
volume compression [volumename]
```

If no option is given, the command displays whether compression is currently enabled or disabled for the specified volume.

Example usage:

```
volume compression vol1
```

Change compression status for a volume:

```
volume compression [volumename] ?enable|disable?
```

To specify a volume, use its name. Enabling compression may improve storage efficiency but could affect performance, depending on workload and hardware.

Disabling compression will stop further data from being compressed, but existing compressed data will remain compressed until it is rewritten.

Example usage:

```
volume compression vol1 enable
volume compression vol1 disable
```

5.5 Configure Erasure Coding with a RAID Layout

Instead of relying on hardware RAID controllers that protect data at a drive level and computing RAID on the disks themselves, PanFS implements patented client-based per-file object erasure coding. Implemented in software, erasure coding is used across distributed Storage Nodes with RAID layouts, resulting in increased data reliability, fault tolerance, and performance.

> **NOTE:** PanFS provides unique file-level flexibility enabling files, including a single file, in the same StorageSet, volume, or even directory to have different erasure-coded RAID levels and other specific layout properties.

Because of the way data is distributed with PanFS erasure coding, reliability increases with scale rather than decreasing, as is the case with traditional storage systems. In addition, VDURA uses a distributed hot-sparing model for added recoverability, where the space needed to recover from a failure is distributed among all of the Storage Nodes in the StorageSets, rather than being centralized on a single spare node.

There is no restriction to the setting of a specific RAID layout for files created in a volume or directory. That is, the user is free to set the RAID layout on each file as it is created, no matter what the default value might be, or whether that default is inherited from the volume or from the directory.

> **NOTE:** A file can be seen as a single virtual object that is sliced into multiple component objects, where one component object is stored on each Storage Node. Components are created only as needed, as the file is written. See Overview of the VDURA System for more information.

Users have three RAID levels available: RAID 6+, RAID 5+, and RAID 10+. As described in the Extended File System Availability (EFSA) section below, the "+" after the RAID level indicates that the redundancy of directories exceeds the redundancy of files in the volume by one. Thus, RAID 5+ and RAID 10+ have triple-replicated directories, and RAID 6+ has quadruple-replicated directories.

In addition, small virtual objects are simply double-replicated in RAID 5+ and RAID 10+ and triple-replicated in RAID 6+ for efficiency. As an object extends in size, it is automatically converted to either a RAID 5+ or 6+ layout, depending on configuration preferences.

Thus, you can mix RAID 10+, RAID 5+, and RAID 6+ levels, as well as any legacy volume layouts, together in the same StorageSet. When the system suffers failures, the availability status of each file is evaluated independently for reconstruction.

> **NOTE:** Volumes that existed in a realm prior to PanFS OS 6.0.0 continue to use their current redundancy for directories; they cannot be converted to a RAID+ volume level.

5.5.1 Extended File System Availability (EFSA)

Extended File System Availability (EFSA) is a proprietary VDURA feature that provides increased availability of a RAID-level volume for reading and writing when the RAID-level protection of the volume has been exceeded. With EFSA, the directory tree is replicated multiple times, allowing the volume to

remain fully intact even when a failure exceeds the RAID level of that volume. VDURA notates EFSA-enabled volumes with a "+" after the RAID level.

For example, with four or more Storage Nodes, RAID 6+ directories are quadruple replicated, while the directories in all other volume configurations are triple replicated. Thus with EFSA, PanFS RAID 6+ level volumes with four or more Storage Nodes remain available even after three failures, while all other volumes configurations remain available after two failures.

PanFS volumes that have exceeded their availability levels are listed as degraded in management interface tools. Appropriate alerts indicate to the administrator that the volume has exceeded its availability level and requires attention; the PanFS system does not attempt to automatically reconstruct volumes that have exceeded their fault tolerance levels. If reconstruction is in progress when a RAID+ volume exceeds its fault tolerance level, it is stopped. If any of the above volume types suffer additional failures, they become unavailable.

EFSA is only available in PanFS RAID 6+, RAID 5+, and RAID 10+ volume layouts; it is not supported on legacy volume layouts.

Your system must have at least four Storage Nodes to use EFSA. With four or more nodes, the system will have quadruple-replicated directories.

5.5.2 RAID 6+

RAID 6+ is the default RAID layout for all volumes. RAID 6+ volumes provide a balance between performance, capacity overhead, and RAID protection. RAID 6+ volumes provide protection against one additional Storage Node failure than RAID 6 or against three simultaneous whole-Storage-Node failures. RAID 6+ also provides approximately a 150x improvement in reliability over RAID 5+. In the case of a catastrophic failure (e.g., three simultaneous Storage Node failures in RAID 6+ layout with a minimum of four nodes), EFSA keeps the file system online and accessible, except for the small fraction of specific files affected by the three failures. In most cases, it is possible for PanFS to provide the storage administrator with a list of affected files, including full path names. This limits disaster recovery to restoring a specific list of damaged files, as opposed to the entire file system. If the user sets a lower-reliability RAID layout on any file(s) in this volume, there may be data loss even with two VPOD failures.

5.5.3 RAID 5+

RAID 5+ layout volumes are optimized for performance and capacity overhead. With RAID 5+ and EFSA, volumes have protection against one additional Storage Node failure than with RAID 5 or two Storage Node failures.

5.5.4 RAID 10+

Users can also create a volume that uses a RAID 10+ layout as the default RAID level. RAID 10+ combines striping and mirroring to provide high performance for applications that perform small random writes, while at the same time providing support for EFSA and resiliency against a single disk failure. Unlike RAID 6+ or RAID 5+, which require more operations to perform writes smaller than one stripe unit, RAID 10+ always performs only two I/O operations regardless of the size of the write. However, with a 100 percent overhead, data stored with RAID 10+ uses more space for redundancy than data stored with RAID 5+ or RAID 6+.

As a result, a RAID 10+ level yields significant performance gains for small random write workloads and for n-to-1 write workloads that are not page- or stripe-aligned. However, this comes with a significant decrease in performance for sequential read and write workloads.

5.6 Horizontal Parity

VDURA offers horizontal parity for end-to-end data redundancy and integrity. Horizontal parity is essentially the traditional use of RAID, where either data is duplicated on multiple independent disk drives or some extra amount of parity data is stored, allowing the data to be rebuilt in the face of a specified number of failures. RAID 10+, RAID 6+, and RAID 5+ layouts are examples. This concept is similar to erasure coding with RAID levels, where data is duplicated on multiple networked Storage Nodes.

As in a traditional RAID system, data is distributed across multiple independent disks. When one or more of those disks fails, or when a single sector on one of those disks cannot be read, the data can be recreated by reading the data and parity from each of the other disks, reconstructing the missing data and/or parity, and writing the missing data/parity out to an available disk not already holding data/parity for the file.

5.7 Multi-Level Erasure Coding

As part of overall system functionality, PanFS uses per-file erasure coding, which is performed at the client. The client determines the number of RAID stripe components necessary, either using cluster-wide defaults or another specification for that file object.

Multi-Level Erasure Coding™ (MLEC) refers to a second level of erasure coding that is done inside a VPOD. Each VPOD consists of enough disks to be able to stripe the data given to it for storage across these disks. Using RAID 5, the VPOD will calculate the local data and parity components of a stripe and will write them as needed. This allows a VPOD to be able to recover from a physical drive failure and continue operating by calculating contents of the missing drive.

MLEC can be configured for use during StorageSet creation or when storage is first being activated; it cannot be adjusted afterwards. However, since MLEC is set per StorageSet, StorageSets with different MLEC settings can coexist in the same realm.

Enabling this second level of RAID (erasure coding) will consume some amount of performance and will reduce usable capacity by one drive per VPOD.

Failed drives will have an indicator LED turned on. The failed drive can be pulled and replaced with a new drive. Reconstruction internal to the VPOD will then automatically start.

5.8 User and Group Quotas

> **IMPORTANT:** Do not confuse user and group quotas with volume quotas. Volume quotas are assigned to each volume when it is created as a way to track or limit the volume size. User and group quotas are assigned to limit or track usage by users and groups.

Because storage is usually a shared resource between multiple users and groups, administrators need to be able to define how the resource is divided between various users. User and group quotas accomplish this by allowing the administrator to limit the amount of storage space and/or number of files a specific user or group may use per volume. Using the following options, the administrator can also define what will occur if a user exceeds the quota limit:

- **Assign a soft quota limit**: When a soft quota limit is set, an email alert is sent to listed recipients when the soft quota has been reached, but the user or group can continue to create files and/or use storage space beyond the soft quota limit.
- **Assign a hard quota limit**: When a hard quota limit is set, an email alert is sent to listed recipients when the hard quota has been reached, and the user or group is prohibited from creating files or using storage space beyond the hard quota limit.
- **Track usage only**: Instead of setting quota limits, the administrator can configure either system level or volume level tracking to only monitor usage. The results are available as reports.

User and group quotas track the number of files and directories and the capacity consumed by those files and directories for each user and group. The capacity includes the RAID parity overhead used by object RAID for the files and directories in the volume. Parity overhead will vary depending on the volume layout you choose (RAID 6+, RAID 5+, or RAID 10+) and the types and sizes of files that exist in the volume.

The system-level user quota feature is disabled by default, since there is a slight decline in performance when user quotas are enabled, even if quotas are set to **Tracked**.

Before user quota limits or monitoring can be activated for a volume or system, you must enable them on that volume or system from the **Policy Configuration** page.

5.8.1 Configure User and Group Quotas

When enabled, **User and Group Quotas Enforcement Policy** can be configured either to enforce a policy or simply to track user quotas for reporting purposes. When the **User and Group Quotas Enforcement Policy** configuration is set at the system level, all volumes will inherit the system default settings unless overrides are specified.

The root volume cannot be configured to enforce user quotas. It is automatically set to **Disabled** and cannot be changed.

Quota limits cannot be modified while the user quotas policy is disabled. If you accidentally update the limits in these circumstances, you should deploy a temporary limit (for example, by updating the soft quota), and then redeploy the actual limit.

Group quota hard byte limits can be overrun. It is possible to successfully change group ownership of a set of files to another group, thereby placing the new group over its group quota hard byte limit.

Click **Storage > User & Group Quotas > Policy Configuration**.

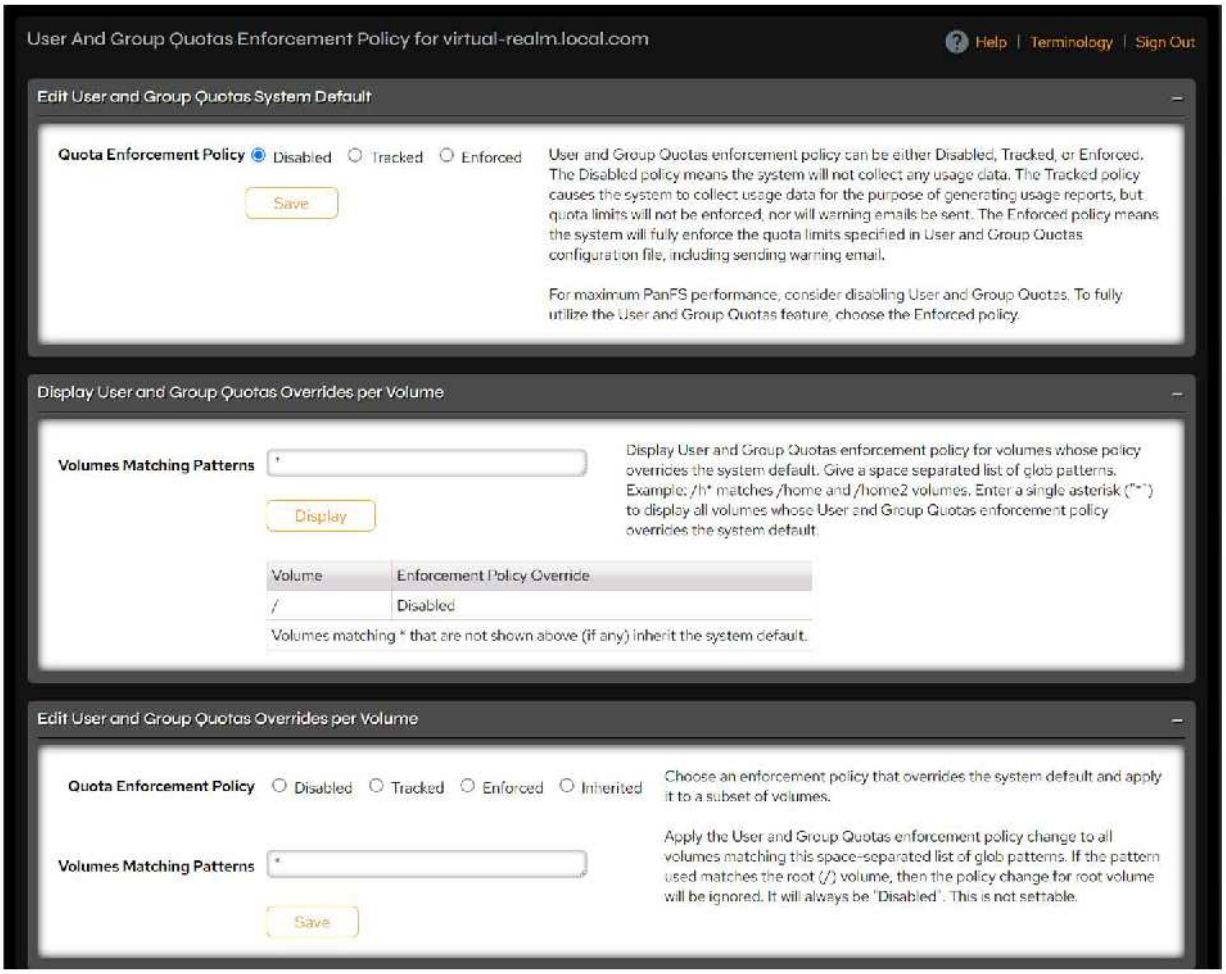


Figure 5-15. User and Group Quotas Enforcement Policy page.

Edit Default Quota Enforcement Policy

Edit the **User and Group Quotas System Default** field's **Quota Enforcement Policy,** if desired, and click **Save**.

- **Disabled** will disable user quotas on all volumes in the system. You may then enable user quotas on specific volumes in the **Overrides** section. For maximum performance, consider disabling user and group quotas.
- **Tracked** collects usage data but does not enforce any limits. You may then modify this setting on specific volumes in the **Overrides** section.
- **Enforced** enables user quotas on all volumes in the system. You may then disable user quotas on specific volumes in the **Overrides** section.

Refine Volume-Level Quota Settings

You can override the system default settings above with volume-level settings. For example, to apply user quotas to only certain volumes, set the system default to either **Enforced** or **Disabled**, then specify the volumes you wish to include or exclude in the **Edit User and Group Quotas Overrides per Volume** text box.

Display current volume user quota settings by entering the volumes you want to display in the **Display User and Group Quotas Overrides per Volume** text box and clicking the **Display** button. The current policy settings will be listed for volumes matching the specified pattern. To change these settings, enter the volumes you want to change in the **Edit User and Group Quotas Overrides per Volume** text box, select the radio button for the appropriate setting, and click **Save**.

The **Tracked** option collects usage data, but does not enforce any limits. Choosing the **Inherited** option causes the specified volumes to inherit the system default setting. The **Inherited** option is useful if you have at some point chosen to override the system default setting on a volume and then want to return that volume to the system default.

Set Individual User and Group Quotas

Click **Storage > User & Group Quotas > Limits Configuration**.

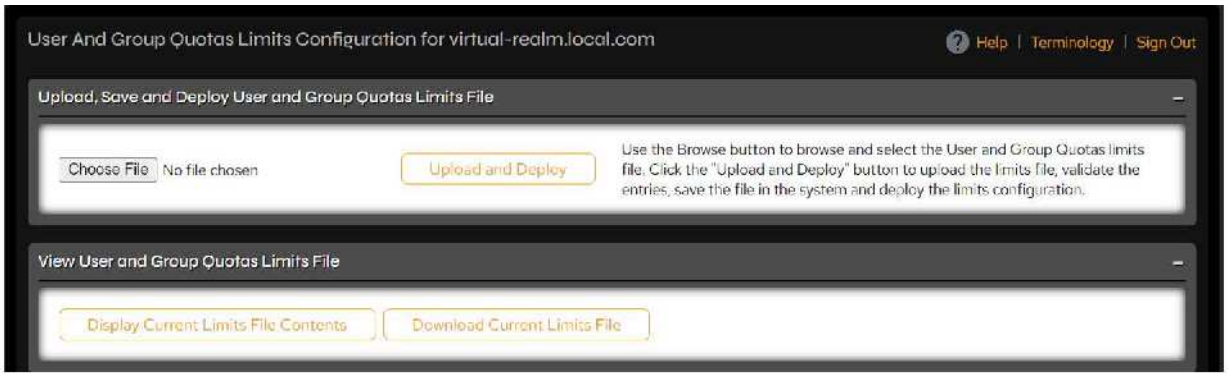


Figure 5-16. User and Group Quotas Limits Configuration page.

You have three options to enter user and group quota information:

- Upload a preconfigured limits file of user and group quotas by browsing to the limits file and clicking the **Upload and Deploy** button. See the table below for a description of the syntax, rules, and options for the **User and Group Quotas Limits** file.

- Click the **Display Current Limits File Contents** button to show the current quota file in the text box. If there is none, the syntax required for the quota limits file is displayed. Make your changes using the syntax, rules, and options described in the table below. Click **Save and Deploy Changes** to implement the quota limits.
- Click **Download Current Limits File** to download an editable version of the file. Make your changes using the syntax, rules, and options described in the table below and save it locally when finished. Browse to it in the **Upload, Save and Deploy User and Group Quotas Limits File** box, then click **Upload and Deploy**.

Example limits file:

```
user aadams /Vol1 100M 100M 0 0 admin@yourcompany.com =
user ssmith /Vol1 50M 1T 0 0 admin@yourcompany.com =
user ddoe /Vol1 0 0 2 3 admin@yourcompany.com =
group users /Vol2 10M 20M 5 10 admin@yourcompany.com =
```

Syntax:

<User/Group> <ID> <Volume> <SoftByteLimit> <HardByteLimit> <SoftFileLimit> <HardFileLimit> <SQViolationEmail> <HQViolationEmail>

Syntax rules:

- One entry per line.
- Separate entries for user quotas and group quotas.
- Separate entries for system level and volume level defaults.
- All fields except email address fields are required. If email address field is missing, it will default to **none**.
- Delimiter is whitespace.
- Any field containing a space must be nested in double quotes (“ ”).
- For a multiple line entry, use the escape “\” character at end of line.

The table below explains how to use the fields.

Field Option	Usage
<User/Group>	• User: keyword to set quota for user • Group: keyword to set quota for group

Field Option	Usage
<ID> Example: jjones	• User or group ID • An asterisk (*****) indicates all users or groups. Use this to set volume-level defaults. • If the value is a text string and prefixed with a domain (e.g., **domain1\user_1**), it will be treated as a Windows user or group name. Otherwise, it will be treated as a Unix user or group name. • If the value starts with **S-**, it will be considered as a Windows SID. • If the value starts with either **uid:** or **gid:** followed by a number, it will be treated as a Unix user or group ID. • If the value starts with either **user_sid:** or **group_sid:** followed by a SID, it will be treated as a Windows user or group SID.
<Volume> Example: /Vol1	• Volume name • An asterisk (*****) indicates all volumes. This will set the system default. • If the volume is an asterisk, the ID field must also be an asterisk. You cannot set a default for a user across all volumes. This must be set with an entry for each volume.

Field Option	Usage
<SoftByteLimit> <HardByteLimit> <SoftFileLimit> <HardFileLimit> Example: 100M 100M 0 0	• Use a hyphen (**-**) to indicate **inherit** and a zero (**0**) to indicate **unlimited**. • Use up to two decimal places for byte limits. Use only whole numbers for file limits. Use a minimum of 10 M for byte limits. • Use suffixes K, M, G, T or P to indicate kilo-, mega-, giga-, tera-, or peta-bytes (power of 10): ex: 800K, 200M, 10G 12.8T etc. **K** suffix will multiply the number by 1,000 **M** suffix will multiply the number by 1,000,000 **G** suffix will multiply the number by 1,000,000,000 **T** suffix will multiply the number by 1,000,000,000,000 **P** suffix will multiply the number by 1,000,000,000,000,000 • For quota bytes, the minimum is 10M and the maximum is approximately 9223 PB. Use a zero (**0**) for unlimited, which is approximately 9223 PB. • For quota files, the minimum is 1 and the maximum is approximately 9223 P. Use a zero (**0**) to indicate **unlimited**, which is approximately 9223 P files.

Field Option	Usage
<SQViolationEmail> <HQViolationEmail> Example: admin@yourcompany.com	Email address for **soft_quota_violation_notification**. Email address for **hard_quota_violation_notification**. • Comma-separated list of email IDs or may be empty. • Use **$user** for Unix email address of user. • Use $LDAP@<yourcompany.com> for email address to be fetched from LDAP. • Use a hyphen (**-**) to inherit from volume. • If email ID contains spaces, they must be nested in double quotes (" "). • Use an equals character (**=**) character by itself as the field value if the email list entered for the other email field is to be used for this also. Do not use an equals character in both email fields in the same entry. • Use the word **none** to explicitly indicate no email is to be sent.

Field Option	Usage
Additional Examples	• Set system default quota and email address list for users: `user * * 100G 110G 500K 600K $LDAP@<yourcompany.com>,admin@<yourcompany.com> -` • Set system default quota and email address list for groups: `group * * 500G 550G 800K 2M - -` • Set default quota for users for /home1 volume: `user * /home1 100G 110G 200K 300K $user $user` • Set default quota for groups for /home2 volume: `group * /home2 0 0 600K 700K - -` **Note**: Zero (**0**) is used to indicate unlimited limit for soft and hard byte limits. • Set user specific quota for usr_1 in /home2 volume: Unix login name usr_1: `user usr_1 /home2 20G 25G 200K 250K admin@company.com \ guest@company.com,admin@company.com` Windows login name: `user mydomain\usr_1 /home2 20G 25G 200K 250K admin@company.com \ guest@company.com,admin@company.com`

The system will automatically validate the **User and Group Quotas Limits** file before it deploys it. If the file is valid, it will be saved and, if possible, deployed. If the file contains errors, they will be displayed ten at a time so they can be corrected and the changes deployed.

Administrators may not be able to deploy a file if a previous deploy operation is in progress or other failures are reported. If the file was previously validated by the system, it has already been saved so there is no need to upload the file again.

If the deploy operation has failed to deploy the **User and Group Quotas Limits** file, the user will see a **Redeploy** button. When pressed, the **Redeploy** button will deploy a previously saved **User and Group Quotas Limits** file.

Edit Quota Thresholds

By default, the system clears the soft quota violations as soon as storage space and/or the number of files are under the soft quota limits. However, hard quota violations are cleared when they are sufficiently under the hard quota limits (default setting is five percent).

You can configure the thresholds for soft and hard quota limits through the CLI in service mode as follows:

- Hard quota threshold percentage for storage space consumption:

```
[pancli_service] pan_rm_configdb_util -s
uint16:smm_clear_hard_quota_thrshld:5
```

- Hard quota threshold percentage for number of files:

```
[pancli_service] pan_rm_configdb_util -s
uint16:smm_clear_hard_quota_files:5
```

- Soft quota threshold percentage for storage space consumption:

```
[pancli_service] pan_rm_configdb_util -s
uint16:smm_clear_soft_quota_thrshld:0
```

- Soft quota threshold percentage for number of files:

```
[pancli_service] pan_rm_configdb_util -s
uint16:smm_clear_soft_quota_files:0
```

5.8.2 Common User Quota Limit Configurations

The table below displays some common user quota limit configurations.

Configuration	**Steps**
No user or group quotas and no tracking	On the **User & Group Quotas > Policy Configuration** page, select **Disabled** as the system default.

Configuration	Steps
Collect usage data for tracking purposes only	On the **User & Group Quotas > Policy Configuration** page, select **Tracked** as the system default.
Apply the same user and group quota limits to all users and groups on all volumes	On the **User & Group Quotas > Policy Configuration** page, select **Enforced** as the system default. On the **User & Group Quotas > Limits Configuration** page, create a **User and Group Quotas Limits** file.
Enforce user or group quotas only for a few volumes on the system	On the **User & Group Quotas > Policy Configuration** page, select **Disabled** as the system default. Under **Edit User and Group Quotas Overrides per Volume**, enter the volumes on which you want to enforce quotas in the **Volumes Matching Patterns** text box, then select **Enforced** or **Tracked** as the **Quota Enforcement Policy**. On the **User & Group Quotas > Limits Configuration** page, create a **User and Group Quotas Limits** file specifying the limits for the volumes where user quotas will be enforced. **NOTE**: A limits file does not need to be created for **Tracked** volumes.

Configuration	Steps
Enforce user or group quotas for most volumes on the system, but disable user quotas on a few volumes	On the **User & Group Quotas > Policy Configuration** page, select **Enforced** or **Tracked** as the system default. Under **Edit User and Group Quotas Overrides per Volume**, enter the volumes on which you want to disable quotas in the **Volumes Matching Patterns** text box, then select **Disabled** as the **Quota Enforcement Policy**. On the **User & Group Quotas > Limits Configuration** page, create a **User and Group Quotas Limits** file specifying the limits for the volumes where user quotas will be enforced. **NOTE**: A Limits file does not need to be created for **Tracked** volumes.

5.8.3 Monitor User and Group Quotas

Click **Storage > User & Group Quotas > Usage Report**.

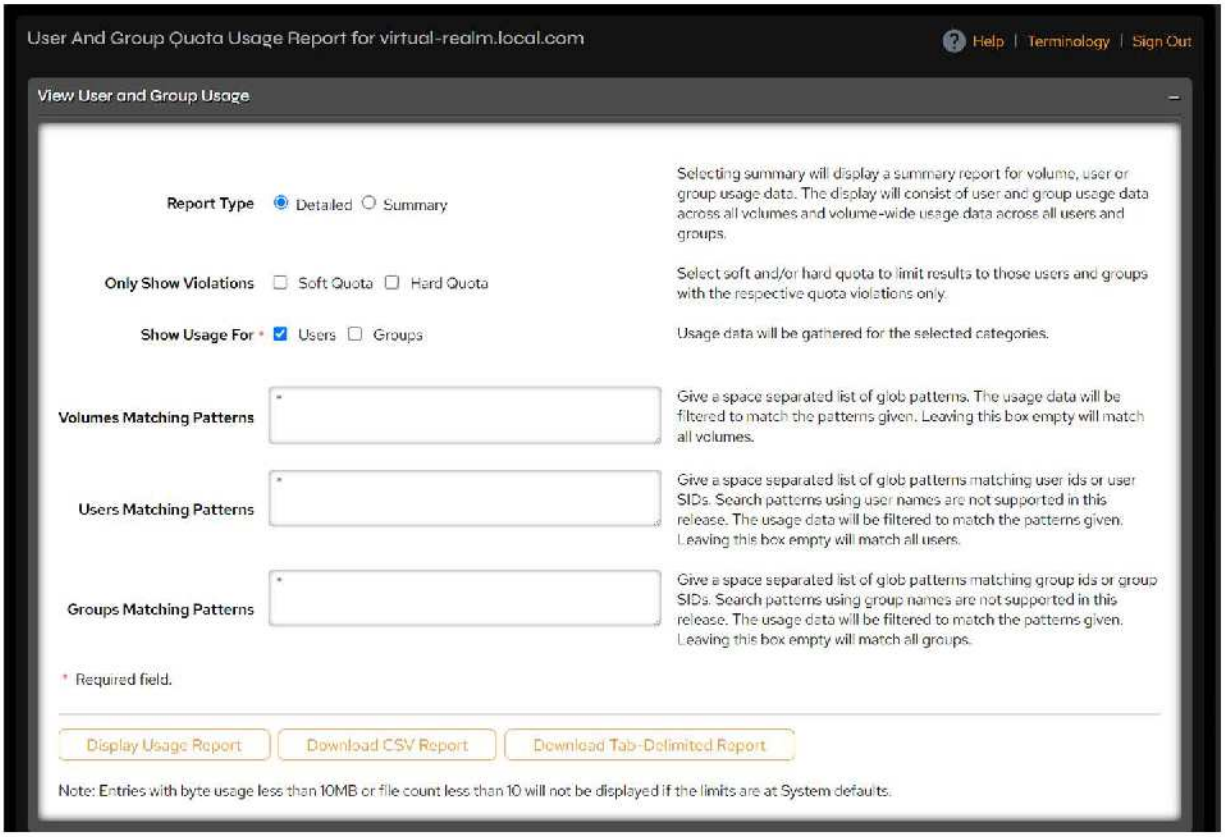


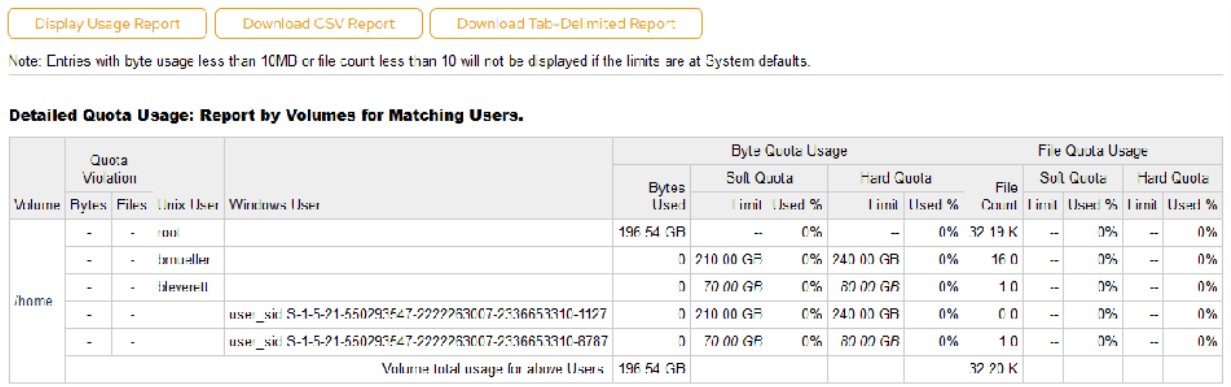
Figure 5-17. User and Group Quota Usage Report page.

Select the controls for the **Quota Usage Report** you wish to display.

Control	Description
Report Type	Selecting **Summary** displays a summary report by volume, user, or group usage data. The display consists of user and group usage data across all volumes and volume-wide usage data across all users and groups. Selecting **Detailed** displays a detailed report by volume for matching users or group usage data. The report includes byte and file quota usage data, including bytes used and file count, soft and hard quota limits and percent used, and quota violation information by user and volume.
Only Show Violations	Select soft and/or hard quota to limit results to those users and groups with the respective quota violations only. Select both to show both. To show all usage data regardless of whether there are any quota violations, remove any check marks.
Show Usage For	Usage data will be displayed for the selected categories.
Volumes Matching Patterns	Enter a space-separated list of glob patterns. The usage data will be filtered to match the patterns given. Leaving this box empty will match all volumes.
Users Matching Patterns	Enter a space-separated list of glob patterns matching user IDs or user SIDs. Search patterns using user names are not supported in this release. The usage data will be filtered to match the patterns given. Leaving this box empty will match all users.

Control	Description
Groups Matching Patterns	Enter a space-separated list of glob patterns matching group IDs or group SIDs. Search patterns using group names are not supported in this release. The usage data will be filtered to match the patterns given. Leaving this box empty will match all groups.

View the **Quota Usage Report** by clicking the **Display Usage Report** button located at the bottom of the **Quota Usage Report** screen.


Display Usage Report | Download CSV Report | Download Tab-Delimited Report

Note: Entries with byte usage less than 10MB or file count less than 10 will not be displayed if the limits are at System defaults.

Detailed Quota Usage: Report by Volumes for Matching Users.

Volume	Quota Violation			Windows User	Byte Quota Usage					File Quota Usage				
					Bytes Used	Soft Quota		Hard Quota		File Count	Soft Quota		Hard Quota	
	Bytes	Files	Unix User			Limit	Used %	Limit	Used %		Limit	Used %	Limit	Used %
/home	-	-	root		196.54 GB	--	0%	--	0%	32.19 K	--	0%	--	0%
	-	-	[illegible]		0	210.00 GB	0%	240.00 GB	0%	16.0	--	0%	--	0%
	-	-	[illegible]		0	70.00 GB	0%	80.00 GB	0%	1.0	--	0%	--	0%
	-	-		user_sid S-1-5-21-550293547-2222263007-2336653310-1127	0	210.00 GB	0%	240.00 GB	0%	0.0	--	0%	--	0%
	-	-		user_sid S-1-5-21-550293547-2222263007-2336653310-8787	0	70.00 GB	0%	80.00 GB	0%	1.0	--	0%	--	0%
				Volume total usage for above Users	196.54 GB					32.20 K				

Figure 5-18. Usage Report output.

- Click the **Download CSV Report** button to download the specified Quota Usage Report to your computer in a CSV file format.
- Click the **Download Tab-Delimited Report** button to download the specified Quota Usage Report to your computer in a tab-delimited file format.

If either the system or the root volume is offline, requests for User Quota Usage Reports will be rejected.

5.9 Enclosure Migration

Use the instructions below to migrate from one VDURA enclosure to another.

5.9.1 Before You Begin

Before beginning a migration, be aware of the following considerations:

- At least one new hardware enclosure must be installed before data can be migrated.
- Multiple Storage Nodes can be drained at the same time. However, for every Storage Node that will be drained, there must be an equivalent migration/spare Storage Node to receive the data.

- Migrations in different StorageSets are independent of each other.
- Migration cannot be used to reduce the number of Storage Nodes in a StorageSet or to move a volume from one StorageSet to another.
- It is possible to migrate more than one enclosure at a time if more than one enclosure has been added.

Take Realm Offline

1. Take the realm offline:

```
[pancli] syscontrol offline
Are you sure you want to take the system offline? (yes/no): yes
System offline received.
```

Taking the realm offline will pause all volumes until the system is brought back online.

2. Verify the realm is offline:

```
[pancli] syscontrol
The system is offline
The system is locked. New enclosures will not be automatically
discovered.
```

Power Off Nodes

Power off every node in the enclosure chassis to be replaced. Use the command below for each node in the enclosure to be replaced.

```
[pancli] syscontrol node {name or IP} halt
```

Example:

```
[pancli] syscontrol node VCH-1,7 halt
```

If a node in the enclosure cannot be reached over the network, or the command above does not work, connect to the node out-of-band. Log in and run halt:

```
[pancli] support servicemode yes
[pancli_service] halt -p
```

Once each node is powered off, proceed to the next step.

> **NOTE:** If you are unsure about how to remove/insert enclosures and their components, consult the *VDURA V5000 Certified Platform Installation and Service Guide*.

Empty Enclosure

Remove all components from the enclosure following normal hardware replacement procedures:

1. Remove the power cords from both power supplies.
2. Remove all network cables. Best practice is to label where each belongs (switch and port) to aid re-installation.
3. Remove each power supply, if needed.
4. Remove each node, labeling each with its appropriate physical slot number.
5. Remove the old enclosure from the rack.

Install Replacement Enclosure

Install the replacement enclosure into the rack and re-install the enclosure nodes and components into it.

1. Insert the new enclosure chassis into the rack.
2. Insert each power supply into the replacement enclosure, if needed.
3. Insert each node into its corresponding slot number in the replacement enclosure.
4. Insert the network cables in their appropriate node ports. Insert the power cables into their respective power supplies.

Monitor Nodes

Monitor the nodes in the enclosure to make sure they come online.

Once nodes power up and report the replacement enclosure to the realm president Director, the president will assign the replacement enclosure a new name. Find the new name:

```
[pancli] sysmap enclosures allcolumns
```

In this example, the failed enclosure was VCH-1. The realm has named the replacement enclosure VCH-5, which is the only enclosure shown in the sample output.

```
[pancli] sysmap enclosures allcolumns
NAME  SERIAL             STATUS NODES (ON/OFFLINE) STORAGESET STATIC LINK
                 DIRECTOR STORAGE              AGGREGATION
VCH-5 27b31362400001001 good   1/0      10/0      Set 1      disabled
```

Monitor the node state:

```
[pancli] sysmap nodes enclosure {enclosure name}
```

Example:

```
[pancli] sysmap nodes enclosure VCH-5
```

If all nodes do not come online, contact VDURA Technical Support before proceeding.

Rename Replacement Enclosure

Rename the replacement enclosure to that of the failed enclosure. Fill in the new enclosure's current and desired name:

```
[pancli] syscontrol enclosure {current name} rename {desired
name}
```

Example:

```
[pancli] syscontrol enclosure VCH-5 rename VCH-1
```

Verify the results:

```
[pancli] sysmap enclosures allcolumns
NAME  SERIAL             STATUS NODES (ON/OFFLINE) STORAGESET STATIC LINK
                 DIRECTOR STORAGE              AGGREGATION
VCH-1 27b31362400001001 good   1/0      10/0      Set 1      disabled
```

Bring Realm Online

Bring the realm back online:

```
[pancli] syscontrol online
System online received.
```

Verify the results:

```
[pancli] syscontrol
The system is online.
The system is locked. New enclosures will not be automatically
discovered.
```

Verify Online Status

Make sure all volumes are online.

```
[pancli] volume list
```

> **NOTE:** If the enclosure to be replaced contains only OSDs that are not in the StorageSet containing the root volume, it may be possible to take the volumes in that StorageSet offline and leave volumes in other StorageSets online. Contact VDURA Technical Support for this modified procedure.

> **IMPORTANT:** The following steps assume that you have installed the proper number of hardware enclosures using the procedures described above.

Add Migration Hardware to the StorageSet

From the **Storage > StorageSets** page, select the StorageSet to which you want to add the new enclosure. On the **StorageSet Detail** page, click the **Add Storage** button.

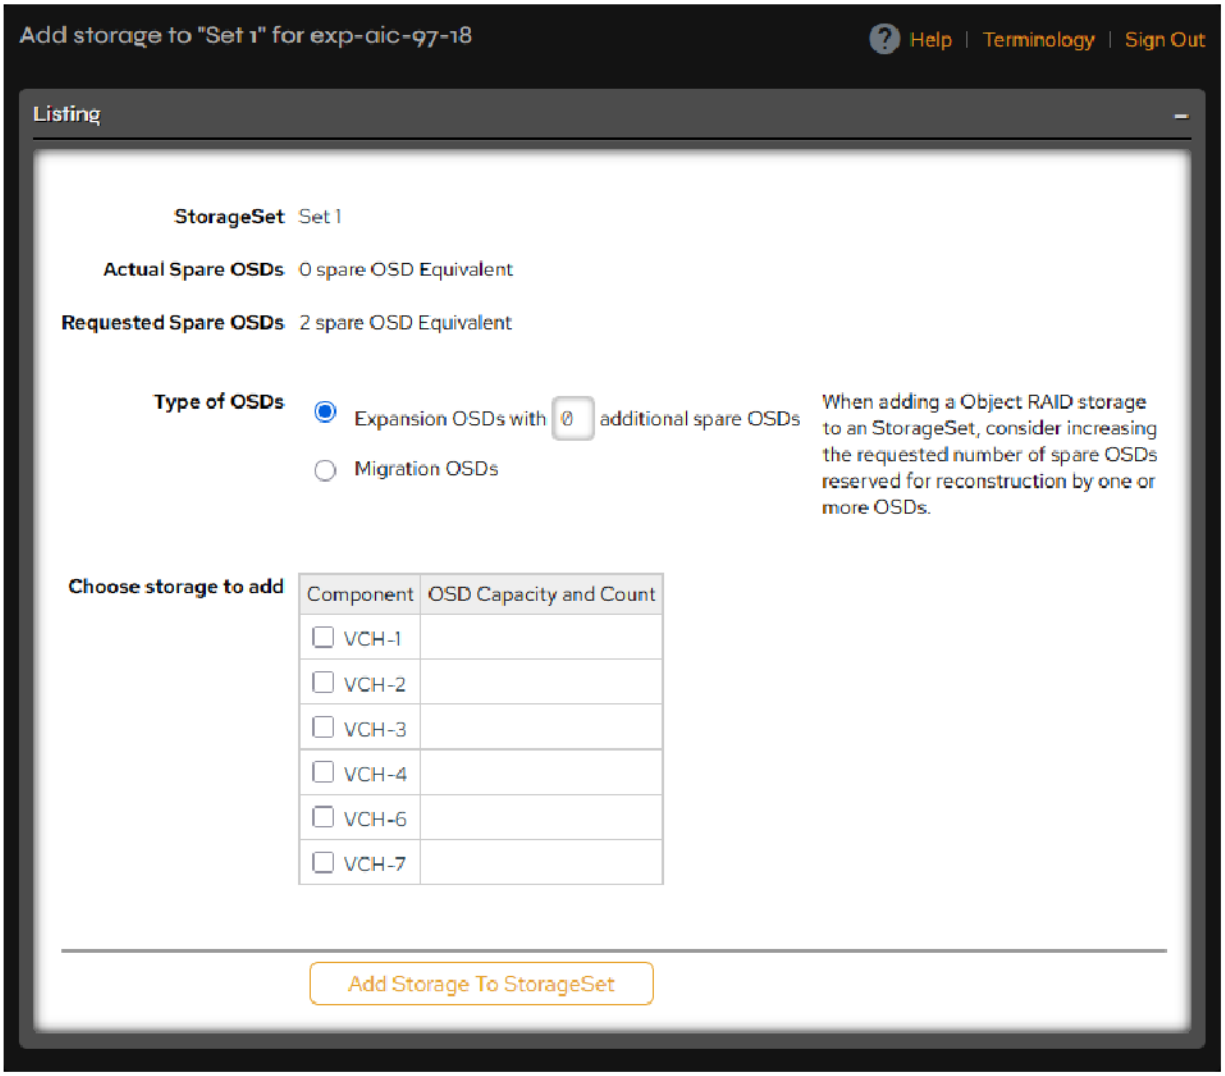


Figure 5-19. Add Storage page.

On the **Add Storage** page:

1. Under **Type of OSDs**, select **Migration OSDs**.
2. Next to **Choose Storage to Add**, check the box with the storage you will use for the migration.
3. Click the **Add Storage to StorageSet** button. You can verify that the enclosure was added to the correct StorageSet from the **StorageSet Details** page. Once sufficient hardware has been added, the migration process can begin.

Drain Storage Nodes

Drain the Storage Nodes in the enclosure from which you are migrating the data. From the **Details** section of the **Hardware > Hardware Management > Hardware Listing** page, click the name of the enclosure you want to migrate. At the bottom of its page, click the **Start Draining OSDs** button. You can monitor the progress of the storage drain from the **Hardware > Drain Status** page.

Move Director Services

After or during the drain, migrate the volume management (FM/SM) and realm management (RM) services to the new Director(s). From the **Storage > Volumes > Service Ownership** page, move the volume services to the new Director via the **Move A Volume Service Manually** field, then click **Move All Volume Services from Director**. If the Director also had realm management services, add the new Director to the repset and remove the old Director by going to the **Hardware > Realm Management** page, adding the new Director and clicking **Remove** in the **Remove from Replication Set** column in the row of the old Director.

Once the Storage Nodes have completed the drain, click the **Clear Record of Node** button from the OSD listed in the **Drained OSDs** section of the **Hardware > Hardware Management > OSD Drain Status** page. At this point, you can either safely remove the drained enclosure or you can reinitialize it by going to the **Hardware > Hardware Management > Hardware Listing** page and selecting the link for the enclosure you have drained. On the **Enclosure Management** page, click the **Reinitialize System** button.

The system will lock as part of the migration process. A message will appear stating that when the enclosure no longer appears, the user must unlock the system. Unlock the system from the **Hardware > Hardware Management > Hardware Listing** page by clicking the **Unlock System** button at the bottom of the page.

5.10 Realm Management

A set of Directors can be configured to host critical system services and replicate system configuration data. Replicated services include the primary system control service, the management console, and supporting services like DHCP and delegated DNS.

An odd number of Directors, greater than two, is required for redundancy. A three-Director configuration provides a single-fault tolerant system realm. A five-Director configuration provides a two-fault tolerant system realm.

When adding or removing Directors from the realm, it can take as long as five minutes for this information to propagate to all Directors in the cluster. Other actions during this intermediate stage, such as changing other system settings or failure of the president Director, may lead to hard-to-diagnose errors.

Click **Hardware > Realm Management**.

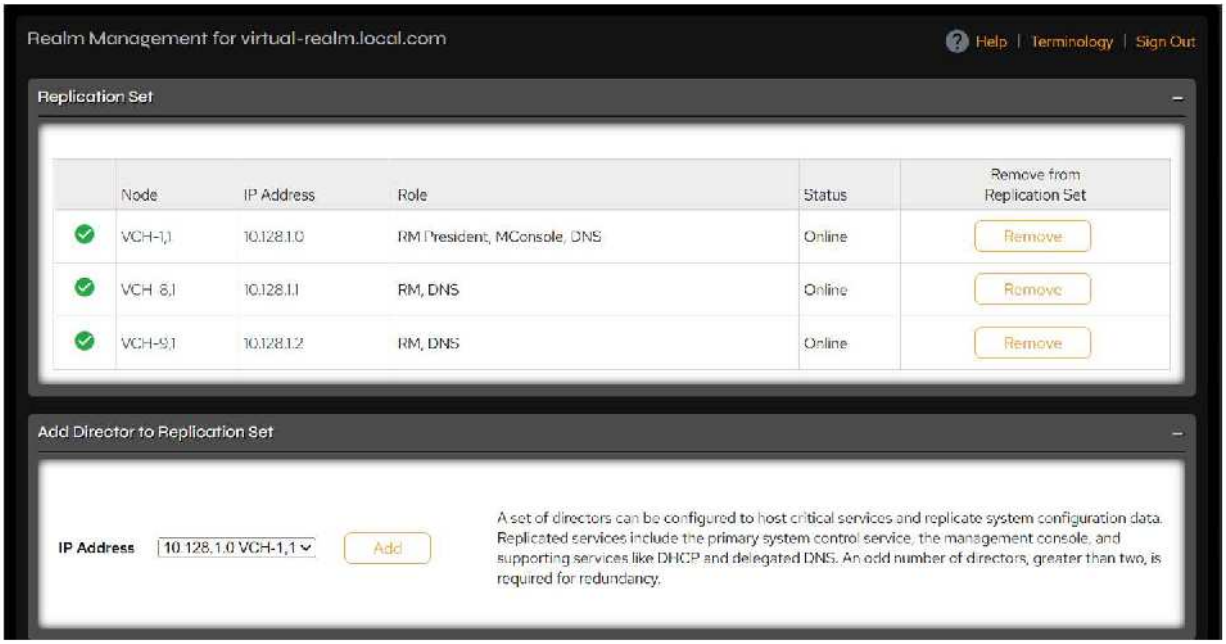


Figure 5-20. Realm Management page.

The **Realm Management** page includes the following information and controls:

Replication Set Table

Field	Description
Node	Enclosure containing the Director that is a member of realm management.
IP Address	IP address of the Director.
Role	Role taken by the Director in the realm.
Status	Status of the Director.
Remove from Replication Set	In the row containing the Director you want to remove, click **Remove** to remove it from the realm.

Add Director to Replication Set Table

Field	Description
IP Address	List of the IP addresses of all Directors in the system. Select a Director from this list to add it to the realm.
Add button	After you have selected a Director from the list, click the **Add** button to add it to the realm.

5.10.1 Add a Director to the Repset

Use the IP address drop-down list to select a Director you want to add to the repset. Click the **Add** button to add the Director to the repset.

If you do not have at least three Directors configured as members of the repset, a message will appear at the top of the screen telling you to add additional repset members.

Remove a Director from the Repset

In the row containing the Director you want to remove, click **Remove** in the **Remove from Replication Set** column.

6. Security Enhanced Linux in VDURA Systems

This chapter explains the VDURA implementation of Security Enhanced Linux (SELinux). This chapter includes the following topics:

- Terminology
- SELinux Overview
- PanFS and MAC
- Configuration and Use
- Important Considerations

6.1 Terminology

Access Control Mechanism (ACM) is a mechanism that can allow or prevent access by a user or a running process to a file system object such as a file or directory, based upon some criteria. Traditional UNIX permissions bits are a simple form of ACM, while Access Control Lists (ACLs) are a more sophisticated form.

Discretionary Access Control (DAC) is a class of security mechanisms for objects in a file system, such as files, directories and devices, that allows the owner of an object to decide who can access that object. Access to a file system object is granted or revoked at the discretion of the owner of the object. Traditional UNIX file permissions and ACLs are discretionary.

Mandatory Access Control (MAC) is a class of security mechanisms that restricts the level of control that users or running processes have over the file system objects that they create. Unlike in a DAC implementation, where users have full control over their own objects, MAC centralizes control of access decisions for those objects through policies defined by the system's administrator. It does this by adding additional labels, or categories, to all file system objects and checking whether a given user or process has been granted access to those categories by the system's administrator before they can interact with those objects.

Multi-Category Security (MCS) is an enhancement to SELinux that allows users to label files with categories. Only users who have been granted access to this category by the system's administrator can access files labeled with the category, assuming the existing DAC and TE rules also permit access.

Multi-Level Security (MLS) is a specific MAC security scheme. The "multi-level" concept comes from the defense community's security classifications of Confidential, Secret, and Top Secret forming an increasingly strict hierarchy. MLS is only relevant if a storage system is asked to store content from different levels in the security hierarchy at the same time. It is intended to prevent a user or process that is only allowed to access content in a lower security level from accessing anything in a higher security level.

Role-Based Access Control (RBAC) allows the system's administrator to establish roles for users and file system objects based on business-functional requirements or similar criteria instead of controlling access purely by permissions associated with each user. For example, an "Operator" role might have permission to backup the data on a storage system but not to change that system's configuration, while an "Administrator" role might have permission to review and change the configuration of the storage system but not to access the data it contains. Each user account on such a storage system might have one or more associated roles it could perform.

Type Enforcement (TE) is an access control mechanism found in MAC systems, where access decisions are made based on a **subject-access-object** set of rules that answer the question: "May <subject> do <action> to <object>?". For example: "May a web server (subject) read (action) files in user X's home directory (object)?".

6.2 SELinux Overview

Security Enhanced Linux (SELinux) refers to a Linux system that has been configured to enable a Linux Security Module (LSM) that is built into the Linux kernel. The SELinux subsystem in the kernel is driven by a security policy which is controlled by the system's administrator and loaded at system boot time. All security-relevant, kernel-level access operations on the system are intercepted by SELinux and examined in the context of the loaded security policy. SELinux answers the question: "May <subject> do <action> to <object>?". For example: "May a web server (subject) read (action) files in user X's home directory (object)?"

SELinux implements Mandatory Access Control (MAC). MAC requires that every process and system resource have a special security label called an SELinux context. (You may also see this referred to as an SELinux label, security label, or MAC label.) An SELinux context is an identifier that abstracts away the details of the security policy and focuses on the security properties of the entity. The SELinux policy uses these contexts in a series of rules that define how users and processes can interact with each other and various system resources. By default, the policy does not allow any interaction unless a rule explicitly grants access.

> **NOTE:** The VDURA SELinux implementation should not be confused with its other access control mechanisms. PanFS offers two different Discretionary Access Control mechanisms: traditional UNIX permissions bits and Access Control Lists (ACLs). This guide also provides information on those two mechanisms.

6.3 PanFS and MAC

Mandatory Access Control (MAC) systems bind security attributes such as SELinux security context labels to objects within a system. These attributes are used with other information in the system to make access control decisions. Today, MAC systems have been mainstreamed in modern operating systems such as Linux, Windows, macOS, and FreeBSD.

The PanFS storage system is MAC-aware and operates in Limited Server Mode, as defined in NFSv4.2 terminology found in IETF RFC 7204. PanFS will store and transmit labels but does not make any assumptions regarding their format or content and will not itself enforce security policies on file system objects. PanFS depends upon the client's operating system to interpret and enforce the policies that the security labels represent. It will inform client systems when a security label changes for a file the client has open or if it has a copy of that file's metadata. The client must apply its own policy to file system objects, using the security labels obtained from PanFS.

The following set of features is supported on Linux-based client systems with Linux kernel versions of at least 2.6.25:

- PanFS stores and retrieves MAC labels as requested by a client.
- MAC labels are coherent across multiple client systems just like file system data and metadata. PanFS will notify clients of changes to the MAC label for a file if that client has that file open or has a copy of the metadata for that file.
- PanFS supports mount options that are specific to SELinux (**context**, **rootcontext**, **fscontext**, **defcontext**) that set some default parameters used by SELinux.
- The same DirectFlow Client operates in both SELinux/non-SELinux environments.
- PanFS supports older DirectFlow Clients without MAC labeling support at the same time as it supports updated DirectFlow Clients that do support MAC labels. Customers are not forced to upgrade their client systems before they can upgrade to a PanFS version that supports SELinux.

6.4 Configuration and Use

SELinux requires security contexts be associated with file system objects to decide whether access is allowed or not, as defined by the security policy. A security context is stored as a string in an Extended Attribute (an **xattr**) with a predefined and fixed name of **security.selinux**.

Within SELinux, a security context is represented as a variable-length string that includes either three or four pieces of information: the SELinux user, role, type, and an optional MCS/MLS security level or range. These items are separated by colon characters:

```
user:role:type[:range]
```

- **User**: This string represents SELinux's identity for a given user. All SELinux policy and MAC mechanisms use this string, which may or may not match the user's normal login ID. Login IDs from LDAP or other identity management systems are mapped to SELinux user strings via a configuration file on the client system.
- **Role**: A role is associated with one or more object types the SELinux user can access. Roles determine what types of objects SELinux users can access.
- **Type**: When a type is associated with a process, it defines what processes or domains an SELinux user can access. When a type is associated with an object, it defines what access permissions the SELinux user has to that object.
- **Range or Level**: This feature is only applicable if the policy supports MCS or MLS. The entry can consist of either:
  - A single security level that contains a sensitivity level and zero or more categories, such as **s0** or **s1:c0, s7:c10.c15**.
  - A range that consists of two security levels (low and high) separated by a dash, such as **s0 - s15:c0.c1023**.

> **IMPORTANT:** SELinux policy rules are checked after DAC rules. SELinux policy rules are not used if DAC rules deny access first, which means that no SELinux denial is logged if the traditional DAC rules prevent the access attempt.

6.5 Important Considerations

To implement a secure environment using SELinux and PanFS, there are two configuration steps that a system administrator must take.

First, DirectFlow Client access must be restricted to only those IP addresses listed in the DFC exports file in PanFS. That step will prevent client systems that are not using SELinux or that do not respect the security labels from accessing the file system. Allowing other access paths would bypass the SELinux security policies entirely.

Second, remember that PanFS does not take part in access control decisions, but rather supports the SELinux security module in the client's kernel with label storage and retrieval. Since clients are responsible for enforcing security policy and access control based on MAC labels, it is imperative that all clients that can access PanFS use the same security policies. Specifically, they must all use the same sets of SELinux user, role, type, and range strings, and they must all have the same policies for how those strings are interpreted and applied.

7. Manage Storage Hardware

This chapter describes how to use PanActive Manager to manage VDURA hardware at the system, enclosure, and node level. It includes the following topics:

- Monitor and Manage the VDURA System
- Monitor and Manage Enclosures
- Monitor and Manage Nodes
- Network and NIC Status

For information on monitoring and managing hardware with the command-line interface, see the *VDURA PanCLI Reference Guide.* For information on installing hardware and the network configurations possible for your VDURA system, see the *VDURA V5000 Certified Platform Installation and Service Guide*.

7.1 Monitor and Manage the VDURA System

You can use PanActive Manager to monitor VDURA hardware at the system level. You can also perform a number of system-level tasks, including taking the system on- or offline, powering it down, and rebooting it.

Click **Hardware** > **Hardware Management** > **Hardware Listing**.

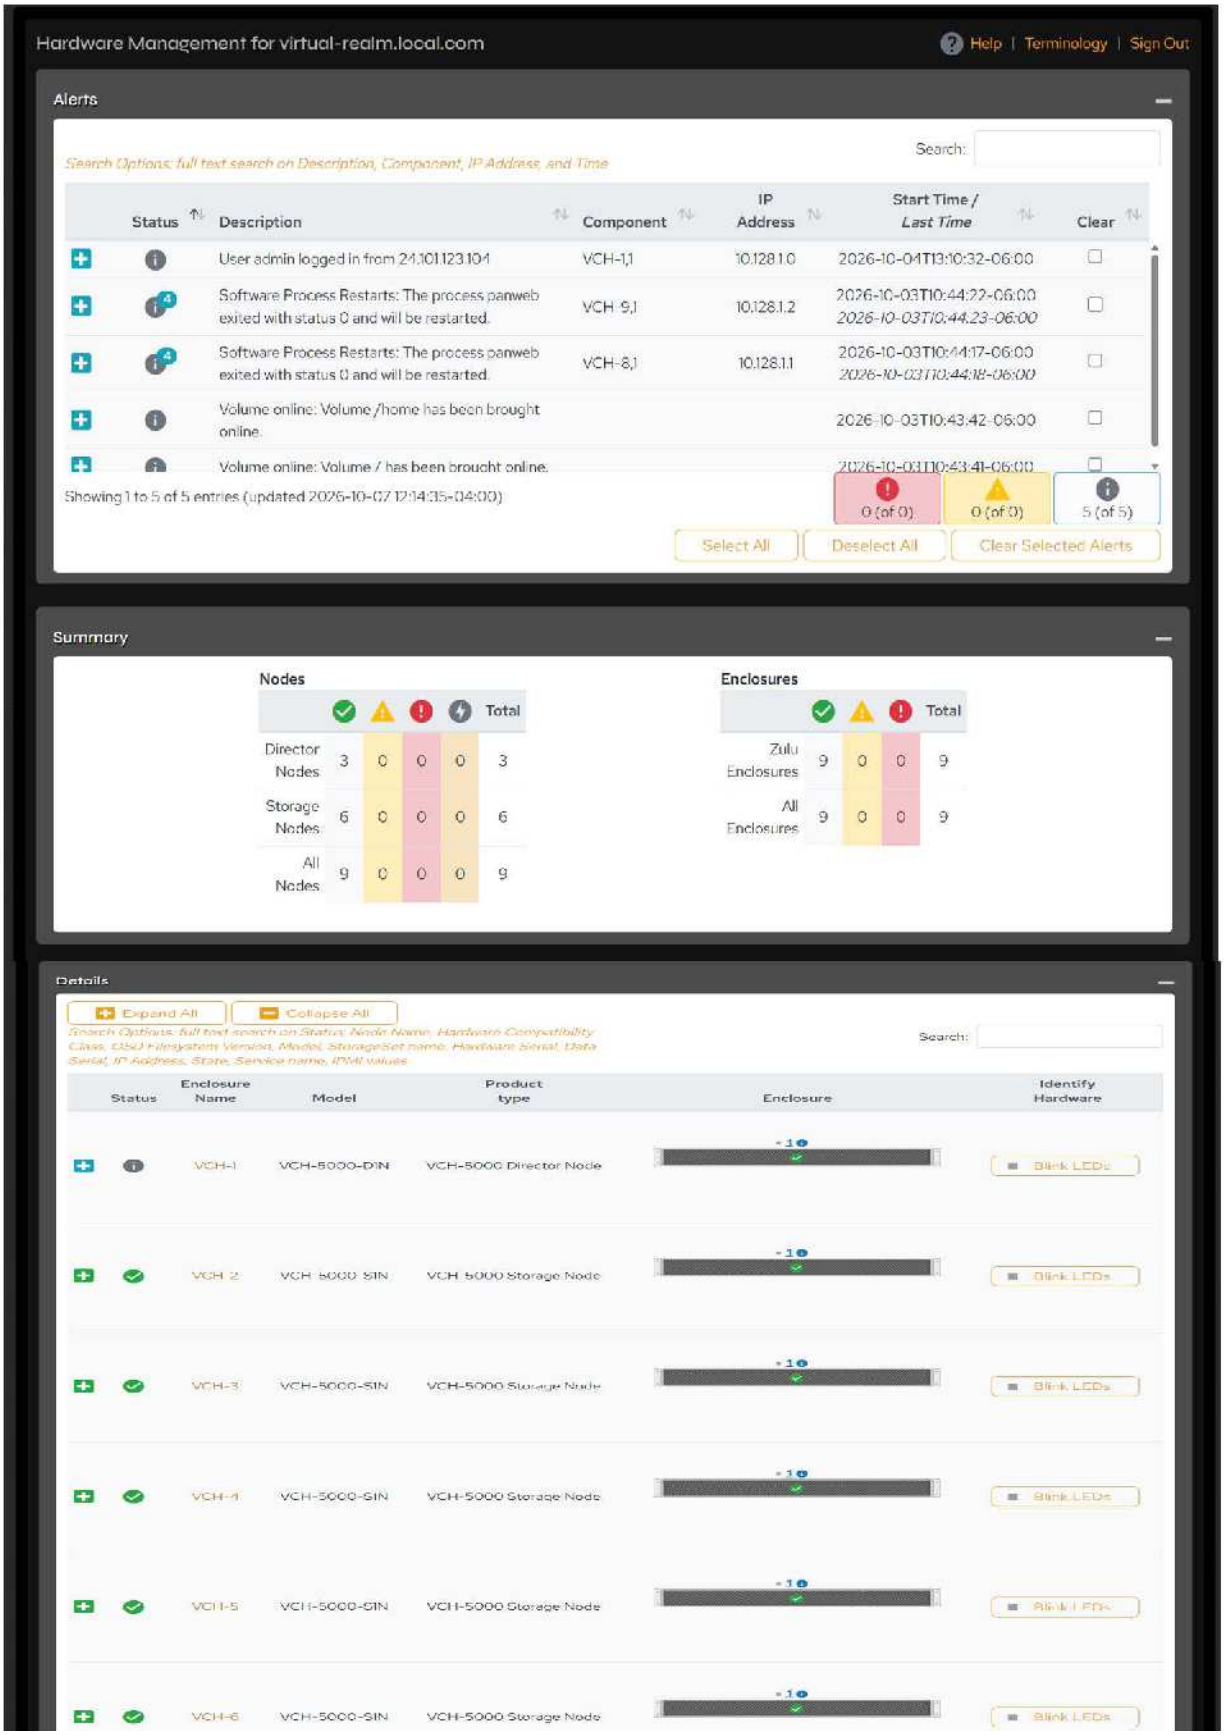

Hardware Management for virtual-realm.local.com
Help | Terminology | Sign Out
Alerts
Search Options: full text search on Description, Component, IP Address, and Time
Search:
Status
Description
Component
IP Address
Start Time / Last Time
Clear
User admin logged in from 24.101.123.104
VCH-1,1
10.128.1.0
2026-10-04T13:10:32-06:00
Software Process Restarts: The process panweb exited with status 0 and will be restarted.
VCH-9,1
10.128.1.2
2026-10-03T10:44:22-06:00
2026-10-03T10:44:23-06:00
Software Process Restarts: The process panweb exited with status 0 and will be restarted.
VCH-8,1
10.128.1.1
2026-10-03T10:44:17-06:00
2026-10-03T10:44:18-06:00
Volume online: Volume /home has been brought online.
2026-10-03T10:43:42-06:00
Volume online: Volume / has been brought online.
2026-10-03T10:43:41-06:00
Showing 1 to 5 of 5 entries (updated 2026-10-07 12:14:35-04:00)
0 (of 0)
0 (of 0)
5 (of 5)
Select All
Deselect All
Clear Selected Alerts
Summary
Nodes
Total
Director Nodes 3 0 0 0 3
Storage Nodes 6 0 0 0 6
All Nodes 9 0 0 0 9
Enclosures
Total
Zulu Enclosures 9 0 0 9
All Enclosures 9 0 0 9
Details
Expand All
Collapse All
Search Options: full text search on Status, Node Name, Hardware Compatibility Class, OS/Flashstack Version, Model, StorageSet name, Hardware Serial, Data Serial, IP Address, State, Service name, IDNR values
Search:
Status
Enclosure Name
Model
Product type
Enclosure
Identify Hardware
VCH-1 VCH-5000-DIN VCH-5000 Director Node
VCH-2 VCH-5000-SIN VCH-5000 Storage Node
VCH-3 VCH-5000-SIN VCH-5000 Storage Node
VCH-4 VCH-5000-SIN VCH-5000 Storage Node
VCH-5 VCH-5000-SIN VCH-5000 Storage Node
VCH-6 VCH-5000-SIN VCH-5000 Storage Node
Blink LEDs



Figure 7-1. Hardware Management page.

The system-level **Hardware Management** page includes the following status tables and controls.

Statuses & Controls	Description
Alerts	**Search**: Use this field to filter alerts by **Description**, **Component**, **IP Address**, and **Time** fields. Clear the box to return to the full view. **Status** icon meanings: • **OK/Online**: The component is online and functioning normally. • **Warning**: A problem is developing that might impair system availability or data integrity. • **Error**: The system is experiencing difficulty that might imminently impair system availability or data integrity. **Description**: Displays summarized error messages. Click the **plus** icon to the left of the **Status** icon to display a detailed description. **Component**: Click the link to display the node detail page for the troubled node. **IP Address**: IP address of the problem node. **Start Time/Last Time**: Date and time of the error. **Clear**: Check boxes individually or click the **Select All/Deselect All** buttons. Click the **Clear Selected Alerts** button to clear selected alerts.

Statuses & Controls	Description
Summary	**Nodes**: Shows a summary of the status of all nodes in the realm. **Enclosures**: Shows a summary of the status of all enclosures in the realm.

Statuses & Controls	Description
Details	Provides a total count of nodes by status and role and enclosures by hardware class and status. **Search:** Use this field to filter alerts by **Description**, **Component**, **IP Address**, and **Time** fields. Clear the box to return to the full view. **Status** icon meanings: • **OK/Online**: The component is online and functioning normally. • **Warning**: A problem is developing that might impair system availability or data integrity. • **Error**: The system is experiencing difficulty that might imminently impair system availability or data integrity. • Click the **plus** icon to the left of the **Status** icon to display a detailed description. **Enclosure Name**: Click a link to go to the item's **Enclosure Management** page. **Model**: Displays the model of each enclosure. **Enclosure**: Provides a graphical representation of the contents and status of each enclosure in the realm. Click an image to go to its **Node Management** page. **Identify Hardware**: Click the **Blink LEDs** button to flash the status lights of the nodes in the enclosure.

Statuses & Controls	Description
Take System Offline/Bring System Online	**Online**: Activate all file services. **Offline**: Disable client access to data but leave the system active so you can manage it with PanActive Manager. For more information, see Manage System-Level Hardware.
Lock System/Unlock System	**Lock**: Lock the system. Some processes, such as migration, require a locked system. **Unlock**: Unlock a locked system. If the system is locked, this status will appear in the **System Status** box in the left pane. For more information, see Lock/Unlock the System.

Statuses & Controls	Description
System Shutdown	Shut down the system. Once the shutdown process is complete, you can turn off the power supplies. For more information, see Shut Down the System. **NOTE**: You must power down and restart hardware from its physical location. **IMPORTANT INFORMATION REGARDING SHUTDOWN/RESTART:** Powering on a JBOD must be done by pressing its power button. If you need to power on a JBOD remotely, VDURA can provide instructions for connecting the JBOD BMC. **At Installation** When installing a V5000 system that includes JBODs, each JBOD must be powered on before powering on its attached Storage Node. VDURA recommends powering on all JBODs, then waiting at least 60 seconds before powering on the Storage Nodes. **During Operation** If the realm is shut down from either PanActive Manager or PanCLI, JBODs will not shut down automatically and must be powered off by hand. JBODs can be powered off by holding down the JBOD enclosure power button located on the front of the unit. Storage Nodes should be powered off prior to shutting down the JBODs. When powering on the realm after a shutdown, JBODs should be powered on before their

Statuses & Controls	Description
	associated Storage Nodes because the nodes must be able to see their JBOD drives as part of the power-on process. VDURA recommends powering on all JBODs, then waiting at least 60 seconds before powering on the Storage Nodes. Use PanActive Manager or PanCLI to power down Storage and Director Nodes; do not remove power at the PDU.
System Reboot	Reboot all hardware components. Hardware and software will come online in a few minutes. For more information, see Reboot the System.
Battery Override	This is a legacy system control and is not used with V5000 hardware.

7.1.1 Manage System-Level Hardware

Several system-wide hardware functions can be controlled from the **Hardware Management** page.

Take the System Offline/Bring the System Online

This procedure disables client access to data but leaves the system active so it can be accessed through PanActive Manager. When the system is brought online again, all file services are reactivated.

From the **Hardware Management** page **Controls** section, click the **Take System Offline**/**Bring System Online** button. You will be prompted to confirm the offlining action. Click **Yes** to continue.

Lock/Unlock the System

Some processes, such as enclosure migration, require that the system be locked before the procedure can be performed.

From the **Hardware Management** page **Controls** section, click the **Lock System** button (or **Unlock System** if the system is locked).

When the system is locked, a message link appears in left pane **System Status** box. Click that link to go to the **Hardware Management** page where you can unlock the system.

Shut Down the System

Shutting down the system powers down the Storage Nodes.

From the **Hardware Management** page **Controls** section, click the **System Shutdown** button. Select **Yes** when asked if you want to shut down the system.

Once the shutdown process is complete, you can turn off the power supplies. All hardware must be powered down and restarted from its physical location.

Reboot the System

Rebooting the system will reboot all hardware components. Hardware and software will come back online in a few minutes.

From the **Hardware Management** page **Controls** section, click the **System Reboot** button. You will be prompted to confirm your choice. Select **Yes** to continue.

7.2 Monitor and Manage Enclosures

You can use PanActive Manager to monitor VDURA hardware at the enclosure, edit VDURA system names, and drain all Storage Nodes in an enclosure.

Click **Hardware** > **Hardware Management** > **Hardware Listing**. In the **Details** section, click the name of the enclosure you want to work with to display its **Enclosure Management** page.

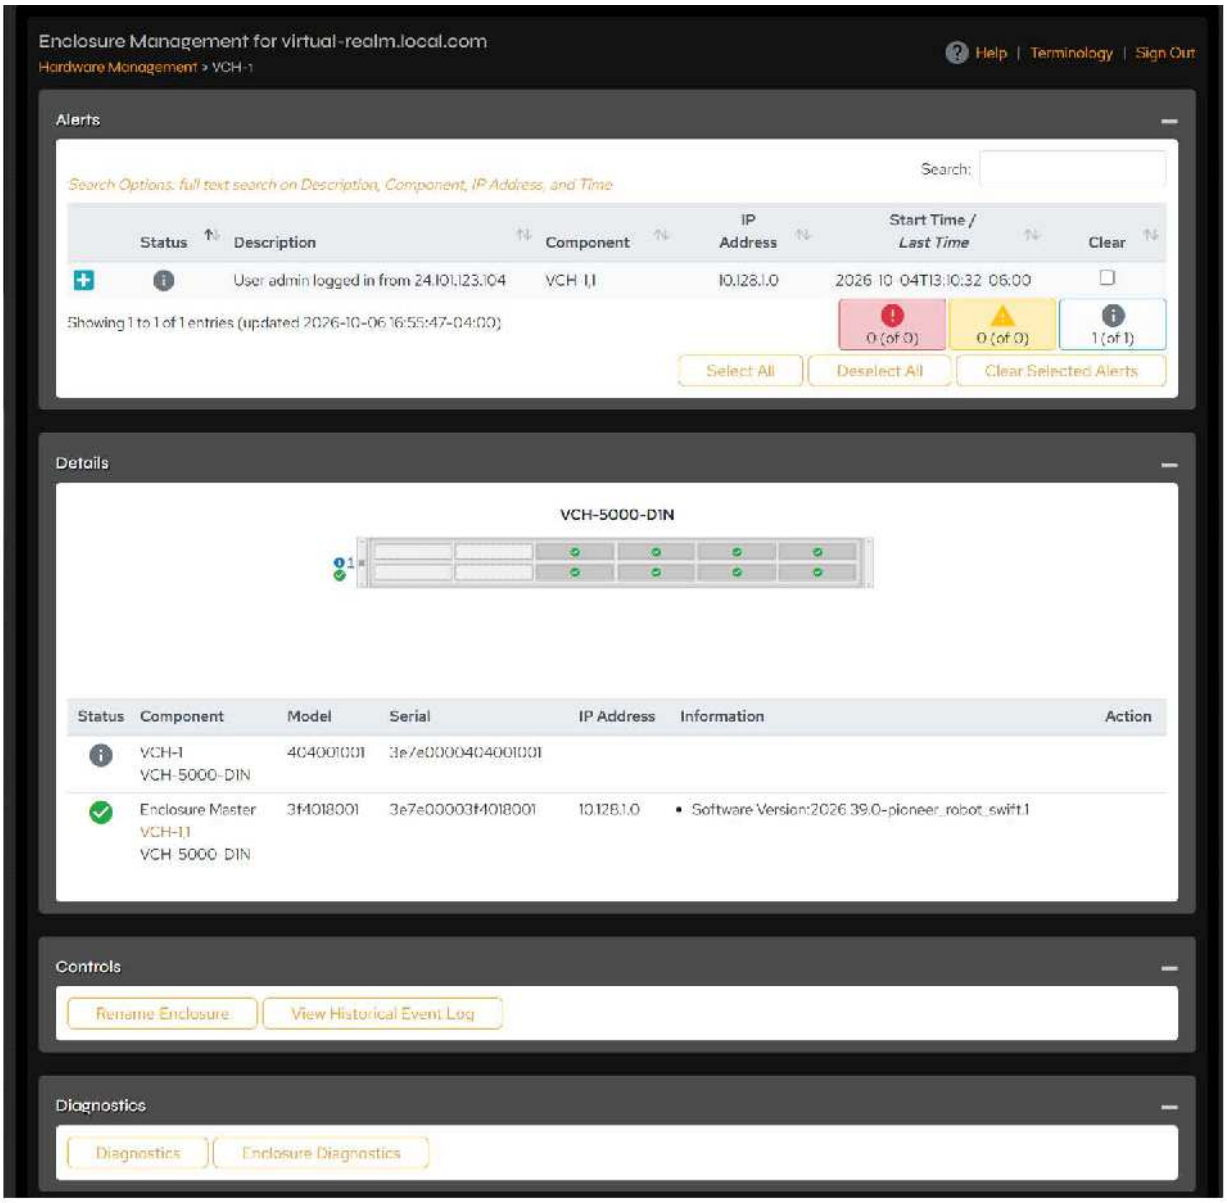


Figure 7-2. Enclosure Management Details, Director Node.

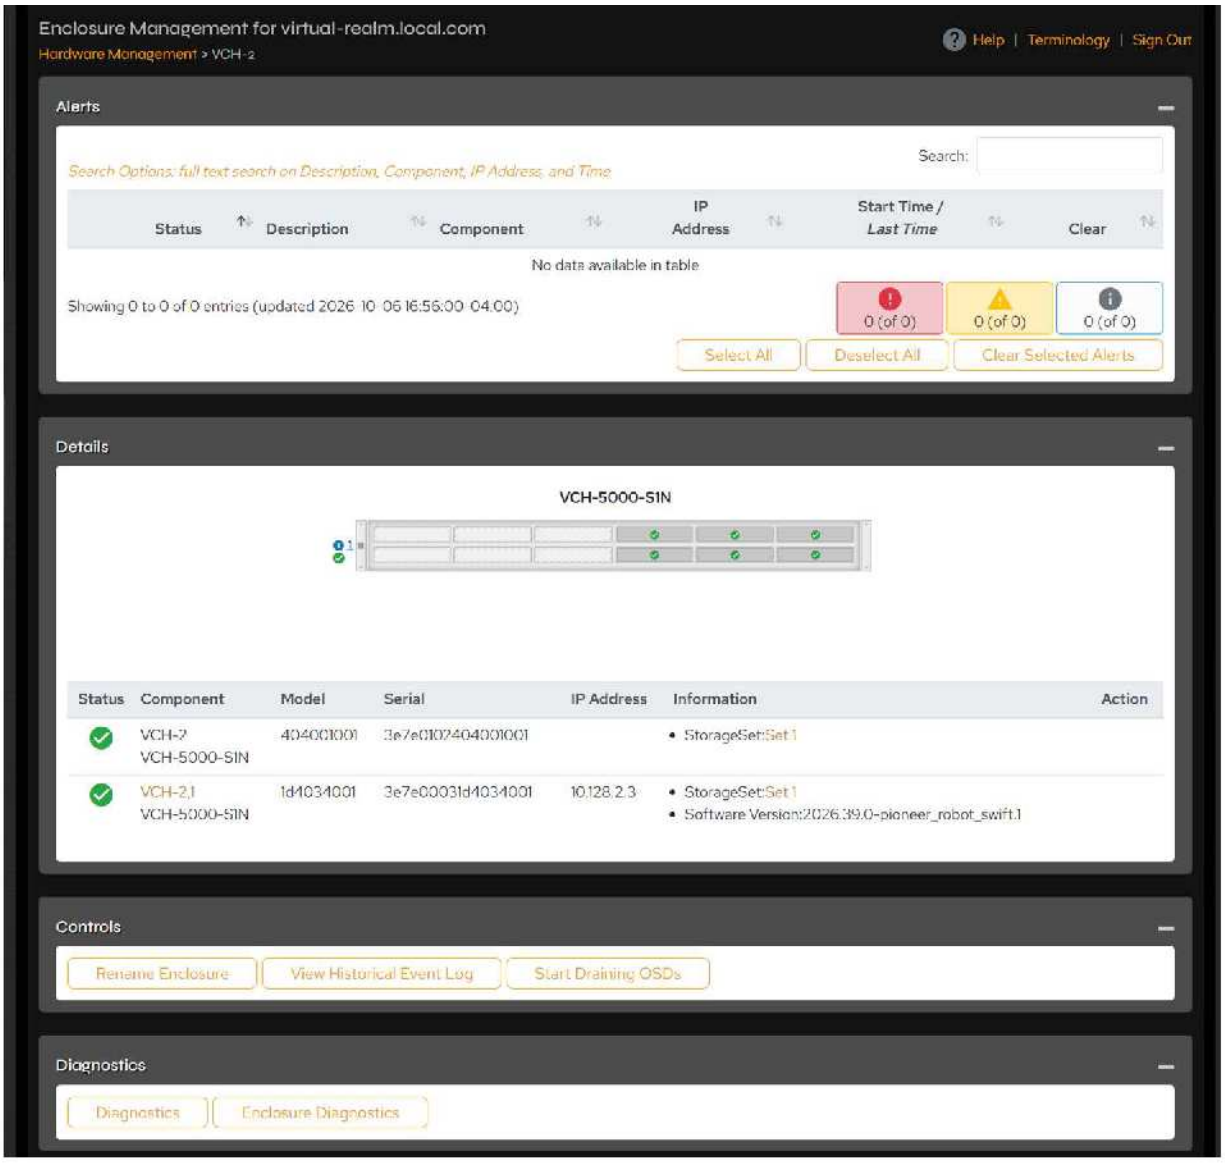


Figure 7-3. Enclosure Management Details, Storage Node.

The **Enclosure Management** page includes the following status tables, controls, and diagnostics.

Statuses & Controls	Description
Alerts	**Search:** Use this field to filter alerts by **Description**, **Component**, **IP Address**, and **Time** fields. Clear the box to return to the full view. **Status** icon meanings: • **Online**: The item is online and functioning normally. • **Warning**: A problem is developing that might impair system availability or data integrity. Click a link to go to an enclosure's management page. • **Error**: The system is experiencing difficulty that might imminently impair system availability or data integrity. **Description**: Error messages appear in summary. Click the **plus** icon to the left of the **Status** icon to display a detailed description. **Component**: Click the link to display the node detail report for the troubled node. **IP Address**: The IP address of the problem node. **Start Time/Last Time**: Date and time of the error. **Clear**: Check boxes individually or click the **Select All/Deselect All** buttons. Click the **Clear Selected Alerts** button to clear selected alerts.

Statuses & Controls	Description
Details	**Hardware diagrams**: Hover mouse icon over a node to view its name. Click a node to open its **Node Management** page. **Status** icon and node graphic LED meanings: • **OK/Online**: The component is online and functioning normally. • **Warning**: A problem is developing that might impair system availability or data integrity. • **Error**: The system is experiencing difficulty that might imminently impair system availability or data integrity. **Component**: Name of the component. If component is a node, click the link to go to its **Node Management** page. **Model:** Displays the model of each component. **Serial**: Serial number of the component. **IP Address**: IP address of the node. **Information**: Displays relevant facts about the component. If component is a node, click the **StorageSet** link to go to its **StorageSet Detail** page. **Action**: If component is a blade or node, click the **Blink LEDs** button to flash its status light.

Statuses & Controls	Description
Controls	**Rename Enclosure**: Click to rename the enclosure. **View Historical Event Log**: Click to display the **Historical Event Log** page for the enclosure. **Start Draining OSDs**: This button begins the draining of data from all eligible Storage Nodes in the enclosure. Draining data will move all data to other Storage Nodes. This process can take several days. To drain all data from this enclosure, type **YES** and click **OK** in the dialog box that opens. This type of drain is used when migrating data from an old enclosure to a new one. For more information, see Enclosure Migration. To drain a single storage node, use the procedure in Drain a Storage Node instead. **Clear Record of Enclosure**: Click to remove all nodes within the specified enclosure.
Diagnostics	**Diagnostics**: Click to open the **Launch Diagnostics** page where you can select diagnostics to perform. **Enclosure Diagnostics:** Click to run enclosure-level diagnostic tests.

7.3 Monitor and Manage Nodes

Each Director and Storage Node has a **Node Management** page from which you can monitor and manage the hardware at the node level.

Click **Hardware > Hardware Management > Node Listing** to display a list of all Director and Storage Nodes in your system.

In the example below, nodes VCH-1,1, VCH-8,1, and VCH-9,1 are Directors. Each Storage Node is listed as its own enclosure (VCH-2,1 through VCH-7,1).

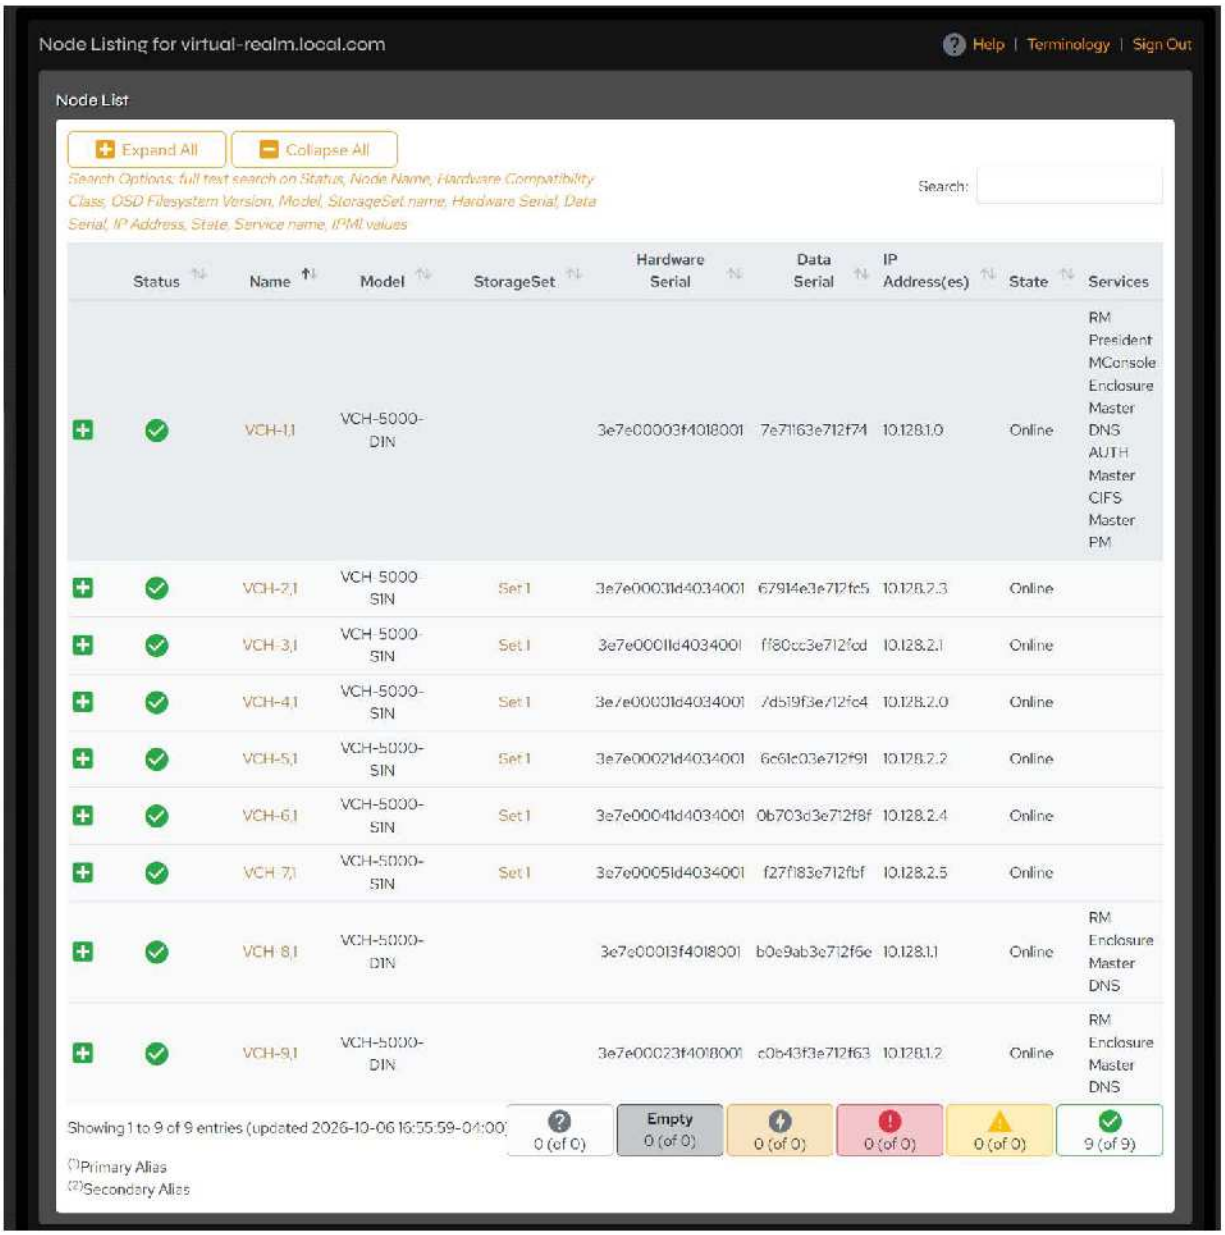


Figure 7-4. Node Listing page.

This page provides high-level information about the nodes in your system, including node status, name, model, StorageSet, hardware and data serial numbers, IP addresses, state, and services provided by Directors. A footnote below the table identifies the primary alias (1) and secondary alias (2) addresses.

To open a specific node, click its **Name link** in the table. You can also open **Node Management** for a specific node from the **Hardware Management** page by clicking the icon for the node in the enclosure's graphical representation.

7.3.1 Manage Node-Level Hardware

The **Node Management** page provides detailed information about errors and node statistics, including drive assets and NIC status. From this page, you can also perform node procedures, set networking preferences, and gather information about the node.

From the **Node Listing** page, click the name of a node to view its **Node Management** page.

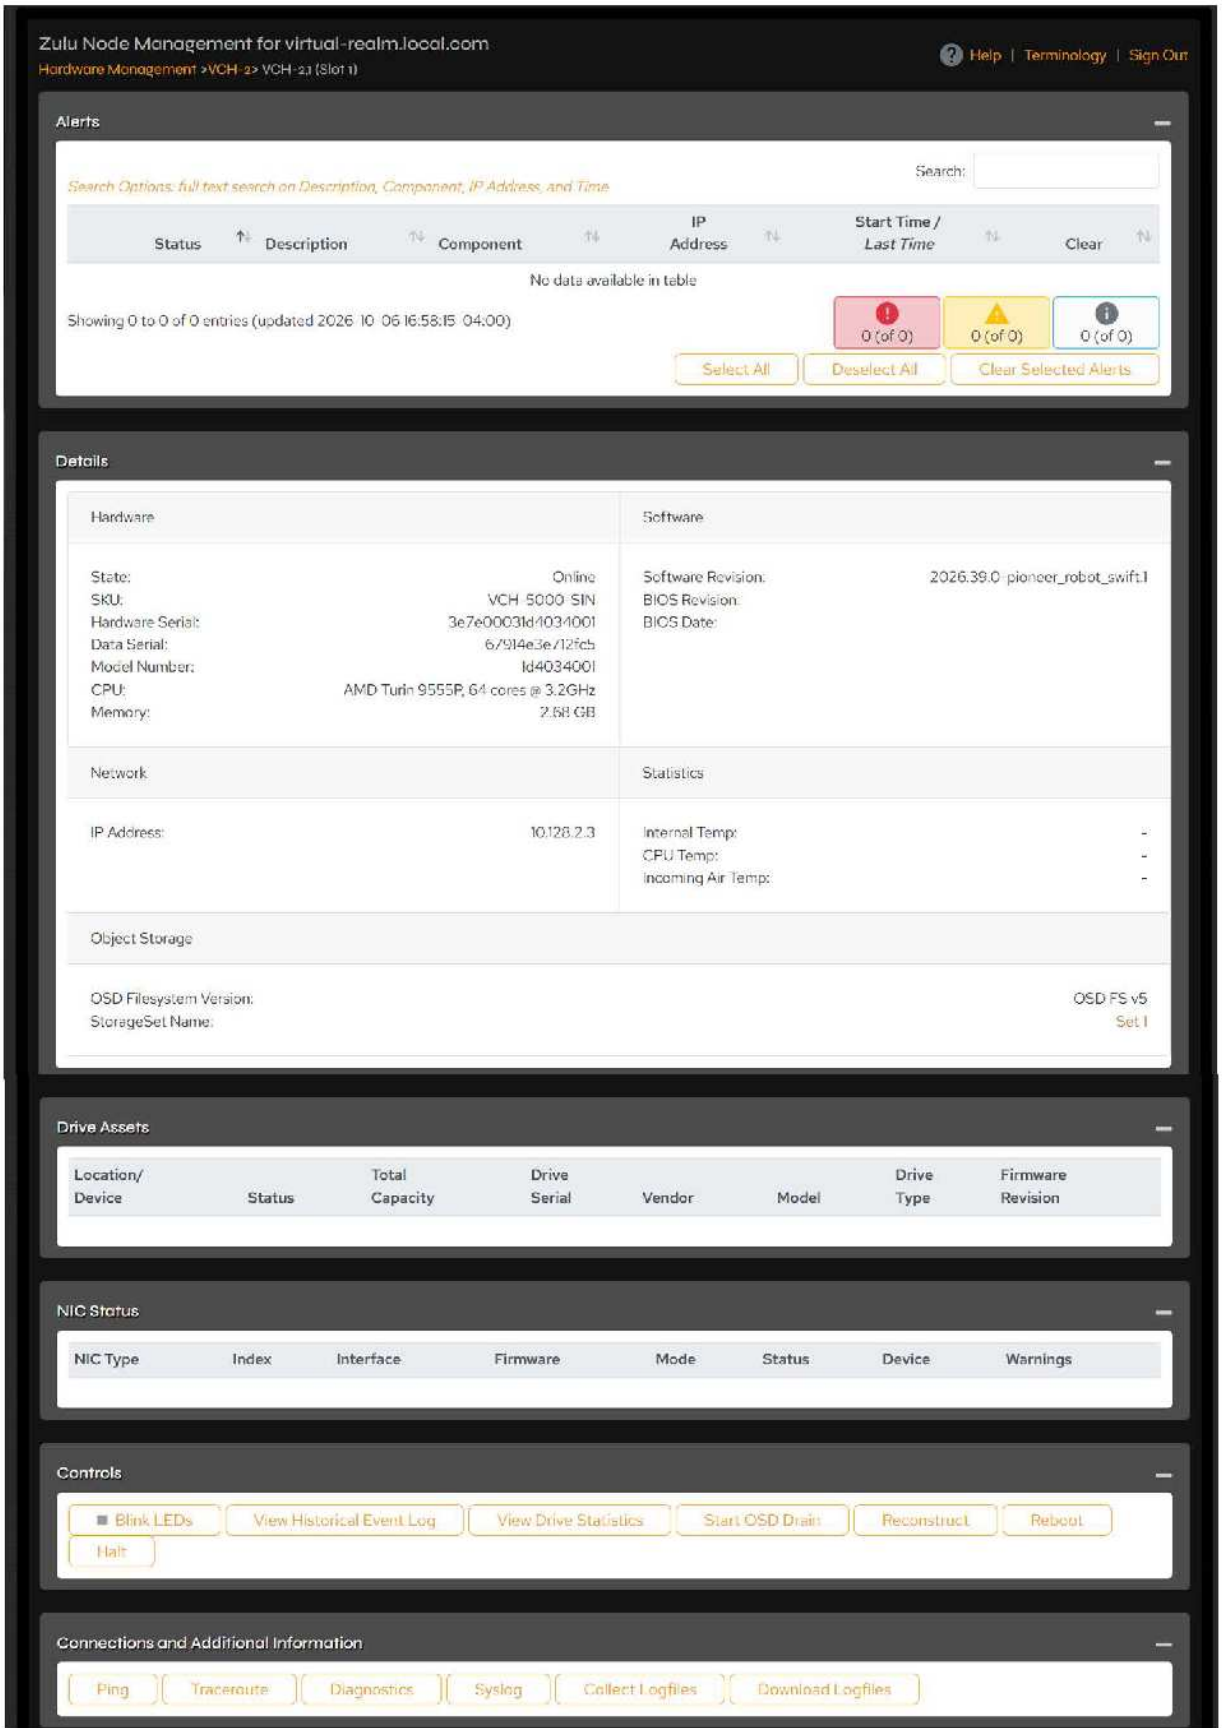

Zulu Node Management for virtual-realm.local.com
Hardware Management >VCH-2> VCH-2.1 (Slot 1)
Help | Terminology | Sign Out
Alerts
Search Options: full text search on Description, Component, IP Address, and Time
Search:
Status
Description
Component
IP Address
Start Time / Last Time
Clear
No data available in table
Showing 0 to 0 of 0 entries (updated 2026-10-06 16:58:15 -04:00)
0 (of 0)
0 (of 0)
0 (of 0)
Select All
Deselect All
Clear Selected Alerts
Details
Hardware
State: Online
SKU: VCH-5000-SIN
Hardware Serial: 3e7e00031d4034001
Data Serial: 67914e3e712fcb
Model Number: Id4034001
CPU: AMD Turin 9555P, 64 cores @ 3.2GHz
Memory: 2.68 GB
Software
Software Revision: 2026.39.0-pioneer_robot_swift.1
BIOS Revision:
BIOS Date:
Network
IP Address: 10.128.2.3
Statistics
Internal Temp: -
CPU Temp: -
Incoming Air Temp: -
Object Storage
OSD Filesystem Version: OSD FS v5
StorageSet Name: Set 1
Drive Assets
Location/ Device
Status
Total Capacity
Drive Serial
Vendor
Model
Drive Type
Firmware Revision
NIC Status
NIC Type
Index
Interface
Firmware
Mode
Status
Device
Warnings
Controls
Blink LEDs
View Historical Event Log
View Drive Statistics
Start OSD Drain
Reconstruct
Reboot
Halt
Connections and Additional Information
Ping
Traceroute
Diagnostics
Syslog
Collect Logfiles
Download Logfiles

Figure 7-5. Node Management page.

The following tables and controls appear on the **Node Management** page:

Statuses & Controls	Description
Alerts	**Search**: Use this field to filter alerts by **Description**, **Component**, **IP Address**, and **Time** fields. Clear the box to return to the full view. **Status** icon meanings: • **Online**: The item is online and functioning normally. • **Warning**: A problem is developing that might impair system availability or data integrity. • **Error**: The system is experiencing difficulty that might imminently impair system availability or data integrity. **Description**: Error messages appear in summary. Click the **plus** icon to the left of the **Status** icon to display a detailed description. **Component**: Click the link to display the node detail report for the troubled node. **IP Address**: The IP address of the problem node. **Start Time/Last Time**: Date and time of the error. **Clear**: Check boxes individually or click the **Select All/Deselect All** buttons. Click the **Clear Selected Alerts** button to clear selected alerts.

Statuses & Controls	Description
Details	Provides a variety of statistics and information about the node, including state, physical characteristics, serial and part numbers, IP address, primary and secondary aliases (Directors only), object storage details (Storage Nodes only), hardware temperatures, and additional hardware and software statistics. The following list describes the possible node states: • **Joining**: The node is performing the initial handshake to join the system. This state occurs following reboot. • **Online**: The node is online. • **Needs Upgrade**: The node has a different version of software from the rest of the system and must be upgraded before it can be used. This is typically a transient state for new nodes, as the system will automatically attempt to upgrade the new off-version node to the version of the rest of the system. If the attempt to upgrade an off-version node fails, it remains in this state awaiting remedy. • **Needs Reinitialization**: The node already has data on it that might have come from a different system. • **Offline**: The node is offline. • **FSCK**: The node is performing a file system check (**fsck**). This state is transient and usually happens following a reboot. • **Not Responding**: The node cannot be contacted. The cause is unknown. The node may be rebooting or temporarily disconnected. • **Not Responding, Lost**: A node has been physically replaced or reinitialized but not removed from the software, so the system is looking for the old one. See Remove a Lost Node for more information. • **Failed**: Data from a failed node is getting or

Statuses & Controls	Description
	has been reconstructed from parity. The node should be removed and its records cleared from the system. • **Rejected**: The node has been rejected from the system and powered down. A new node can be rejected if it has an invalid serial number or if its capacity differs from the rest of the enclosure.

Statuses & Controls	Description
Details, cont.	• **Draining**: The transparent drain operation moves all Storage Node data to other storage nodes and removes it from the system. • **Needs Repair**: Needs attention. Call VDURA support for the best course of action. • **Upgrading**: If a new node is added to the system with a different software version, the system will automatically upgrade the new off-version node to the version of the rest of the nodes in the system. • **Clock Sync**: The node is performing clock synchronization. This transient state follows reboot. • **Booting**: Node is rebooting. • **Shutdown**: Node is shutting down. • **Secure Erase**: Node is undergoing secure erasure of data, as requested by the admin. Node state modifiers include: • **Not Activated**: The node has joined the system but is not yet active. For Storage Nodes, this means the admin needs to add the node to a StorageSet. For Directors, this means the admin must activate it. • **Draining**: Data is being moved from this node to other nodes. Once this action is completed, the node should be removed from the system. • **Drain suspended**: Drain activity was requested, but has been temporarily suspended.
Drive Assets	Provides information about the drives, including location, status, total capacity, serial number, vendor, model, drive type, and firmware revision. For usage information, click the **View Drive Statistics** button in the **Controls** section.

Statuses & Controls	Description
NIC Status	Provides information about the network interface cards, including NIC type, index, interface, mode, status, device type, and warnings.

Statuses & Controls	Description
Controls	The buttons described below only appear if the function can be performed. • **Blink LEDs/Stop LEDs**: Click to flash the status lights on a node when you need to locate it in a machine room. Click **Stop LEDs** to stop the flashing. • **Clear Record of Node:** Use this button on a Storage Node that has completed draining (**Drained** state) or has completed reconstruction (**Failed** state). Clearing is also appropriate for removing a rejected node or a new node that still has data from another system (**Needs Reinit** state) without erasing the existing data. • **Halt**: Click to halt activity on individual nodes. For the procedure to halt a node, see Halt a Node. • **Reboot**: Click to reboot individual nodes. For the procedure to reboot a node, see Reboot a Node. • **Reconstruct:** Click to start reconstruction of a Storage Node from parity. This action causes all volumes in that node's StorageSet to be in **Degraded** mode until reconstruction of those volumes completes. The node will immediately transition to the **Failed** state, and the admin should subsequently clear the record of the node. For more information, see Reconstruct Failed Storage Nodes in a StorageSet. • **Reinitialize Node**: This control appears only when reinitialization is required. If a new node is added to the system but still has data on it from a previous system (**Needs Reinit** state), it must be reinitialized before it can be brought online. • **Remove from Replication Set**: Click to remove a node from the software that has been

Statuses & Controls	Description
	physically replaced in the enclosure. See Remove a Lost Node.
Controls, cont.	• **Start OSD Drain/Stop OSD Drain**: This option allows users to drain or stop draining all data from a Storage Node. The transparent drain operation does not require the corresponding StorageSet to be offline and does not put the StorageSet into a degraded operation. For more information, see Drain a Storage Node. • **View Drive Statistics**: Click to view data on all drives in the realm. • **View Historical Event Log**: Click to display the Historical Event Log page for the enclosure.
Networking Preferences	These controls are specifically for Director Nodes. • **Enable/Disable Aliases:** Use this button to decide whether a Director should have IP aliases assigned to it. By default, all Directors can receive aliases. It may be useful to disable alias assignment for a handful of Directors if they are busy with other specific system tasks. For more information about IP aliases, see IP Aliases.

Statuses & Controls	Description
Connections and Additional Information	The buttons described below only appear if the function can be performed. • **Ping**: Use this button to attempt a handshake from the node to a specified IP address with two retry attempts. • **Traceroute**: Use this button to show IP addresses for up to five hops from the node to a specified IP address with a two second wait for response. • **Diagnostics**: Opens the **Launch Diagnostics** page, from which you can select diagnostic tests to run or results to view. • **Syslog**: Generates a system log file for the node. • **Collect Logfiles**: Click to generate all log files for the node. • **Download Logfiles**: Downloads the logs to your local host. You can then send the log files to VDURA Technical Support for troubleshooting.

7.3.2 Node Management Tasks

You can perform the following tasks from the **Node Management** page.

Halt a Node

Nodes can safely be removed after they are halted.

Click **Hardware** > **Hardware Management** > **Node Listing**, then click the name of the node you want to halt. The Node Management page opens.

In the **Controls** section, click the Halt button to start the halt process. A warning dialog box asks you to confirm your request. Click **OK** to confirm the request. The system will display a message at the top of the screen indicating that the halting activity is underway.

> **NOTE:** This operation can also be performed using the PanCLI **halt** command.

Reboot a Node

Click **Hardware** > **Hardware Management** > **Node Listing**, then click the name of the node you want to reboot. The Node Management page opens.

In the **Controls** section, click the Reboot button to start the process. A warning dialog box asks you to confirm your request. Click **OK** to confirm. The system will display a message at the top of the screen indicating that the rebooting activity is underway.

> **NOTE:** This operation can also be performed using the PanCLI **reboot** command.

Drain a Storage Node

Users can drain (migrate) all data from a Storage Node. The transparent drain operation does not require the corresponding StorageSet to be offline and does not put the StorageSet into a degraded state of operation. This feature allows users to proactively drain a Storage Node when hardware failures are predicted, but the node is still in an **Online** state.

Draining a Storage Node will move all its data to other nodes and will remove it from the system. Once the drain operation is complete, it can be physically removed from the system.

To drain a node, follow these steps:

From **Hardware** > **Hardware Management** > **Node Listing**, click the name of the node you want to drain. In the **Controls** section of the **Node Management** page that opens, click the **Start OSD Drain** button. A message will appear asking you to confirm that you want to drain. Click **Yes**.

You can monitor the drain status either from the **Details** section of the **Node Management** page or from the **Hardware > Hardware Management > Drain Status page** (see Monitor or Cancel a Drain).

If progress seems to stop after draining more than 95 percent of the data, it is safe to physically remove the node. Failure to drain completely is not uncommon and is caused by damaged objects that remain on the node. The system will reconstruct whatever remaining data it was unable to drain. Instead of removing the node, you can also click the **Reconstruct** button to trigger a reconstruction. For more information on reconstructing storage nodes, see Reconstruct Failed Storage Nodes in a StorageSet.

The storage node may appear in the **Not Responding, Lost** state during reconstruction. Once reconstruction completes, the system will remove the node automatically; no manual intervention is needed.

Monitor or Cancel a Drain

You can monitor the status of nodes being drained or cancel an in-progress drain from the **Hardware > Hardware Management > Drain Status** page. This page also presents a list of previously drained or reconstructed nodes in the system.

7.3.3 Remove a Lost Node

When a Director has been physically replaced but has not been removed from the software, it will appear in the **Not Responding, Lost** state, as the system is looking for the "lost" node. To remove the node from the software in this specific case, click the **Remove from Replication Set** button, then use the **Clear Record of Node** button to remove the director from PanActive Manager.

Storage Nodes undergoing reconstruction after a drain may also appear as **Not Responding, Lost**. Once reconstruction completes, the system will remove the node automatically; no manual intervention is needed.

7.3.4 Disk Firmware

VDURA system administrators can update node drive firmware using PanActive Manager. VDURA recommends planning firmware updates in advance. If necessary, contact VDURA Technical Support for consultation or assistance with firmware updates.

Update Disk Firmware

Click **Hardware** > **Hardware Management** > **Drive Listing** to display information about the system nodes and disks, as well as recent disk firmware updates.



Figure 7-6. Drive Listing page with Update Firmware field.

Select the drives you wish to update by marking either the **Storage Drives** box, the **Director Drives** box, or both.

Click the **Start Serial Disk Firmware Update** button. The **Start Serial Disk Firmware Update** button will toggle to the **Cancel Disk Firmware Update** button.

Firmware updates will occur one node at a time. The system will wait for a node to return to the **Online** state before starting to update the next one. The system may wait as long as 20 minutes for a node to come back online. If the node does not come back online in this time, the system will not continue with further firmware updates, and a warning will be sent to the administrator. When a node is unavailable, a "Node is unavailable" error message will appear on the **Hardware Management** page.

Cancel a Disk Firmware Update

Cancel an in-progress firmware update by clicking the **Cancel Disk Firmware Update** button.

If a disk firmware update is canceled by the user, the system will finish processing the update of any node that is in progress. As soon as the firmware update on that node has finished, the update will stop.

Restart a Disk Firmware Update

A canceled disk firmware update may also be restarted by performing the procedure Update Disk Firmware. If the update was canceled and started again, the system will go through each node to check its firmware version, but no updates will be performed on disks that have already been updated to the current version of firmware.

7.3.5 Node Temperature

To monitor the temperatures of nodes, click **Hardware > Hardware Management > Temperature**. You can view internal, CPU, and incoming air temperatures from this page.

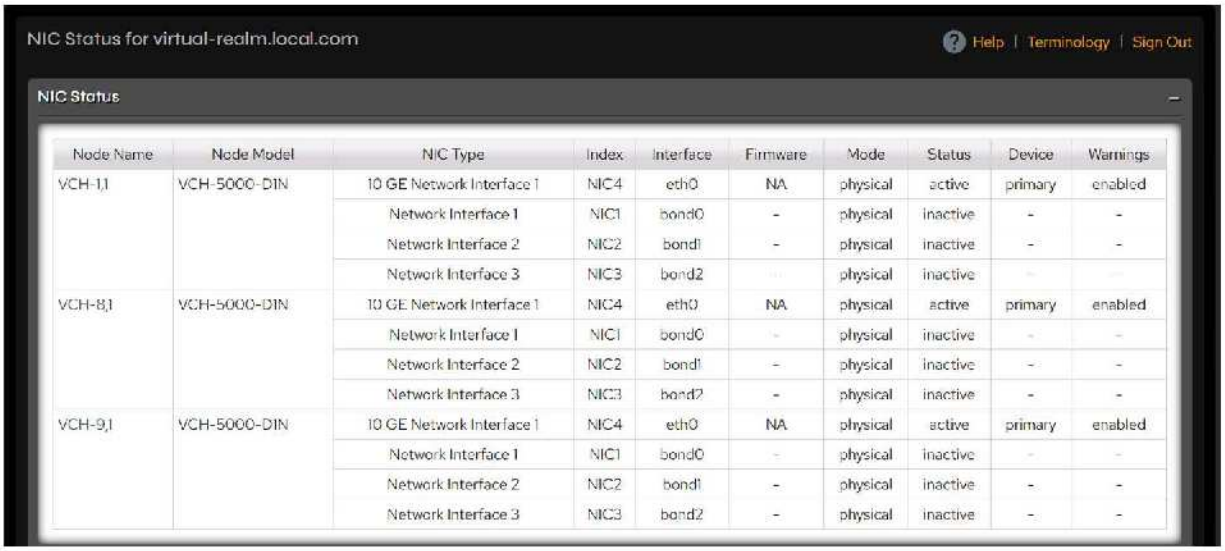
7.3.6 Replace a Node

The procedures to replace Director and Storage Nodes are documented in the *VDURA V5000 Certified Platform Installation and Service Guide.* It is not recommended to replace nodes without the assistance of a qualified VDURA representative.

7.4 Network and NIC Status

You can view the NIC status of any specific node via that node's **Node Management** page.

To obtain detailed information about the NIC status of all nodes in the system, navigate to **Network > Network Status > NIC Status**.



Node Name	Node Model	NIC Type	Index	Interface	Firmware	Mode	Status	Device	Warnings
VCH-1,1	VCH-5000-D1N	10 GE Network Interface 1	NIC4	eth0	NA	physical	active	primary	enabled
		Network Interface 1	NIC1	bond0	-	physical	inactive	-	-
		Network Interface 2	NIC2	bond1	-	physical	inactive	-	-
		Network Interface 3	NIC3	bond2	-	physical	inactive	-	-
VCH-8,1	VCH-5000-D1N	10 GE Network Interface 1	NIC4	eth0	NA	physical	active	primary	enabled
		Network Interface 1	NIC1	bond0	-	physical	inactive	-	-
		Network Interface 2	NIC2	bond1	-	physical	inactive	-	-
		Network Interface 3	NIC3	bond2	-	physical	inactive	-	-
VCH-9,1	VCH-5000-D1N	10 GE Network Interface 1	NIC4	eth0	NA	physical	active	primary	enabled
		Network Interface 1	NIC1	bond0	-	physical	inactive	-	-
		Network Interface 2	NIC2	bond1	-	physical	inactive	-	-
		Network Interface 3	NIC3	bond2	-	physical	inactive	-	-

Figure 7-7. NIC Status page.

The following table explains the NIC status fields and their values.

Field	Description
Node Name	Link to node's management page.
Node Model	Brief description of hardware.
NIC Type	Type of NIC.
Index	Identification of NIC for tracking purposes. This number is referred to in Historical Event Log and system log messages.
Interface	Type of NIC interface.
Firmware	Firmware version.

Field	Description
Mode	Operating mode of the NIC: • **Failover**: This NIC is part of "LAG Network Interface 1/2" and is operating in failover mode. When a LAG is operating in failover mode, only one NIC will be used for communication. • **Physical**: This NIC is not part of any LAG and is operating standalone. • **No Warn**: The system admin has chosen not to be warned about any error seen on this NIC. • **Disabled**: This interface is not intended to be used. For example, if storage-failover/secondary-alias is not configured, then a second interface on the nodes will be disabled.
Status	Operational state of interfaces: • **Active**: Link on the interface is up. • **Inactive**: Link on the interface is down.
Device	Type of IP address configured on this interface: • **Primary**: Configured with primary address and primary alias. • **Secondary**: Configured with secondary alias.
Warnings	Displays whether or not warnings are enabled for the NIC.

8. Monitor Performance

This chapter describes how to use PanActive Manager to monitor performance and generate reports for your VDURA system. It includes the following sections:

- Display Performance Status
- Create Reports

8.1 Display Performance Status

You can use PanActive Manager to display summaries of current system performance:

- System-At-A-Glance
- Storage Statistics
- Drive Statistics
- Director Statistics
- DirectFlow Clients

8.1.1 System-At-A-Glance

Click **Status & Reports > System Monitor > System-At-A-Glance**.

The **System-At-A-Glance** page is the default page for the **Status & Reports** tab and provides a capsule of data about the entire system. You can also click the **System Status** icon in the left pane of any page to view **System-At-A-Glance**.

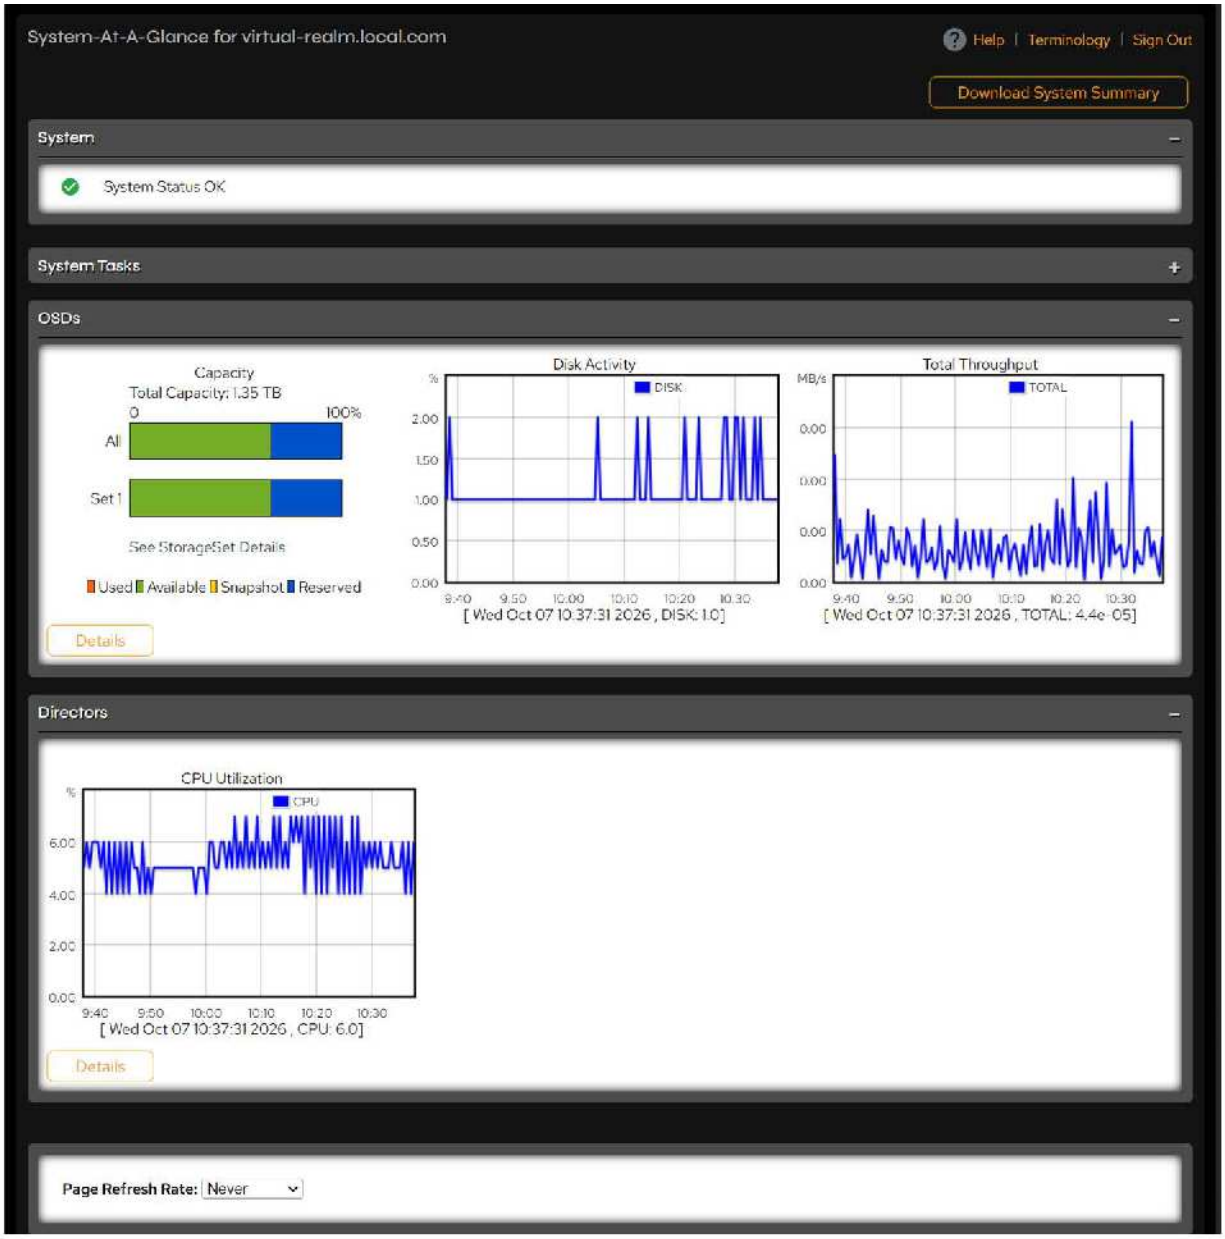


Figure 8-1. System-At-A-Glance page.

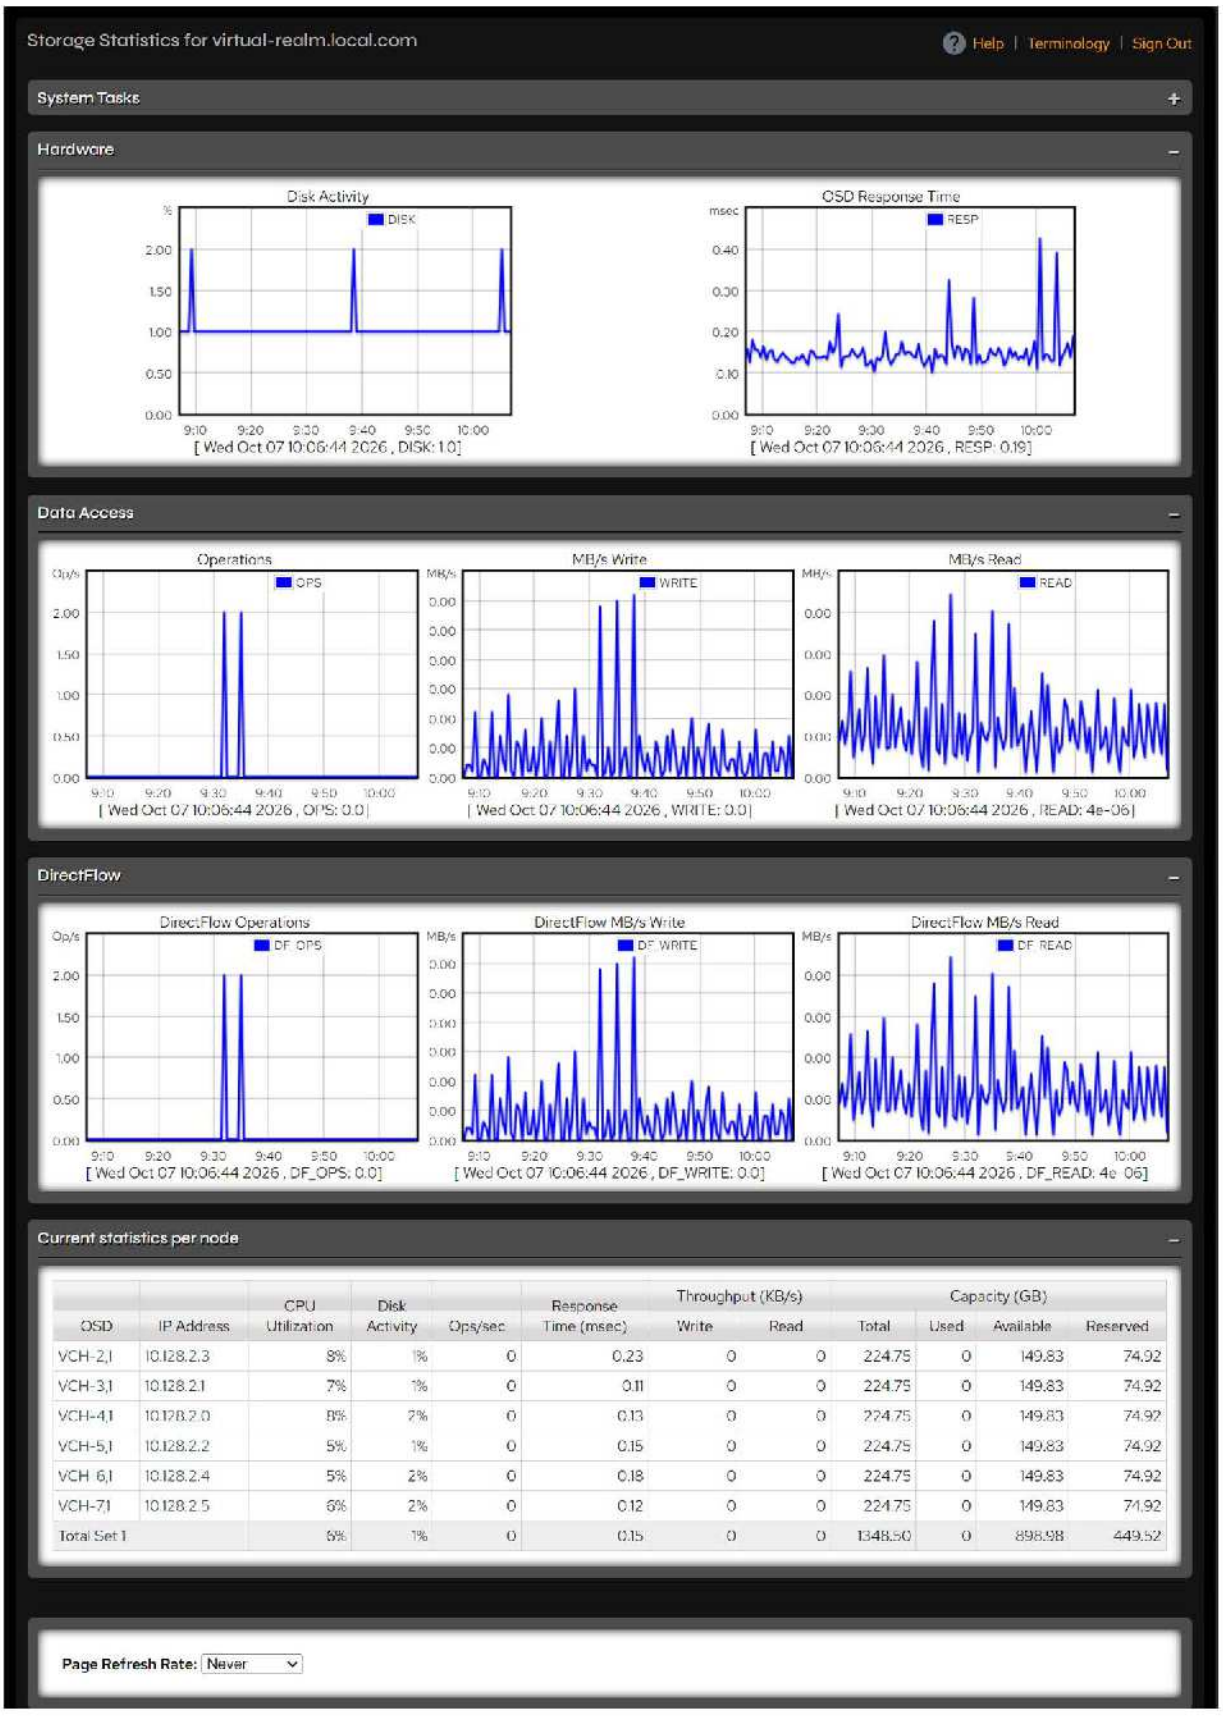
The **System-At-A-Glance** page provides an aggregate status for the VDURA system.

Metric	Description
Download System Summary	Click to download a system summary in TXT format.
Hardware	Summary of hardware errors. Click the icon(s) for details. For a description of error messages, see VDURA System Error Messages.
Volumes	Storage warnings and errors. Click the icon(s) for details. For a description of error messages, see VDURA System Error Messages.
StorageSets	StorageSet warnings and errors. Click the icon(s) for details. For a description of error messages, see VDURA System Error Messages.
System Tasks	Click the plus icon in the right corner of the System Tasks heading to display any system tasks that are in progress, such as diagnostic processes. Overhead processes can explain unidentifiable spikes in performance and use.

Metric	Description
OSDs	Statistics are averages of all Storage Nodes during the past hour. To display detailed information about individual nodes, click the **Details** button. VDURA uses some disk space to store internal data, such as object metadata. This overhead is responsible for the fact that the available node capacity is lower than its total disk capacity. **Capacity: Distribution of storage capacity for all StorageSets combined and for each StorageSet.** • **Used**: Disk space allocated to data objects. • **Available**: Disk space not allocated to data objects. • **Snapshot**: Disk space used by snapshots. • **Reserved**: Space reserved for reconstruction. **See StorageSet Details**: Click the link to display the **StorageSet Listing** page. **Disk Activity**: Percent of time the disk is busy with storage operations. **Total Throughput**: MB of data read from or written to disk per second.
Directors	Statistics are averages of all Director Nodes during the past hour. To display detailed information about individual nodes, click the **Details** button. **CPU Utilization**: Percent of time the CPU is busy with storage operations.

8.1.2 Storage Statistics

Click **Status & Reports > System Monitor > Storage Statistics**.



OSD	IP Address	CPU Utilization	Disk Activity	Ops/sec	Response Time (msec)	Throughput (KB/s) Write	Throughput (KB/s) Read	Capacity (GB) Total	Capacity (GB) Used	Capacity (GB) Available	Capacity (GB) Reserved
VCH-2,1	10.128.2.3	8%	1%	0	0.23	0	0	224.75	0	149.83	74.92
VCH-3,1	10.128.2.1	7%	1%	0	0.11	0	0	224.75	0	149.83	74.92
VCH-4,1	10.128.2.0	8%	2%	0	0.13	0	0	224.75	0	149.83	74.92
VCH-5,1	10.128.2.2	5%	1%	0	0.15	0	0	224.75	0	149.83	74.92
VCH-6,1	10.128.2.4	5%	2%	0	0.18	0	0	224.75	0	149.83	74.92
VCH-7,1	10.128.2.5	6%	2%	0	0.12	0	0	224.75	0	149.83	74.92
Total Set 1		6%	1%	0	0.15	0	0	1348.50	0	898.98	449.52

Figure 8-2. Storage Statistics page.

The **Storage Statistics** page provides a performance summary of Storage Nodes in your system.

Metric	Description
System Tasks	Click the plus icon in the right corner of the System Tasks heading to display any system tasks that are in progress, such as diagnostic processes. Overhead processes can explain unidentifiable spikes in performance and use.
Hardware (aggregated for all Storage Nodes)	**Disk activity**: Percent of time the disk is busy with storage operations. **OSD Response Time**: Response time of the OSDs in milliseconds.
Data Access (aggregated for all Storage Nodes)	**Operations**: Number of operations per second. **MB/s Write**: MB of data written to disk per second. **MB/s Read**: MB of data read from disk or cache per second.
DirectFlow (aggregated for all Storage Nodes)	**DirectFlow Operations**: Number of DirectFlow operations per second. **DirectFlow MB/s Write**: MB per second of data written to disk via DirectFlow access. **DirectFlow MB/s Read**: MB per second of data read from disk or cache via DirectFlow access.

Metric	Description
Current statistics per node	**OSD**: Click to open the **Node Management** page for the node. **IP Address**: Click to open the **Node Management** page for the node. **CPU Utilization**: Percent the CPU is busy with storage operations. **Disk Activity**: Percent the disk is busy with storage operations. **Ops/sec**: Number of storage operations per second. **Response Time**: Response time in milliseconds. **Throughput**: KB per second of data written to and read from disk, shown in separate Write and Read columns. **Capacity**: Total, used, available, and reserved capacity for the node, in GB. VDURA uses some disk space to store internal data, such as object metadata. This overhead is responsible for the available node capacity being lower than the total disk capacity.

8.1.3 Drive Statistics

The **Drive Statistics** page provides a summary of drive statistics for your system. For each node, information is included about status, node and StorageSet, drive number, drive name, drive type, drive serial number, LUN, total capacity, data capacity used, metadata capacity used, total capacity used, free capacity, total inodes, reallocation count, reads since boot, and writes since boot. Use the Search box to filter the list. For SSDs, remaining lifetime is also listed.

Click **Status & Reports > System Monitor > Drive Statistics**.

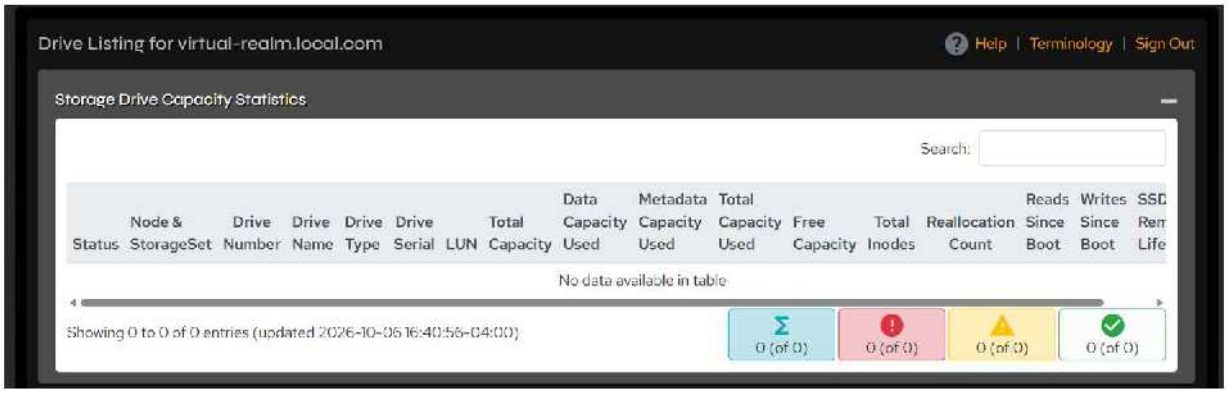


Figure 8-3. Drive Listing page with capacity statistics.

> **NOTE:** Capacity information is reported using base 10, not base 2; i.e. KB = 1000, not 1024.

> **NOTE:** For capacity expansion nodes, the drives associated with the node in the JBOD are listed with the node's internal drives.

8.1.4 Director Statistics

Click **Status & Reports > System Monitor > Director Statistics**.

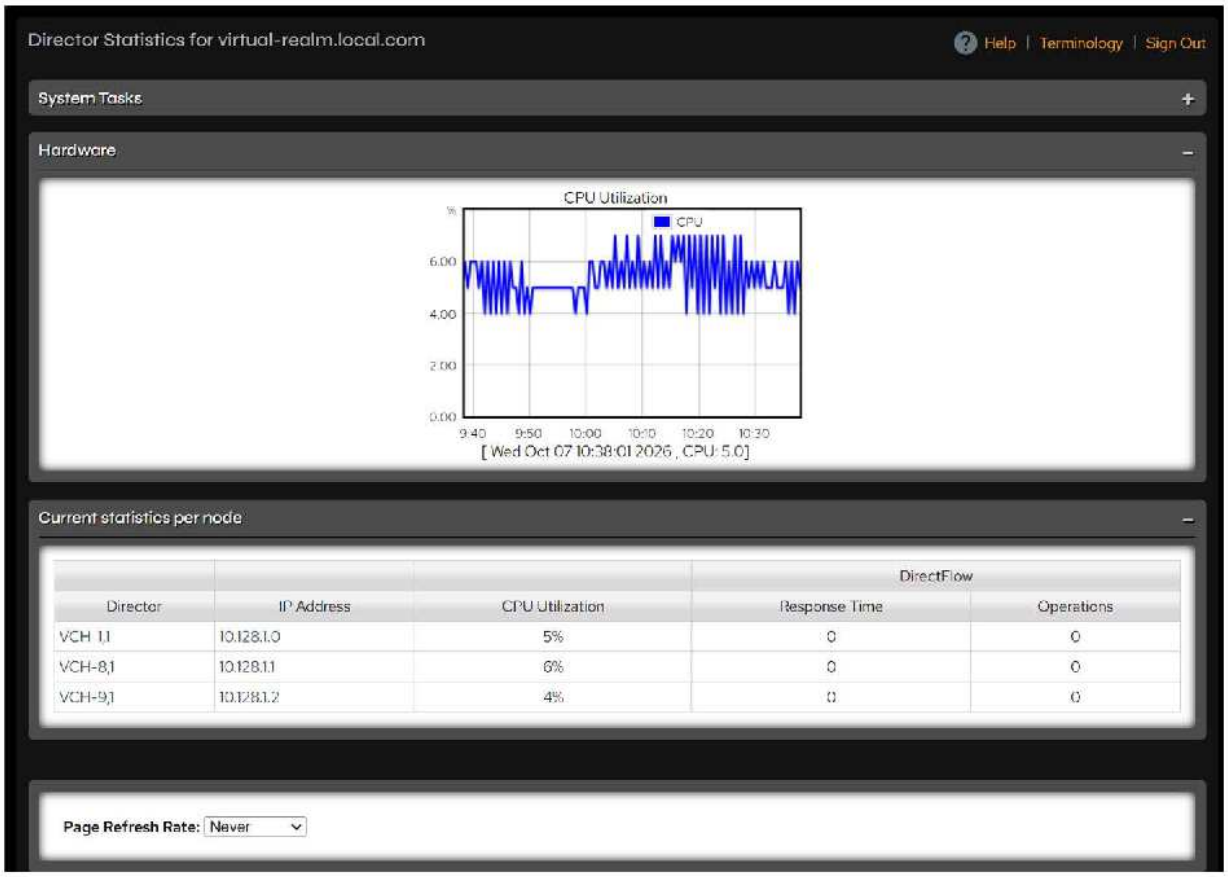


Figure 8-4. Director Statistics page.

The **Director Statistics** page provides the following information for Directors in your system.

Metric	Description
System Tasks	Click the **plus** icon in the right corner of the **System Tasks** heading to display any system tasks that are in progress, such as diagnostic processes. Overhead processes can explain unidentifiable spikes in performance and use.
Hardware (aggregate for all Director Nodes)	**CPU Utilization**: Percent of time the CPU is busy with storage operations.

Metric	Description
Current statistics per node (last 30 seconds)	**Director**: Click to open the **Node Management** page for the node. **IP Address**: Click to open the **Node Management** page for the node. **CPU Utilization**: Percent of time the CPU is busy with storage operations. The page displays the following DirectFlow statistics: **Response Time**: Response time in milliseconds for requests. **Operations**: Number of DirectFlow operations per second.

8.1.5 DirectFlow Clients

You can monitor abnormal or excessive usage patterns of DirectFlow clients. This feature is useful for determining which clients have been most active recently.

Click **Status & Reports > System Monitor > DirectFlow Clients**.

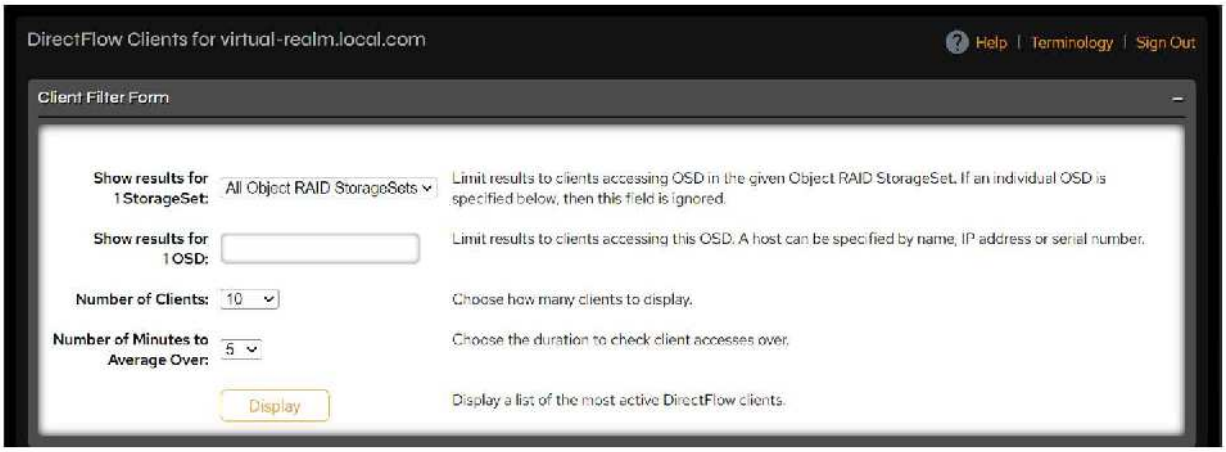


Figure 8-5. DirectFlow Clients page.

Using the Client Filter Form, select:

- The StorageSet to show results for, or an OSD name, IP address, or serial number to show results for a single OSD (the StorageSet selection is ignored if you enter an OSD).
- The number of clients to display (10, 20, 50, 100, 200, 500, or 1024).
- The number of minutes to average over (5, 10, 15, 30, or 60).

Click the **Display** button to display a list of the most active DirectFlow clients. Results appear below the **Client Filter Form**.

8.2 Create Reports

PanActive Manager provides the following reports that summarize current and historical system information:

- Performance Report
- System Statistics Report
- Capacity Report
- Asset Report
- Volume Report

VDURA uses industry standard metrics, such as operations per second, throughput, and response time, to measure the performance of the system. Users must be familiar with these performance metrics to interpret and use the reports.

8.2.1 Report Intervals

Most reports allow you to select a time interval for the report period. One method of specifying the interval is to select an integer number of hours, days, weeks, or months; the report will contain data for the most recent period selected (for example, the last six hours). A second method is to enter the start and end times of the interval. You can use any of the following date/time formats:

- A date in one of these forms:
  - *mm/dd/yy* (06/05/25)
  - *monthname dd*, *yyyy* (June 05, 2025)
  - *dd monthname yyyy* (05 June 2025)
  - *day, dd monthname yyyy* (Thursday, 05 June 2025)
  - *yyyymmdd* (20250605)
  - *yyyy-mm-dd* (2025-06-05)

The default year is the current year. If no time is specified, midnight is used.

- The time in the form *hh:mm:ss* or *hh:mm* followed by am or pm (2:30pm). If no am or pm is specified, *hh* is interpreted on a 24-hour clock. If no date is specified, the current date is used.
- A combination of date and time in any of the above forms (2:00pm 06/05/25).
- A date and time expressed in ISO 8601 point-in-time specification, such as *yyyymmdd*T*hhmmss* (“20250605T023000”), *yyyymmdd hhmmss*, or *yyyymmdd*T*hh*:*mm*:*ss*.

Export or Print Reports

Most reports allow you to export the report data as a comma-separated value (CSV) file, and some also offer a printable view. After generating the report, click the Export button to download the CSV file or, where available, the Printable View button to print the report.

> **TIP:** Depending on your browser settings, your browser might attempt to display the CSV file directly. If the CSV data does not display properly in your browser, use your browser's save-as feature to save the file.

8.2.2 Performance Report

The System Performance Report and Director Performance Report features have been deprecated. They have not yet been removed from PanActive Manager, so the Performance page is still displayed, but these reports are no longer supported.

The Performance page is available at **Status & Reports** > **Reports** > **Performance**.

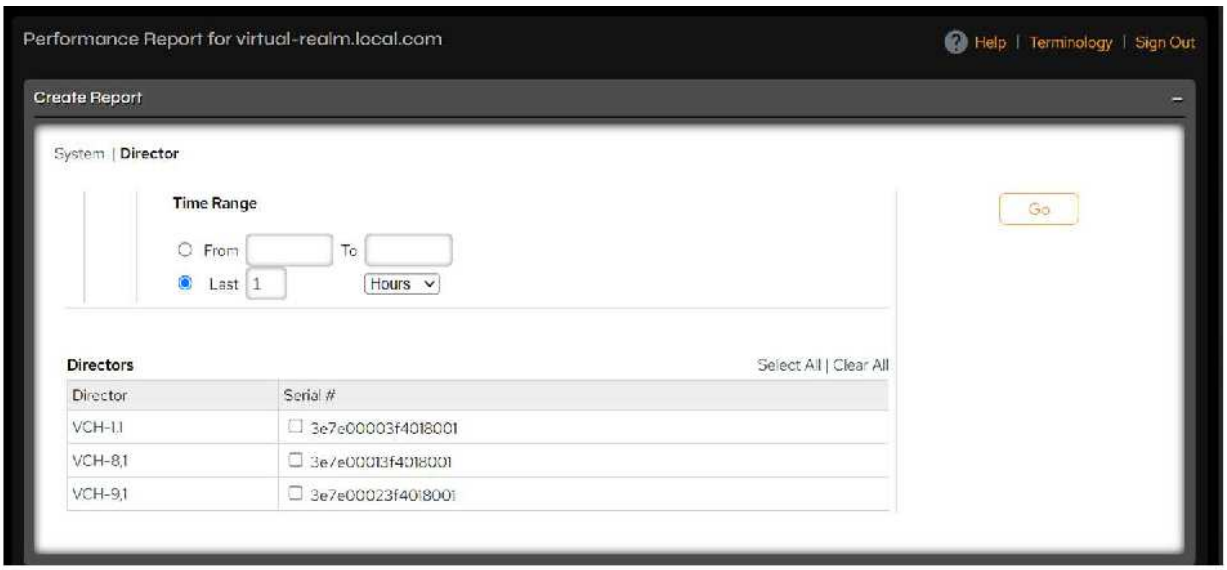


Figure 8-6. Performance Report page for a Director.

8.2.3 System Statistics Report

Use PanActive Manager to create a **System Statistics** report that can include the following metrics:

- **CPU Utilization**: Percent of time the CPU is busy with storage operations.
- **Disk Activity**: Percent of time the disk is busy with storage operations.
- **Operations**: Total number of operations per second.
- **OSD Response Time**: Response time in milliseconds of OSDs.
- **Total Throughput**: MB of data read from or written to disk or cache per second.
- **MB/s Write**: Total MB of data written to disk per second.
- **MB/s Read**: Total MB of data read from disk or cache per second.
- **DirectFlow Operations**: Number of DirectFlow operations per second.

- **DirectFlow MB/s Write**: MB of data written to disk per second.
- **DirectFlow MB/s Read**: MB of data read from disk or cache per second.

Click **Status & Reports > Reports > System Statistics**.

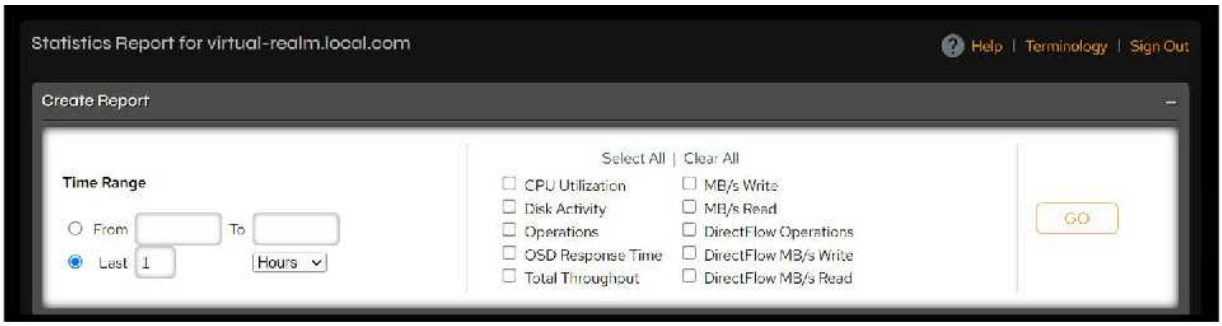


Figure 8-7. System Statistics Report page.

Select statistics information and time interval for your report. For information on allowable date/time formats, see Report Intervals.

Click **Go** to display your report.

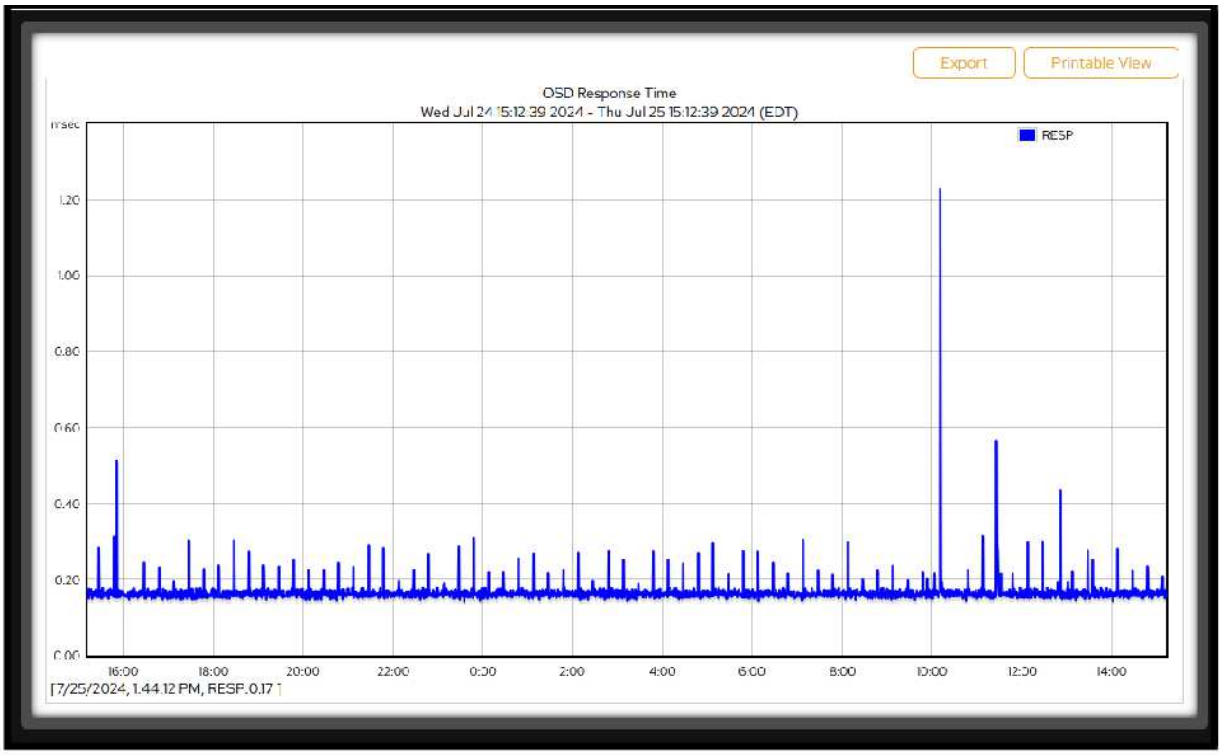


Figure 8-8. System Statistics – OSD Response Time output.

8.2.4 Capacity Report

Use PanActive Manager to create **System Capacity** and **StorageSet Capacity** reports including the following metrics:

- **Total Space**: Total capacity of the system or a StorageSet.
- **Used Space**: Utilized capacity of the system or a StorageSet.
- **Reserved Space**: Capacity reserved for use by reconstruction (related to the current spare count in each StorageSet).
- **Snapshot Space**: Capacity consumed by snapshots.

> **NOTE:** The VDURA operating system uses some disk space to store internal data, such as object metadata. This overhead is responsible for the fact that the available node capacity is lower than its total disk capacity.

Click **Status & Reports > Reports > Capacity**.



Figure 8-9. Capacity Report page for a StorageSet.

Make the following selections:

- **System** | **StorageSet**: Click to select report level.
- **StorageSet Name**: For StorageSet report, select a StorageSet.
- **Time Range**: For information on allowable date/time formats, see Report Intervals.

Click **Go** to display the report.

> **NOTE:** All capacity information is reported using base 10, not base 2; i.e., KB = 1000, not 1024.

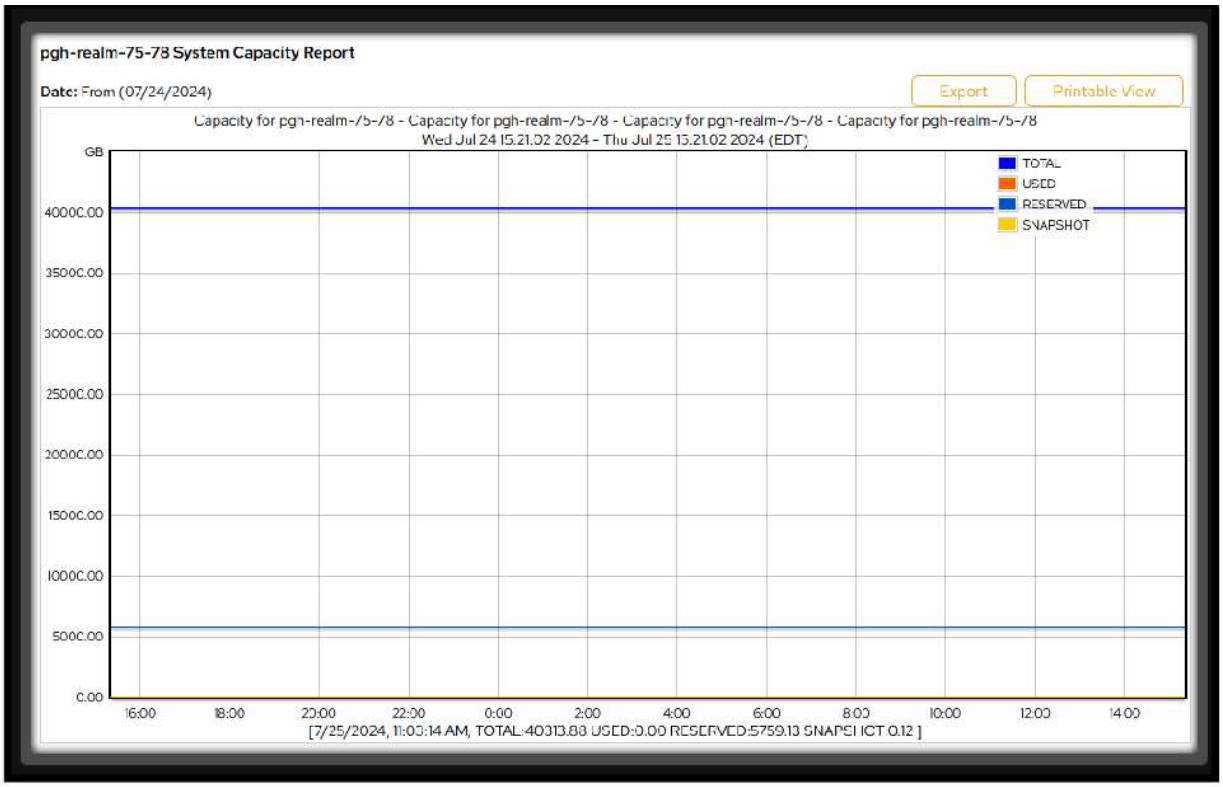


Figure 8-10. System Capacity Report output.

8.2.5 Asset Report

Use PanActive Manager to create a report of the hardware assets in your system.

Click **Status & Reports** > **Reports** > **Asset Listing** or **Hardware** > **Hardware Management** > **Asset Report**.

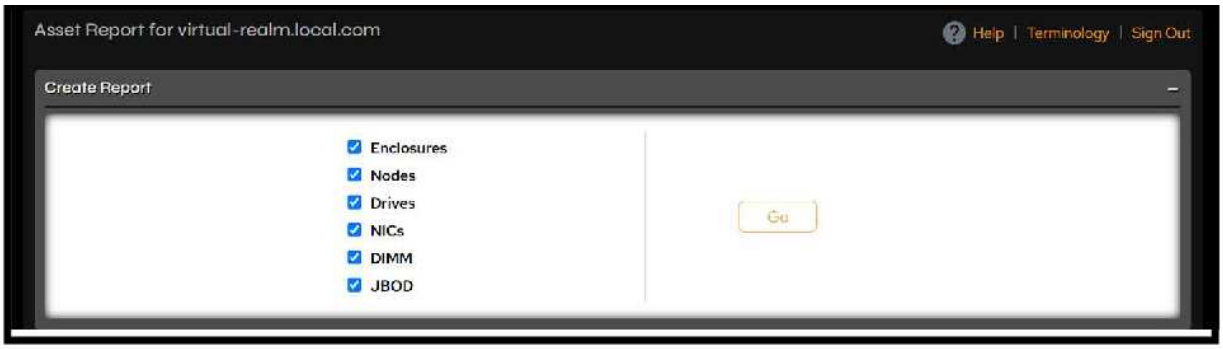


Figure 8-11. Asset Report page.

Select the assets to display (Enclosures, Nodes, Drives, NICs, DIMM, and JBOD), then click **Go** to display the report.

8.2.6 Volume Report

Use PanActive Manager to generate a report of the system's space allocation.

Click **Status & Reports** > **Reports** > **Volume Listing**.

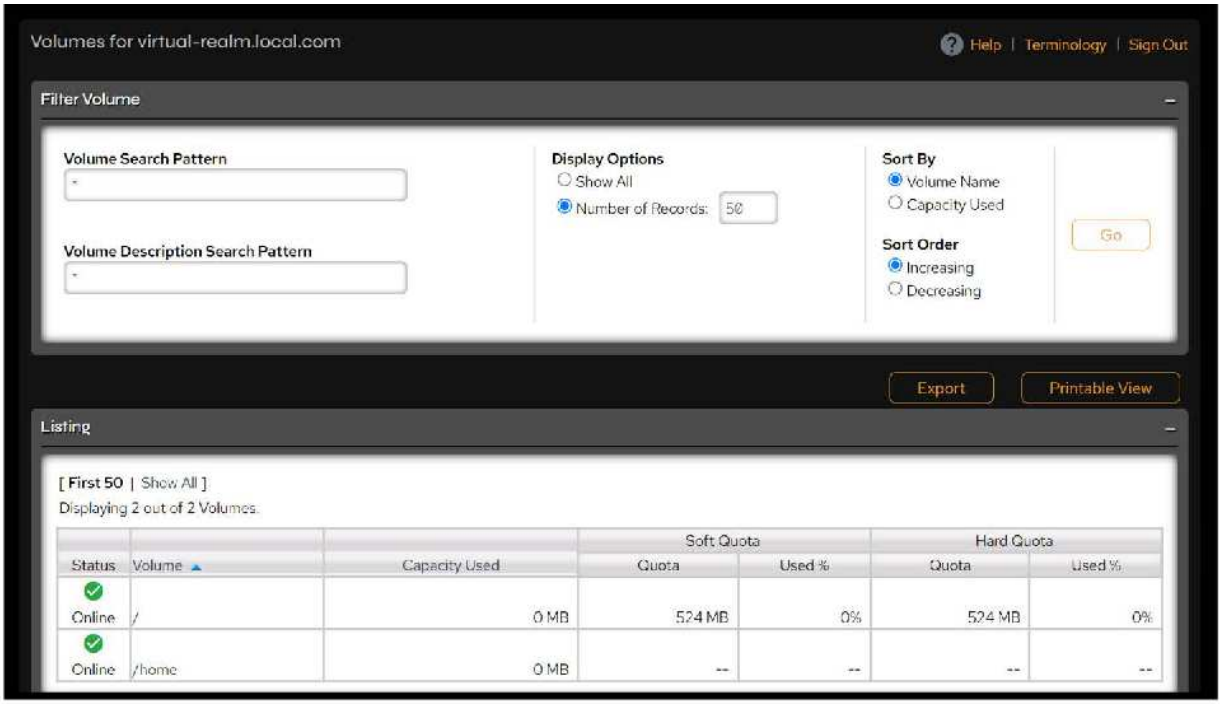


Figure 8-12. Volumes Report page.

Enter search strings and select display and sort options, then click **Go** to display the report.

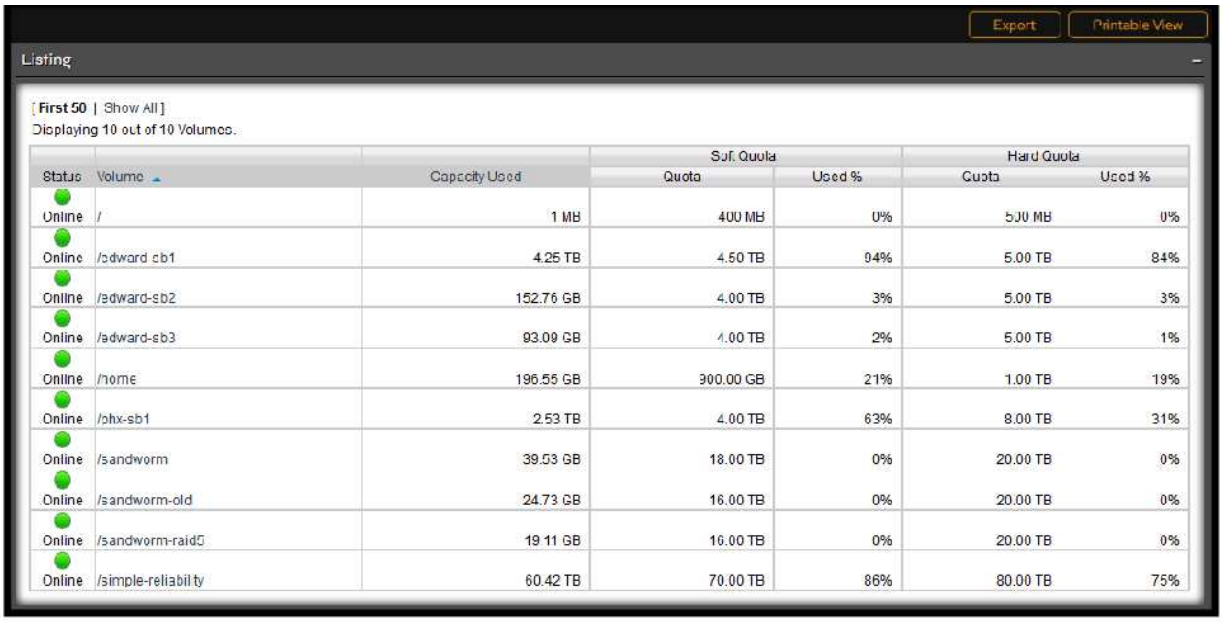


Figure 8-13. Volume Report output.

9. Upgrade VDURA Software

This chapter describes how to use PanActive Manager to upgrade your VDURA software. It includes the following topics:

- Upgrade System Software
- Upgrade and Uninstall DirectFlow Client Software

9.1 Upgrade System Software

VDURA provides two methods for accessing system software upgrades. PanActive Link automatically provides upgrade files. For instructions on activating PanActive Link, see Configure PanActive Link. If PanActive Link is not enabled, software can be downloaded from **support.VDURA.com** to a directory on your VDURA system for installation or to a remote host, then uploaded to your system and installed.

9.1.1 Before You Begin

The realm should be manually taken offline before starting an upgrade and should be manually brought back online when the upgrades for Directors, VPOD Containers, and all VPODs are complete. If this is not done, PanFS may attempt to start volume services while VPODs are updating, which can lead to extra errors in event log, alerts, and potentially false triggering of reconstruction.

From the **Hardware Management** page **Controls** section, click the **Take System Offline/Bring System Online** button. You will be prompted to confirm the action. Click **Yes** to continue.

The overall sequence is: take the system offline, install the upgrade, verify that all nodes have been upgraded, then bring the system back online.

> **NOTE:** For information on upgrading system software with the command-line interface, see the **upgrade** command in the *VDURA PanCLI Reference Guide*.

> **NOTE:** A VPOD Container is the part of a Storage Node that represents its hypervisor role and main OS (see VPODs). VPODs will be off-version upgraded after VPOD Container and Director upgrades are complete.

Access Available Upgrades

This feature has been deprecated in PanActive Manager as of PanFS 12.0.0 but has not yet been removed from the UI. Contact VDURA Technical Support to obtain upgrade images.

Click **Advanced Tools > Software Management > Upgrade PanFS**.

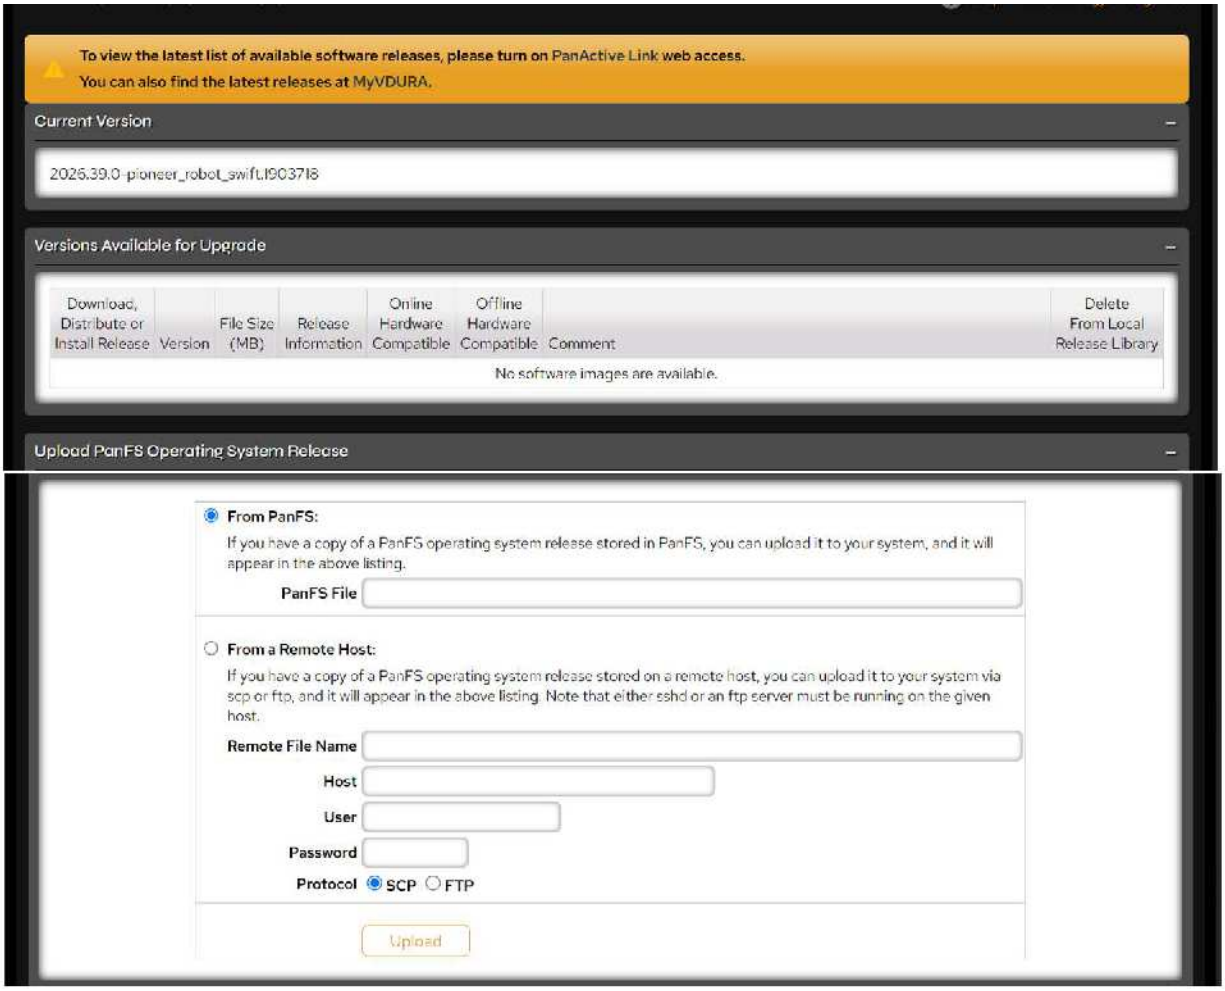


Figure 9-1. PanFS Operating System Upgrade page.

If PanActive Link is enabled, the **Versions Available for Upgrade** table will automatically be populated with available upgrades. Pay particular attention to the columns that indicate compatibility with your hardware. If the upgrade is not compatible with existing hardware, the column provides a link to a table of software versions that are compatible.

If PanActive Link is disabled, either enable it from the **Configuration > Notification > PanActive Link** page or go to the **Downloads** page of **support.VDURA.com**, download a compatible software version to the VDURA system or a remote host. Use the **Upload PanFS Operating System Release** options on

the **PanFS Operating System Upgrade** page to upload the software. The new version will then appear in the **Versions Available for Upgrade** table.

> **NOTE:** Remote access via FTP or SCP requires that the remote system be running an FTP or SSH service.

Install the Upgrade

Once you have identified a compatible software upgrade in the **Versions Available for Upgrade** table, click its **Download** button to download it to a single node in the system. This step can be completed while the system is online.

When the download is successfully completed, the **Download** button becomes a **Distribute** button. Click the **Distribute** button to distribute the upgrade package to all other nodes in the system.

If the distribution is unsuccessful, an error is displayed on the **PanFS Operating System Upgrade** page and a control appears in the **Versions Available** table allowing you to redistribute the software to all nodes or specific ones identified by IP address. Select the nodes to which you want the software redistributed, then click **Redistribute to Selected** at the bottom of the list.

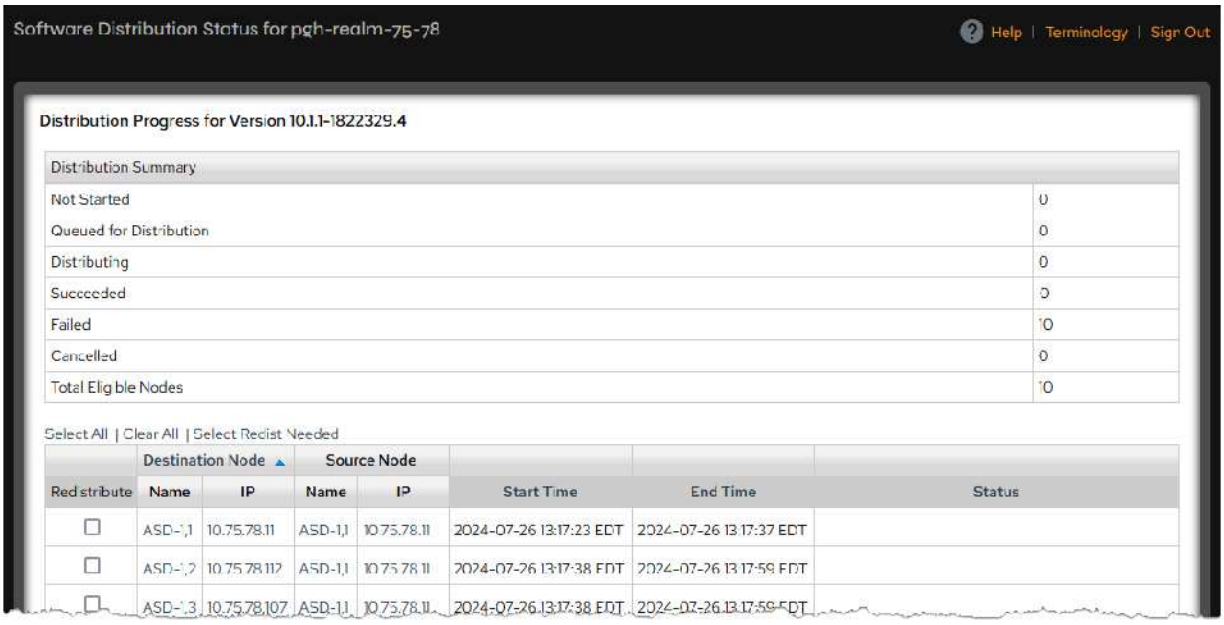


Figure 9-2. Software Distribution Status with Redistribute boxes.

You can also go to the **Distribution Status** page to view the status of the nodes and select ones to redistribute. This step can be completed while the system is online.

Once the image has successfully been distributed, the **Distribute** button becomes an **Install** button. You are now ready to unpack, verify, and install the software. This step requires the system to be offline.

When the installation is successful on all nodes, the system will reboot and all nodes will come back online. If the installation fails, the system will revert to the current version of the software.

Check Upgrade Status

The status of the software distribution can be checked from **Advanced Tools > Software Management > Distribution Status**.

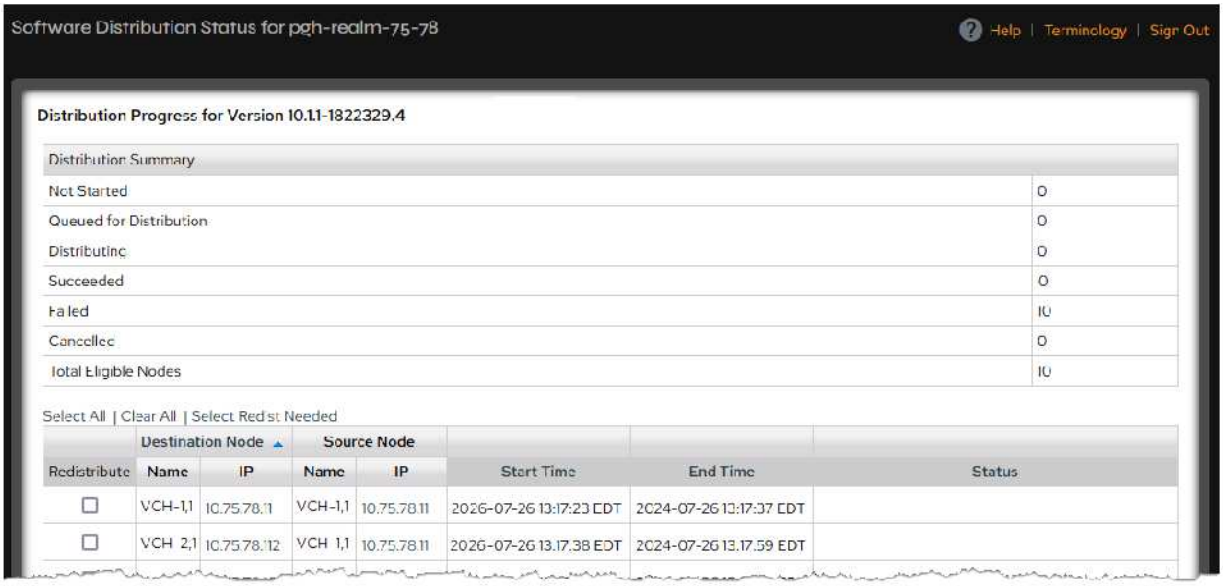


Figure 9-3. Software Distribution Status page with Distribution Summary at top.

View Upgrade History

Click **Advanced Tools > Software Management > Version History** to see the recent history of software upgrades and installation dates.



Figure 9-4. PanFS Operating System Version History page.

Roll Back to a Previous Version

The previous version of the VDURA software is retained on the system in the event a problem requires you to back out of an upgrade. This version will be listed in the **Versions Available for Upgrade** table on the **PanFS Operating System Upgrade** page. In general, VDURA supports downgrading by one release. However, always check with VDURA for any limitations or exceptions before initiating a downgrade.

To roll back to an earlier version, follow the same instructions used to upgrade to a newer version. Since the software is already on the system, start with the distribution step.

Upgrade Times

Upgrade and downgrade times will vary, but in general, the following holds true:

- The download and distribute phases take time, but the system is fully online during these steps. Since both download and distribution are just single-click steps, there is also practically no administrative time required.
- The time it takes to download a release depends on the method. Downloading to each node directly from the VDURA support portal takes longer than copying the file (downloaded from the support portal separately) into a volume, then downloading from that volume.
- Distribution time depends on the number of nodes in the system. VDURA uses an exponential distribution method to hasten this process: one node gets the image and sends it to two other nodes, then each of those three nodes send a copy to two additional nodes, then those nine nodes each then send to two additional nodes, etc.

When the install button is pushed, the system does some work to validate packages, etc., before it goes offline and reboots. At this point, clients will experience the system hanging. This step will take about ten minutes to complete before coming back online.

9.2 Upgrade and Uninstall DirectFlow Client Software

For instructions on installing and configuring the DirectFlow Client, see DirectFlow Client Support.

9.2.1 Upgrade DirectFlow Client Software

Complete the following steps to upgrade the DirectFlow Client software while retaining your current **/etc/fstab** entries for VDURA mount points. This procedure is for RPM files only. Customers with Debian-based systems are advised to create a backup copy of **/etc/fstab** prior to reinstalling or upgrading the DirectFlow Client.

- Download the appropriate RPM for the client from **support.VDURA.com** and copy the RPM to the client. If you are unsure of the exact name of your RPM file, contact VDURA Technical Support before upgrading the DirectFlow Client.
- Run **umount /panfs** to unmount the VDURA file system.
- Run **rmmod /panfs** to unload the DirectFlow kernel module.
- Run **rpm -ivh [panfs*.rpm]**, where **[panfs*.rpm]** is the name of the RPM file. For example, using the command with a file named **panfs-5.14.0-**

  570. 12.1.el9_6.x86_64-11.1.1-1853317.62.rhel_9_x86_64.rpm** looks like this:

  ```
  rpm -ivh panfs-5.14.0-570.12.1.el9_6.x86_64-11.1.1-
  1853317.62.rhel_9_x86_64.rpm
  ```

This updates the DirectFlow Client software and re-installs the VDURA kernel module.

- Use **mount** to mount your VDURA mount points.

Uninstall DirectFlow Client Software

If you have a previous version of the DirectFlow Client software on your system, you can uninstall it with the following procedure:

- Run **umount /panfs** to unmount the VDURA file system.
- Run **rmmod /panfs** to unload the DirectFlow kernel module.
- Run **rpm -e [panfs*.rpm]**, where **[panfs*.rpm]** is the name of the old version, to uninstall the DirectFlow Client software.

10. End-to-End Encryption

This chapter provides information about and instructions for managing VDURA's end-to-end encryption implementation.

> **NOTE:** End-to-end encryption works only on the platform tenant, and only in a realm that does not use multi-tenancy. It cannot be used with tenants other than the platform tenant.

It includes the following topics:

- Introduction and Overview
- Key Management
- DirectFlow Client Behavior
- Gateway Protocols
- Configuration Considerations
- Failure Modes
- Performance Considerations
- Upgrades and Downgrades
- Encryption Configuration Procedures

10.1 Introduction and Overview

To support a more secure environment for processing data, the DirectFlow Client supports an end-to-end encryption feature. This enables data to be encrypted and decrypted at the client, protecting user data from a range of attacks. VDURA's encryption keeps user data secure once it leaves the client and not available at any point past that.

The security of user data is a top concern for storage systems. It is important to understand the threats that these mechanisms are protecting against. VDURA's encryption implementation protects user data from:

- Snooping on the wire.
- Theft of any components of the system (servers, storage media, JBODs).
- Hacking into the storage software or device firmware.

10.1.1 What is Encrypted

The primary goal of this feature is to encrypt user data (file contents). This means that every byte of user data that passes from a DirectFlow Client to Storage Nodes (VPODs) is encrypted. Note that the protocol data units that transport user data to the Storage Nodes are not encrypted. RPC messages between the DirectFlow Client and the realm are also unencrypted.

How Data is Encrypted

The DirectFlow Client uses the AES-XTS cryptographic algorithm. This algorithm is FIPS-140 approved with key size of 256 bits. Data is encrypted and decrypted within the context of the DirectFlow Client. The Storage Nodes do not have the keys and do not attempt to decrypt data.

Metadata Encryption

Currently, metadata is not encrypted. Since metadata must be interpreted by the system, the goal of keeping data encrypted once it leaves the client would become compromised, as more parts of the system would need keys to allow for access to and interpretation of the data.

10.2 Key Management

Clients require keys to allow them to encrypt and decrypt data. These keys are managed by an external key management system (KMS). Directors must be able to create and delete keys from the KMS, and the DirectFlow Client must be able to retrieve keys.

10.2.1 Role-Based Authentication

The KMS grants keys to DirectFlow Clients to mutually authenticate with each other. To do this, each DirectFlow Client node requires a digital certificate to present to the KMS. Creation and provision of this certificate is managed by the customer. This certificate must be installed on the DirectFlow Client node by the customer for PanFS software to use. A DirectFlow Client with a certificate has the authority to retrieve any encryption key from the KMS for any encrypted volume.

10.2.2 Encryption Domain

VDURA's encryption implementation provides encryption at PanFS volume-level granularity. Encryption can be enabled when the volume is created and cannot be changed for the lifetime of the volume. A volume with encryption enabled will have a Volume Data Key (VDK) which will be used to encrypt and decrypt all file data in the volume. This VDK cannot be changed, as that would require re-encryption of all data in the volume, which is impractical.

VDURA system administrators can create and delete encrypted volumes through the PanCLI interface via an option on the command line when a volume is created. Since encryption is an optional feature, administrators can continue to create standard, or unencrypted, volumes.

10.2.3 KMS Integration

VDURA's encryption implementation implements a KMIP protocol (version 1.4 or better).

10.3 DirectFlow Client Behavior

A client connected to the VDURA system requires two pieces of software for encryption. First is the PanFS Client, which is a kernel module that is loaded and runs the client side of the file system. This is the software that performs encryption and decryption of user data. The second piece is a userspace module that allows for connection to the KMIP server. This module retrieves keys on behalf of the DirectFlow Client. To do that, it requires credentials to present to the KMS.

10.4 Gateway Protocols

While VDURA offers the DirectFlow Client for a high-performance, parallel data path, S3 is also available to access data. To enable S3, Directors act as a gateway hosting the service. S3 accesses data through a PanFS Client that runs on the Director, called the Gateway Client. As of PanFS 12.0.0, S3 configuration is done through PanCTL.

The Gateway Client currently does not support encryption. Therefore, no gateway protocols are supported. For S3, if encryption is enabled, then S3 support is disabled.

10.5 Configuration Considerations

Several configuration options must be set to allow a user to manage and control a system with end-to-end encryption enabled.

10.5.1 External KMS Configuration

The external KMS is configured via PanCLI. Configuration requires the following information:

- IP address or fully qualified domain name
- Port
- KMIP client credentials
- Certificates

See KMIP Server Configuration below for instructions on using PanCLI commands to configure the KMIP server.

10.5.2 Key Management

Each encrypted volume has a volume Data Encryption Key (DEK). The DEK remains active in the KMS for the lifetime of the volume and its value cannot be changed. When a DirectFlow Client needs a volume's encryption key, it requests and receives the DEK directly.

10.5.3 Encrypted Volume Configuration

When creating a volume in PanActive Manager, there is an advanced configuration item that allows you to select **Encrypt Volume**. This option is dependent on a KMS being configured; if there is no KMS, then this will not be visible.

10.5.4 Client Configuration

Credentials for each DirectFlow Client must be provided by the customer from a Certificate Authority. A client with proper credentials can then retrieve any encrypted volume's key from the KMS.

Mount Options

DirectFlow Clients must be told the hostname or IP address of the KMS, as well as other relevant information to contact the KMS, via a mount option. See Client Interface Configuration below for mount options to explicitly indicate that an encrypted volume is being mounted.

10.5.5 Secure Status

By querying the current state of the VDURA system via the **syscontrol encryption** and **volume list showencryption** PanCLI commands, the administrator can see what parts of the system are encrypted.

10.5.6 Mounting PanFS

When the **kmip-config-file** mount option is specified (see Client Interface Configuration), the DirectFlow Client will deny PanFS global mounts (/panfs) and realm root mounts. Attempting to do a global mount or mounting any path inside the realm root volume will result in an error returned for the mount attempt; only direct volume mounts will be allowed. Attempting to mount a volume that is encrypted without explicitly providing the **kmip-config-file** mount option will also result in a failure.

10.6 Failure Modes

Several failure modes are possible when end-to-end encryption is enabled on a VDURA system:

- KMS network connection issues.
- Lost or destroyed keys.

- Compromised keys.
- Accessing data with no key.
- Accessing data with the wrong key.

10.6.1 KMS Network Connection Issues

If the KMS is unavailable, then the data will not be available. Failure to communicate with the KMS or failure to retrieve keys will result in the DirectFlow Client returning an error to the application.

10.6.2 Lost or Destroyed Keys

If keys are lost or destroyed, then data in an encrypted volume will be permanently inaccessible.

10.6.3 Compromised Keys

No provision is made to deal with compromised keys.

10.6.4 Reading Data without a Key

If a DirectFlow Client attempts to read data without a key, an error will be returned in a manner similar to a lost or destroyed key scenario.

10.6.5 Reading Data with the Wrong Key

AES-XTS does not allow for detection of an incorrect key used for decryption. Therefore, a client using the wrong key to decrypt data will result in random data being returned to the application.

10.7 Performance Considerations

There are two main points to keep in mind when considering the implications of encryption and decryption on performance: the time it takes for the DirectFlow Client to retrieve a key relevant to the data and CPU overhead.

- **Key retrieval overhead**: Users should expect key retrieval overhead time will be generally minimal. Instances where it is not should be infrequent.
- **CPU overhead**: Various software algorithms must be run to encrypt and decrypt data. These algorithms consume CPU and memory cycles on the client systems.

10.8 Upgrades and Downgrades

Downgrading a system with encrypted volumes to a version that does not support encryption will result in volume data being permanently inaccessible.

10.9 Encryption Configuration Procedures

Administrators must use PanCLI to configure and interact with the end-to-end encryption feature; the PanActive Manager GUI does not offer any visibility into the feature.

10.9.1 KMIP Server Configuration

Use the PanCLI **kmip** command for KMIP server configuration.

The following information should be set with the **kmip config set** command:

- KMIP server host (IP address or hostname)
- KMIP server port
- CA certificate (location of root certificate PEM file)
- Client certificate (location of client certificate PEM file; used by Director Nodes)
- Client private key (location of client private key PEM file)
- KMIP protocol version to be used by Director Nodes (KMIP version number)

Below is an example series of **kmip config set** configuration commands. All files (such as **/director.pem**) should be uploaded to the Director Node by the PanFS administrator prior to running.

```
kmip config set host 10.70.25.173
kmip config set port 5696
kmip config set ca /RootCertificate.pem
kmip config set client /director.pem
kmip config set key /director.key
kmip config set kmip_version 1.4
kmip config show
kmip config validate
kmip config save
```

Display the active KMIP server configuration:

```
kmip server show
```

Delete the active KMIP server configuration:

```
kmip server delete
```

10.9.2 Realm-Wide Encryption Configuration

Use the PanCLI **syscontrol encryption** command to change realm-wide, default end-to-end encryption behavior when new volumes are created.

The system-wide encryption setting defines default encryption status for new volumes in the realm. By default, new volumes follow this realm-wide encryption setting; enabling encryption does not affect existing non-encrypted volumes. However, if there are encrypted volumes in the system, it is not possible to turn this setting off. Encryption cannot be enabled if S3 is enabled. A valid KMIP configuration should be set as a prerequisite to enabling encryption.

Enable system-wide volume encryption for new volumes:

```
syscontrol encryption enable
```

Disable system-wide volume encryption:

```
syscontrol encryption disable
```

View the current encryption setting:

```
syscontrol encryption
```

10.9.3 Volume-Level Encryption Configuration

Create a Volume with/without Encryption

If system-wide encryption is off by default, use the PanCLI **volume create** command to create a volume with encryption enabled:

```
[pancli] volume create encryption on
```

If system-wide encryption is on by default, use the PanCLI **volume create** command to create a volume with encryption disabled:

```
[pancli] volume create encryption off
```

Show Volume Encryption Status

Use the PanCLI **volume list** command to show a volume's encryption state:

```
[pancli] volume list showencryption

E2EE   Space   Soft   Hard
Volume   StorageSet  RAID           State   Used   Quota   Quota  Status
/        Set 1     Object RAID6+  --      0 MB  524 MB  524 MB  Online
Jencvol  Set 1     Object RAID6+  On      0 MB      --      --  Online
/home    Set 1     Object RAID6+  Off     0 MB      --      --  Online
Displayed 3 out of 3 volumes in the system.
```

10.9.4 Client Interface Configuration

When mounting the DirectFlow Client, credentials must be set to allow KMS access. For the credential requirements, see Client Configuration.

Configuring a KMIP server on a DirectFlow Client assumes you have the DirectFlow Client PanFS RPMs already installed and PanFS is not mounted yet.

1. Copy the required two certificate files and one key file to a location of your choice on the DirectFlow Client. In this example, **/tmp** is used:

   ```
   cp /path/to/certs/RootCertificate.pem /tmp
   cp /path/to/certs/dfc_client.pem /tmp
   cp /path/to/certs/dfc_client.key /tmp
   ```

2. Create a KMS config file named **/tmp/kmip_config**. This file is passed to mount so the mount binary knows the KMS config. Below is an example of file contents:

```
config_version = 1

[global]
primary-server = server:dev

[server:dev]
host = <kms_server_ip>
port = 5696
version = 1.4

connect-timeout = 300
reconnect-retries = 5
ca-cert-path = /tmp/RootCertificate.pem
client-cert-path = /tmp/dfc_client.pem
client-key-path = /tmp/dfc_client.key
```

3. Create a directory to mount under. You cannot use **/panfs** because global mounts are not allowed:

```
mkdir /mnt/enc_vol
```

4. Locally mount an encrypted volume, specifying the KMS config file:

```
mount -t panfs -o kmip-config-file=/tmp/kmip_config panfs://<realm_ip>/enc_vol /mnt/enc_vol
```

11. Kubernetes CSI Driver for PanFS

This chapter explains how to configure and use CSI with PanFS. It contains the following sections:

- The PanFS CSI Driver
- Installation, Configuration, and Diagnostics

11.1 The PanFS CSI Driver

Modern containerized workloads often require persistent storage that remains available beyond the lifecycle of individual pods. Kubernetes provides a standardized mechanism for managing such storage through the Container Storage Interface (CSI), an open, community-driven specification that allows any compliant storage system to integrate seamlessly with the Kubernetes ecosystem.

With CSI, Kubernetes takes full responsibility for orchestrating the storage lifecycle: creating volumes when **PersistentVolumeClaims** (PVCs) are submitted, attaching them to the appropriate worker nodes, mounting them into workload pods, and later detaching or deleting them as needed. This eliminates manual intervention at the storage backend, enabling developers and operators to define storage requirements declaratively while letting the cluster handle the details.

In this model, the CSI driver serves as the bridge between Kubernetes and the underlying storage system, translating Kubernetes CSI API calls into the native operations required by the backend. Importantly, when using the PanFS/DFC CSI driver, there is no need to pre-create volumes manually within the PanFS realm; the driver will dynamically create and configure them on demand, based on the parameters defined in each PVC.

By implementing a CSI driver for PanFS storage, Kubernetes workloads can take full advantage of the performance, scalability, and reliability of the backend without custom scripts or manual configuration, aligning the storage experience with Kubernetes' native workflows and best practices.

The PanFS CSI Driver integrates Kubernetes with the PanFS/DFC storage backend, enabling seamless, Kubernetes-native dynamic provisioning, attachment, and lifecycle management of volumes without manual creation in the PanFS realm.

11.2 Installation, Configuration, and Diagnostics

The most recent files and documentation can be found at the links below.

> **IMPORTANT:** GitHub Personal Access Token (PAT) is required for Kubernetes to authenticate with GitHub container registry. Follow the instructions described in GitHub documentation to create a PAT with at least read permissions.

Basic information to get started:

https://github.com/panasasinc/panfs-container-storage-interface-oss

PanFS CSI Driver Architecture Overview:

https://github.com/panasasinc/panfs-container-storage-interface-oss/blob/main/docs/Overview.md

CSI Driver Installation - Quick Start:

https://github.com/panasasinc/panfs-container-storage-interface-oss#csi-driver-installation---quick-start

PanFS CSI Driver & StorageClass Deployment Guide:

https://github.com/panasasinc/panfs-container-storage-interface-oss/tree/main/deploy/k8s

PanFS CSI Driver Diagnostic Commands Reference:

https://github.com/panasasinc/panfs-container-storage-interface-oss/blob/main/docs/Diagnostic.md

PanFS CSI Driver Upgrade Guide:

https://github.com/panasasinc/panfs-container-storage-interface-oss/blob/main/docs/Upgrade.md

Kernel Module Management (KMM) Guide:

https://github.com/panasasinc/panfs-container-storage-interface-oss/blob/main/docs/KMM.md

PanFS CSI Driver Volume Encryption Guide:

https://github.com/panasasinc/panfs-container-storage-interface-oss/blob/main/docs/Encryption.md

SELinux Configuration Reference for Kubernetes CSI Drivers:

https://github.com/panasasinc/panfs-container-storage-interface-oss/blob/main/docs/SELinux.md

PanFS CSI Driver Usage Guide:

https://github.com/panasasinc/panfs-container-storage-interface-oss/blob/main/docs/Usage-Guide.md

12. Remote Direct Memory Access in VDURA Systems

This chapter explains the VDURA implementation of Remote Direct Memory Access (RDMA). This chapter includes the following topics:

- Introduction
- Prerequisites
- Enable or Disable RDMA on the Realm
- Mount the RDMA Client

12.1 Introduction

PanFS can be configured to move bulk read/write data between a client and Storage Nodes over RDMA instead of TCP. Each client-to-node session keeps its TCP connection for control traffic, but adds an RDMA connection for data. Currently, VDURA systems provide one cluster-wide on/off switch for a single RDMA device per client and two client mount options. There are no other tunable parameters.

RDMA is turned off by default and must be enabled in two places: once cluster-wide on the realm, and once per client at mount time.

12.2 Prerequisites

To use RDMA with your VDURA system, you must have an RDMA-capable device on the client, e.g., an NVIDIA/Mellanox adapter (device names such as mlx5_0 and mlx5_1). The RDMA device must sit on the network interface that carries the storage traffic. If it is on a different interface, RDMA mounts fail even though the device exists.

Check what the client sees:

```
rdma link
```

The output should resemble the following, where the device name is listed before the "/":

```
link mlx5_0/1 state ACTIVE physical_state LINK_UP netdev ens1f0
```

In this example, **mlx5_0** is the device name.

12.3 Enable or Disable RDMA on the Realm

RDMA is disabled by default on a freshly initialized realm. Follow the steps below to enable it.

1. Enter service mode:

```
[pancli] support servicemode yes
```

2. Issue the command to enable RDMA:

```
[pancli_service] mconsole_feature RDMA on
```

If you wish to disable RDMA, use the **off** option instead:

```
[pancli_service] mconsole_feature RDMA off
```

3. Verify the current state:

```
[pancli_service] mconsole_feature
```

This command lists every feature; look for the RDMA row.

4. Leave service mode:

```
[pancli_service] support servicemode no
```

If you are logged in to a bash shell on a Director instead of the PanCLI shell, you can run a command in service mode directly, for example: **pancli -s -c mconsole_feature RDMA on**

Be aware of the following points when enabling or disabling RDMA.

- The only relevant RDMA options are **on** and **off**. For more information, see the fifth point in Known Limitations.
- Setting the state to what it is already results in a “No change to RDMA needed.” message.
- Changing the state restarts services and the file system is briefly unavailable. Wait for the realm to come back online before mounting volumes or creating new ones.

12.4 Mount the RDMA Client

Use **rdma-mode** to mount the PanFS file system. Optionally, you can select a specific RDMA device with **rdma-dev**.

12.4.1 Mount Options

RDMA options behave identically for volume mounts, realm-root mounts (**panfs:// [IP]/**), and global mounts (**panfs://[IP]:global**).

Option	Description
`rdma-dev [RDMA device name, e.g. mlx5_0]`	By default, this is not set; the client selects a device. When **rdma-mode** is set to off, **rdma-dev** is accepted but ignored.

Option	Description
`rdma-mode off auto enforce`	Any **rdma-mode** value other than the three listed below will fail the mount. **off**: • **Off** is the default setting. • If RDMA is available, TCP is used and no RDMA connection is made. • If RDMA is unavailable, or **rdma-dev** names a device that does not exist, TCP is used. • If the realm RDMA feature is turned off, TCP is used. **auto**: • If RDMA is available, RDMA is used. • If RDMA is unavailable, the mount succeeds but falls back to TCP with no error. • If the realm RDMA feature is turned off, the mount succeeds and uses TCP. **enforce**: • If RDMA is available, RDMA is used. • If RDMA is unavailable, the mount fails. • If the realm RDMA feature is turned off, the mount fails. For more information, see the third point in Known Limitations.

Examples:

```
mount -t panfs -o rdma-mode=auto panfs://<director_ip>/
<volume> /mnt/<volume>

mount -t panfs -o rdma-mode=enforce,rdma-dev=mlx5_0 panfs://
<director_ip>/<volume> /mnt/<volume>

mount -t panfs -o rdma-mode=auto panfs://<director_ip>/ /mnt/
panfs
```

12.4.2 Error Messages

You may see several error messages if RDMA is incorrectly configured.

- **If there is an invalid mode value**:

```
mount.panfs: invalid rdma-mode 'test'; valid values are
'auto', 'off', or 'enforce'
```

- **If enforce is used with no usable RDMA**:

```
mount.panfs: Invalid argument: inspect options and try again.
/var/log/messages may have more details about the mount
failure.
```

See **/var/log/messages** or **dmesg** for detailed information:

```
pan_rdma_device_init: no RDMA-capable device found on any
network interface
RDMA: init failed, mode=enforce, aborting, 0x7 (Invalid
argument)
```

- **If a mount inherits settings from an earlier PanFS mount, mount.panfs prints a warning beginning with:**

```
warning: inheriting the rdma-mode=… setting from existing
PanFS mount; per-mount RDMA options are not supported
```

For more information, see the first point in Known Limitations.

12.4.3 Known Limitations

Be aware that there are currently several limitations to this implementation.

- RDMA settings are per client, not per mount. The first PanFS mount on a client fixes **rdma-mode** and **rdma-dev** for the whole client. Every later mount inherits them and silently ignores its own values. To change the setting, unmount every PanFS file system on that client first.
- Volumes with a stripe unit smaller than 64 KiB fall back to TCP for data I/O, even with **rdma-mode=enforce**.

- Currently, **rdma-mode=enforce** succeeds even when the realm RDMA feature is off, and subsequent writes hang. Ideally, it should fail the mount.
- Automatic device selection is unreliable on some hardware realms; **rdma-mode=enforce** fails unless **rdma-dev** is used.
- Currently, **mconsole_feature** treats anything that is not **on** as **off**. For example, trying to use **mconsole_feature RDMA enable** results in a "Turned RDMA off." message instead of reporting an invalid value.
- Toggling the RDMA feature interrupts the file system. Operations issued immediately afterwards can fail. For example volume creation may fail resulting in an "rmm: the realm root volume is not online" error.

13. Troubleshoot Your System

This chapter describes how to use the diagnostic and troubleshooting tools in PanActive Manager to troubleshoot your system. It includes the following sections:

- System Registration
- View Events
- Perform Network Diagnostic Tests
- Perform System Diagnostics
- Troubleshooting

13.1 System Registration

If you did not complete the registration of your system during installation, you can register using PanActive Manager.

Navigate to **Configuration > Registration** to fill out the registration form in PanActive Manager or confirm registration by phone.

- Select the **Register by Phone** (1-888-726-2727) radio button and click **Save**.
- Select the **Register Now** radio button. Fill in all required fields on the form and click **Save**.

> **NOTE:** Once registration is completed, the **Registration** item will no longer be visible in the **Configuration** menu.

13.2 View Events

PanActive Manager's **Historical Event Log** allows the user to view historical events such as system errors, warnings, system changes, and administrator actions.

> **NOTE:** For information on viewing events with the command-line interface, see the **eventlog** command in the *VDURA PanCLI Reference Guide*.

Events displayed in the **Historical Event Log** are categorized as follows:

Icon	Status	Description
	Information	Error recovery messages, system changes, and administrator actions.
	**Warning**	Developing problems that might impair system availability or data integrity. For information on setting email recipients for warning messages, see Set Status Recipients.
	**Critical**	Problems that are an imminent risk to system availability or data integrity. For information on setting email recipients for critical messages, see Set Alert Recipients.

For a description of VDURA information and error messages, see VDURA System Error Messages.

Click **Status & Reports > Historical Event Log**.

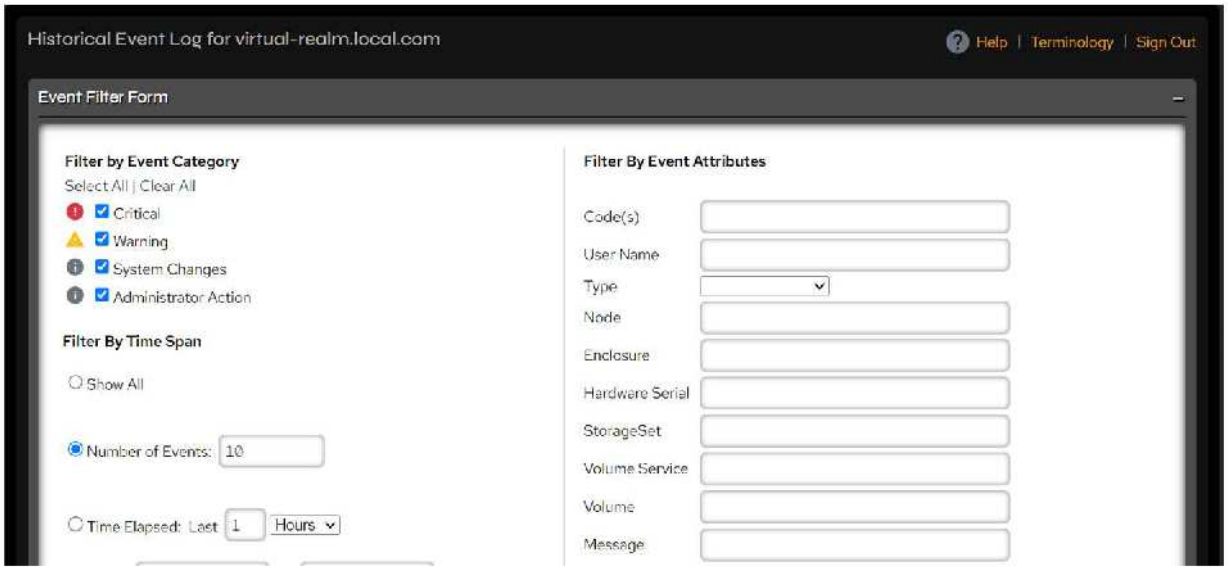


Figure 13-1. Historical Event Log page.

Use the filter options to customize the view:

Filter	Description
Filter by Event Category	**Critical**: View critical errors. Errors are critical if system availability or data integrity is at risk. **Warning**: View error warnings. Warnings may indicate a developing problem. **System Changes**: View system changes such as reconfigurations. **Administrator Action**: View administrator actions such as logins.

Filter	Description
Filter By Time Span	**Show All**: View errors, warnings, system changes, and administrator actions. **Number of Events**: Enter a number of events to display beginning with the most recent. **Time Elapsed**: Enter a number of the past hours, days, or weeks for which to display events. **From/To**: Enter a date interval, using one of the formats described in Report Intervals.

Filter	Description
Filter By Event Attributes	**Code(s)**: Limit results to events with the given event code. Multiple code options may be used to show multiple codes. **User Name**: Limit results to events attributed to the given user. **Type**: Limit results to events for the given component type. **Node**: Limit results to events for the given node. **Enclosure**: Limit results to events for the given enclosure. **Hardware Serial**: Limit results to events that occurred on the given piece of hardware. The serial can identify a node, enclosure, etc. **StorageSet**: Limit results to events for the given StorageSet. **Volume Service**: Limit results to events attributed to the given volume service. **Volume**: Limit results to events for the given volume. **Message**: Limit results to events whose message matches the given glob pattern.

Click **Search** to display the customized list of events.

Click the icon in the **Event Category** column to display detailed error messages. Click the linked names of components, enclosures, or IP addresses to display the **Enclosure Management** or **Node Management** page for the affected hardware. For more information, see Manage Storage Hardware.

To export all **Historical Event Log** data, click the **Export** button.

13.2.1 Monitor Login Activity with the Historical Event Log

By checking to view the **Administrator Action Event Category,** you can view both successful and unsuccessful login attempts. This information is useful for administrators responsible for monitoring storage enclosure access and can help determine if someone is trying to compromise the storage system through an unauthorized entry.

As a login security feature, the **syslog_watcher** daemon will report a message to PanActive Manager when there are ten failed login attempts over a five minute timespan or 30 failed login attempts over a 60 minute period. After the first report, **syslog_watcher** will report additional problems as often as once every 60 minutes. Each message will include the number of failures, the types (**ssh**/**http**/**https**), and the number of minutes between the first and last message. For each such report, the system will trigger a current **Administrator Action** item for the node on which it was reported. That item will expire in 60 minutes or can be cleared manually. If the node already has an **Administrator Action** item for failed logins, the new info message will replace the old one. All reports from **syslog_watcher** are logged as info items in the **Historical Event Log** linked to the reporting node.

The **Warning** icon will appear when the system detects too many recent login failures. Clicking the warning icon displays a message alerting the user to the excessive login failures.

If the user has enabled the **Warn About Failed Logins** feature in **Configuration > Security > Administrator**, then the system will also trigger a realm-wide warning alert that will send email and pager messages as well as create an **Historical Event Log** warning entry. The realm-wide warning allows the administrator to clear the event once rather than for each node.

13.3 Perform Network Diagnostic Tests

PanActive Manager includes an interface for network diagnostic tests you can use to gather and process information about network failures.

Click **Network > Network Status > Network Diagnostics**.



Figure 13-2. Network Diagnostics page.

Click the **Start New Test** button to test that the following components are configured and responsive: SMTP mailserver, default router, DNS server, NIS server, LDAP server, NTP server, secure web proxy server, and remote syslog collector. Once the test is completed, the page will show the name/ID, status, and test results for each of the above components (if configured), including response time.

Click the **Connections** controls to perform specific tests:

Control	Description
Netperf	Turns the **netperf** daemon on/off on all nodes in the system. **Netperf** is a benchmark that can be used to measure various aspects of networking performance. Its focus is on bulk data transfer and request/response performance using either TCP or UDP and the Berkeley Sockets interface. When enabled, the **netperf** server will run on port 4005.
Ping	Pings from the node to a specified IP address with two retry attempts.
Traceroute	Shows IP addresses for up to five hops from the node to a specified IP address; waits two seconds for a response.

Use the **Clear All Test Results** button to clear all network diagnostic test results.

13.4 Perform System Diagnostics

PanActive Manager includes an interface for system diagnostic tests used to gather and process information about your system. This information can be used by VDURA Technical Support to help diagnose problems with your system. Typically, you will not run these diagnostic tests unless directed to by VDURA Technical Support.

13.4.1 Select a Diagnostic Test

Click **Hardware > Hardware Diagnostics > Select Diagnostic**. This page contains a list of tests that can be performed against the hardware, approximate run time for each test, and detailed test descriptions.

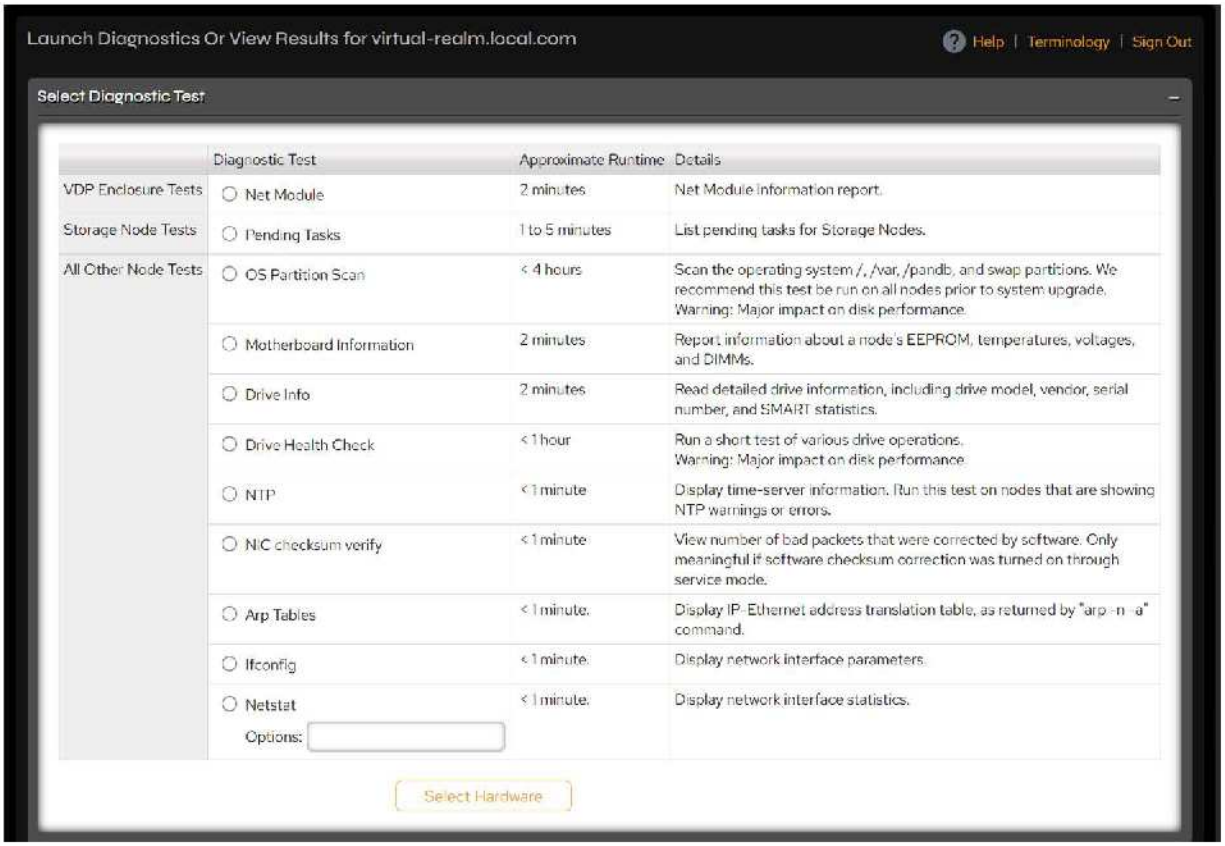


Figure 13-3. Launch Diagnostics Or View Results page.

Select a diagnostic to run, then click the **Select Hardware** button to choose hardware.

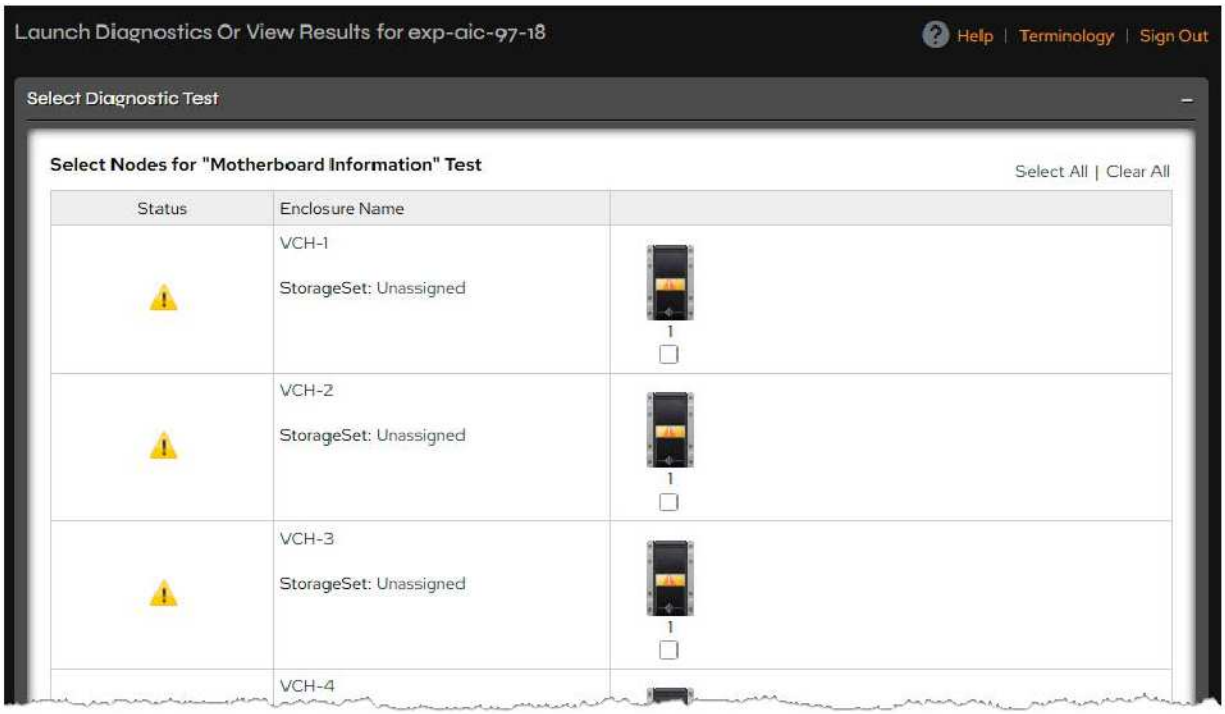


Figure 13-4. Hardware selection fields.

Use the check boxes next to the nodes to select which ones to test, then click the **Launch New Diagnostic** button. When complete, the system will display the test results.

You can view any existing tests on selected hardware by clicking the **View Results for Diagnostics** button on the **Select Hardware page.** If there are no test results for the hardware or test type you have chosen, you can use the **Run Test Now** link from the **Diagnostic Results** screen.

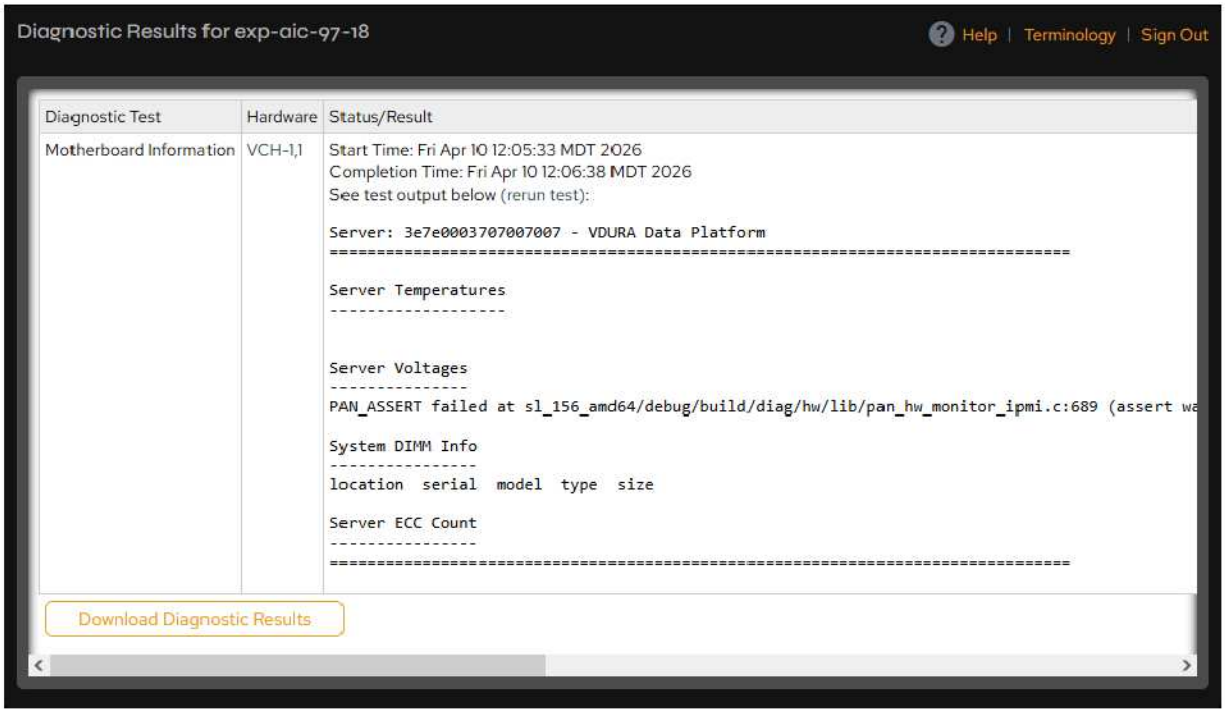


Figure 13-5. Diagnostic Results page.

13.4.2 Download a Diagnostic Test

Once the test has completed and the results are displayed, you can download the results by clicking the **Download Diagnostic Results** button at the bottom of the **Diagnostic Results** page. A compressed file containing your report will be downloaded to the location you provide.

13.5 Troubleshooting

PanActive Manager detects system and hardware errors, summarizes them in reports, collects them in the **Historical Event Log**, and distributes them to system administrators in email alerts and status summaries. These messages provide information about self-healing corrective actions and actions the system administrator can take to prevent or correct the problems.

In some cases, such as a shutdown caused by power failure, PanActive Manager messages are unavailable. Correct any obvious hardware issues, reboot the system, and review the event listings to be informed of self-healing processes that have been implemented or for guidance on system administrator actions.

If problems persist, contact VDURA Technical Support.

PanActive Manager includes networking troubleshooting tools that can be used to gather log and process information about network failures. Use the **Collect State Information** and **Enable Diagnostic Tracing** tools to troubleshoot problems with the help of VDURA Technical Support.

You can also use the **panfs_trace** CLI command to collect DirectFlow Client trace data.

> **NOTE:** If you prefer to troubleshoot using a command-line interface, see the *VDURA PanCLI Reference Guide* for details on additional VDURA command-line interface commands.

13.5.1 Collect State Information

State information can be used to help understand what happened just before a failure or problem occurred.

Click **Advanced Tools > Trace Collection**.

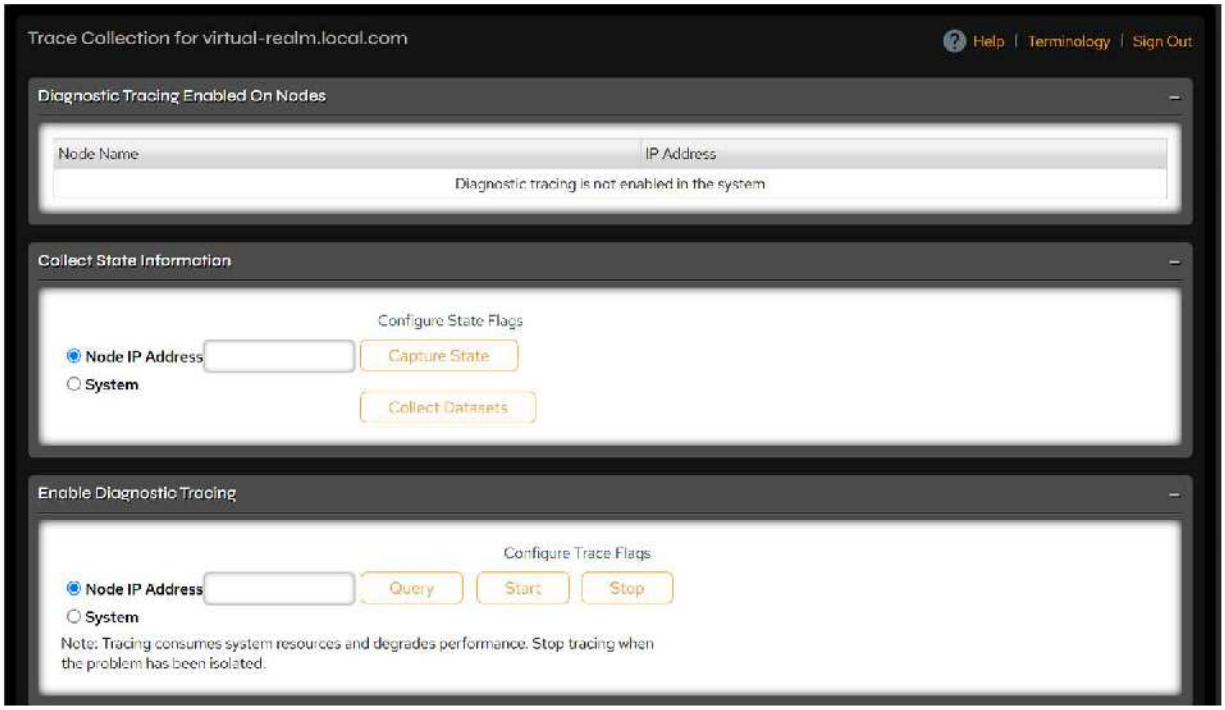


Figure 13-6. Trace Collection page.

Fill in the **Node IP Address** field for the node experiencing problems. For VDURA system problems, select **System**. If instructed by VDURA Technical Support staff, click **Configure State Flags**.

Click **Capture State** to save a compressed state information file to your local system. Click **Collect Datasets** to write datasets for all Directors in the system or given host to a local disk. Send the file(s) to VDURA Technical Support to analyze failures.

13.5.2 Enable Diagnostic Tracing

Use **Diagnostic Tracing** to help recreate the cause of system problems, such as the following:

- Unresponsive or slow system
- Incorrect data (wrong file size, premature end of file, wrong attributes, wrong time stamp, etc.)
- Functionality issues such as an attempt to remove a file without proper permissions or a file that cannot be removed even with permission
- Components crashing with similar debug information

Click **Advanced Tools > Trace Collection**.

Fill in the **Node IP Address** field for the node experiencing problems. For VDURA system problems, select **System**. If instructed by VDURA Technical Support, click **Configure Trace Flags** and follow further instructions.

Click **Start** to begin diagnostic tracing and recreate the problem you want to trace.

Click **Stop** to end diagnostic tracing and save the tracing file to your local host. Send this file to VDURA Technical Support to analyze failures.

If you are not sure if tracing is currently enabled, click **Query**. PanActive Manager will indicate whether or not tracing is currently enabled.

13.5.3 Collect DirectFlow Client Traces

The DirectFlow Client writes trace data for some operations into a circular trace buffer in memory. As more trace data is added to the trace buffer, the oldest data is pushed out of the buffer. VDURA Technical Support may request a **panfs_trace** command to collect this trace data from the trace buffer for troubleshooting. Follow the instructions in Knowledge Base article KB-300915, "How do I collect long-term DirectFlow Client traces?" to use the CLI **panfs_trace** command.

14. PanMove Sync

PanMove Sync is a Linux-based client-side parallel data replication tool that allows customers to replicate and migrate data across a LAN. This chapter explains how to configure and use PanMove Sync. It contains the following sections:

- PanMove Sync Basics
- Configuration and Use
- Logging and Error Handling
- Limitations

14.1 PanMove Sync Basics

PanMove Sync is a Linux-based client-side parallel data replication tool that allows customers to replicate and migrate data across a LAN. It is included as part of every **panfs-apps.rpm** package. Use cases for PanMove Sync include:

- One-time movement of data from one VDURA volume to another VDURA volume
- One-time movement of data from third-party storage to a VDURA volume
- One-time movement of data from a VDURA volume to third-party storage
- Replication of data between two or more file systems for disaster recovery or latency reduction purposes

As shown in the following illustration, the existing **rsync** open source utility is leveraged as part of PanMove Sync to perform the actual synchronization of files and directories. Because Linux systems administrators are most likely familiar with **rsync** and its options, PanMove Sync follows **rsync's** syntax and usage where possible.

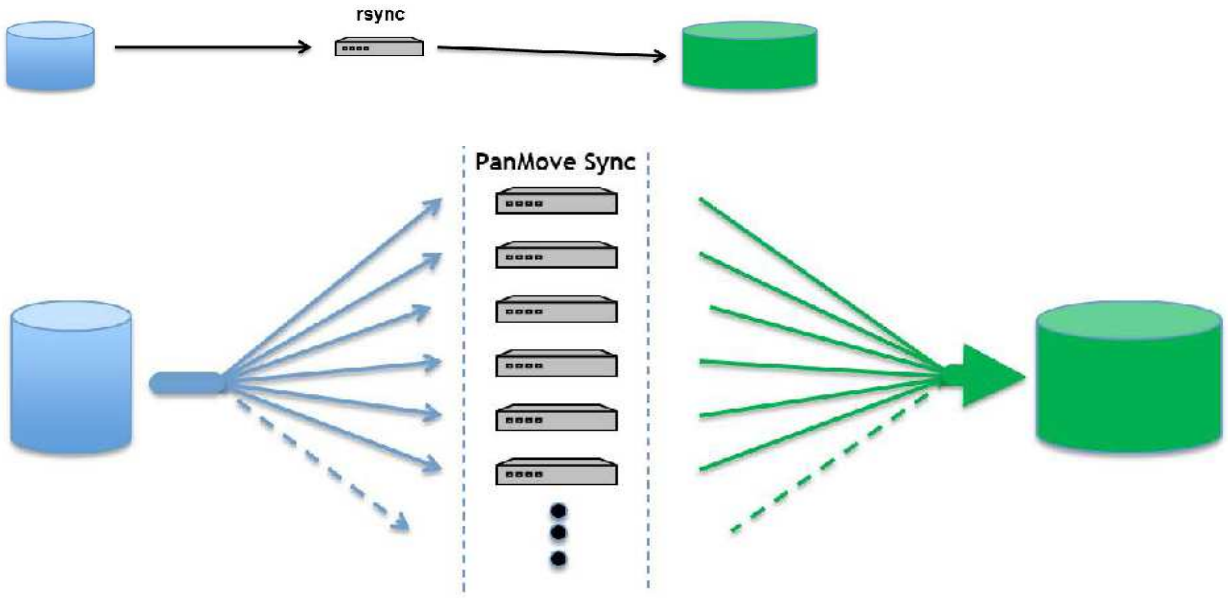


Figure 14-1. The rsync utility as part of PanMove Sync.

PanMove Sync has several advantages over **rsync**, particularly when working with large datasets. PanMove Sync offers:

- **Scalable performance**: PanMove Sync scales with multiple clients (tested with up to 64 worker processes).
- **Can be interrupted and resumed**: PanMove Sync periodically saves the replication and can resume from the last saved point.
- **Fast start**: PanMove Sync begins copying data virtually immediately. For very large datasets, **rsync** can take hours analyzing the directory structure before moving any data.
- **Automatic load-balancing**: PanMove Sync automatically redistributes workload evenly across the available workers. Thus, faster workers are automatically given more work.

14.1.1 Operation

PanMove Sync copies an existing directory tree to a new location, using a parallel, multi-client process that preserves permissions, attributes, ownership, sparseness, and hardlink relationships of the original files.

You initiate operation by starting the PanMove Sync manager, which controls the overall process; see Start PanMove Sync Manager for more information. The PanMove Sync manager can perform the replication operation independently or can call up to 64 worker processes to assist using the **–client-list=<file>** option.

> **NOTE:** More than one worker process can be run on a single worker machine. The number is controlled through the **–workers-per-client** option. The default number is 4, however you may try changing from the default to maximize performance.

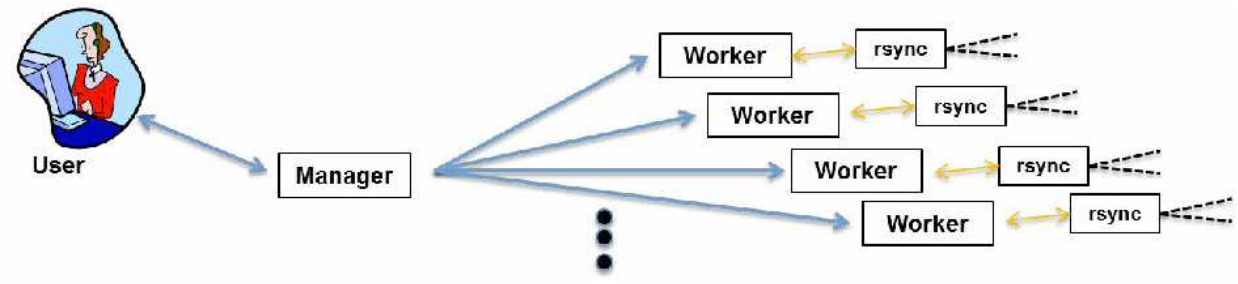


Figure 14-2. Worker processes.

During operation, PanMove Sync determines the directories and/or files to be synchronized and queues instructions for the workers to handle the copying. Unless given an explicit list of items to replicate (using **–files-from, –exclude-from,** or **–include-from** options), PanMove Sync begins with the top-level source directory and traverses the tree in a depth-first manner. For each directory visited, the manager queues an instruction for one of the workers to replicate the directory and its contained files. In this way, the manager visits the entire source directory tree structure, while the tree content is replicated by the workers. The workers operate in parallel (independent of each other) to create and copy directories, files, and links.

PanMove Sync may be interrupted and resumed without starting over by using the **–resume** option. During operation and when terminating in a controlled fashion, the manager periodically saves the current state of the operation to a temporary directory, so that if the operation is terminated before completion, it can be later resumed from the last saved state. By default, the temporary directory is located in the replication destination directory. Use the **–temp-dir** option to direct PanMove Sync to place the temporary directory in a different location.

In addition, if a replication operation runs to completion but fails to replicate some files due to errors, the operation can be retried using the **–retry** option to attempt to replicate the failed files without performing the entire replication operation over.

Once the replication operation is complete, the manager will output to the console a summary of the results. The results can also be copied to a designated log file using the –**log-file=<file>** option. These results include the number of files replicated, total number of bytes transferred, and the number of files and/or directories that failed to replicate. If any files or directories failed to replicate, the manager sends out a notice instructing the user to examine the log for the names of failed items, and also reminding the user that the **–retry** option can be used to retry the failed items.

> **NOTE:** If the **–temp-dir-del** option is used and set to **always**, the temporary directory is deleted and the **–retry** option is not possible.

14.1.2 Compatibility

PanMove Sync is supported on all versions of the Linux operating system supported by DirectFlow. A complete list of supported client distributions and kernels can be found at **support.VDURA.com** on the **Download > DirectFlow Client for Linux** page.

The utility also supports third-party file storage, providing that the file systems are accessible to the clients running PanMove Sync.

All PanMove Sync workers must run the same PanMove Sync version as the manager. If the manager sees that the worker does not have an acceptable version, it tells the worker to terminate.

All PanMove Sync workers must run **rsync** version 2.6.8 or later to ensure proper operation.

Best performance is obtained when using PanFS mounts for both source and destination directories.

> **NOTE:** You can check client compatibility when you first set up PanMove Sync. See Verify Client Setup for more information.

14.1.3 Operation Requirements

The operation of PanMove Sync assumes the following points are true. If any one of these requirements is not met, PanMove Sync may fail to replicate all directories and files as expected.

- The source directory must be mounted on all clients at the same mountpoint.
- The destination directory must be mounted on all clients at the same mountpoint.
- All clients must have an appropriate version of the rsync utility available.
- The **panfs-apps.rpm** is installed on all clients.
- The user has the proper permissions on the source and destination directories to replicate the files.
- The source directory is not being modified during operation.

- No other clients will be writing to the destination directory at the same time that PanMove Sync is running.
- The destination directory is not contained within the source directory. Otherwise, an infinite recursion may result which will consume all space on the source directory volume.
- PanMove Sync is able to reach workers using the IP address or hostname and user ID provided via SSH. This is necessary to auto-start workers. We recommend key-based SSH access between PanMove Sync and workers.
- There is sufficient space on the destination volume or another volume shared between manager and worker clients, to hold temporary transaction data.

14.2 Configuration and Use

The PanMove Sync manager controls the copying process by traversing the directory tree, distributing instructions to worker clients, providing control, and reporting status during the processing.

This section contains the following sections:

- Before Using PanMove Sync Manager
- Start PanMove Sync Manager
- Verify Client Setup
- Create a Client-List File
- Specify the Number of Processes to Start on Each Client
- Replicate or Exclude Specific Files
- Resume an Interrupted Replication
- Retry Replication of Failed Files
- Terminate a Replication Process
- Manually Start a Worker
- Specify a Log File Location
- Example Scenarios

14.2.1 Before Using PanMove Sync Manager

Verify the following two points before starting the PanMove Sync manager:

- Before using the PanMove Sync manager, make sure that **panfs-apps** is installed on the master node and each worker client. The **panfs-apps** RPM is available at **support.VDURA.com** on the **Downloads tab** > **DirectFlow Client** page. While on this page, search for **panfs-apps** and select the correct version for your system.

- The master node must have root SSH keys configured to itself and to each worker node. The SSH keys must not require pass phrases, so that they can be used without a password or pass phrase prompt. Test SSH access from the master node to itself:

```
# ssh localhost
```

Test SSH keys to any worker nodes:

```
# ssh root@{node name or IP}
```

The SSH command should succeed without a password prompt.

14.2.2 Start PanMove Sync Manager

To start the PanMove Sync manager, enter the following command:

```
pan_sitesync [options] <src_dir> <dst_dir>
```

- **<src_dir>** is source directory to replicate data from. This can be any arbitrary directory.
- **<dst_dir>** is destination directory to replicate data to. This can be any arbitrary directory not residing within the source directory tree.

The table below explains command options.

Option	Description
`-h, --help`	Display a brief help page describing PanMove Sync options.
`--addr=<IP>`	Use the specified IP address for worker communications. This option allows selection of the host address and the attached subnet for exchanging messages with workers. Generally, this is only needed with multi-homed hosts when there is no routing between attached subnets. By default, PanMove Sync will use the first routable interface address.

Option	Description
`--client-list=<FILE>`	Specify a file with the list of clients to start as workers. The total number of workers will be the number of clients specified times the value of the **workers-per-client** option. Changes to this file that occur once processing has started are ignored until the next restart or resume. See Create a Client-List File for more details about the content of this file. If not specified, workers are created on the manager's host machine as the current user, using the localhost as the client. The worker executable and a proper version of **rsync** must be found in the default PATH for this to work properly.
`--delete`	Delete files and directories that exist only on destination, for any directory being replicated.
`-n, --dry-run`	Operate as usual, except do not copy any directories or files. Instead, PanMove Sync will determine which files to copy and obtain statistics about the replication. A dry-run operation may be interrupted and resumed just like a regular operation.
`--eta`	Perform a dry-run in order to gather statistics before starting replication. During replication, PanMove Sync will compute estimates for completion (ETA) which are included in progress messages. This command implies **–progress.** Use of this option will extend the time it takes to perform a replication, since it adds a dry-run before starting updates.

Option	Description
`--exclude-from=<FILE>`	Specify a file that contains **rsync** exclude patterns (one per line). Blank lines in the file and lines starting with ";" or "#" are ignored. **FILE** cannot be "-" (i.e., the list cannot be read from standard input as with regular **rsync**).
`--files-from=<FILE>`	Specify optional file containing an exact list of files and directories to be replicated. Unlike **–include-from,** this file cannot contain wildcards. Each directory or file listed must be specified with a path relative to the source directory. All directory names must end with a slash delimiter. When using this option, directory recursion is disabled.
`-I, --ignore-times`	Do not skip files that match size and time. This forces all files to be updated regardless of the last-changed time or file size.
`--include-from=<FILE>`	Specify a file that contains **rsync** include patterns (one per line). Blank lines in the file and lines starting with ";" or "#" are ignored. **FILE** cannot be "-" (i.e., the list cannot be read from standard input as with regular **rsync**).
`--log-date`	Add the current date to log timestamps (dd-MM-yy, HH:MM:SS). Default is HH:MM:SS.
`--log-file=<FILE>`	Specify the path and name of the PanMove Sync log file to create. Default is "-" to write log output to **stdout.** The log file size is usually in the 5-10 MB range, but if verbose logging is requested, the size can become much larger.

Option	Description
`--log-msec`	Extend log timestamps to millisecond accuracy (HH:MM:SS.mmm). Default is HH:MM:SS.
`--modify-window=<NUM>`	Compare last-modification times with reduced accuracy; allows a difference of **NUM** seconds to be considered equal.
`--no-itemize`	Skip generating **ITEM** log messages. Normally when any verbose options are specified, PanMove Sync will produce a log message for every file and directory examined or updated. Use of this option suppresses these itemized messages while still allowing other verbose log messages to appear. This can greatly reduce log size and processing overhead if such individual messages are not needed.
`--no-xattrs`	Omit the **-X** parameter to **rsync** (i.e., do not copy or update xattrs on the destination). PanFS ACLs are stored as extended attributes, so they will not be preserved when running PanMove Sync with the **–no-xattrs** option.
`--port-m=<NUM>`	Set the port number to use for the manager **XML-RPC** service. Default is 9980.
`--port-w=<NUM>`	Set the base port number to use for the worker **XML-RPC** service. Default is dynamic.
`--progress`	Periodically show a progress message to **stdout.** The message will look similar to this: `[10:15:33] NOTICE: Processed 0 files (0B) in 2 directories in 00h:00m:04s`

Option	Description
`-r, --resume`	Resume a previous incomplete operation. If **–resume** is specified, a valid temporary directory must exist and contain the state from a previous operation. Similarly, if **–resume** is not given, the temporary directory must not already exist. If a migration has been begun and interrupted, you can only resume the migration using the version with which you began. For example, if you began a migration using PanFS 5.5.x, you cannot resume that migration using PanFS 6.0.x or later.
`--retry`	Only used to restart a replication which completed with one or more file/directory errors. Only the failed items will be retried, and no other options specifying which files to replicate (e.g., **include**/**exclude**, **files-from**, etc.) are allowed. If a migration ends with failures, you can only retry the migration using the version with which you began. For example, if you began a migration using PanFS 5.5.x you cannot retry that migration using PanFS 6.0.x or later.
`--save-period=<MINUTES>`	Specify how long to wait in minutes between periodic updates of the saved state in the temporary directory. Increasing this value may cause more reprocessing when resuming PanMove Sync. Default value is 5.
`--size-only`	Skip files that match in size regardless of whether the source and destination have matching last-modified times.
`--stats-only`	Only used to display statistics of an interrupted replication without resuming the operation.

Option	Description
`--sync-size=<SIZE>`	Set the threshold for processing a single directory by multiple clients. Larger values allow for greater use of single-client processing, which is more efficient but may reduce the overall performance due to decreased use of parallel processing. Default value is 5 GB.
`--sync-tax=<SIZE>`	Control file batch sizes for replication. Larger values will decrease the number of files handled at one time, which is less efficient but helps distribute workloads more evenly across workers. Default value is 1 MB. The value of sync-size divided by the value of sync-tax sets the maximum number of files to be handled at one time by a single worker. If this ratio is set too high, excessive delays and/or errors during processing may occur.
`--temp-dir`	Set the temporary directory used for status and communications. This directory allows the replication to be interrupted and resumed later without starting over or losing all progress or performance statistics. Default is **<dest_dir>/.pps_temp**.

Option	Description
`--temp-dir-del=<WHEN>`	Specify whether to delete the temporary directory when PanMove Sync exits. Choices are: • **always**: always delete the temporary directory • **never**: never delete the temporary directory • **success**: delete upon successful completion only Default value is **success.** When using **–temp-dir-del=always** and PanMove Sync is interrupted, the temp directory will only be removed if PanMove Sync is interrupted with **Ctrl-C** from the command line or **kill -SIGINT.** It will be left intact if PanMove Sync is killed with **SIGKILL** (**kill -9**).
`--test-workers`	Do not perform replication. Instead, test the ability of the manager to start the specified workers, and test each worker's access to the file system directories. All other options and arguments are ignored except **–client-list**.
`--timeout=<SECONDS>`	Set the **rsync** I/O timeout. In order to prevent worker processes from running indefinitely when waiting for blocked I/O, the timeout default is not zero (infinite) as it usually is for **rsync**. Therefore, on congested file systems or networks where I/O timeouts may be more likely, this option can be used to increase the timeout duration and reduce or eliminate **rsync** timeout errors. Default value is 300.
`-u, --update`	Skip files that are newer on the destination.

Option	Description
`--use-rsync=<FILE>`	Tell PanMove Sync to use the specified executable as **rsync.** Default is **rsync** in standard search path.
`-v, --verbose`	Increase the amount of information output to the log. By default, PanMove Sync outputs only warnings and errors. A single **-v** will add information about what files are being transferred and a brief summary at the end. Two **-v** options (**–vv)** will add additional information about PanMove Sync operation. Three or more **-v** options are only needed for debugging PanMove Sync. **CAUTION!** Even a single **-v** will cause PanMove Sync to generate **ITEM** log messages for every file handled. This can have adverse effects, such as greatly increasing log size and network traffic. In extreme cases, the excess network traffic can cause the temporary loss of one or more workers. Use the **–no-itemize** option to suppress such messages.
`--version`	Display current version of PanMove Sync. All other arguments are ignored.
`-w, --workers-per-client=<NUM>`	Specify the number of worker processes to start on each worker client. The PanMove Sync utility has been tested running up to 64 processes. Default value is 4.
The following options are passed as-is to **rsync.** See the **rsync** man page for more information.	

Option	Description
`--delete`	
`--exclude-from`	
`--ignore-times`	
`--include-from`	
`--modify-window`	
`--size-only`	
`--update`	

14.2.3 Verify Client Setup

When setting up PanMove Sync for the first time, use the **–test-workers** option to verify your client setup. This option checks:

- NAS access at all clients
- SSH access from PanMove Sync to workers
- RPC connectivity to all workers
- PanMove Sync and **rsync** availability and versions

14.2.4 Create a Client-List File

At startup, PanMove Sync may be given a file containing a list of available worker clients by using the **–client-list=<file>** option. Any valid host can be used as a worker and listed in this file, including the host on which PanMove Sync is run.

With no client-list, PanMove Sync uses **localhost** with the current user ID, SSH keys for authentication, and expects executables to be in the user's PATH. An empty **client-list** can also be used if it is desirable to start workers manually.

Each line in the **client-list** file takes the following form:

```
IP_address_or_hostname userID <pw> <pan_sitesync_worker>
<rsync>
```

- **userID**: the ID for SSH to use to access the client machine. A single dot can be used for the **userID** to represent the current user at the time PanMove Sync is started (see examples below). The **userID** (or current user if a dot is specified) must have at least read access to the source directory and read+write+execute access to the destination directory.
- **<pw>**: optional password in open text. If using key-based SSH, leave this field empty.
- **<pan_sitesync_worker>**: optional executable for each worker client.
- **<rsync>**: optional **rsync** executable to use.

Each optional item is enclosed in angle brackets; these angle brackets must always be present. Paths to the executables may be included, as well, as either absolute or relative to the **userID** home directory. For any executable not specified, the default executable names are used, and it is assumed they will be found in the default search path for the SSH **userID**.

A few examples of **client-list** files:

- ssh to 169.132.32.10 using the current **userID** and key-based access and run **pan_sitesync_worker** and **rsync** from default PATH:

```
169.132.32.10 . <> <> <>
```

- ssh to the host as the current user, with key-based access and run **pan_sitesync_worker** from root PATH and use specific **rsync**:

```
169.132.32.10 . <> <> </test/bin/rsync>
```

- ssh to the host as **rnwp** with a password and run **/global/utilities/pan_sitesync_worker**:

```
169.132.32.4 rnwp <password> </global/utilities/
pan_sitesync_worker> <>
```

- ssh to the host as ssikah with a password and run **~/tbin/pan_sitesync_worker**:

```
169.132.32.4 ssikah <password> <tbin/pan_sitesync_worker> <>
```

- ssh to the host as root with key-based access and run **pan_sitesync_worker** from root PATH and use specific **rsync**:

```
169.132.32.10 root <> <> </test/bin/rsync>
```

> **NOTE:** The **client-list** provides PanMove Sync with available workers. However, the distribution of replication work to multiple workers depends on the data layout and parameters such as **sync_size** and **sync_tax.** In some cases (e.g., smaller datasets), PanMove Sync may not need to use all available workers.

14.2.5 Specify the Number of Processes to Start on Each Client

PanMove Sync can support up to 64 worker processes running across the hosts specified in your client list. You can specify the number of processes to run on a worker by using the **–workers-per-client= <NUM>** option. For example, the following command tells PanMove Sync to start two worker processes for each entry found in the client list file **client_list.txt**:

```
pan_sitesync --workers-per-client=2 --client-list=client_list.txt
src_dir dst_dir
```

In this example, if the client list file has five clients specified, a total of 10 worker processes would be started.

14.2.6 Replicate or Exclude Specific Files

If an operation is started with the **–files-from <file>** option, the manager does not traverse the source directory tree. Instead, it simply iterates the specified replication list, generating work orders for each item (file or directory) it finds.

The example below replicates the directory tree from **/vol1/src** to **/vol3/archive3** using the workers that are defined in the file **/usr/foo/sync_workers**:

```
pan_sitesync --client-list=/usr/foo/sync_workers /vol1/src /
vol3/archive3
```

If either the **–include-from <file>** or **–exclude-from <file>** option is used, PanMove Sync will traverse the source directory as usual, but will include or exclude the files stipulated.

If PanMove Sync is initially started with either the **–include-from <file>, –exclude-from <file>,** or **–files-from <file>** options, the corresponding files are copied to the temporary directory. If the operation is interrupted and restarted, the cached versions of the files are used. To change the contents of the **include-from, exclude-from,** or **files-from** options, you will need to rerun the operation with the updated files.

To create a file for the **files-from, include-from,** or **exclude-from** option, list the path of each file to be included or excluded on a separate line.

> **NOTE:** Only files can be excluded, not subdirectories.

14.2.7 Resume an Interrupted Replication

If a replication is interrupted before completion, it can be resumed at the point where it was interrupted as long as PanMove Sync was not started with the **–temp-dir-del=always** option.

The example below resumes an interrupted replication from **/vol1/src** to **/vol3/archive3**:

```
pan_sitesync --resume /vol1/src /vol3/archive3
```

If a migration has been begun and interrupted, you can only resume the migration using the version with which you began. For example, if you began a migration using PanFS 5.5.x, you cannot resume that migration using PanFS 6.0.x or later.

14.2.8 Retry Replication of Failed Files

If a replication completes successfully, but with failed files or directories, PanMove Sync sends a notice instructing the user to examine the log identifying the failed files. Those files or directories can be retried using the **–retry** option.

The example below retries replication of the files that failed during the **/vol1/src** to **/vol3/archive3** replication:

```
pan_sitesync --retry /vol1/src /vol3/archive3
```

Retry works from an internally saved list of failed files. No other options are needed with the **–retry** option.

If a migration ends with failures, you can only retry the migration using the version with which you began. For example, if you began a migration using PanFS 5.5.x, you cannot retry that migration using PanFS 6.0.x or later.

14.2.9 Terminate a Replication Process

If the replication process must be interrupted, use the **Ctrl-C** key combination or the system command **kill -SIGINT <procID>** in order to ensure an orderly termination and save the latest progress. If the replication is interrupted with the **SIGKILL** command, no termination handling will occur.

14.2.10 Manually Start a Worker

The PanMove Sync worker is the worker client tool and is usually started automatically by the PanMove Sync manager. You can also start a worker directly as follows:

```
pan_sitesync_worker [options] <managerURL>
```

The **<managerURL>** parameter is the hostname or IP address and port where PanMove Sync RPC services can be found. The manager must be initialized before starting a worker. For example:

```
pan_sitesync_worker http://10.75.1.134:9980
```

Available options include:

Option	Description
`--index=<n>`	The index of this worker process as defined by the PanMove Sync manager. This option is not normally used when starting a worker manually.
`--log-level=<level>`	The level of verbosity in the worker log. This option is not normally used when starting a worker manually.

Option	Description
`--port=<n>`	Port number for worker's RPC service. Default value is **dynamic.**
`--report-rsync-version=<x.y.z>`	Force worker to report its **rsync** version as **x.y.z**. This option is not normally used when starting a worker manually.
`--signature=<n>:`	The worker's unique signature provided by PanMove Sync. This option is not normally used when starting a worker manually.
`--spawn`	Start worker as background process. This option is not normally used when starting a worker manually.
`--use-rsync=<FILE>`	The **rsync** executable for worker to use. The default is **rsync** on the standard search path.
`--version`	Report the version of PanMove Sync worker and exit.

Manually started workers can be added to an in-progress replication or removed from participation at any time.

14.2.11 Specify a Log File Location

By default, all messages are sent to **stdout.** However, you can specify a log file to hold messages by using the **–log-file=<file>** option. The example below replicates the directory tree from **/vol1/src** to **/vol3/archive3** and sends all log output including periodic progress messages to **/usr/foo/psync.log**:

```
pan_sitesync --progress --log-file=/usr/foo/psync.log /vol1/
src /vol3/archive3
```

For information on the contents of the log file, see Logging and Error Handling.

14.2.12 Example Scenarios

The following command replicates the directory tree from **/vol1/src** to **/vol3/archive3** using **localhost** workers. Note that if files are found in **/vol3/archive3** that do not exist in **/vol1/src,** they are left untouched.

```
pan_sitesync /vol1/src /vol3/archive3
```

The following command replicates the directory tree from **/vol1/src** to **/vol3/archive3** and displays the name of each file replicated. In this case, if any files or directories exist in **/vol3/archive3** that are not found in **/vol1/src,** they are deleted.

```
pan_sitesync -v --delete /vol1/src /vol3/archive3
```

The following command replicates the directory tree from **/vol1/src** to **/vol3/archive3** using the workers that are defined in the file **/usr/foo/sync_workers**.

```
pan_sitesync --client-list=/usr/foo/sync_workers /vol1/src /
vol3/archive3
```

The following command replicates the directory tree from **/vol1/src** to **/vol3/archive3,** including only those files stipulated in **/usr/foo/myincludes**.

```
pan_sitesync --include-from=/usr/foo/myincludes /vol1/src /
vol3/archive3
```

14.3 Logging and Error Handling

PanMove Sync creates a log file for every operation. Log entries are labeled with a timestamp, originator, and one of several severity levels, in this format:

```
<timestamp> <severity> [workerID] : <message>
```

Severity levels are as follows:

- **NOTICE**: Important user information (always shows)
- **STAT**: Operational statistic (always shows)

- **FATAL**: Unexpected error, terminating (always shows)
- **ERROR**: Unexpected error, operation may fail (always shows)
- **ITEM**: File or directory being replicated (**-v** and above, unless **–no-itemize** specified)
- **WARN**: Unexpected event, operation proceeding (**–vv** and above)
- **INFO**: Generally used for debugging PanMove Sync operation (**–vvv** and above)
- **ALL**: Low-level information used for debugging only (**–vvvv**)

Below are some example log file entries:

```
[11:35:52] INFO: Completing: WorkQueue now has 24 WorkOrders
[11:35:53] ERROR [w(0):compute-5-6]: FILE: SubDir1/file0.dat:
ERROR: Update missed (rsync error)
[11:36:07] STAT: Elapsed operational time: 00h:00m:17s
[11:36:07] STAT: Number of directories visited: 30
[11:36:07] STAT: Files updated: 129
```

14.4 Limitations

The following table lists known limitations with this release of PanMove Sync:

Issue	Description
symlinks	If a symlink was constructed using a full path to another location in the directory being replicated, such links may end up broken.
Non-default layouts	Non-default PanFS layouts are not preserved. In particular, the PanFS attributes **agg_type_files** and **agg_layout_info_files** are not preserved across the replication. If you wish to have child files that use non-default layouts, you will need to exclude any directories containing such files from the process and perform the copy of those directories and their contents manually.
CW mode setting	Copying CW mode setting is not handled. This must be done manually after the data replication process completes.

Issue	Description
Timestamp synchronization	On some versions of Linux, it is not possible for **rsync** to synchronize the timestamp of a replicated symlink. In such cases, the replicated symlink has a timestamp equal to its creation time instead of the original symlink's timestamp.
Excluding replication of subdirectories	Unlike **rsync,** PanMove Sync does not support excluding replication of subdirectories by naming them in an exclude file. Only files can be specified for exclusion in an exclude file.
Extremely long pathnames	PanMove Sync is unable to replicate directories which have a path longer than the OS maximum pathlength (usually 256 characters on Linux). Unlike regular **rsync,** this limit applies even to files within deeply nested subdirectories, whose full pathname exceeds this limit.

15. System Limits and Requirements

This chapter describes general limits and requirements of VDURA systems.

Limit	Description
Smallest V5000 system	3 Directors and 3 Storage Nodes
Largest system	• No enforced limit • VDURA has tested up to 1300 OSDs
Maximum number of files or directories	No enforced limit
Maximum file size	• 32-bit Linux DirectFlow: 16 TB • 64-bit Linux DirectFlow: 8192 PB
Maximum number of volumes	• Per Director: no limit • Per StorageSet: no limit • Per realm: Default of 1200, up to 4096
Maximum number of Realm Managers (RMs) in repset	5
Maximum StorageSet size	No enforced limit
Maximum number of entries per directory	• 1,000,000 (contact VDURA if more are needed) • 250,000 (recommended)
Maximum number of clients	5000
Maximum path length	4096 bytes
Maximum file name length	• 512 bytes • May be limited further by the client operating system
Minimum physical client memory	1 GB
Maximum number of mounted realms	1024 total global and local mounts

Limit	Description
Maximum number of IP aliases	• 100 per NIC per Director, but no more than 512 per realm • This is configured when the realm is created
Maximum number of hard links	255

16. VDURA System Error Messages

This chapter provides a numerical list of VDURA system error messages.

Error messages are categorized as follows:

Icon	Status	Description
	Info	Error recovery messages or configuration changes are typically logged to the **Events** page. For information on viewing event messages see View Events.
	**Warning**	Developing problems that might impair system availability or data integrity are logged to the events page and emailed to system administrators. For information on viewing event messages see View Events. For information on setting recipients for warning messages, see Set Status Recipients.
	Critical	Problems that are an imminent risk to system availability or data integrity are logged to the events page and emailed to system administrators. For information on viewing event messages see View Events. For information on setting recipients for critical messages, see Set Alert Recipients.

The following table is a list of error messages and PanActive Manager error handling.

Error Code	Category	Error Description	Recommended Action
8405	Warning	Power supply 1 failed: Power supply 1 has failed (PS1 serial %%SERIAL%%, shelf serial %%SHELFID%%). Power supply 1 is on right when looking at rear of shelf.	Contact VDURA Technical Support.
8406	Warning	Power supply 1 unsupported version: Power supply 1 is an unsupported version (PS1 serial %%SERIAL%%, shelf serial %%SHELFID%%). Power supply 1 is on right when looking at rear of shelf.	Contact VDURA Technical Support.
8407	Critical	Incompatible Power Supplies: Power supplies in shelf are incompatible (shelf serial %%SHELFID%%). This configuration is not supported.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8408	Warning	Power supply 1 not found: Power supply 1 is missing (shelf serial %%SHELFID%%). Power supply 1 is on right when looking at rear of shelf.	Check if power supply 1 is inserted properly.
8409	Critical	Power supply 1 over maximum temperature: Power supply 1 temperature is above threshold (PS1 serial %%SERIAL%%, shelf serial %%SHELFID%%). Power supply 1 is on right when looking at rear of shelf.	Ambient temperature is too high. Adjust room temperature.
8410	Warning	Power supply 1 temperature high: Power supply 1 exceeds optimal temperature (PS1 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Ambient temperature is too high. Adjust room temperature.

Error Code	Category	Error Description	Recommended Action
8411	Info	Power supply 1 under temperature: Power supply 1 is under optimal temperature (PS1 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Ambient temperature is too low. Adjust room temperature.
8412	Warning	Power supply 1, AC power fail: AC power disconnected from power supply 1 (PS1 serial %%SERIAL%%, shelf serial %%SHELFID%%). Power supply 1 is on right when looking at rear of shelf.	Check that power supply is plugged in.
8413	Warning	Power supply 1 failed during the monthly battery test (power supply module serial %%SERIAL%%). %%DESCRIPTION%%	Contact VDURA Technical Support.
8414	Warning	A 950W power supply was replaced with an 800W power supply in PS1 (shelf serial %%SHELFID%%). This scenario can cause damage to the 800W power supply.	Insert a 950W power supply into this location.

Error Code	Category	Error Description	Recommended Action
8415	Warning	Power supply 2 failed: Power supply 2 has failed (PS2 serial %%SERIAL%%, shelf serial %%SHELFID%%). Power supply 2 is on left when looking at rear of shelf.	Contact VDURA Technical Support.
8416	Warning	Power supply 2 unsupported version: Power supply 2 is an unsupported version (PS2 serial %%SERIAL%%, shelf serial %%SHELFID%%). Power supply 2 is on left when looking at rear of shelf.	Contact VDURA Technical Support.
8417	Warning	Node disk 9 encountered disk errors.	Contact VDURA Technical Support.
8418	Warning	Power supply 2 not found: Power supply 2 is missing (shelf serial %%SHELFID%%). Power supply 2 is on left when looking at rear of shelf.	Check if power supply 2 is inserted properly.

Error Code	Category	Error Description	Recommended Action
8419	Critical	Power supply 2 over maximum temperature: Power supply 2 temperature is above threshold (PS2 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Ambient temperature is too high. Adjust room temperature.
8420	Warning	Power supply 2 temperature high: Power supply 2 exceeds optimal temperature (PS2 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Ambient temperature is too high. Adjust room temperature.
8421	Info	Power supply 2 temperature high: Power supply 2 exceeds optimal temperature (PS2 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Ambient Temperature is too low. Adjust room temperature.

Error Code	Category	Error Description	Recommended Action
8422	Warning	Power supply 2, AC power fail: AC power disconnected from power supply 2 (PS2 serial %%SERIAL%%, shelf serial %%SHELFID%%). Power supply 2 is on left when looking at rear of shelf.	Check that power supply is plugged in.
8423	Warning	Power supply 2 failed during the monthly battery test (power supply module serial %%SERIAL%%). %%DESCRIPTION%%	Contact VDURA Technical Support.
8424	Warning	A 950W power supply was replaced with an 800W power supply in PS2 (shelf serial %%SHELFID%%). This scenario can cause damage to the 800W power supply.	Insert a 950 W power supply into this location.
8425	Critical	Battery temperature over maximum: Shelf battery temperature is above critical threshold (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8426	Critical	Battery failed: Shelf battery has failed (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.
8427	Warning	Battery unsupported version: Shelf battery is an unsupported version (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.
8428	Warning	Battery failure predicted: Shelf battery is predicted to fail (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.
8429	Critical	Battery not found: Shelf battery is missing (shelf serial %%SHELFID%%).	Check if battery module is inserted properly.
8430	Warning	Battery temperature high: Shelf battery exceeds optimal temperature (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	Ambient temperature is too high. Adjust room temperature.

Error Code	Category	Error Description	Recommended Action
8431	Warning	Battery low: Shelf battery charge is low (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	Wait until fully charged. It can take up to 6 hours to fully charge battery. If the problem persists, contact VDURA Technical Support for further assistance regarding this message.
8432	Info	Battery under temperature: Shelf battery is under optimal temperature (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	Ambient temperature is too low. Adjust room temperature.
8433	Warning	Node disk 10 encountered disk errors, contact VDURA Technical Support.	Contact VDURA Technical Support.
8434	Warning	Enclosure over power budget: The power draw of nodes exceeds the power supply capacity.	Contact VDURA Technical Support for steps to bring the shelf within power budget.

Error Code	Category	Error Description	Recommended Action
8435	Critical	Net Module 1 failed: Net Module 1 has failed (switch 1 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.
8436	Warning	Net Module 1 unsupported version: Net Module 1 is an unsupported version (switch 1 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.
8437	Warning	Node disk 11 encountered disk errors, contact VDURA Technical Support.	Contact VDURA Technical Support.
8438	Critical	Net Module 1 not found: Net Module 1 is missing (shelf serial %%SHELFID%%).	Check if switch 1 is inserted properly.
8439	Warning	Node disk 12 encountered disk errors, contact VDURA Technical Support.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8440	Warning	Net Module 1 checksum error: Checksum error on Device U6 (switch 1 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.
8441	Warning	Net Module 1 checksum error: Checksum error on Device U8 (switch 1 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.
8442	Warning	Node disk 13 encountered disk errors.	Contact VDURA Technical Support.
8443	Critical	Net Module 1 over maximum temperature: Switch 1 temperature is above threshold (switch 1 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Ambient temperature is too high. Adjust room temperature.
8444	Warning	Net Module 1 temperature high: Net Module 1 exceeds optimal temperature (switch 1 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Ambient temperature is too high. Adjust room temperature.

Error Code	Category	Error Description	Recommended Action
8445	Info	Net Module 1 is under optimal temperature (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Ambient temperature is too low. Adjust room temperature.
8446	Warning	Net Module 1 is resetting (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Switch management interface did not respond. Please wait 90 seconds for reset to complete.
8447	Warning	Management interface of Net Module 1 is not responding (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Switch management interface is not responding. Please reset the switch when system is not busy and check that the switch firmware is up-to-date.
8448	Warning	Net Module 1 firmware is out of date (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Switch firmware is out of date; upgrade is recommended. Contact VDURA Technical Support for further assistance.

Error Code	Category	Error Description	Recommended Action
8449	Warning	Net Module 1 firmware is unsupported (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Switch firmware is not supported. Downgrading firmware to a supported version is recommended. Contact VDURA Technical Support for further assistance.
8450	Critical	Node is unavailable: Node is unavailable due to a software failure. Administrative action may be needed to reconstruct this node.	Contact VDURA Technical Support to repair this node. If the node cannot be repaired, go to the Node Detail page and use the Reconstruct control to start reconstruction of data.
8451	Critical	Node exceeds maximum temperature: Node temperature exceeds threshold. Shutdown is imminent.	Adjust room temperature and check fans.
8452	Critical	Node shutdown: Node has shutdown automatically due to critically high temperature.	Adjust room temperature and check fans.

Error Code	Category	Error Description	Recommended Action
8453	Critical	AC power failed: Shelf %%SHELFID%% is running on backup battery power. Shutdown is imminent.	Go to the Hardware page and click the System Shutdown button.
8454	Warning	Node disk nearly full: Node's disk space that is available for user data has reached %%PERCENTFULL%%% of capacity.	Add more storage to the StorageSet.
8455	Critical	Node disk full: Node's disk space that is available for user data has reached 100% of capacity.	Add more storage to the StorageSet.
8456	Warning	Node temperature high: Node exceeds optimal operating temperature.	Ambient temperature is too high. Adjust room temperature.
8457	Warning	Node memory errors: Number of memory errors in node are abnormally high.	Contact VDURA Technical Support for a replacement node.
8458	Warning	Node is now offline.	
8459	Warning	Node needs to be reinitialized: Node has unknown data on it and must be reinitialized to be used in this system.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8460	Warning	Node is unavailable.	Node may be rebooting. Check for any network problems and be sure that the node is properly inserted in the slot.
8461	Warning	OSD out of service: Reconstruction will begin on the volumes in the StorageSet.	Wait until reconstruction completes on all volumes in the StorageSet and then replace the OSD.
8462	Warning	No compatibility info for Net Module 1 firmware (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	There is no compatibility info for the switch firmware. Downgrading firmware to a supported version is recommended. Contact VDURA Technical Support for further assistance.
8463	Critical	Net Module 1 firmware is incompatible (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Upgrade switch firmware to a supported version. Contact VDURA Technical Support for further assistance.

Error Code	Category	Error Description	Recommended Action
8464	Warning	Node disk 14 encountered disk errors, contact VDURA Technical Support.	Contact VDURA Technical Support.
8465	Info	Node removed: Node has been removed from the system.	No specific user action is required.
8466	Info	Node under temperature: Node is under optimal operating temperature.	Ambient temperature is too low. Adjust room temperature.
8467	Info	Node discovered by the system.	No specific user action is required.
8468	Info	Node is now online.	No specific user action is required.
8469	Info	Node has been moved from its previous enclosure and/or slot location.	No specific user action is required.
8470	Info	Node needs upgrade: Node software is inconsistent with system software.	The system will automatically attempt to apply its software to this node. No action is needed.
8471	Info	Node is performing a fsck.	Wait until fsck completes.

Error Code	Category	Error Description	Recommended Action
8472	Warning	Node unsupported version: Node hardware is not supported by current version of software (hardware serial %%BLADESERIAL%%, shelf serial %%SHELFID%%, slot %%SLOTID%%).	Check that the hardware serial on this node is valid.
8473	Warning	Node filesystem nearly full: The local filesystem %%FILESYSTEM%% on this node is %%PERCENTFULL%%% full.	
8474	Info	Node filesystem no longer full: The filesystem %%FILESYSTEM%% on node is %%PERCENTFULL%%% full. There is enough space now.	No specific user action is required.
8475	Info	Node disk no longer full: Node's disk space that is available for user data has reached %%PERCENTFULL%%% of capacity. There is enough space now.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
8476	Warning	Node disk 9 SMART predicted failure.	Contact VDURA Technical Support.
8477	Warning	Node disk 10 SMART predicted failure.	Contact VDURA Technical Support.
8480	Critical	Node disk 0 encountered disk errors.	Contact VDURA Technical Support.
8481	Warning	Node disk 0 SMART predicted failure.	Contact VDURA Technical Support.
8482	Warning	Node disk 0 is an unsupported version.	Contact VDURA Technical Support.
8483	Warning	Node disk 0 is missing from the node.	Contact VDURA Technical Support.
8484	Info	Node disk 0 is rebuilding to mirror disk 1.	No action is needed.
8485	Warning	Net Module 1 firmware needs upgrade but upgrade firmware file not available (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Switch firmware is not supported. Downgrading firmware to a supported version is recommended. Contact VDURA Technical Support for further assistance.

Error Code	Category	Error Description	Recommended Action
8486	Warning	Net Module 1 firmware is newer than tested version (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Contact VDURA Technical Support to obtain the firmware file and for further assistance.
8487	Warning	Net Module 1 detected uplink on a port that should not be an uplink (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Check that the enclosure switch is cabled correctly. Uplink should be sfp6, cx4, or lagg of ge.
8488	Warning	Net Module 1 is not using trigger logic - potential performance problem (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	If the warning persists, contact VDURA Technical Support.
8489	Info	Net Module 1 is experiencing packet flooding - performance may be reduced (switch 1 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	No specific user action is required.
8490	Critical	Node disk 1 encountered disk errors.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8491	Warning	Node disk 1 SMART predicted failure.	Contact VDURA Technical Support.
8492	Warning	Node disk 1 is an unsupported version.	Contact VDURA Technical Support.
8493	Warning	Node disk 1 is missing from the node.	Contact VDURA Technical Support.
8494	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
8495	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
8496	Info	Node disk 1 is rebuilding to mirror disk 0.	No specific user action is required.
8497	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
8498	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
8499	Warning	Node disk 12 SMART predicted failure.	Contact VDURA Technical Support.
8500	Warning	Node disk 13 SMART predicted failure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8501	Info	Storage added to StorageSet %%SETNAME%%.	No specific user action is required.
8502	Info	StorageSet affinity changed for shelf %%SHELFID%%.	No specific user action is required.
8503	Critical	StorageSet full: StorageSet %%SETNAME%% is at 100 % capacity, further writes or snapshots may not be allowed.	Add more storage to the StorageSet.
8504	Warning	Node disk 14 SMART predicted failure.	Contact VDURA Technical Support.
8505	Warning	Node disk 9 is an unsupported version.	Contact VDURA Technical Support.
8506	Warning	OSDs are not online in StorageSet %%SETNAME%%. See node and volume status for more details.	Check node and volume status.
8507	Warning	Node disk 10 is an unsupported version.	Contact VDURA Technical Support.
8508	Warning	StorageSet nearly full: %%PERCENTFULL%%% of space available to user in StorageSet %%SETNAME%% is full.	Add more storage to the StorageSet.

Error Code	Category	Error Description	Recommended Action
8509	Critical	StorageSet has no spare space: Reconstruction cannot be started on StorageSet %%SETNAME%% as there are no reserved OSDs in it. Add more storage to the StorageSet as soon as possible.	Add more storage to the StorageSet.
8510	Info	StorageSet created: StorageSet %%SETNAME%% is created.	No specific user action is required.
8511	Info	StorageSet deleted: StorageSet %%SETNAME%% is deleted.	No specific user action is required.
8513	Info	StorageSet capacity imbalance detected: The distribution of used space on OSDs in StorageSet %%SETNAME%% is not balanced. The system is working to rebalance the StorageSet automatically.	Turn on capacity balancing for this StorageSet. If capacity balancing is already turned on, the system will automatically balance the distribution of used space in this StorageSet.

Error Code	Category	Error Description	Recommended Action
8515	Warning	OSDv4 block allocation policy configuration is not consistent for all nodes in StorageSet %%SETNAME%%	
8516	Warning	StorageSet %%SETNAME%% uses OSDv4 block allocation policy that will be deprecated in the future releases	
8517	Warning	Node disk 11 is an unsupported version.	Contact VDURA Technical Support.
8518	Warning	Node disk 12 is an unsupported version.	Contact VDURA Technical Support.
8519	Warning	Node disk 13 is an unsupported version.	Contact VDURA Technical Support.
8520	Warning	Node disk 14 is an unsupported version.	Contact VDURA Technical Support.
8521	Warning	Volume degraded: One OSD is not responding. Performance on volume %%VOLNAME%% is degraded.	Minimize accessing this volume for reads and writes. The performance of the system will be slower than expected. Reconstruction of data will start soon.

Error Code	Category	Error Description	Recommended Action
8522	Warning	Volume offline: Volume %%VOLNAME%% has been taken offline- unavailable for administrative purpose.	
8523	Critical	Volume down: Volume %%VOLNAME%% is not accessible.	
8524	Info	Volume online: Volume %%VOLNAME%% has been brought online.	No specific user action is required.
8525	Info	Volume created: Volume %%VOLNAME%% has been created.	No specific user action is required.
8526	Info	Volume deleted: Volume %%VOLNAME%% has been deleted.	No specific user action is required.
8527	Warning	Volume has damaged files: Volume %%VOLNAME%% has damaged files.	Contact VDURA Technical Support.
8528	Warning	Volume offline: Volume %%VOLNAME%% has been taken offline for administrative purpose.	

Error Code	Category	Error Description	Recommended Action
8529	Warning	Inconsistent parity equation: Object %%OBJID%% (offset %%OFFSET%%, length %%LENGTH%%) has an inconsistent parity equation among its component objects.%%REPORTED_BY%%%%TRANSIENT% %	Contact VDURA Technical Support.
8531	Info	Volume %%VOLNAME%% has been set read- only.	No specific user action is required.
8532	Warning	Node disk 9 is missing from the node.	Contact VDURA Technical Support.
8533	Info	Volume background cleanup started: Background cleanup on the volume %%VOLNAME%% has begun.	No specific user action is required.
8534	Info	Volume background cleanup stopped: Background cleanup on the volume %%VOLNAME%% has stopped.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
8535	Critical	File System Recovery start: File System Recovery (FSRC) on the volume %%VOLNAME%% has begun.	The system is checking and repairing the volume. During this process, the volume is offline.
8536	Warning	File System Recovery recommended: File System Recovery (FSRC) recommended to repair inconsistencies in volume %%VOLNAME%%.	The system has discovered an unexpected file system inconsistency. To repair inconsistencies, click the volume name, then start FSRC using the Recover Volume control. The volume will be offline while FSRC repairs it. Note that depending on the size of the volume, FSRC may take several hours to complete. Until the volume is repaired, some files or directories may be inaccessible.

Error Code	Category	Error Description	Recommended Action
8537	Info	File System Recovery end: File System Recovery (FSRC) on the volume %%VOLNAME%% has completed. Fixed %%FIXED%% objects during recovery.	The system has finished checking and repairing the volume. To examine any damaged files, click the volume name and follow the link to Damaged Files in Volume.
8538	Info	Volume background cleanup completed: Background cleanup on the volume %%VOLNAME%% has completed. Reconstructed %%RECONED%% objects, removed %%REMOVED%% stale components.	No specific user action is required.
8539	Critical	Volume fault-tolerance exceeded: Volume %%VOLNAME%% has too many failed OSDs.	

Error Code	Category	Error Description	Recommended Action
8540	Critical	File System Recovery failed: File System Recovery (FSRC) on the volume %%VOLNAME%% failed to recover volume. Fixed %%FIXED%% objects, found %%FENCED%% damaged objects, and deleted %%DELETED%% objects during recovery.	An error occurred while checking and repairing the volume. The system is unable to bring the volume online. To start FSRC again, click the volume name, then click the Recover Volume control. If FSRC is unable to complete, contact VDURA Technical Support for assistance.
8541	Critical	File System Recovery failed: File System Recovery (FSRC) on the volume %%VOLNAME%% failed to recover volume.	One or more data fragments have been removed from file system but are still occupying space in the system. Contact VDURA Technical Support for assistance to recover them.
8542	Info	Volume reconstruction for drain: The system has started reconstructing data in volume %%VOLNAME%% to complete the drain process.	Wait until volume comes back online.

Error Code	Category	Error Description	Recommended Action
8543	Info	Volume rescue completed: Volume rescue on the volume %%VOLNAME%% has completed. Fixed %%FIXED%% objects, and deleted %%DELETED%% objects during rescue.	The system has finished rescuing the volume. Run PanCLI volume rescue report in service mode to download a detailed report.
8544	Warning	Node disk 10 is missing from the node.	Contact VDURA Technical Support.
8545	Warning	Node disk 11 is missing from the node.	Contact VDURA Technical Support.
8546	Warning	Node disk 12 is missing from the node.	Contact VDURA Technical Support.
8547	Warning	Node disk 13 is missing from the node.	Contact VDURA Technical Support.
8548	Warning	Node disk 14 is missing from the node.	Contact VDURA Technical Support.
8549	Info	Volume reconstruction stop: reconstruction of data in volume %%VOLNAME%% has stopped due to configuration or state changes.	Check node and volume status for more details.

Error Code	Category	Error Description	Recommended Action
8550	Warning	StorageSet has no spare space: StorageSet %%SETNAME%% has requested %%NUMDESIRED%% spare storage nodes but currently has no spares. This will not have any effect upon reconstruction of data that may be in progress, but the system will not be able to reconstruct or drain any additional storage nodes until another storage node is added.	Add more storage to the StorageSet.
8551	Critical	StorageSet spare failure: StorageSet %%SETNAME%% does not have enough space to reserve space for reconstruction. Requested %%NUMDESIRED%% MB but only reserved %%NUMRESERVED%% MB.	

Error Code	Category	Error Description	Recommended Action
8552	Warning	Volume reconstruction began: Reconstruction of data in volume %%VOLNAME%% began.	Wait until reconstruction completes.
8553	Critical	Reconstruction failed: Encountered error during reconstruction of data in volume %%VOLNAME%%.	Contact VDURA Technical Support.
8554	Info	Volume reconstruction complete: Reconstruction of data in volume %%VOLNAME%% is complete.	No specific user action is required.
8555	Info	StorageSet spare success: Successfully reserved requested space of %%NUMDESIRED%% MB in StorageSet %%SETNAME%%.	No specific user action is required.
8556	Critical	Volume reconstruction failed: some objects could not be reconstructed.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8557	Critical	Reconstruction failed: There was not enough space or not enough storage nodes to complete reconstruction in volume %%VOLNAME%%. Add more storage to the StorageSet.	Add more storage to the StorageSet.
8558	Info	StorageSet spare success: StorageSet %%SETNAME%% contains %%NUMDESIRED%% spare storage nodes for reconstruction.	No specific user action is required.
8559	Critical	Volume rescue reconstruction began: rescue reconstruction of data in volume %%VOLNAME%% began.	Wait until rescue reconstruction completes.

Error Code	Category	Error Description	Recommended Action
8560	Critical	Hard Quota Violation: Hard quota limit reached on volume %%VOLNAME%%. No further writes allowed in this volume.	Try one or more of the following: • Increase hard quota for the volume. • Delete unneeded files from the volume. Note that in an environment with snapshots, deleting files may not actually free space. • Delete the oldest snapshot and verify that enough space is available from the Storage tab. • Add more storage to the StorageSet.

Error Code	Category	Error Description	Recommended Action
8561	Warning	Soft Quota Violation: Soft quota limit reached on volume %%VOLNAME%%.	Try one or more of the following: • Increase soft quota for the volume. • Delete unneeded files from the volume. Note that in an environment with snapshots, deleting files may not actually free space. • Delete the oldest snapshot and verify that enough space is available from the Storage tab. • Add more storage to the StorageSet.
8564	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
8566	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
8567	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
8568	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8571	Warning	Software - Unexpected result: System software encountered unexpected behavior in module %%CODEMODULE%%: %%MESSAGE%%. (version %%VERSION%%, %%FILENAME%%:%%LINENUMBER%%)	
8572	Info	Software Process Failed: The process %%MODULENAME%% has failed with exit code %%EXITCODE%% and will be restarted.	The system has restarted this application automatically, so no action is needed. If you see many such messages, please report a bug with VDURA Technical Support.
8573	Info	Software Process Restarts: The process %%MODULENAME%% exited with status %%EXITCODE%% and will be restarted.	The system has restarted this application automatically, so no action is needed. If you see many such messages, please report a bug with VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8574	Info	DHCP Leases Low: A node failed to join the system as there are no IP Addresses available. Increase the range of IP Addresses.	Add more IP addresses to address range. Go to the Network page and select the IP Ranges page to make changes.
8575	Info	System online: The system is now online.	No specific user action is required.
8576	Info	System going offline-unavailable: The system is being brought offline. Data services will be turned off.	No specific user action is required.
8577	Warning	System offline-unavailable: Data services off.	
8578	Info	System rebooting: System is preparing for a reboot.	No specific user action is required.
8579	Info	System power down: System is powering down due to a failed node.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
8580	Warning	System is locked for new nodes: A new node was not allowed to join the system as the system is locked.	To let the node join, go to the Hardware page and unlock the system. If the new node is intended for a different system on this subnet, no action is required.
8581	Warning	IP address %%IP%% appears to be in use by a machine that is not part of this storage cluster.	To fix this problem, either edit the system IP ranges not to include any IPs that are taken by machines outside of your system or change the outside machine to use an IP address that is outside of the system's reserved IP ranges.
8582	Warning	DirectFlow Client with IP Address %%IPADDR%% clock skew of %%CLOCKSKEW%% seconds exceeds %%MAXCLOCKSKEW%% second limit	
8583	Warning	System reboot failed.	
8584	Warning	System power down failed.	

Error Code	Category	Error Description	Recommended Action
8585	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
8586	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
8587	Info	Generic informational message.	This is an informational message. No action is needed.
8588	Warning	Node has a warning condition the cause of which is either a kernel crash or unknown.	No specific user action is required. Contact VDURA Technical Support if the event reported causes any malfunction.
8589	Critical	Node is in error condition the cause of which is unknown.	No specific user action is required. Contact VDURA Technical Support if the event reported causes any malfunction.
8590	Info	Syslog warning: slow syncs seen by iopathv4	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
8591	Info	If the drift increases to 10,000 milliseconds, DirectFlow clients will degrade to read-only mode.	This is an informational message. However, it does suggest that the node generating the message is having problems contacting time servers or with its clock.
8592	Warning	This will cause DirectFlow clients to degrade to read-only mode.	Clock drift can be caused by loss of contact with your external time sources.
8593	Info	If the drift increases to 10,000 milliseconds, it may cause volumes on the StorageSet that this node is a part of to become inaccessible.	Clock drift can be caused by loss of contact with your external time sources.
8594	Warning	This will cause volumes on the StorageSet that this node is a part of to become temporarily inaccessible.	Clock drift can be caused by loss of contact with the realm president or local hardware or software failures.

Error Code	Category	Error Description	Recommended Action
8595	Info	ntpd monitor: Clock skew from time server %%IP%% of %%NUMMS%% milliseconds has dropped below the threshold of %%NUMMSTHRESHOLD%% milliseconds.	No action is needed. The time skew between the external time source and the node is improving.
8596	Warning	If this continues, DirectFlow clients will degrade to read-only mode.	Entering local clock sync can be caused by loss of contact with the time servers configured. You can use tools like ntptrace and ping to diagnose and rectify them.
8597	Warning	If this situation persists, this node's clock may drift from the clocks on the system, this node may go offline. Please check to see that this node re-synchronizes within the hour.	Nodes going into local time sync like this can be caused by loss of contact with the realm president or local hardware or software failures.
8598	Info	This event relates to being in local sync.	No action is needed. Whatever caused nodes to move to local sync has been rectified.

Error Code	Category	Error Description	Recommended Action
8599	Info	Timeout after 1 hour+10 min, as syslog_watcher will wait for 1 hour before sending a repeat message.	Collect the /var/log/ security file from this host to check the origin of failed login attempts.
8602	Warning	Failed to send email via SMTP server %%SMTPSERVER%% to %%EMAILADDRESSES%% with subject %%SUBJECT%%: %%MESSAGE%%.	Check the connectivity to SMTP server as well as SMTP server security settings.
8603	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
8604	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
8605	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
8606	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8607	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
8608	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
8609	Info	Node disk 9 is rebuilding to mirror disk 1.	No specific user action is required.
8610	Info	Node disk 10 is rebuilding to mirror disk 1.	No specific user action is required.
8612	Info	Node disk 11 is rebuilding to mirror disk 1.	No specific user action is required.
8613	Info	Node disk 12 is rebuilding to mirror disk 1.	No specific user action is required.
8614	Info	Node disk 13 is rebuilding to mirror disk 1.	No specific user action is required.
8615	Info	Node disk 14 is rebuilding to mirror disk 1.	No specific user action is required.
8620	Info	Upgrading message.	No specific user action is required.
8621	Info	Upgrading message.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
8622	Info	Upgrading message.	No specific user action is required.
8623	Info	Upgrading message.	No specific user action is required.
8624	Info	Upgrading message.	No specific user action is required.
8625	Critical	Upgrading message.	
8626	Info	Upgrading message.	No specific user action is required.
8627	Critical	Upgrading message.	
8628	Critical	Upgrading message.	
8629	Critical	Upgrading message.	
8630	Critical	Upgrading message.	
8631	Critical	Upgrading message.	
8632	Critical	Upgrading message.	
8633	Info	Upgrading message.	No specific user action is required.
8634	Critical	Upgrading message.	
8635	Warning	Upgrading message.	
8636	Info	Upgrading message.	No specific user action is required.
8637	Info	Upgrading message	No specific user action is required.
8638	Info	Upgrading message.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
8639	Info	Upgrading message.	No specific user action is required.
8640	Info	Upgrading message.	No specific user action is required.
8641	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
8642	Critical	Node disk %%DRIVENUMBER%% encountered disk errors.	Contact VDURA Technical Support.
8643	Warning	Node disk %%DRIVENUMBER%% SMART predicted failure.	Contact VDURA Technical Support.
8644	Warning	Node disk %%DRIVENUMBER%% is missing from the node.	Contact VDURA Technical Support.
8645	Warning	Node disk %%DRIVENUMBER%% is an unsupported version.	Contact VDURA Technical Support.
8646	Info	Node disk %%DRIVENUMBER%% is rebuilding to mirror disk 1.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
8647	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
8650	Info	System is coming up.	Wait a few minutes.
8651	Info	Initial system startup has completed.	Check the health of the system.
8652	Info	Volume service moved: Service %%SERVICEID%% moved from Director %%FROMHWSERIAL%% to Director %%TOHWSERIAL%%.	No specific user action is required.
8653	Warning	Volume service moved leaving no backup: Service %%SERVICEID%% moved from Director %%FROMHWSERIAL%% to Director %%TOHWSERIAL%%.	To ensure the volume service gets a backup location, make sure all directors are online.
8654	Info	Volume service backup enabled: Service %%SERVICEID%% now has backup Director %%TOHWSERIAL%%.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
8655	Warning	Volume service backup disabled: Service %%SERVICEID%% does not have a backup.	To ensure the volume service gets a backup location, make sure all directors are online.
8656	Info	System going offline: The system is being brought offline. Data services will be turned off.	No specific user action is required.
8657	Warning	System offline: Data services off.	
8658	Info	System going offline for upgrade: The system is being brought offline for upgrade. Data services will be turned off.	No specific user action is required.
8659	Warning	System offline for upgrade: Data services off.	
8661	Info	Alias %%IP_ALIAS%% has been assigned.	No specific user action is required.
8662	Warning	Alias %%IP_ALIAS%% cannot be assigned: %%MESSAGE%%	Check the status of all directors and validate IP alias configuration.

Error Code	Category	Error Description	Recommended Action
8663	Info	Volume service removed: Service %%SERVICEID%% manages no volumes and has been removed.	No specific user action is required.
8664	Warning	Detected too many recent login failures. Check event log for details about failures on individual hosts.	Collect the /var/log/ security file from affected hosts to check the origin of failed login attempts.
8665	Warning	Detected one or more AWOL clients. Check event log for client details.	Check the health of the clients reported as AWOL.
8670	Warning	This battery module is nearing the end of its useful 3.5 year life.	The battery module can continue to be used but should be replaced as soon as reasonably possible. Please open a Service Request with a list of battery hardware serial numbers impacted. The VDURA battery module is a consumable item but may be covered under the service entitlement you have purchased for this enclosure.

Error Code	Category	Error Description	Recommended Action
8684	Critical	Node's Small File Filesystem has reached 90% or more of capacity.	Delete data or add additional storage.
8685	Warning	Node's Small File Filesystem capacity is more than 80% full.	Delete data or add additional storage.
8686	Info	Node's Small File Filesystem has reached 50% or more of capacity.	No specific user action is required.
8687	Info	Node's Small File Filesystem is below 50% capacity.	No specific user action is required.
8688	Critical	Node's Large File Filesystem has reached 90% or more of capacity.	Delete data or add additional storage.
8689	Warning	Node's Large File Filesystem capacity is more than 80% full.	Delete data or add additional storage.
8690	Info	Node's Large File Filesystem has reached 50% or more of capacity.	No specific user action is required.
8691	Info	Node's Large File Filesystem is below 50% capacity.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
8697	Critical	Node's Metadata Filesystem has reached 90% or more of capacity.	Contact VDURA Technical Support.
8698	Warning	Node's Metadata Filesystem capacity is more than 70% full.	Contact VDURA Technical Support.
8699	Info	Node's Metadata Filesystem has reached below 70% of capacity.	No specific user action is required.
8800	Critical	JBOD (JBOD serial number: %%SERIAL%%) failed	Contact VDURA Technical Support.
8801	Critical	JBOD (JBOD serial number: %%SERIAL%%) not found.	Contact VDURA Technical Support.
8802	Warning	JBOD %%SERIAL%% predicted failure.	Contact VDURA Technical Support.
8803	Critical	IO module %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) failed.	Contact VDURA Technical Support.
8804	Critical	IO module %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) not found.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8805	Warning	IO module %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) predicted failure.	Contact VDURA Technical Support.
8806	Critical	JBOD fan %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) failed.	Contact VDURA Technical Support.
8807	Critical	JBOD fan %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) not found.	Contact VDURA Technical Support.
8808	Warning	JBOD fan %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) predicted failure.	Contact VDURA Technical Support.
8809	Critical	JBOD power supply %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) failed.	Contact VDURA Technical Support.
8810	Critical	JBOD power supply %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) not found.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8811	Warning	JBOD power supply %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) predicted failure.	Contact VDURA Technical Support.
8812	Warning	JBOD power supply %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) failed during the monthly battery test.	Contact VDURA Technical Support.
8813	Warning	JBOD power supply %%INDEX%% (JBOD serial number %%JBOD_SERIAL%%) AC power fail.	Contact VDURA Technical Support.
8900	Warning	Power supply 1 fan failed: Enclosure fan has failed in power supply 1 (fan serial %%SERIAL%%, shelf serial %%SHELFID%%). Power supply 1 is on right when looking at rear of shelf.	Contact VDURA Technical Support.
8901	Warning	Battery fan failed: Enclosure fan has failed in battery module (fan serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8902	Warning	Power supply 2 fan failed: Enclosure fan has failed in power supply 2 (fan serial %%SERIAL%%, shelf serial %%SHELFID%%). Power supply 2 is on left when looking at rear of shelf.	Contact VDURA Technical Support.
8903	Critical	There are 2 or more fans failed or missing in the enclosure. System will shut down automatically in a few minutes if this condition is not cleared up.	Contact VDURA Technical Support.
8950	Critical	Net module 2 has failed (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Contact VDURA Technical Support.
8951	Warning	Net Module 2 unsupported version (switch 2 serial, enclosure serial %%SHELFID%%).	Contact VDURA Technical Support.
8952	Critical	Net module 2 not found (enclosure serial %%SHELFID%%).	Check if switch 2 is inserted properly.

Error Code	Category	Error Description	Recommended Action
8953	Critical	Net module 2 over maximum temperature (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Ambient temperature is too high. Adjust room temperature.
8954	Warning	Net Module 2 exceeds optimal temperature (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Ambient temperature is too high. Adjust room temperature.
8955	Info	Net Module 2 under optimal temperature (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Ambient temperature is too low. Adjust room temperature.
8956	Warning	Net Module 2 checksum error: Checksum error on Device U6 (switch 2 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.
8957	Warning	Net Module 2 checksum error: Checksum error on Device U8 (switch 2 serial %%SERIAL%%, shelf serial %%SHELFID%%).	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
8958	Warning	Net Module 2 is resetting (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Switch management interface did not respond. Please wait 90 seconds for reset to complete.
8959	Warning	Management interface of Net Module 2 is not responding (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Switch management interface is not responding. Please reset the switch when system is not busy and check that the switch firmware is up-to-date.
8960	Warning	Net Module 2 firmware is out of date (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Switch firmware is out of date; upgrade is recommended. Contact VDURA Technical Support for further assistance.
8961	Warning	Net Module 2 firmware is unsupported (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Switch firmware is not supported. Downgrading firmware to a supported version is recommended. Contact VDURA Technical Support for further assistance.

Error Code	Category	Error Description	Recommended Action
8962	Warning	No compatibility info for Net Module 2 firmware (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	There is no compatibility info for the switch firmware. Downgrading firmware to a supported version is recommended. Contact VDURA Technical Support for further assistance.
8963	Critical	Net module 2 firmware is incompatible (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Upgrade switch firmware to a supported version. Contact VDURA Technical Support for further assistance.
8964	Warning	Net Module 2 firmware needs upgrade but upgrade firmware file not available (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Contact VDURA Technical Support.
8965	Warning	Net Module 2 firmware is newer than tested version (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Consider downgrading switch firmware to a supported version. Contact VDURA Technical Support for further assistance.

Error Code	Category	Error Description	Recommended Action
8966	Warning	Net Module 2 detected uplink on a port that should not be an uplink (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	Check that the enclosure switch is cabled correctly. Uplink should be sfp6, cx4, or lagg of ge.
8968	Info	Net Module 2 is experiencing packet flooding - performance may be reduced (switch 2 serial %%SERIAL%%, enclosure serial %%SHELFID%%).	If the warning persists, contact VDURA Technical Support.
8975	Info	Battery test started. Any change in LED pattern on this shelf can be safely ignored during the test. The test will take up to one minute (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	No specific user action is required.
8976	Info	Battery test completed successfully (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
8977	Critical	Battery test failed: battery drained too quickly (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	
8978	Info	Transient error during Battery test: battery test aborted (battery serial %%SERIAL%%, shelf serial %%SHELFID%%).	Battery test will automatically run again later.
9050	Warning	Node low battery: Node is waiting for shelf battery to get charged.	
9051	Info	Node online fsck active: Online fsck is running on the node.	No specific user action is required.
9052	Info	Node online fsck inactive: Online fsck stopped running on the node.	No specific user action is required.
9053	Info	Node got address: Node assigned IP Address by the system.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
9062	Info	Node clock sync: Node is performing NTP clock synchronization.	Wait up to 2 minutes for clock synchronization to complete. If clock synchronization lasts longer than 2 minutes, make sure at least one of the NTP servers is responding.
9063	Warning	One or more of this node's drives have failed.	Contact VDURA Technical Support.
9064	Warning	One or more of this node's drives may be damaged.	Contact VDURA Technical Support.
9065	Info	The quality of this node's drives cannot be determined at this time. The system will continue testing the drives until a determination can be made.	No action is needed.
9066	Info	Tests for this node's drives have passed.	No specific user action is required.
9067	Info	Node not ready to serve data: No active object storage service.	No specific user action is required.
9068	Info	Hardware activated.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
9069	Info	Performing secure erase.	No specific user action is required.
9070	Info	Hardware is booting.	No specific user action is required.
9071	Info	Hardware is shutting down.	No specific user action is required.
9072	Info	Timeout after 1 hour+10 min, as syslog_watcher will wait for 1 hour before sending a repeat message.	Check the health of this client.
9073	Warning	Found %%NUMECC%% ECC errors in the past %%TIMEPERIOD%%.	Node may need to be drained or replaced depending on severity of ECC correctable error rate.
9074	Warning	Found %%NUMERR%% %%ERRTYPE%% errors on drive %%DRIVENAME%%.	Node may need to be drained or replaced depending on severity of ATA sector error rate.
9075	Warning	Node is degraded: Controller is reconstructing this OSD.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9079	Warning	%%NUMERR%% IPS communication timeout errors occurred to OSD %%IP%% in the past %%TIMEPERIOD%%.	Collect logs for the OSD and contact VDURA Technical Support.
9080	Info	Hardware deactivated	No specific user action is required.
9081	Info	Node not ready to serve data: Network error.	No specific user action is required.
9082	Info	Node not ready to serve data: Polling/DB state is not reconciled.	No specific user action is required.
9083	Info	Node not ready to serve data: Configuration is not valid.	No specific user action is required.
9100	Info	OSD has successfully drained.	
9101	Info	OSD drain status: %%BYTES_REMAINING%% bytes (in %%OBJECTS_REMAINING%% objects) out of %%TOTAL_BYTES%% (and %%TOTAL_OBJECTS%% objects) still need to be drained.	Check the drain status on the OSD Detail page to see whether drain is blocked.

Error Code	Category	Error Description	Recommended Action
9102	Warning	Configuration message.	Correct the NIS configuration.
9103	Info	Configuration message.	No specific user action is required.
9104	Warning	Configuration message.	Make sure the NIS servers is responding.
9105	Info	Configuration message.	No specific user action is required.
9120	Warning	Lost connection to LDAP service could not be reestablished.	Check the LDAP server connectivity. To retry, disable and re-enable LDAP. Also make sure that LDAP is in the Name Resolution Order list.
9122	Warning	Connection to LDAP service reestablished.	
9126	Warning	Lost connection to NIS service could not be reestablished.	Check the NIS server connectivity. To retry, disable and re- enable NIS User/Group Name Resolution. Also make sure that NIS is in the Name Resolution Order list.
9127	Warning	Connection to NIS service reestablished.	

Error Code	Category	Error Description	Recommended Action
9150	Info	User Quota Violation Hard (bytes): Limit reached on volume %%VOLNAME%% for %%QUOTATYPE%%: %%QUOTAIDNAME%% (Id: %%QUOTAID%%) Limit = %%QUOTALIMIT%%. No further writes allowed for this %%QUOTATYPE%% in this volume unless it has at most %%QUOTATHRSHLD%% of data.	Increase hard quota bytes limit. Delete unneeded files from volume to keep usage below limit.
9151	Info	User Quota Violation Soft (bytes): Limit reached on volume %%VOLNAME%% for %%QUOTATYPE%%: %%QUOTAIDNAME%% (Id: %%QUOTAID%%) Limit = %%QUOTALIMIT%%.	Increase soft quota bytes limit. Delete unneeded files from volume to keep usage below limit.

Error Code	Category	Error Description	Recommended Action
9152	Info	User Quota Violation Hard (files): Limit reached on volume %%VOLNAME%% for %%QUOTATYPE%%: %%QUOTAIDNAME%% (Id: %%QUOTAID%%) Limit = %%QUOTALIMIT%%. No further file creates allowed for this %%QUOTATYPE%% in this volume.	Increase hard quota file limit. Delete unneeded files from volume to keep usage below limit.
9153	Info	User Quota Violation Soft (files): Limit reached on volume %%VOLNAME%% for %%QUOTATYPE%%: %%QUOTAIDNAME%% (Id: %%QUOTAID%%) Limit = %%QUOTALIMIT%%.	Increase soft quota file limit. Delete unneeded files from volume to keep usage below limit.
9170	Info	User quota usage calculation in progress for volume %%VOLNAME%%. Quota usage data may not be current.	

Error Code	Category	Error Description	Recommended Action
9171	Info	User quota usage calculation is pending for volume %%VOLNAME%%. Quota usage data may not be current.	No specific user action is required.
9172	Info	User quota usage calculation completed for volume %%VOLNAME%%.	No specific user action is required.
9182	Warning	Node NIC 1 failed: %%DESCRIPTION%%	Contact VDURA Technical Support.
9183	Warning	Node NIC 2 failed: %%DESCRIPTION%%	Contact VDURA Technical Support.
9184	Warning	Node NIC 3 failed: %%DESCRIPTION%%	Contact VDURA Technical Support.
9185	Warning	Node NIC 4 failed: %%DESCRIPTION%%	Contact VDURA Technical Support.
9186	Warning	Node NIC 5 failed: %%DESCRIPTION%%	Contact VDURA Technical Support.
9187	Warning	Node NIC 6 failed: %%DESCRIPTION%%	Contact VDURA Technical Support.
9188	Warning	Node NIC 7 failed: %%DESCRIPTION%%	Contact VDURA Technical Support.
9189	Warning	Node NIC 8 failed: %%DESCRIPTION%%	Contact VDURA Technical Support.
9190	Warning	Node NIC 9 failed: %%DESCRIPTION%%	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9202	Warning	Node NIC 1 transmission error: %%DESCRIPTION%%	Check cabling.
9203	Warning	Node NIC 2 transmission error: %%DESCRIPTION%%	Check cabling.
9204	Warning	Node NIC 3 transmission error: %%DESCRIPTION%%	Check cabling.
9205	Warning	Node NIC 4 transmission error: %%DESCRIPTION%%	Check cabling.
9206	Warning	Node NIC 5 transmission error: %%DESCRIPTION%%	Check cabling.
9207	Warning	Node NIC 6 transmission error: %%DESCRIPTION%%	Check cabling.
9208	Warning	Node NIC 7 transmission error: %%DESCRIPTION%%	Check cabling.
9209	Warning	Node NIC 8 transmission error: %%DESCRIPTION%%	Check cabling.
9210	Warning	Node NIC 9 transmission error: %%DESCRIPTION%%	Check cabling.
9222	Warning	Node NIC 1 slow transmission: %%DESCRIPTION%%	Check 10 GE switch configuration.

Error Code	Category	Error Description	Recommended Action
9223	Warning	Node NIC 2 slow transmission: %%DESCRIPTION%%	Check 10 GE switch configuration.
9224	Warning	Node NIC 3 slow transmission: %%DESCRIPTION%%	Check 10 GE switch configuration.
9225	Warning	Node NIC 4 slow transmission: %%DESCRIPTION%%	Check 10 GE switch configuration.
9226	Warning	Node NIC 5 slow transmission: %%DESCRIPTION%%	Check 10 GE switch configuration.
9227	Warning	Node NIC 6 slow transmission: %%DESCRIPTION%%	Check 10 GE switch configuration.
9228	Warning	Node NIC 7 slow transmission: %%DESCRIPTION%%	Check 10 GE switch configuration.
9229	Warning	Node NIC 8 slow transmission: %%DESCRIPTION%%	Check 10 GE switch configuration.
9230	Warning	Node NIC 9 slow transmission: %%DESCRIPTION%%	Check 10 GE switch configuration.
9242	Warning	Node NIC 1 not full duplex: %%DESCRIPTION%%	Check switch configuration.
9243	Warning	Node NIC 2 not full duplex: %%DESCRIPTION%%	Check switch configuration.

Error Code	Category	Error Description	Recommended Action
9244	Warning	Node NIC 3 not full duplex: %%DESCRIPTION%%	Check switch configuration.
9245	Warning	Node NIC 4 not full duplex: %%DESCRIPTION%%	Check switch configuration.
9246	Warning	Node NIC 5 not full duplex: %%DESCRIPTION%%	Check switch configuration.
9247	Warning	Node NIC 6 not full duplex: %%DESCRIPTION%%	Check switch configuration.
9248	Warning	Node NIC 7 not full duplex: %%DESCRIPTION%%	Check switch configuration.
9249	Warning	Node NIC 8 not full duplex: %%DESCRIPTION%%	Check switch configuration.
9250	Warning	Node NIC 9 not full duplex: %%DESCRIPTION%%	Check switch configuration.
9262	Warning	Node NIC 1 not active: %%DESCRIPTION%%	Check cabling.
9263	Warning	Node NIC 2 not active: %%DESCRIPTION%%	Check cabling.
9264	Warning	Node NIC 3 not active: %%DESCRIPTION%%	Check cabling.

Error Code	Category	Error Description	Recommended Action
9265	Warning	Node NIC 4 not active: %%DESCRIPTION%%	Check cabling.
9266	Warning	Node NIC 5 not active: %%DESCRIPTION%%	Check cabling.
9267	Warning	Node NIC 6 not active: %%DESCRIPTION%%	Check cabling.
9268	Warning	Node NIC 7 not active: %%DESCRIPTION%%	Check cabling.
9269	Warning	Node NIC 8 not active: %%DESCRIPTION%%	Check cabling.
9270	Warning	Node NIC 9 not active: %%DESCRIPTION%%	Check cabling.
9282	Critical	Node NIC 1 wrong broadcast address: %%DESCRIPTION%%	Configure the broadcast address correctly.
9283	Critical	Node NIC 2 wrong broadcast address: %%DESCRIPTION%%	Configure the broadcast address correctly.
9284	Critical	Node NIC 3 wrong broadcast address: %%DESCRIPTION%%	Configure the broadcast address correctly.
9285	Critical	Node NIC 4 wrong broadcast address: %%DESCRIPTION%%	Configure the broadcast address correctly.

Error Code	Category	Error Description	Recommended Action
9286	Critical	Node NIC 5 wrong broadcast address: %%DESCRIPTION%%	Configure the broadcast address correctly.
9287	Critical	Node NIC 6 wrong broadcast address: %%DESCRIPTION%%	Configure the broadcast address correctly.
9288	Critical	Node NIC 7 wrong broadcast address: %%DESCRIPTION%%	Configure the broadcast address correctly.
9289	Critical	Node NIC 8 wrong broadcast address: %%DESCRIPTION%%	Configure the broadcast address correctly.
9290	Critical	Node NIC 9 wrong broadcast address: %%DESCRIPTION%%	Configure the broadcast address correctly.
9302	Warning	Node NIC 1 not active: %%DESCRIPTION%%	Try rebooting the node.
9303	Warning	Node NIC 2 not in use: %%DESCRIPTION%%	Try rebooting the node.
9304	Warning	Node NIC 3 not in use: %%DESCRIPTION%%	Try rebooting the node.
9305	Warning	Node NIC 4 not in use: %%DESCRIPTION%%	Try rebooting the node.
9306	Warning	Node NIC 5 not in use: %%DESCRIPTION%%	Try rebooting the node.
9307	Warning	Node NIC 6 not in use: %%DESCRIPTION%%	Try rebooting the node.

Error Code	Category	Error Description	Recommended Action
9308	Warning	Node NIC 7 not in use: %%DESCRIPTION%%	Try rebooting the node.
9309	Warning	Node NIC 8 not in use: %%DESCRIPTION%%	Try rebooting the node.
9310	Warning	Node NIC 9 not in use: %%DESCRIPTION%%	Try rebooting the node.
9322	Warning	Node NIC 1 has no external connectivity: %%DESCRIPTION%%	Check cable and switch configuration.
9323	Warning	Node NIC 2 has no external connectivity: %%DESCRIPTION%%	Check cable and switch configuration.
9324	Warning	Node NIC 3 has no external connectivity: %%DESCRIPTION%%	Check cable and switch configuration.
9325	Warning	Node NIC 4 has no external connectivity: %%DESCRIPTION%%	Check cable and switch configuration.
9326	Warning	Node NIC 5 has no external connectivity: %%DESCRIPTION%%	Check cable and switch configuration.
9327	Warning	Node NIC 6 has no external connectivity: %%DESCRIPTION%%	Check cable and switch configuration.
9328	Warning	Node NIC 7 has no external connectivity: %%DESCRIPTION%%	Check cable and switch configuration.

Error Code	Category	Error Description	Recommended Action
9329	Warning	Node NIC 8 has no external connectivity: %%DESCRIPTION%%	Check cable and switch configuration.
9330	Warning	Node NIC 9 has no external connectivity: %%DESCRIPTION%%	Check cable and switch configuration.
9362	Warning	Node NIC 1 not in use: %%DESCRIPTION%%	Try rebooting the node.
9363	Warning	Node NIC 2 not in use: %%DESCRIPTION%%	Try rebooting the node.
9364	Warning	Node NIC 3 not in use: %%DESCRIPTION%%	Try rebooting the node.
9365	Warning	Node NIC 4 not in use: %%DESCRIPTION%%	Try rebooting the node.
9366	Warning	Node NIC 5 not in use: %%DESCRIPTION%%	Try rebooting the node.
9367	Warning	Node NIC 6 not in use: %%DESCRIPTION%%	Try rebooting the node.
9368	Warning	Node NIC 7 LACP not active: %%DESCRIPTION%%	Check switch configuration.
9369	Warning	Node NIC 8 LACP not active: %%DESCRIPTION%%	Check switch configuration.
9370	Warning	Node NIC 9 LACP not active: %%DESCRIPTION%%	Check switch configuration.

Error Code	Category	Error Description	Recommended Action
9400	Warning	Power supply 1 failed.	Contact VDURA Technical Support.
9401	Warning	Power supply 1 has unsupported version.	Contact VDURA Technical Support.
9402	Warning	Power supply 1 not found.	Check whether power supply 1 is inserted properly.
9403	Warning	Power supply 1 failure predicted.	Contact VDURA Technical Support.
9404	Warning	Power supply 1 AC power fail: AC power disconnected from power supply 1.	Check that power supply 1 is plugged in.
9410	Warning	Power supply 2 failed.	Contact VDURA Technical Support.
9411	Warning	Power supply 2 has unsupported version.	Contact VDURA Technical Support.
9412	Warning	Power supply 2 not found.	Check whether power supply 2 is inserted properly.
9413	Warning	Power supply 2 failure predicted.	Contact VDURA Technical Support.
9414	Warning	Power supply 2 AC power fail: AC power disconnected from power supply 2.	Check that power supply 2 is plugged in.
9420	Warning	IPMI card failed.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9421	Warning	IPMI card not responding correctly.	Contact VDURA Technical Support.
9422	Warning	IPMI card not found.	Check whether IPMI card is inserted properly.
9423	Warning	Non-volatile memory system failure is predicted.	Contact VDURA Technical Support.
9424	Warning	Non-volatile memory system has failed.	Contact VDURA Technical Support.
9430	Warning	Fan 1 failed.	Contact VDURA Technical Support.
9431	Warning	Fan 1 speed too slow.	Contact VDURA Technical Support.
9435	Warning	Fan 2 failed.	Contact VDURA Technical Support.
9436	Warning	Fan 2 speed too slow.	Contact VDURA Technical Support.
9440	Warning	Fan 3 failed.	Contact VDURA Technical Support.
9441	Warning	Fan 3 speed too slow.	Contact VDURA Technical Support.
9445	Warning	Fan 4 failed.	Contact VDURA Technical Support.
9446	Warning	Fan 4 speed too slow.	Contact VDURA Technical Support.
9450	Warning	Fan 5 failed.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9451	Warning	Fan 5 speed too slow.	Contact VDURA Technical Support.
9455	Warning	Fan 6 failed.	Contact VDURA Technical Support.
9456	Warning	Fan 6 speed too slow.	Contact VDURA Technical Support.
9460	Warning	Fan 7 failed.	Contact VDURA Technical Support.
9461	Warning	Fan 7 speed too slow.	Contact VDURA Technical Support.
9470	Warning	Fibre channel card in port 1 failed.	Contact VDURA Technical Support.
9471	Warning	Fibre channel port 1 link is slow.	Check cabling and SFP+ modules.
9472	Warning	Fibre channel port 1 link is down.	Check cabling and SFP+ modules.
9473	Warning	Fibre channel card in port 2 failed.	Contact VDURA Technical Support.
9474	Warning	Fibre channel port 2 link is slow.	Check cabling and SFP+ modules.
9475	Warning	Fibre channel port 2 link is down.	Check cabling and SFP+ modules.
9500	Critical	Node Disk 2 encountered disk errors, contact Panasas Technical Support.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9501	Warning	Node disk 2 SMART predicted failure.	Contact VDURA Technical Support.
9502	Warning	Node disk 2 is an unsupported version.	Contact VDURA Technical Support.
9503	Warning	Node disk 2 is missing from the node.	Contact VDURA Technical Support.
9504	Warning	SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
9505	Warning	SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
9506	Info	The hardware serial of node has changed.	No specific user action is required.
9600	Warning	Battery Fan Module 1 failed: Fan Module 1 of battery module serial %%SERIAL%% in enclosure serial %%SHELFID%% has failed. Fan Module 1 is on the left when looking down on the mother board of battery module from the rear of the enclosure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9601	Warning	Battery Fan Module 2 failed: Fan Module 2 of battery module serial %%SERIAL%% in enclosure serial %%SHELFID%% has failed. Fan Module 2 is in the middle when looking down on the mother board of battery module from the rear of the enclosure.	Contact VDURA Technical Support.
9602	Warning	Battery Fan Module 3 failed: Fan Module 3 of battery module serial %%SERIAL%% in enclosure serial %%SHELFID%% has failed. Fan Module 3 is on the right when looking down on the mother board of battery module from the rear of the enclosure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9603	Warning	Battery Fan Module 1 not found: Fan Module 1 of battery module serial %%SERIAL%% in enclosure serial %%SHELFID%% not found. Fan Module 1 is on the left when looking down on the mother board of battery module from the rear of the enclosure.	Contact VDURA Technical Support.
9604	Warning	Battery Fan Module 2 not found: Fan Module 2 of battery module serial %%SERIAL%% in enclosure serial %%SHELFID%% not found. Fan Module 2 is in the middle when looking down on the mother board of battery module from the rear of the enclosure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9605	Warning	Battery Fan Module 3 not found: Fan Module 3 of battery module serial %%SERIAL%% in enclosure serial %%SHELFID%% not found. Fan Module 3 is on the right when looking down on the mother board of battery module from the rear of the enclosure.	Contact VDURA Technical Support.
9609	Warning	Node Fan Module 1 failed: Fan Module 1 of node serial %%SERIAL%% in enclosure serial %%SHELFID%% has failed. Fan Module 1 is on the left when looking down on the mother board of node from the rear of the enclosure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9610	Warning	Node Fan Module 2 failed: Fan Module 2 of node serial %%SERIAL%% in enclosure serial %%SHELFID%% has failed. Fan Module 2 is in the middle when looking down on the mother board of node from the rear of the enclosure.	Contact VDURA Technical Support.
9611	Warning	Node Fan Module 3 failed: Fan Module 3 of node serial %%SERIAL%% in enclosure serial %%SHELFID%% has failed. Fan Module 3 is on the right when looking down on the mother board of node from the rear of the enclosure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9612	Warning	Node Fan Module 1 not found: Fan Module 1 of node serial %%SERIAL%% in enclosure serial %%SHELFID%% not found. Fan Module 1 is on the left when looking down on the mother board of node from the rear of the enclosure.	Contact VDURA Technical Support.
9613	Warning	Node Fan Module 2 not found: Fan Module 2 of node serial %%SERIAL%% in enclosure serial %%SHELFID%% not found. Fan Module 2 is in the middle when looking down on the mother board of node from the rear of the enclosure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9614	Warning	Node Fan Module 3 not found: Fan Module 3 of node serial %%SERIAL%% in enclosure serial %%SHELFID%% not found. Fan Module 3 is on the right when looking down on the mother board of node from the rear of the enclosure.	Contact VDURA Technical Support.
9624	Warning	Node's metadata filesystem capacity is nearly full: %%PERCENTFULL%%% of metadata filesystem space has been used.	Delete data or add additional storage.
9625	Critical	Node's metadata filesystem space has reached 100% of capacity. No additional data can be written to the node.	Delete data or add additional storage.
9626	Info	Node metadata filesystem is no longer full: Node's metadata filesystem reached %%PERCENTFULL%%% of capacity.	No specific user action is required.

Error Code	Category	Error Description	Recommended Action
9627	Warning	StorageSet metadata filesystem capacity is nearly full: %%PERCENTFULL%%% of metadata filesystem space in StorageSet %%SETNAME%% has been used.	Contact VDURA Technical Support.
9628	Critical	StorageSet %%SETNAME%% is at 100% metadata filesystem capacity. No additional data can be written to the StorageSet %%SETNAME%%.	Contact VDURA Technical Support.
9629	Info	StorageSet metadata filesystem capacity is no longer full: StorageSet %%SETNAME%% metadata filesystem reached %%PERCENTFULL%%% of capacity.	No specific user action is required.
9630	Warning	KMIP server 1 (%%KMIP_ADDR%%) failed: %%RC%%	Contact VDURA Technical Support.
9631	Warning	KMIP server 2 (%%KMIP_ADDR%%) failed: %%RC%%	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9640	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9641	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9642	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9643	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9644	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9645	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9646	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9647	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9648	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9649	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9650	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9651	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9652	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9653	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9654	Critical	Failed to unlock %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9655	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9656	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9657	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9658	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9659	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9660	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9661	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9662	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9663	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9664	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9665	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9666	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9667	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9668	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.
9669	Warning	Failed to configure %%DRIVE%% drive: %%RC%%	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
9870	Warning	Client certificate for KMIP 1 will soon expire.	Update the client certificate for KMIP 1 until the expiration date.
9871	Warning	Root certificate for KMIP 1 will soon expire.	Update the client certificate for KMIP 1 until the expiration date.
9872	Warning	Client certificate for KMIP 2 will soon expire.	Update the client certificate for KMIP 2 until the expiration date.
9873	Warning	Root certificate for KMIP 2 will soon expire.	Update the root certificate for KMIP 2 until the expiration date.
9874	Critical	Client certificate for KMIP 1 expired.	Update the client certificate.
9875	Critical	Root certificate for KMIP 1 expired.	Update the root certificate.
9876	Critical	Client certificate for KMIP 2 expired.	Update the client certificate.
9877	Critical	Root certificate for KMIP 2 expired.	Update the root certificate.
9902	Critical	PAN_MGMT_EV_NODE_DRIVES_MISMATCH	Place drives into proper location.
9903	Warning	PAN_MGMT_EV_NODE_DRIVES_UNEXPECTED	Place drives into proper location.

Error Code	Category	Error Description	Recommended Action
10000	Warning	Power supply 3 failed (power supply module serial %%SERIAL%%). %%DESCRIPTION%%	Contact VDURA Technical Support.
10001	Warning	Power supply 3 unsupported version. %%DESCRIPTION%%	Contact VDURA Technical Support.
10002	Warning	Power supply 3 not found in enclosure with serial %%SHELFID%%. %%DESCRIPTION%%	Check if power supply module 3 is inserted properly.
10003	Critical	Power supply 3 over maximum temperature (power supply module serial %%SERIAL%%, enclosure serial %%SHELFID%%). %%DESCRIPTION%%	Ambient temperature is too high. Adjust room temperature.
10004	Warning	Power supply 3 temperature high (power supply module serial %%SERIAL%%, enclosure serial %%SHELFID%%). %%DESCRIPTION%%	Ambient temperature is too high. Adjust room temperature.

Error Code	Category	Error Description	Recommended Action
10005	Info	Power supply 3 under temperature (power supply module serial %%SERIAL%%, enclosure serial %%SHELFID%%). %%DESCRIPTION%%	Ambient temperature is too low. Adjust room temperature.
10006	Warning	Power supply 3, AC power fail: AC power disconnected from power supply module 3 (power supply module serial %%SERIAL%%, enclosure serial %%SHELFID%%). %%DESCRIPTION%%	Check that power supply module is plugged in.
10007	Warning	Power supply 3 failed during the monthly battery test (power supply module serial %%SERIAL%%). %%DESCRIPTION%%	Contact VDURA Technical Support.
10008	Warning	A 950W power supply module was replaced with an 800W power supply module 3 (enclosure serial %%SHELFID%%). This scenario can cause damage to the 800W power supply module. %%DESCRIPTION%%	Insert a 950W power supply module into this location.

Error Code	Category	Error Description	Recommended Action
10100	Warning	Power supply 4 failed (power supply module serial %%SERIAL%%). %%DESCRIPTION%%	Contact VDURA Technical Support.
10101	Warning	Power supply 4 unsupported version. %%DESCRIPTION%%	Contact VDURA Technical Support.
10102	Warning	Power supply 4 not found in enclosure with serial %%SHELFID%%. %%DESCRIPTION%%	Check if power supply module 4 is inserted properly.
10103	Critical	Power supply 4 over maximum temperature (power supply module serial %%SERIAL%%, enclosure serial %%SHELFID%%). %%DESCRIPTION%%	Ambient temperature is too high. Adjust room temperature.
10104	Warning	Power supply 4 temperature high (power supply module serial %%SERIAL%%, enclosure serial %%SHELFID%%). %%DESCRIPTION%%	Ambient temperature is too high. Adjust room temperature.

Error Code	Category	Error Description	Recommended Action
10105	Info	Power supply 4 under temperature (power supply module serial %%SERIAL%%, enclosure serial %%SHELFID%%). %%DESCRIPTION%%	Ambient temperature is too low. Adjust room temperature.
10106	Warning	Power supply 4, AC power fail: AC power disconnected from power supply module 4 (power supply module serial %%SERIAL%%, enclosure serial %%SHELFID%%). %%DESCRIPTION%%	Check that power supply module is plugged in.
10107	Warning	Power supply 4 failed during the monthly battery test (power supply module serial %%SERIAL%%). %%DESCRIPTION%%	Contact VDURA Technical Support.
10108	Warning	A 950W power supply module was replaced with an 800W power supply module 4 (enclosure serial %%SHELFID%%). This scenario can cause damage to the 800W power supply module. %%DESCRIPTION%%	Insert a 950W power supply module into this location.

Error Code	Category	Error Description	Recommended Action
10902	Warning	Power supply 3 fan failed: enclosure fan has failed in power supply module 3 (power supply module serial %%SERIAL%%, enclosure serial %%SHELFID%%). Power supply 3 is on bottom left when looking at rear of enclosure.	Contact VDURA Technical Support.
10903	Warning	Power supply 4 fan failed: enclosure fan has failed in power supply module 4 (power supply module serial %%SERIAL%%, enclosure serial %%SHELFID%%). Power supply 4 is on top left when looking at rear of enclosure.	Contact VDURA Technical Support.
12000	Warning	Node disk 3 encountered disk errors.	Contact VDURA Technical Support.
12001	Warning	Node disk 3 SMART predicted failure.	Contact VDURA Technical Support.
12002	Warning	Node disk 3 is an unsupported version.	Contact VDURA Technical Support.
12003	Warning	Node disk 3 is missing from the node.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
12004	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
12005	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
12006	Info	Node disk 3 is rebuilding to mirror disk 1.	No specific user action is required.
12200	Info	Node disk 1 is rebuilding to mirror disk 0.	No specific user action is required.
12020	Warning	Node disk 4 encountered disk errors.	Contact VDURA Technical Support.
12021	Warning	Node disk 4 SMART predicted failure.	Contact VDURA Technical Support.
12022	Warning	Node disk 4 is an unsupported version.	Contact VDURA Technical Support.
12023	Warning	Node disk 4 is missing from the node.	Contact VDURA Technical Support.
12024	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
12025	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
12026	Info	Node disk 4 is rebuilding to mirror disk 1.	No specific user action is required.
12040	Warning	Node disk 5 encountered disk errors.	Contact VDURA Technical Support.
12041	Warning	Node disk 5 SMART predicted failure.	Contact VDURA Technical Support.
12042	Warning	Node disk 5 is an unsupported version.	Contact VDURA Technical Support.
12043	Warning	Node disk 5 is missing from the node.	Contact VDURA Technical Support.
12044	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
12045	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
12046	Info	Node disk 5 is rebuilding to mirror disk 1.	No specific user action is required.
12060	Warning	Node disk 6 encountered disk errors.	Contact VDURA Technical Support.
12061	Warning	Node disk 6 SMART predicted failure.	Contact VDURA Technical Support.
12062	Warning	Node disk 6 is an unsupported version.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
12063	Warning	Node disk 6 is missing from the node.	Contact VDURA Technical Support.
12064	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
12065	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
12066	Info	Node disk 6 is rebuilding to mirror disk 1.	No specific user action is required.
12080	Warning	Node disk 7 encountered disk errors.	Contact VDURA Technical Support.
12081	Warning	Node disk 7 SMART predicted failure.	Contact VDURA Technical Support.
12082	Warning	Node disk 7 is an unsupported version.	Contact VDURA Technical Support.
12083	Warning	Node disk 7 is missing from the node.	Contact VDURA Technical Support.
12084	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
12085	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.

Error Code	Category	Error Description	Recommended Action
12086	Info	Node disk 7 is rebuilding to mirror disk 1.	No specific user action is required.
12100	Warning	Node disk 8 encountered disk errors.	Contact VDURA Technical Support.
12101	Warning	Node disk 8 SMART predicted failure.	Contact VDURA Technical Support.
12102	Warning	Node disk 8 is an unsupported version.	Contact VDURA Technical Support.
12103	Warning	Node disk 8 is missing from the node.	Contact VDURA Technical Support.
12104	Warning	Node SSD wearout warning: Drive is nearing failure.	Contact VDURA Technical Support.
12105	Warning	Node SSD available reserve space low: Drive is nearing failure.	Contact VDURA Technical Support.
12106	Info	Node disk 8 is rebuilding to mirror disk 1.	No specific user action is required.