



Vantage User Guide

The web-based management console for VDURA storage realms

© 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 12.0.0-2026-10-19-11:59:32

1. Introduction	1
1.1 Important Note Regarding VDURA System GUI Options	1
1.2 What Vantage Is	1
1.3 How Vantage Relates to panctl and PanAPI	1
1.4 Audience and Scope of This Guide	2
2. Getting Started	3
2.1 Accessing the UI	3
2.2 Signing In	3
2.3 Layout Overview	4
3. Common UI Features in Vantage	7
3.1 Table Toolbar	7
3.2 Status Icons	7
3.3 Required Fields and Tooltips	8
3.4 Confirm Before Destroy	8
3.5 Page Not Found	8
4. Dashboard	9
4.1 Node/StorageSet/Volume Summary Cards	9
4.2 Total Capacity	11
4.3 Alerts	11
4.4 System Performance	12
5. Status Center	14
5.1 Alerts Panel	14
5.2 Events Panel	15
5.3 Audit Log Panel	17
6. File System	18
6.1 StorageSets	18
6.1.1 StorageSets List Page	18
6.1.2 StorageSet Detail Page	20
6.1.3 Edit Spare VPODs Modal	24
6.1.4 Add Storage Modal	25
6.1.5 Delete StorageSet Confirmation	27
6.1.6 Create StorageSet Page	27
6.2 Creating and Managing StorageSets	28
6.2.1 Create a StorageSet	28
6.2.2 Add Storage to a StorageSet	29

6.2.3 Modify a StorageSet	29
6.2.4 Delete a StorageSet	29
6.3 Volumes	30
6.3.1 Volumes List Page	30
6.3.2 Volume Detail Page	32
6.3.3 Delete Volume Confirmation	35
6.3.4 Create Volume Page	35
6.4 Creating and Managing Volumes	39
6.4.1 Create a Volume	39
6.4.2 Delete a Volume	40
7. Realm Nodes	41
7.1 Director Nodes Table	41
7.2 Storage Nodes Table	42
7.3 Node Detail Page	44
8. Hardware	48
8.1 Hardware List Page	48
8.2 Chassis Detail Page	49
8.3 Parent Field on Realm Node Detail	53
8.4 Hardware on Virtual Realms	54
9. Tenants	55
9.1 Tenants List Page	55
9.2 Tenant Detail Page	56
9.2.1 Delete Tenant Confirmation	62
9.2.2 Finishing Tenant Setup	63
9.2.3 Networking Tab	65
9.2.4 Create Network Configuration Page	68
9.3 Create Tenant Page	71
9.4 Acting as a Tenant (Context Switching)	72
9.4.1 Give a Tenant Its Own Isolated Storage	78
9.5 Creating and Managing Tenants	78
9.5.1 Create a Tenant	79
9.5.2 Assign a StorageSet to a Tenant	79
9.5.3 Modify a Tenant's StorageSet Assignment	79
9.5.4 Modify a Tenant	79
9.5.5 Delete a Tenant	79

9.5.6 Create a Network Configuration for a Tenant	80
10. Identity and Access	81
10.1 Users List Page	81
10.2 User Detail Page	82
10.2.1 Edit Email Modal	83
10.2.2 Reset Password Modal	84
10.2.3 Manage Roles Modal	85
10.2.4 Delete User Confirmation	86
10.3 Create User Page	86
10.4 Creating and Managing Users	87
10.4.1 Create a User	87
10.4.2 Change a User's Email	88
10.4.3 Reset a User's Password	88
10.4.4 Change a User's Roles	88
10.4.5 Delete a User	88
10.5 Service Accounts List Page	89
10.6 Service Account Detail Page	90
10.6.1 API Tokens Tab	91
10.6.2 Issue Token Modal	92
10.6.3 Token Issued Modal	93
10.6.4 Manage Roles Modal	94
10.6.5 Delete Service Account Confirmation	95
10.7 Create Service Account Page	95
10.8 Creating and Managing Service Accounts	96
10.8.1 Create a Service Account	96
10.8.2 Issue an API Token	96
10.8.3 Revoke an API Token	97
10.8.4 Change a Service Account's Roles	97
10.8.5 Delete a Service Account	97
11. Account, Preferences, and Help	98
11.1 Account Menu and My Account	98
11.2 Theme	101
11.3 Get Help	101
11.4 About	103
11.5 PanActive Manager	104
12. Security Enhanced Linux in VDURA Systems	105

12.1 Terminology	105
12.2 SELinux Overview	106
12.3 PanFS and MAC	106
12.4 Configuration and Use	107
12.5 Important Considerations	108
13. End-to-End Encryption	109
13.1 Introduction and Overview	109
13.1.1 What is Encrypted	109
13.1.2 How Data is Encrypted	109
13.1.3 Metadata Encryption	110
13.2 Key Management	110
13.2.1 Role-Based Authentication	110
13.2.2 Encryption Domain	110
13.2.3 KMS Integration	110
13.3 DirectFlow Client Behavior	110
13.4 Gateway Protocols	111
13.5 Configuration Considerations	111
13.5.1 External KMS Configuration	111
13.5.2 Key Management	111
13.5.3 Encrypted Volume Configuration	112
13.5.4 Client Configuration	112
13.5.5 Secure Status	112
13.5.6 Mounting PanFS	112
13.6 Failure Modes	112
13.6.1 KMS Network Connection Issues	113
13.6.2 Lost or Destroyed Keys	113
13.6.3 Compromised Keys	113
13.6.4 Reading Data without a Key	113
13.6.5 Reading Data with the Wrong Key	113
13.7 Performance Considerations	113
13.8 Upgrades and Downgrades	113
13.9 Encryption Configuration Procedures	114
13.9.1 KMIP Server Configuration	114
13.9.2 Realm-Wide Encryption Configuration	115
13.9.3 Volume-Level Encryption Configuration	115
13.9.4 Client Interface Configuration	116

14. S3	118
14.1 Concepts	118
14.2 Logging In with panctl	119
14.3 Deployment Options	120
14.3.1 Dedicated Tenant (Recommended)	120
14.3.2 Platform Tenant	120
14.3.3 Choices Within a Tenant	120
14.4 Prerequisites	121
14.4.1 Tenant and Storage	121
14.4.2 Tenant Subnet	121
14.4.3 Attaching Client Hosts to the Tenant Subnet	121
14.4.4 Volume	122
14.5 Enabling S3	123
14.5.1 Realm Switch	123
14.5.2 Tenant Configuration	123
14.5.3 Settings Reference	124
14.6 Gateway Instances	125
14.6.1 Choose Nodes	125
14.6.2 Create and Start	126
14.6.3 Endpoints	126
14.6.4 States and Health	127
14.6.5 Stop, Start, Delete	127
14.7 S3 Users	128
14.8 Buckets	129
14.9 TLS Certificates	130
14.9.1 Tenant-Provided Certificate (Recommended)	130
14.9.2 Realm-Generated Certificate (Default)	133
14.10 Aggregator and Named Access (DNS)	133
14.10.1 Aggregator (Experimental)	133
14.10.2 Named Access with Your Own DNS (external)	133
14.10.3 Realm-Published Names (panfs, Experimental)	135
14.11 Client Access	135
14.11.1 AWS CLI	136
14.11.2 Python (boto3)	137
14.11.3 rclone	137
14.11.4 Connecting to a Realm-Generated Certificate	138

14.12 Supported S3 Features	138
14.12.1 Bucket Policies	140
14.12.2 Bucket ACLs	140
14.12.3 Public Access	141
14.12.4 Browser Uploads (POST Object)	142
14.12.5 Object Lock	143
14.12.6 Not Implemented	145
14.12.7 Server-Side Encryption	145
14.13 S3 and File Access to the Same Data	146
14.14 Monitoring and Logs	146
14.15 Day-to-Day Operations	147
14.16 Removing S3 from a Tenant	148
14.17 S3 on the Platform Tenant	148
14.18 Troubleshooting	149
15. Kubernetes CSI Driver for PanFS	153
15.1 The PanFS CSI Driver	153
15.2 Installation, Configuration, and Diagnostics	153
16. Troubleshooting	155
17. VDURA System Error Messages	157

1. Introduction

1.1 Important Note Regarding VDURA System GUI Options

Vantage is VDURA's new GUI to manage V5000 storage hardware. As of this release, it does not yet fully cover all aspects of system management. Users wishing to use a GUI for system administration tasks not yet covered in Vantage can switch to the legacy PanActive Manager GUI via the link at the bottom of the left nav bar. Conversely, there is new functionality in Vantage that is not available in PanActive Manager. Users in PanActive Manager can switch to Vantage with the VANTAGE UI link in the upper-right corner of the page.

The next several PanFS releases will continue to add functionality to Vantage until PanActive Manager can be fully deprecated.

The VDURA Data Platform User Guide, which covers PanActive Manager, is available on the VDURA documentation portal.

1.2 What Vantage Is

Vantage is VDURA's web-based management console for a VDURA storage realm. It is the graphical counterpart to `panctl`: where `panctl` is a command-line client for scripting and automation, Vantage is a browser UI for day-to-day visual administration: dashboards, guided create/edit forms, tables you can filter and sort, and confirmation dialogs before anything destructive happens.

1.3 How Vantage Relates to `panctl` and PanAPI

Vantage and `panctl` are independent clients of the same backend REST service (**PanAPI**/ `vdura-api`). Neither of them wraps another, and none is more authoritative; they both read and write the same underlying state:

- A StorageSet, volume, tenant, or user created in Vantage is immediately visible to `panctl` (e.g. `panctl storagesets list`), and vice versa.
- Vantage's **Audit Log** panels (see 5.3, 9.2, and 10.2) show the same records as `panctl audit list`.
- Vantage's **Events** panels show the same feed as `panctl events list`.
- The **Alerts** table in Status Center (see 5) shows the same active-alert set as `panctl alerts list`.

Vantage does not expose every `panct1` capability, and `panct1` does not expose every UI convenience (for example, the Dashboard's visual capacity charts). Treat them as complementary front ends to the same realm.

1.4 Audience and Scope of This Guide

This guide is written for VDURA storage administrators who will use Vantage to monitor and configure a realm: reviewing health, managing StorageSets and volumes, and managing tenants and users. It walks through every page, panel, button, and modal in the UI, explaining what each screen is telling you and what each control does.

2. Getting Started

2.1 Accessing the UI

Vantage is served directly from the realm over HTTPS; there is nothing to install. Point a browser at the realm's address (an IP address or hostname), e.g.:

```
https://<realm-address>/
```

Certificate warning on first visit: VDURA realms ship with a self-signed TLS certificate by default. The first time any browser visits a given realm's UI, it will show a certificate-trust warning. This is expected; click through it (or add a permanent exception) to proceed. You will not see this warning again on subsequent visits from the same browser.

2.2 Signing In

Visiting the realm address without a valid session redirects to a dedicated **Login** page.




Figure 2.1. The Login page: Vantage wordmark, centered Username/Password form, Sign in button, and footer links.

Vantage uses the same credentials as `panct1` and PanAPI: the username and password of a platform user, such as the `admin` account. The page shows the Vantage wordmark top-left and a centered card with **Username** and **Password** fields and a **Sign in** button. The **Username** field is selected when the page loads: its label moves up onto the field's border and the prompt "Enter username" appears inside it.

Three links run along the bottom of the page, above the copyright line:

Link	Description
Documentation	Opens the VDURA documentation page on vdura.com .
Get Support	Opens the VDURA technical support page on vdura.com .
Switch To Legacy UI	Opens PanActive Manager, the legacy UI, at the same realm address (see 11.5).

If an existing session expires, you're returned to this same page.

Note: PanActive Manager, the legacy UI (see [11.5](#)), maintains a **separate** session from Vantage. Signing in to one does not sign you in to the other.

2.3 Layout Overview

Every page in Vantage shares the same frame:

Links at the top right of each page provide basic realm and user information.

- Status icon: Displays overall realm health (a green checkmark or a warning/error icon). Click it to navigate to [Status Center](#) and view a full list of alerts and events. Its tooltip reads "All systems healthy · Click to view Status Center" when nothing is wrong.
- Realm name: Displays the name of the realm you are viewing (e.g. `virtual-realm.local.com`).
- User dropdown menu: Displays the username currently logged in (e.g. "AD" for the built-in `admin` user). Click the dropdown menu to access links to your account information (see [11.1](#)) or to sign out.
- System time: Displays the current system time, shown in UTC.

Left sidebar: a collapsible icon rail. By default it shows icons only; clicking any icon (or the rail itself) expands it to show text labels. The items, top to bottom, are:

Icon	Function
Vantage logo	Click to expand/collapse the left menu horizontally to view icons only or icons with names.

Icon	Function
Dashboard	Click to display overall realm health and performance (see 4).
Status Center	Click to display realm alerts and events (see 5).
File System	Click to select from two categories (see 6): StorageSets (view, create, modify, and delete StorageSets) or Volumes (view, create, modify, and delete volumes).
Realm Nodes	Click to display all Director and Storage Nodes in the realm (see 7). Click a node name to view its details.
Hardware	Displays realm hardware details on physical VDURA hardware; disabled in a virtual realm (see 8).
Tenants	Click to view, create, modify, and delete realm tenants (see 9).
Identity & Access	Click to select from two categories (see 10): Users (view, create, modify, and delete several types of realm admin and observer users) or Service Accounts (view, create, modify, and delete the accounts that scripts and integrations use to authenticate).
Get Help	Click to access help resources (see 11.3).
About	Click to view information such as Vantage version and links to the VDURA website, software terms and conditions, and third-party licenses (see 11.4).
UI Mode	Click to toggle light or dark UI mode (see 11.2).
Switch to Legacy UI	Click to switch to the PanActive Manager legacy UI (see 11.5).

Navigating to a new top-level page collapses the sidebar back down to icons; expand it again as needed.




Figure 2.2. Full left sidebar expanded with text labels, including the greyed-out Hardware item and the open Identity & Access menu.

3. Common UI Features in Vantage

Two sets of features recur across nearly every page.

3.1 Table Toolbar

Every data table in Vantage (Alerts, Events, Audit Log, Volumes, Storage Nodes, etc.) shares the same toolbar, typically in the top-right corner of the panel:

Icon	Function
Funnel	Opens a filter builder: pick a Column , an Operator (e.g. "is"), and a Value , with a "+ Add filter" link to add more conditions and a "Remove all" link to clear them.
Three vertical bars	Opens a column-visibility picker: a checkbox per column, plus "Show/Hide All" and "Reset" controls.
Magnifying glass	Expands into an inline free-text search box.

Pagination sits at the bottom of each table: a **Rows per page** dropdown, an `x-y of z` counter, and previous/next arrows.

3.2 Status Icons

The same four status icons recur on every summary tile and status badge:

Icon	Meaning
 green check	Healthy/active/online
 red circle	Critical (something is broken)
 amber triangle	Warning
 gray clock	Pending/not yet provisioned/no storage assigned

An icon is gray while its count is zero and takes its color when something is in that state. For example, a realm with two failed Director Nodes shows a red circle with the count 2 on the Director Nodes tile, and a healthy realm shows it gray with 0.

3.3 Required Fields and Tooltips

Required form fields are marked with a trailing gold asterisk (*). Fields with extra context have a small (?) icon next to their label; clicking it opens a short tooltip explanation (these do not appear on hover in this build; you must click them).

3.4 Confirm Before Destroy

Every destructive action in Vantage (deleting a StorageSet, a volume, a user, or a service account) requires a confirmation modal that restates the action and explicitly warns it cannot be undone. None of them require re-typing the resource's name to confirm; a single click on the modal's own confirm button executes the action immediately, so don't treat the modal as a safety net against a mis-click.

3.5 Page Not Found

If you open an address that does not match a Vantage page, for example a mistyped URL or an old bookmark, Vantage shows a "404" page reading "Page not found" with the message "The page you're looking for doesn't exist or has been moved." Click **Back to Dashboard** to return to the Dashboard, then navigate with the sidebar.




Figure 3.1. The Page not found screen.

4. Dashboard

The Dashboard is Vantage's landing page and the first thing a user sees upon logging in. It gives a realm-wide, at-a-glance health summary and is laid out as a grid of cards.




Figure 4.1. The whole Dashboard page, with all four card rows visible.

4.1 Node/StorageSet/Volume Summary Cards

The top row has three cards:

Card	Contents	(?) Tooltip
Director Nodes/Storage Nodes	Combined card with two counts side by side, each broken into 4 status buckets (see below)	"Status overview of director and storage nodes in the realm. VDPs are counted separately on the Realm Nodes page."
StorageSets	Total count + 4 status buckets	"Status overview of storage sets in the realm."
Volumes	Total count + 4 status buckets	"Status overview of volumes in the realm."

Each count row uses four consistent status icons, used throughout Vantage (see also 3.2):

Icon	Meaning
✓ green check	Healthy/active/online
⊖ red circle	Critical (something is broken)
⚠ amber triangle	Warning
🕒 gray clock	Pending/not yet provisioned/no storage assigned

Freshly initialized realm example: Before any storage is provisioned, the StorageSets card shows a single pending StorageSet (clock icon, not a green check), Volumes reads **0**, and Total Capacity shows the realm’s full raw capacity as Unassigned.




Figure 4.2. Dashboard on a freshly initialized realm: StorageSets pending, Volumes 0, capacity fully Unassigned.

This example also shows what alerts look like; two of the three Director Nodes are **Critical** (red), the top bar shows a red alert badge with a count of 3, and the Alerts panel lists three critical alerts (for example “PanFS root filesystem is not mounted” and “director node not in the realm-manager ReplicaSet”). See 3.2 for the status icons.

4.2 Total Capacity

A card showing realm-wide storage capacity as a horizontal bar plus a legend:

Legend item	Meaning
Used	Occupied space on StorageSets
Available	Free space on StorageSets
Unassigned	Capacity on Storage Nodes that are not part of any StorageSet yet, still usable for a new StorageSet or to grow an existing one

Below the bar, each StorageSet gets its own row with a mini progress bar and a `used / total` value (e.g. `Set 1: 5.60 MB / 898.99 GB`).




Figure 4.3. Dashboard with the Total Capacity (?) tooltip open.

4.3 Alerts

A card headed **Alerts** with a count, four filter chips (**All/Critical/Warning/Info**, each showing its own count), and an external-link icon that jumps to the Alerts table in Status Center. When there is nothing active it simply reads “No alerts” (plain text, no icon); this is distinct from the fuller Alerts table on the Status Center page itself, which shows a checkmark icon plus “No active alerts” when clear (see 5.1).

4.4 System Performance

A full-width card with a **Time Period** dropdown (Last 15 min [default], Last 1 hour , Last 6 hours , Last 24 hours) and a chart area. Its tooltip reads: "Disk throughput, IOPS, metadata capacity, and free capacity trends over the selected time period." With real workload history available, it renders four line charts stacked vertically. **Disk Throughput** (MB/s) and **Total IOPS** each have one line per active node/component (e.g. a Director Node and a VPOD), color-coded with a small legend in the chart's top-right corner. **Metadata Capacity** (GiB) and **Metadata Free Capacity** (GiB) follow, each plotted as a single line with no legend.



Figure 4.4. System Performance card showing real Disk Throughput and Total IOPS line charts.



Figure 4.5. The Metadata Capacity and Metadata Free Capacity charts at the bottom of the System Performance card.

Note: On a realm with no historical performance telemetry yet (e.g. immediately after a fresh reset, before any workload has run), this chart instead shows a perpetual loading spinner regardless of the selected time period. This is a property of having no data to plot, not a UI defect.

5. Status Center

Status Center repeats the same node/StorageSet/volume summary cards from the Dashboard and adds two full-width panels: **Alerts**, and a tabbed panel whose tabs are **Events Log** and **Audit Log**.




Figure 5.1. Full Status Center page: summary cards + Alerts panel + Events Log/Audit Log panel.

5.1 Alerts Panel

Column	Description
Severity	Icon-coded severity of the alert
Alert Message	Human-readable description
Component	The realm component the alert concerns
Alert Group	Logical grouping/category
Code	Machine-readable alert code

Additional columns available via the column picker but hidden by default: **Alert Rule**, **IP Address**.

Toolbar (top-right of the panel): severity filter chips (**All/Critical/Warning/Info**, each with a count), a filter-builder icon, a column-visibility icon, and a search icon. See 3.1 for how these work; the same pattern is reused on every table in the product.

Column headers add their own controls: an up/down arrow sorts the table by that column, ascending or descending, and a three-dot menu offers sort order, pin, filter, hide-column, and manage-all-columns options.

When nothing is active, the panel shows a checkmark and “No active alerts”.

5.2 Events Panel

The **Events Log** tab is the default view of the bottom panel.

Column	Description
Time (UTC)	Sortable; defaults to newest-first
Event Message	Human-readable description
Event Group	Logical grouping/category
Code	Machine-readable event code
Component	The realm component the event concerns

This panel uses the same toolbar pattern as Alerts (filter, column-visibility, and search icons). Once the realm has done anything (brought a volume online, enabled a backup service, etc.), this panel populates with real rows: a timestamp, a human-readable message (e.g. “Volume online: Volume /home has been brought online.”), an Event Group, a numeric Code, and the Component it concerns.

The screenshot displays the Vantage interface with two main panels. The top panel, titled 'Alerts', shows a single critical alert: 'PanFS root filesystem is not mounted' with the component 'VCH-1'. The bottom panel, titled 'EVENTS LOG', shows a list of events. The events are all 'Volume service backup enabled' messages for various components like VCH-1, VCH-6, and VCH-8. The events are sorted by time, with the most recent at the top.

Time (UTC) ↓	Event Message	Event Group	Code	Component
15:37:21	Volume service backup enabled: Service 0x0400000000000008(FM) now has backup Director 3e7e00003f4018001.	Volume	8654	VCH-1
15:37:20	Volume service backup enabled: Service 0x0400000000000001a(FM) now has backup Director 3e7e000023f4018001.	Volume	8654	VCH-6
15:37:20	Volume service backup enabled: Service 0x0400000000000001d(FM) now has backup Director 3e7e000023f4018001.	Volume	8654	VCH-6
15:37:20	Volume service backup enabled: Service 0x0400000000000001c(FM) now has backup Director 3e7e000023f4018001.	Volume	8654	VCH-6
15:37:19	Volume service backup enabled: Service 0x0400000000000001b(FM) now has backup Director 3e7e000023f4018001.	Volume	8654	VCH-6
15:37:19	Volume service backup enabled: Service 0x04000000000000021(FM) now has backup Director 3e7e000013f4018001.	Volume	8654	VCH-8
15:37:19	Volume service backup enabled: Service 0x04000000000000024(FM) now has backup Director 3e7e000013f4018001.	Volume	8654	VCH-8
15:37:19	Volume service backup enabled: Service 0x04000000000000023(FM) now has backup Director 3e7e000013f4018001.	Volume	8654	VCH-8
15:37:19	Volume service backup enabled: Service 0x04000000000000022(FM) now has backup Director 3e7e000013f4018001.	Volume	8654	VCH-8
15:37:19	Volume service backup enabled: Service 0x04000000000000005(FM) now has backup Director 3e7e000013f4018001.	Volume	8654	VCH-8

Figure 5.2. Alerts panel filtered to a Critical alert, with the Events Log tab beneath it populated with Volume-component rows.

Note: On a realm with no activity yet (e.g. immediately after a fresh reset), this panel instead shows a perpetual loading skeleton with no rows, even after waiting. Elsewhere in the product (the Tenant detail page, 9.2) the equivalent Events tab shows the same empty/loading behavior under the same no-activity-yet condition, while the Audit Log tabs and panels (5.3, and on the Tenant and User detail pages, 9.2 and 10.2) populate immediately since simply *viewing* a page is itself an audited action. This is a property of having nothing to show yet, not a broken control.

5.3 Audit Log Panel

The **Audit Log** tab, next to Events Log, lists the actions users and service accounts have taken in the realm, newest first. Because viewing a page is itself an audited action, this tab has rows as soon as anyone has signed in.




Figure 5.3. The Audit Log tab of the Status Center bottom panel.

Column	Description
Date & Time (UTC)	Sortable; defaults to newest-first
Actor	Who performed the action, e.g. "admin"
Action	e.g. "View", "List"
Resource	What the action was performed on, e.g. "Volume", "Tenant", or a linked ID
Source IP	The IP address the action originated from
Result	e.g. "Success"

This panel uses the same toolbar pattern as Alerts (filter, column-visibility, and search icons). It shows the same records as `panctl audit list`.

6. File System

The **File System** sidebar item expands to two sub-sections: **StorageSets** and **Volumes**.

6.1 StorageSets

6.1.1 StorageSets List Page

This page lists every StorageSet in the realm as its own card. The top-right contains a search icon and a **+ Create StorageSet** button.




Figure 6.1. StorageSets list page with a single active StorageSet card.

Below a summary strip (count + the same 4-icon status breakdown described in 4.1), each StorageSet card shows:

Card element	Description
Name	e.g. "Set 1"
Status badge (top-right)	Active (green) or No Storage Assigned (gray clock)
Provisioned	Percentage of capacity provisioned
VPODS	x/y online

Card element	Description
Storage Capacity	used / total with a mini progress bar

Clicking a card opens its detail page.

On a freshly initialized realm with no storage provisioned anywhere yet, this list shows a single empty shell and no active StorageSets.




Figure 6.2. StorageSets list on a freshly initialized realm: a single empty shell, no active StorageSets yet.

6.1.2 StorageSet Detail Page

This page shows full details for a single StorageSet. The breadcrumb trail `StorageSets > Set N` is on the left, with the StorageSet’s health status icon beside its name. The name is shown as an inline-editable title; see 6.2.3 to rename it. The top-right contains the **+ Add Storage** and **Delete** actions.




Figure 6.3. Full detail page for the active StorageSet, Volumes tab selected.

Left column:

Card	Field	Meaning
Summary Tiles	TOTAL	Total storage capacity of the StorageSet
Summary Tiles	AVAILABLE	Free space on the StorageSet (% + amount)
Summary Tiles	PROVISIONED	Percent of the StorageSet’s storage that has been provisioned (% + amount left)
Hardware Component	VPODs	Number of VPODs in the StorageSet that are online (x/y Online)

Card	Field	Meaning
Hardware Component	Spare VPODs	OSDs requested and available for reconstruction (Requested N / Available N), with an inline pencil-icon edit to increase this number: see 6.1.3

Right column:

Card	Field	Meaning
Storage Capacity	Used	Utilized capacity of the StorageSet (% of total, with an (i) info icon)
Storage Capacity	Reserved	Capacity reserved for use by reconstruction (% of total)
Storage Capacity	Available	Capacity ready to be consumed by growth in Used capacity, for any volume (% of total)
Storage Capacity	Total	The StorageSet's overall total capacity, shown in the legend alongside the three segments

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

Vantage 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.

Bottom table:

Tab	Column	Meaning
Volumes (n)	Component	Volume path (e.g. / , /home); links to volume detail

Tab	Column	Meaning
Volumes (n)	EC Protection Level	e.g. "Dual Parity (n+2)"
Volumes (n)	Capacity Used	Amount of capacity currently used by the volume
Volumes (n)	Available	Amount of capacity still available to the volume
Volumes (n)	Soft Quota	The soft quota for the volume, if set.
Volumes (n)	Hard Quota	The volume's hard quota, if set: the maximum capacity the volume can consume
Volumes (n)	Status	The volume's status (e.g. "Online")
Storage Nodes (n)	Name	Links to node detail under Realm Nodes
Storage Nodes (n)	IP	The node's IP address
Storage Nodes (n)	Capacity Total	Total capacity of this node
Storage Nodes (n)	Status	The node's health/operational status (see 3.2)




Figure 6.4. Same page with the Storage Nodes tab selected.

Empty StorageSet example: A StorageSet that exists but has no storage assigned yet shows status badge **No Storage Assigned**, VPODs **0/0 Online**, capacity **0 B / 0 B**, and both the Volumes and Storage Nodes tabs read “No rows”. This is the expected appearance of a StorageSet shell before any nodes have been added to it.




Figure 6.5. Detail page for the empty StorageSet ('No Storage Assigned', 0/0 VPODs).

6.1.3 Edit Spare VPODs Modal

Opened via the pencil icon next to Spare VPODs, this modal changes the number of spare VPODs reserved for reconstruction.




Figure 6.6. The Edit Spare VPODs for Set N modal open.

To open it, click the pencil icon on the Spare VPODs line in the Hardware Component field. Title: "Edit Spare VPODs for Set N". Body text: "Number of spare VPODs to reserve for reconstruction. This value can only be increased, not decreased." A numeric stepper labeled "OSD equivalent", then Cancel/Save.

6.1.4 Add Storage Modal

This modal adds storage to a StorageSet.




Figure 6.7. The Add Storage modal showing that there are no unassigned nodes available.

This modal has an **Additional spare OSDs** field at the top (described below) and, under it, a **Component** list: every unassigned Storage Node in the realm as a checkbox row, with Component and Status columns and its own search box. Select the nodes to add and confirm. If the realm has no unassigned nodes, it displays “No unassigned nodes available.” instead of the node picker.

When adding storage, the value in the **Additional spare OSDs** field is added to the StorageSet's existing spare reservation (originally set by **Space Reserved for Reconstruction** on the Create StorageSet page, see 6.1.6). Leaving it at zero keeps the current reservation unchanged; it cannot reduce it. The amount of space reserved for reconstruction determines the number of consecutive VPOD failures the StorageSet can withstand while still providing data redundancy.




Figure 6.8. The Add Storage modal's node picker populated with unassigned Storage Nodes.

Once confirmed, the StorageSet's status flips to **Active**, its VPODs come online, and Total Capacity updates accordingly. Expect a brief burst of critical/warning alerts in the top bar while this happens, clearing on its own within a minute or two; see [Troubleshooting](#) if it doesn't.

Note: Adding storage to a StorageSet automatically provisions the built-in platform tenant's `/` and `/home` volumes on it; no separate Create Volume step is needed for those two. **Create Volume** (see 6.3.4) is for anything beyond them: additional paths or any volume belonging to a tenant other than the built-in default.

6.1.5 Delete StorageSet Confirmation

Click Delete in the delete StorageSet confirmation prompt window to confirm and delete the StorageSet.




Figure 6.9. The delete confirmation modal for a StorageSet.

6.1.6 Create StorageSet Page

Use this page to create a new StorageSet.




Figure 6.10. The full Create StorageSet form.

Field	Notes
Name*	Required text field
Space Reserved for Reconstruction	Numeric stepper, "OSD equivalent", default 2 . Helper: "Number of spare OSDs to reserve for reconstruction. After the StorageSet is created, this value can only be increased, not decreased."
Multi-Level Erasure Coding	This feature provides a second layer of erasure coding and reconstruction inside VPODs. Note that this cannot be turned off once the StorageSet is created.
Buttons	Cancel/Create

6.2 Creating and Managing StorageSets

Use Vantage to create a new StorageSet when new storage devices are added to the system. After the StorageSet is created, you can still increase the number of spares (but never decrease it), either with the pencil icon on the StorageSet Detail page (see 6.1.3) or when you add storage (see 6.2.2). You can also increase the capacity of an existing StorageSet by adding new, unassigned components to it.

Note: For information on configuring and managing StorageSets using the command-line interface, see the `storageset` command in the *VDURA Data Platform CLI Reference Guide*.

6.2.1 Create a StorageSet

1. Click **+ Create StorageSet** on the StorageSets list page (see 6.1.1).
2. Enter a descriptive name for the StorageSet, up to 30 ASCII characters, in the **Name** field.
3. Enter the number of spare VPODs to reserve for reconstruction or use the default of two. The amount of space reserved for reconstruction determines the number of consecutive VPOD failures the StorageSet can withstand while still providing data redundancy, so reserve sufficient space now. After the StorageSet is created, this value can only be increased, not decreased, either with the pencil icon on the StorageSet Detail page (see 6.1.3) or when additional storage is added (see 6.2.2).
4. If desired, turn on **Multi-Level Erasure Coding** for a second layer of erasure coding and reconstruction inside VPODs. This cannot be turned off once the StorageSet is created.
5. Click **Create** (see 6.1.6). The new StorageSet appears on the StorageSets list page.

6.2.2 Add Storage to a StorageSet

1. Click the StorageSet's card on the StorageSets list page (see [6.1.1](#)) to open its detail page. A newly created StorageSet has no storage assigned to it until storage is added.
2. Click **+ Add Storage** (see [6.1.4](#)).
3. Select the available storage to add from the **Component** list.
4. Enter the number of spare OSDs to add to the StorageSet's reservation in the **Additional spare OSDs** field. Leave it at zero to keep the current reservation. As at creation, the total reservation determines how many consecutive VPOD failures the StorageSet can withstand while still providing data redundancy.
5. Click **Add Storage** to return to the StorageSet Detail page.

6.2.3 Modify a StorageSet

See [6.1.2](#) for the StorageSet Detail page.

- To rename it: click the StorageSet's name, enter the new name, and click **Save**.
- To change the spare VPOD reservation: click the pencil icon on the **Spare VPODs** line of the **Hardware Component** card (see [6.1.3](#)).

6.2.4 Delete a StorageSet

1. Open the StorageSet's detail page (see [6.1.2](#)).
2. Click **Delete** (see [6.1.5](#)).
3. Click **Delete** again in the confirmation prompt.

6.3 Volumes

6.3.1 Volumes List Page

This page lists every volume in the realm. The top-right contains a **+ Create Volume** button; the table toolbar has filter, column-visibility, and search icons (see 3.1).


Component	Tenant Name	StorageSet	EC Protection Level	Capacity Used	Available	Soft Quota	Hard Quota	Status
/	platform	Set1	Dual Parity (n+2)	94.20 KB	524.19 MB	524.28 MB	524.28 MB	Online
/home	platform	Set1	Dual Parity (n+2)	77.82 KB	898.98 GB	0 B	0 B	Online

Figure 6.11. The Volumes table with at least two rows.

Column	Description
Component	volume path (e.g. <code>/</code> , <code>/home</code>); links to volume detail
Tenant Name	e.g. "platform"
StorageSet	Links to the owning StorageSet
EC Protection Level	e.g. "Dual Parity (n+2)"
Capacity Used	Amount of capacity currently used by the volume
Available	Amount of capacity still available to the volume
Soft Quota	The soft quota for the volume, if set.

Column	Description
Hard Quota	The volume's hard quota, if set. This is the maximum capacity the volume can consume.
Status	The volume's status (e.g. "Online")

In a realm with more than one tenant, this list shows every tenant's volumes side by side, distinguished by the Tenant Name column.



Component	Tenant Name	StorageSet	EC Protection Level	Capacity Used	Available	Soft Quota	Hard Quota	Status
/	platform	Set 1	Dual Parity (n+2)	159.74 KB	524.12 MB	524.28 MB	524.28 MB	Online
/home	platform	Set 1	Dual Parity (n+2)	77.82 KB	898.98 GB	0 B	0 B	Online
/projects	tenant1	Set 1	Dual Parity (n+2)	77.82 KB	99.99 GB	0 B	100.00 GB	Online

Rows per page: 10 1-3 of 3

Figure 6.12. The realm-wide Volumes list showing volumes from more than one tenant side by side.

6.3.2 Volume Detail Page

This page shows full details for a single volume. The header shows the path and status badge, and the top-right contains a  **Delete** action.




Figure 6.13. Full detail page for one volume.

Card	Field	Meaning
StorageSet	Name	Link to the owning StorageSet
StorageSet	Status	Health status of the owning StorageSet
StorageSet	Provisioned	Percent of the StorageSet's capacity provisioned (% + remaining)
Tenant Info	Tenant ID	The owning tenant's UUID
Tenant Info	Tenant Name	The owning tenant's name
Tenant Info	Status	The owning tenant's status (e.g. "Active")

Card	Field	Meaning
Service Ownership	Volume Service	Which Director's metadata service manages this volume. Provides some control over volume-management load balancing across Directors. Each Director runs a metadata service that manages one or more volumes, and responsibility for a given volume is not spread across multiple Directors. On a realm with several Directors, defining a like number of volumes helps ensure all Directors are assigned some of the metadata workload, which improves performance (e.g. "Automatic").
Service Ownership	Status	Status of the volume's service assignment (e.g. "Online")
Service Ownership	Recovery Priority	Any integer from 1 to 100; a lower number is a higher priority, with 1 the highest. Defaults to 50 if left unchanged (e.g. "1 (High)").

Card	Field	Meaning
Service Ownership	EFSA Mode	Extended File System Availability Mode: the path pursued when an object has exceeded its fault-tolerance level. <code>retry</code> blocks and retries I/O requests until the needed storage devices are back online; <code>file-unavailable</code> returns I/O errors instead. The system default is <code>retry</code> (e.g. "Retry I/O Requests").
Erasure Coding	Protection Level	The volume's erasure-coding scheme: Dual Parity (n+2) , a balance between capacity overhead and performance (the default); Single Parity (n+1) , optimized for performance with less capacity overhead; or Striped Mirror (2x) , optimized for small write performance.
Right Column	SOFT QUOTA	The volume's soft quota, if set.
Right Column	HARD QUOTA	The volume's hard quota, if set: the maximum capacity the volume can consume.
Right Column	Available Capacity	A card with a bar %, a legend of Used, Other Volumes & Data, and Available, the StorageSet Capacity total beneath it, and a help (?) icon.

6.3.3 Delete Volume Confirmation

This modal confirms deletion of a volume.




Figure 6.14. The delete confirmation modal.

6.3.4 Create Volume Page

This page is used for any volume beyond the platform tenant's automatically-provisioned `/` and `/home` (see 6.1.4) or for any volume belonging to another tenant.




Figure 6.15. Top of the Create Volume form: Pathname through Erasure Coding Protection Level radios.

Field	Notes
Pathname*	Text, e.g. /data . The leading / shown in the field is fixed UI chrome, not something to type yourself. A value that already includes / is rejected (“Only letters, numbers, dashes, and underscores are allowed”).
StorageSet Placement*	Required dropdown. Lists every StorageSet with its available capacity, e.g. Set 1: 898.47 GB , Set 2: 0.00 GB (a StorageSet with no assigned storage is still listed, just with 0 GB available). Shows “Storage set is required” if left empty.
Volume Service Placement	Controls which Director’s metadata service manages this volume, giving some control over volume-management load balancing across Directors. Each Director runs a metadata service that manages one or more volumes, and responsibility for a given volume is not spread across multiple Directors. On a realm with several Directors, defining a like number of volumes helps ensure all Directors are assigned some of the metadata workload, which improves performance. Dropdown, default “Automatic Selection by System”; a specific Director can be selected manually instead.
Hard Quota	Numeric + GB unit. Optional when acting as Platform Admin against a tenant with an Unlimited StorageSet quota (the platform tenant, see 9.2); required (with a “Max” quick-fill link next to the field) once you’re acting inside a tenant that has a finite Hard Quota on the StorageSet, see 9.4.
Erasure Coding Protection Level	Radio buttons (see table below).

Erasure Coding Protection Level options:

Option	Description
Dual Parity (n+2)	Balance between capacity overhead and performance (<i>Default</i>)
Single Parity (n+1)	Optimized for performance with less capacity overhead
Striped Mirror (2x)	Optimized for small write performance

Two collapsible sections follow:

Advanced Erasure Coding Configuration

Field	Notes
Stripe Unit	Numeric + unit dropdown (KB/MB).
Stripe Width	Integer.

User & Group Permissions

User and group can be specified as either a UNIX name, UNIX ID, or Windows SID (e.g. `0`).

Row	Name field (default)	Permission dropdown (default)
User	<code>root</code>	<code>all</code>
Group	<code>wheel</code>	<code>read-execute</code>
Others	<i>(none, implicit)</i>	<code>read-execute</code>

Each of the fields above contains the following permission dropdown options: `none` , `read-only` , `write-only` , `execute-only` , `read-write` , `read-execute` , `write-execute` , `all` .

Cancel/Create Volume.

The screenshot shows the 'Create Volume' interface. The 'StorageSet Placement' dropdown is open, displaying two options: 'Set 1' with 124796 GB and 'Set 2' with 0.00 GB. The form includes fields for Pathname, Hard Quota, Erasure Coding Protection Level, and expandable sections for Advanced Erasure Coding Configuration and User & Group Permissions.

Figure 6.16. The StorageSet Placement dropdown open, showing each StorageSet with its available capacity.

The screenshot shows the 'Create Volume' interface with the 'Advanced Erasure Coding Configuration' section expanded. It includes fields for Stripe Unit (set to KB) and Stripe Width. The 'User & Group Permissions' section is also expanded, showing fields for User (root) and Group (wheel).

Figure 6.17. Advanced Erasure Coding Configuration section expanded (Stripe Unit/Stripe Width).




Figure 6.18. User & Group Permissions section expanded, with a permission dropdown open.

6.4 Creating and Managing Volumes

Use Vantage to create a new volume for any path beyond the platform tenant’s automatically-provisioned `/` and `/home` or for any volume belonging to another tenant. This build does not provide an in-place way to edit a volume’s settings after creation.

Note: For information on configuring and managing volumes using the command-line interface, see the `volume` command in the *VDURA Data Platform CLI Reference Guide*.

6.4.1 Create a Volume

This procedure uses the Create Volume page; 6.3.4 describes every field on that page in detail.

1. Click **+ Create Volume** on the Volumes list page (see 6.3.1).
2. Enter the volume’s path in the **Pathname** field. The leading `/` shown in the field is fixed UI chrome; type only the rest of the path.
3. Select the **StorageSet Placement**. Every StorageSet in the realm is listed with its available capacity, even one with none.
4. If desired, choose a specific **Volume Service Placement** instead of the default automatic selection; this cannot be changed after the volume is created.
5. Enter a **Hard Quota** if required (mandatory when acting inside a tenant with a finite StorageSet quota; optional otherwise). Use the **Max** link to fill in the tenant’s full remaining allowance in one click.
6. Choose an **Erasure Coding Protection Level**; this cannot be changed after the volume is created.

7. Optionally expand **Advanced Erasure Coding Configuration** to set a custom Stripe Unit/Stripe Width, and **User & Group Permissions** to set ownership and access beyond the defaults.
8. Click **Create Volume** (see [6.3.4](#)).

6.4.2 Delete a Volume

1. Open the volume's detail page (see [6.3.2](#)).
2. Click  **Delete** (see [6.3.3](#)).
3. Click **Yes, Delete Volume** in the confirmation prompt.

7. Realm Nodes

This page has two summary tiles (**Director Nodes**, **Storage Nodes**) with the same 4-icon breakdown as elsewhere, followed by two tables.

Director Nodes					
Name	Status	IP	Model	Hardware Serial	Data Serial
VCH-1	Online	10.128.1.0	VCH-5000-D1N	3e7e00003f4018001	7e71163e71c90d
VCH-6	Online	10.128.1.2	VCH-5000-D1N	3e7e00023f4018001	c0b43f3e71c91d
VCH-8	Online	10.128.1.1	VCH-5000-D1N	3e7e00013f4018001	b0e9ab3e71c910

Storage Nodes									
Name	Type	Status	IP	Model	StorageSet	Capacity Used	Capacity Total	Hardware Serial	Data Serial
VCH-2	VPOD	Online	10.128.2.1	VCH-5000-S1N	Set 1	1.14 MB	224.74 GB	3e7e0001d40340...	ff80cc3e71c915
VCH-3	VPOD	Online	10.128.2.2	VCH-5000-S1N	Set 1	1.14 MB	224.74 GB	3e7e0002d40340...	6c61c03e71c90b
VCH-4	VPOD	Online	10.128.2.3	VCH-5000-S1N	Set 1	1.14 MB	224.74 GB	3e7e0003d40340...	67914e3e71c90f
VCH-5	VPOD	Online	10.128.2.5	VCH-5000-S1N	Set 1	1.14 MB	224.74 GB	3e7e0005d40340...	f27f183e71c8d0

Figure 7.1. Realm Nodes page displaying both the Director Nodes and Storage Nodes tables.

7.1 Director Nodes Table

Column	Description
Name	Links to node detail
Status	The node's health/operational status (see 3.2)
IP	The node's IP address
Model	e.g. "VCH-5000-D1N"
Hardware Serial	The node's hardware serial number
Data Serial	Matches the <code>dataserial</code> identifier used by <code>panctl</code> nodes

7.2 Storage Nodes Table

Column	Description
Name	Links to node detail
Type	e.g. "VPOD"
Status	The node's health/operational status (see 3.2)
IP	The node's IP address
Model	e.g. "VCH-5000-S1N"
StorageSet	Links to the owning StorageSet
Capacity Used	Amount of capacity currently used on this node
Capacity Total	Total capacity of this node
Hardware Serial	The node's hardware serial number
Data Serial	Matches the <code>dataserial</code> identifier used by <code>panctl</code> nodes

On a freshly initialized realm, every Storage Node shows `-` in the StorageSet column until storage is assigned to a StorageSet. This is a reliable way to confirm a node is genuinely unassigned, rather than just showing 0 capacity.

The screenshot shows the Vantage Realm Nodes interface. At the top, there are summary counts: Director Nodes 3 (1 Online, 2 Critical) and Storage Nodes 6 (all Online). The interface is divided into two main sections: Director Nodes and Storage Nodes.

Director Nodes						
Name	Status	IP	Model	Hardware Serial	Data Serial	
VCH-1J	Online	10.128.1.0	VCH-5000-DIN	3e7e0003f4018001	7e71163e7126ac	
VCH-5J	Critical	10.128.1.1	VCH-5000-DIN	3e7e00013f4018001	b0e9ab3e7126a5	
VCH-8J	Critical	10.128.1.2	VCH-5000-DIN	3e7e00023f4018001	c0b43f3e7126df	

Storage Nodes									
Name	Type	Status	IP	Model	StorageSet	Capacity Used	Capacity Total	Hardware Serial	Data Serial
VCH-2J	VPOD	Online	10.128.2.1	VCH-5000-SIN	-	0 B	-	3e7e00011d40340...	ff80cc3e71272a
VCH-3J	VPOD	Online	10.128.2.4	VCH-5000-SIN	-	0 B	-	3e7e00041d4034...	0b703d3e71270e
VCH-4J	VPOD	Online	10.128.2.0	VCH-5000-SIN	-	0 B	-	3e7e00001d4034...	7d519f3e712705
VCH-6J	VPOD	Online	10.128.2.5	VCH-5000-SIN	-	0 B	-	3e7e00051d4034...	f27f183e7126ce

Figure 7.2. Realm Nodes' Storage Nodes table with every StorageSet column showing a dash.

This example also shows what alerts look like; two of the three Director Nodes show a `Critical` status (see 3.2), while the Storage Nodes are all `Online`.

7.3 Node Detail Page

This page shows full details for a single node. The breadcrumb trail `Realm Nodes > <name>` is on the left. Content differs by node type: Director nodes, Storage Nodes (VPODs) and, on physical hardware, VDP nodes.




Figure 7.3. Detail page for a Director Node (Identity/System Details/Alerts/NIC Info, with Drive Assets below).

Director Node:

Card	Fields
Identity	Type (<code>Director</code>), IP Address, Activation Status (whether the node is configured to be part of the realm’s replication set), Data Serial
System Details	Model, Hardware Serial, CPU, Memory. On physical hardware, an arrow in the card’s corner (Open hardware reference) opens that node’s chassis detail page (8.2).
Alerts (<i>right column</i>)	Expand icon jumps to Status Center

Card	Fields
NIC Info (<i>right column</i>)	Summary line (e.g. "eth0 · Physical · 1× 10GE, 1× Network Interface 1, 1× Network Interface 2, 1× Network Interface 3") plus a table: Interface, Type, Mode, Status. A primary physical interface like <code>eth0</code> typically shows Up ; unused bonded interfaces (<code>bond0</code> / <code>bond1</code> / <code>bond2</code>) show Unknown until configured.
Drive Assets (<i>below the cards</i>)	A table of the node's own drives with the columns Name, Status, Type (SSD or HDD), Model, Vendor, Serial, FW Version, Drive Number and Node Name. The drive count appears beside the card title.

Storage Node (VPOD):

Card	Fields
Identity	Type (<code>VPOD</code>), IP Address, StorageSet Assigned (link), Activation Status (whether the node is configured and available to be used as part of a StorageSet), Data Serial
Alerts (<i>right column</i>)	Expand icon jumps to Status Center
Drives Storage (<i>right column</i>)	Capacity bar with Used/Available legend
Drive Assets (<i>below the cards</i>)	A table of the node's drives with the columns Name, Status, Type, Model, Vendor, Serial, FW Version, Drive Number, Bytes Read, Bytes Written and Node Name. The drive count appears beside the card title.




Figure 7.4. Detail page for a Storage Node/VPOD (Identity/Alerts/Drives Storage/Drive Assets).

Note: On physical hardware, a VPOD's Identity card shows an additional **Parent** field, and the breadcrumb grows a level to match. See [8.3](#).

VDP Node (physical hardware):

Each Storage chassis also runs one VDP node, the chassis's main node. It is reached from the chassis's Linked Node card ([8.2](#)) or from the Parent field of one of its VPODs ([8.3](#)).

Card	Fields
Identity	Type (VDP), IP Address, StorageSet Assigned (link), Activation Status, Data Serial
System Details	Model, Hardware Serial, CPU, Memory, with the same Open hardware reference arrow as a Director node
Alerts (<i>right column</i>)	Expand icon jumps to Status Center
VPODs Storage (<i>right column</i>)	Capacity bar with Used/Available legend for the chassis's VPODs
NIC Info (<i>right column</i>)	Summary line plus a table: Interface, Type, Mode, Status

Card	Fields
VPODs (<i>below the cards</i>)	A table of the chassis's VPOD nodes with the columns Name, Status, Used Capacity and Total Capacity
Drive Assets (<i>below the cards</i>)	A table of the VDP node's own drives with the columns Name, Status, Type, Model, Vendor, Serial, FW Version, Drive Number and Node Name

The screenshot displays the Vantage UI for a VDP Node named VCH-3,1. The interface is divided into several sections:

- Identity:**
 - Type: VDP
 - IP Address: 192.168.200.113
 - StorageSet Assigned: Set 1
 - Activation Status: Activated
 - Data Serial: f24ba45271af6b
- System Details:**
 - Model: VCH-5050-SIN
 - Hardware Serial: 527e0003e4170001
 - CPU: AMD EPYC 9565 72-Core Processor
 - Memory: 384GB
- Alerts:** 4 alerts listed, including "Node NIC 4 not active: 10 GE Network Interface 1 is not active." and "Node NIC 5 not active: 10 GE Network Interface 2 is not active." (1-2 of 4).
- VPODs Storage:** A progress bar showing 100.00% usage. Used: 8.55 GB, Available: 1.32 PB.
- NIC Info:** A table showing network interface details.

Interface	Type	Mode	Status
bond0	Failover - 1x 200 GIB Network Interface 1, 1x 200 GIB Network Interface 2	Failover	Up
pri2	200 GIB Network Interface	Failover	Up
pri3	200 GIB Network Interface	Failover	Up

(1-3 of 9)

Figure 7.5. Detail page for a VDP Node (Identity/System Details/Alerts/VPODs Storage/NIC Info, with the VPODs and Drive Assets tables below).

8. Hardware

On physical VDURA hardware, the **Hardware** sidebar entry is active and opens a list of the realm's chassis. In a virtual realm, this same entry is greyed out and non-interactive instead; see 8.4 for that state.

8.1 Hardware List Page

This page lists every chassis in the realm in a single table.

Column	Description
Name	The chassis model name, e.g. <code>VCH-5050-S1N</code> (Storage) or <code>VCH-5100-D1N</code> (Director). Links to chassis detail
Product Type	The kind of chassis, e.g. "VCH-5050 Storage Node" or "VCH-5100 Director Node"
Status	The chassis's health/operational status (see 3.2)
Identify Hardware	A Blink LED button that blinks a physical LED on that chassis, to help locate it in a rack

Name	Product Type	Status	Identify Hardware
VCH-5050-S1N	VCH-5050 Storage Node	Online	Blink LED
VCH-5050-S1N	VCH-5050 Storage Node	Online	Blink LED
VCH-5050-S1N	VCH-5050 Storage Node	Online	Blink LED
VCH-5100-D1N	VCH-5100 Director Node	Online	Blink LED
VCH-5100-D1N	VCH-5100 Director Node	Online	Blink LED
VCH-5100-D1N	VCH-5100 Director Node	Online	Blink LED

Figure 8.1. Hardware list showing Storage and Director chassis rows.

Above the table are filter, column and search icons; below it are a **Rows per page** control and a page counter.

8.2 Chassis Detail Page

Clicking a chassis **Name** opens its detail page. From top to bottom it shows two cards (Specifications and Linked Node), a row of three temperature cards, a NIC Info card, and a tabbed table.

Card	Fields
Specifications	System Vendor, Variant, System Model, Product Type, Hardware Serial, CPU, Cores/Threads, Memory, BIOS Version, BIOS Date
Linked Node	Role (e.g. Director or VDP), Node (links to that node's Realm Node detail page, 7.3), StorageSet (a link; shown only when the node belongs to a StorageSet, so never on a Director chassis and not on a Storage chassis whose node is Not Activated), Node Status, Activation Status (Activated or Not Activated)
Internal Temp, CPU Temp, Incoming Air Temp	Three cards in a row below the Specifications and Linked Node cards. Each shows one temperature in both Celsius and Fahrenheit.
NIC Info	A summary line naming the primary bond (e.g. its mode and member interfaces), plus a table with the columns Interface, Type, Mode and Status. Rows include bonds, physical interfaces and unused interfaces; Status shows Up, Disabled or Unknown. If no interfaces are detected, the card instead shows a warning titled "No Network Interfaces detected", with the text "This usually indicates a configuration issue; a Storage Node normally has network interfaces."

Below the cards, three tabs (each with a count badge):

Tab	Contents
Base Elements	The individual node board(s) inside this chassis: Base Elements name, Identify Hardware (Blink LED), Status, Software Version, SKU, IP, Model, Serial. A Director chassis shows one row. A Storage chassis shows its VDP node plus one row per VPOD; a Storage chassis whose node is not yet activated shows only the VDP row. Clicking a Base Elements name opens that node's existing Realm Node detail page (7.3) rather than a separate Hardware-specific page.
Power Supplies	The chassis's power supply units, when present
Drive Assets	Every drive in the chassis, one row each, with the columns Name, Identify (a Blink LED button per drive), Status, Type (SSD or HDD), Model, Vendor, Serial, FW Version, Drive Number, Bytes Read, Bytes Written and Node Name (the base element the drive belongs to). The Node Name, Drive Number, Name and Serial headers carry sort-order markers (1 to 4). The table has filter, column and search icons and shows 25 rows per page.




Figure 8.2. Storage chassis detail page (VCH-5050-S1N), top of the page: Specifications, Linked Node, the three temperature cards and the start of the NIC Info table.




Figure 8.3. Director chassis detail page (VCH-5100-D1N), top of the page. A Director's Linked Node card has no StorageSet row.

Interface	Type	Mode	Status
pri2	200 GIB Network Interface 1	Failover	Up
pri3	200 GIB Network Interface 2	Failover	Up
bond1	LAG	Disabled	Disabled
bond2	LAG	Disabled	Disabled
enp163s0f0np0	10GE	Disabled	Disabled
enp163s0f1np1	10GE	Disabled	Disabled
enp164s0	1GE	Disabled	Disabled
enp165s0	1GE	Disabled	Disabled

Base Elements	Identify Hardware	Status	Software Version	SKU	IP	Model	Serial
VCH-3.3	Blink LED	Online	12.0.0-1903681.2	VCH-5050-SIN	192.168.200.107	VCH-5050-SIN	6fbfe26a1d4170001
VCH-3.12	Blink LED	Online	12.0.0-1903681.2	VCH-5050-SIN	192.168.200.96	VCH-5050-SIN	29a790e41d4170002
VCH-3.13	Blink LED	Online	12.0.0-1903681.2	VCH-5050-SIN	192.168.200.97	VCH-5050-SIN	36cb8ce91d4170002
VCH-3.9	Blink LED	Online	12.0.0-1903681.2	VCH-5050-SIN	192.168.200.99	VCH-5050-SIN	eb7bf29d1d4170002
VCH-3.5	Blink LED	Online	12.0.0-1903681.2	VCH-5050-SIN	192.168.200.109	VCH-5050-SIN	0096c0991d4170001
VCH-3.4	Blink LED	Online	12.0.0-1903681.2	VCH-5050-SIN	192.168.200.98	VCH-5050-SIN	d9c3990d1d4170001
VCH-3.10	Blink LED	Online	12.0.0-1903681.2	VCH-5050-SIN	192.168.200.94	VCH-5050-SIN	3e507f2d1d4170002

Figure 8.4. Storage chassis detail page scrolled down to the tabs, with the Base Elements tab open below the NIC Info table.

When a tab has nothing to show, it displays a plain text message instead of an empty table, e.g. “No power supplies to show.” / “No drive assets to show.”

Interface	Type	Mode	Status
bond0	LAG	Failover	Up
pri2	200 GIB Network Interface 1	Failover	Up
pri3	200 GIB Network Interface 2	Failover	Up
bond1	LAG	Disabled	Disabled
bond2	LAG	Disabled	Disabled
enp163s0f0np0	10GE	Disabled	Disabled
enp163s0f1np1	10GE	Disabled	Disabled
enp164s0	1GE	Disabled	Disabled
enp165s0	1GE	Disabled	Disabled

No power supplies to show.

Figure 8.5. Power Supplies tab in its empty state, below the NIC Info table.

Name	Identify	Status	Type	Model	Vendor	Serial	FW Version	Drive Number	Bytes Read	Bytes Written	Node Name
pn0n1	Blink LED	Good	SSD	TS1TMTE712P-VS1	Transcend	J090930015	82B7X5LD	1	—	—	VCH-3,1
sd1	Blink LED	Good	HDD	ST18000NM0012	Seagate	ZVV07QQ5	E002	2	—	—	VCH-3,1
sd2	Blink LED	Good	HDD	ST18000NM0012	Seagate	ZVV07EHW	E002	3	—	—	VCH-3,1
sd3	Blink LED	Good	HDD	ST18000NM0012	Seagate	ZVV06EH5	E002	4	—	—	VCH-3,1
sd4	Blink LED	Good	HDD	ST18000NM0012	Seagate	ZVV07N6G	E002	5	—	—	VCH-3,1
sd5	Blink LED	Good	HDD	ST18000NM0012	Seagate	ZVV07YPC	E002	6	—	—	VCH-3,1
sd6	Blink LED	Good	HDD	ST18000NM0012	Seagate	ZVV07DFR	PAN2	7	—	—	VCH-3,1
sd7	Blink LED	Good	HDD	ST18000NM0012	Seagate	ZVV067ZM	E002	8	—	—	VCH-3,1
sd8	Blink LED	Good	HDD	ST18000NM0012	Seagate	ZVV07WSM	E002	9	—	—	VCH-3,1
sd10	Blink LED	Good	HDD	ST18000NM0012	Seagate	ZVV076CB	E002	11	—	—	VCH-3,1
sd11	Blink LED	Good	HDD	ST18000NM0012	Seagate	ZVV06DSD	E002	12	—	—	VCH-3,1

Figure 8.6. Drive Assets tab of a Storage chassis, showing the table’s columns and sort-order markers.

8.3 Parent Field on Realm Node Detail

On physical hardware, a Storage Node (VPOD)’s Realm Node detail page (7.3) shows an additional **Parent** field in the Identity card, directly after Type. It links to the VDP node that runs on the same chassis, and the breadcrumb trail grows a level to match, e.g. `Realm Nodes > VCH-3,1 > VCH-3,3`. Director nodes have no Parent field and keep the plain `Realm Nodes > <name>` breadcrumb. Neither appears on a virtual realm, since there’s no physical chassis.

To get from a node back to its chassis, use the **Open hardware reference** arrow in the System Details card on the Director or VDP node's page; it opens that chassis's detail page (see 8.2).




Figure 8.7. Realm Node detail page showing the Parent field and chassis-level breadcrumb.

8.4 Hardware on Virtual Realms

In a virtual realm, the **Hardware** sidebar entry is present but greyed out and non-interactive. Hovering it shows the tooltip **"Hardware is unavailable for virtual realms"**, and clicking it does nothing (no navigation occurs).




Figure 8.8. Sidebar expanded with the Hardware item's disabled tooltip visible.

9. Tenants

Note: Multi-tenancy is not supported on InfiniBand. Tenants require an Ethernet network, because tenant isolation is carried by VLAN or VXLAN encapsulation and InfiniBand cannot carry either one. A realm whose clients connect over InfiniBand cannot have tenants created on it. Realms that run only the platform tenant are not affected.

9.1 Tenants List Page

This page lists every tenant in the realm as its own card. The top-right contains a search icon and a **Create Tenant** button.




Figure 9.1. The Tenants list with at least one tenant card.

Each card shows:

Card element	Description
Name	e.g. "PanAPI Platform Tenant (platform)"
Status badge	e.g. <code>Active</code> or <code>Setup Incomplete</code> for a tenant that does not yet have everything it needs to serve data (see 9.2.2)

Card element	Description
Subtext	A capacity summary: one line per StorageSet the tenant can use, showing the StorageSet name, a mini progress bar, and the tenant's capacity used / Hard Quota (e.g. "Set 1 172.03 KB / Unlimited"). A tenant with no assigned StorageSets reads "No StorageSet assignments" instead.

Note: The built-in platform tenant has implicit access to every StorageSet in the realm, so its card shows an "Unlimited" quota against each one (see 9.2). Any other tenant lists only the StorageSets it has been explicitly assigned.




Figure 9.2. The Tenants list with the platform tenant and two newly created tenants.

In this example, one new tenant has been assigned a StorageSet and shows its usage against the 100.00 GB quota, while the other has no StorageSets yet and reads "No StorageSet assignments". Both show **Setup Incomplete**.

9.2 Tenant Detail Page

This page shows full details for a single tenant. The tenant's name is shown as an inline-editable title, and the top-right contains a **Delete Tenant** action (not shown for the built-in platform tenant, which cannot be deleted). A tenant you created also shows a status chip beside its name (**Setup Incomplete**

until setup is finished, see 9.2.2) and two tabs: **Overview**, described here, and **Networking** (see 9.2.3).




Figure 9.3. Info/Description/StorageSets/Volumes cards.

Card	Fields
Info	Short Name, Contact Email
Description	Free-text description. The built-in platform tenant shows a fixed "Built-in platform tenant"; a tenant you created yourself instead shows an inline-editable placeholder, "No description. Click to add."

StorageSets table:

Column	Description
Name	Link to the StorageSet
Tenant Capacity Used	Value + mini bar
Tenant Hard Quota	e.g. "Unlimited" (platform tenant) or a specific size (any tenant you explicitly assigned a quota to, e.g. "100.00 GB")

Column	Description
Total Available on StorageSet	The space still available on the StorageSet, not the tenant's quota
Status	The StorageSet's status (e.g. "Active")
Actions	Only present once a StorageSet has been explicitly assigned to the tenant (see below): pencil (edit quota) and trash (remove assignment) icons

The built-in platform tenant gets implicit, unlimited access to every StorageSet in the realm with no assignment step required. Any other tenant starts out with **"No StorageSets assigned"** and must be given access explicitly before it can place volumes on a StorageSet.




Figure 9.4. A brand-new tenant's detail page: 'No StorageSets assigned' + the + Add StorageSet button.

This card’s header carries a **+ Add StorageSet** button for any tenant other than the built-in platform one, opening an “Add StorageSet” modal. The button is greyed out when every StorageSet in the realm has already been assigned to the tenant.




Figure 9.5. The ‘Add StorageSet’ modal: Name dropdown (with capacities) + required Hard Quota.

Field	Notes
Name*	Dropdown listing every StorageSet in the realm, each annotated with its total capacity (e.g. “Set 1 – 1.34 TB”)
Hard Quota*	Required numeric + GB . Once a StorageSet is selected, a helper line shows how much of it is still unprovisioned, e.g. “Provisioned: 0% (1.34 TB left)”.

Once assigned, the StorageSets table shows the tenant's specific quota (not "Unlimited") against the StorageSet's full total, along with the edit/delete Actions icons.




Figure 9.6. The tenant's StorageSets table after assignment: quota shown, plus edit/delete Actions icons.

Volumes table:

Column	Description
Component	Volume path (e.g. <code>/</code> , <code>/home</code>); links to volume detail
StorageSet	Links to the owning StorageSet
EC Protection Level	e.g. "Dual Parity (n+2)"
Capacity Used	Amount of capacity currently used by the volume
Available	Amount of capacity still available to the volume
Soft Quota	The soft quota for the volume, if set.
Hard Quota	The volume's hard quota, if set: the maximum capacity the volume can consume
Status	The volume's status (e.g. "Online")

This table is the same as the global Volumes list (see [6.3.1](#)), minus the Tenant Name column, since this table is already scoped to the current tenant. A tenant with no volumes shows “No volumes yet”, with the note “Volumes are created in the tenant’s own zone.”


Bottom tabbed panel: Audit Log/Events

This table has the same toolbar pattern as Status Center:

Audit Log column	Description
Date & Time (UTC)	Sortable
Actor	e.g. “admin”
Action	e.g. “View”, “List”
Resource	e.g. “Tenant” or a linked ID like “Volume ID 2”
Source IP	The IP address the action originated from
Result	e.g. “Success”

Events column	Description
Time	When the event occurred
Category	The event’s category/type
Event Message	Human-readable description
Code	Machine-readable event code
User Name	The user associated with the event
IP Address	The IP address associated with the event

The Audit Log tab populates immediately, since simply viewing a page is itself an audited action. The Events tab follows the general behavior described in 5.2.



The screenshot shows the 'AUDIT LOG' tab in a dark-themed interface. The table has columns for Date & Time, Actor, Action, Resource, Source IP, and Result. The data shows various actions like 'Create', 'View', 'List', and 'Update' performed by a user with a long ID, mostly resulting in 'Success'.

Date & Time	Actor	Action	Resource	Source IP	Result
21:18:55	00000000-0000-0000-0000-0...	Create	Volume	127.0.0.1	Failed
21:18:00	00000000-0000-0000-0000-0...	View	Tenant	127.0.0.1	Success
21:18:00	00000000-0000-0000-0000-0...	List	Volume	127.0.0.1	Success
21:17:53	00000000-0000-0000-0000-0...	View	Tenant	127.0.0.1	Success
21:17:53	00000000-0000-0000-0000-0...	List	Volume	127.0.0.1	Success
21:17:53	00000000-0000-0000-0000-0...	View	Tenant ID 2e10785f-0271-42f1-88da-d2a611902d38	127.0.0.1	Success
21:17:53	00000000-0000-0000-0000-0...	View	Roles	127.0.0.1	Success
21:17:53	00000000-0000-0000-0000-0...	Me	Session	127.0.0.1	Success
21:17:11	00000000-0000-0000-0000-0...	View	Roles	127.0.0.1	Success
21:17:11	00000000-0000-0000-0000-0...	List	User	127.0.0.1	Success
21:17:11	00000000-0000-0000-0000-0...	View	Tenant	127.0.0.1	Success
21:17:11	00000000-0000-0000-0000-0...	List	Volume	127.0.0.1	Success
21:15:09	00000000-0000-0000-0000-0...	Update	StorageSet ID 1	127.0.0.1	Success
21:10:13	00000000-0000-0000-0000-0...	Create	Tenant ID 2e10785f-0271-42f1-88da-d2a611902d38	127.0.0.1	Success

Figure 9.7. The bottom Audit Log tab, populated with rows.

9.2.1 Delete Tenant Confirmation

This modal confirms deletion of a tenant. It reads, for example, “Delete”Tenant 1”? This action will permanently remove tenant and all its data. This cannot be undone. Ensure there are no active jobs or users.” and offers **No, Keep It** and **Yes, Delete Tenant** buttons.




Figure 9.8. The delete confirmation modal for a tenant.

A tenant that still has a network configuration cannot be deleted. Vantage shows an error banner reading “Failed to delete tenant: Cannot delete tenant while it still owns subnets: <name>”, with a **Contact Support** button and a close (X) button. Delete the tenant’s network configuration first (see 9.2.3), then delete the tenant.




Figure 9.9. The error shown when deleting a tenant that still owns a network configuration.

9.2.2 Finishing Tenant Setup

A tenant you create does not serve data right away. Until it is set up, its status chip reads **Setup Incomplete**, both beside the tenant name on its detail page and on its card in the Tenants list. Click the chip to open a **Finish setup** checklist of the steps the tenant still needs:

Step	Link	Description
StorageSet assignment	Add StorageSet	Gives the tenant a StorageSet and a Hard Quota (see 9.5.2).
Network configuration	Configure network	Defines the subnet the tenant’s clients use to reach storage (see 9.2.4).




Figure 9.10. The Finish setup checklist opened from the Setup Incomplete chip.

Once both steps are done, the chip changes to **Finishing Setup...** while the realm brings the tenant's network online, and then to **Active**, usually within a minute or two. Reload the page if the chip does not update on its own. A tenant must be **Active** before it can hold volumes. If you delete the tenant's network configuration later, the chip returns to **Setup Incomplete**.




Figure 9.11. A tenant in the Finishing Setup... state.

9.2.3 Networking Tab

The **Networking** tab lists the tenant's network configurations. A network configuration defines the subnet this tenant's clients use to reach storage; a tenant has one network configuration. A tenant with none shows "No network configurations yet" with a **Create Configuration** button, which is also available in the top-right of the tab.




Figure 9.12. The Networking tab of a new tenant, with no network configurations.

Once a network configuration exists, the tab shows it as a card titled with the configuration's name, and **Create Configuration** is greyed out (a tenant has one network configuration). The card has **Edit** and **Delete** actions and three groups of read-only fields:

Group	Fields
Addressing	Network (CIDR), Gateway, Address Range (first and last usable address)
Network Segment	Segment Type (e.g. "VXLAN"), Segment ID, Tunnel UDP Port (shows "Realm default" when you left it empty)
Advanced	MTU (shows "Platform default" when you left it empty), Address Allocation (e.g. "System"), ID




Figure 9.13. The Networking tab showing a network configuration.

Edit opens an “Edit Network Configuration” dialog. The Network is shown for reference but cannot be changed; you can change the **Name**, **Gateway**, **Address Range**, and **MTU**, then click **Save Changes** (or **Cancel**).




Figure 9.14. The Edit Network Configuration dialog.

Delete opens a confirmation reading, for example, "Delete tenant1-net (10.50.0.0/16)? This cannot be undone." with **Cancel** and **Delete** buttons.




Figure 9.15. The Delete Network Configuration confirmation.

9.2.4 Create Network Configuration Page

This page is the form for creating a tenant network configuration. When the realm has no network configuration yet, a banner at the top reads “Realm segment type: VXLAN. This is the first network configuration on the realm. After you create it, the segment type is fixed at VXLAN. Changing the realm to VLAN requires deleting all network configurations first.”




Figure 9.16. The top of the Create Network Configuration form, with the realm segment type banner.

Field	Notes
Name*	Text, 3 to 31 characters: letters, digits, underscore or hyphen
Network*	The subnet in CIDR notation, e.g. 10.0.1.0/24 . The host part must be zero. Cannot be changed after the configuration is created.
Gateway	Optional
Address Range	Optional from and to addresses. Leave empty to make the whole network usable.
Segment ID	In the Network Segment (VXLAN) section. Leave empty to derive one automatically.
Tunnel UDP Port	In the Network Segment (VXLAN) section. Leave empty to use the standard port 4789.

Field	Notes
MTU	In the Advanced section. Leave empty to use the platform default.
Address Allocation	In the Advanced section. Choose System (the system assigns each address itself; the default) or External allocator (an outside DHCP server assigns addresses and the system only records them). Cannot be changed after the configuration is created.

Segment settings cannot be changed after the configuration is created. **Cancel/Create Network Configuration.**

The screenshot displays the 'Create Network Configuration' form. It features a dark-themed interface with a sidebar on the left containing various icons. The main content area is divided into two sections: 'Network Segment (VXLAN)' and 'Advanced'.

Network Segment (VXLAN) Section:

- Address Range:** Two input fields separated by 'to'. Below them is the text: 'Leave empty to make the whole network usable'.
- Segment ID:** A single input field. Below it is the text: 'Leave empty to derive one automatically'.
- Tunnel UDP Port:** A single input field. Below it is the text: 'Leave empty to use the standard port 4789'.
- A warning message at the bottom of this section states: 'Segment settings cannot be changed after the configuration is created'.

Advanced Section:

- MTU:** A single input field. Below it is the text: 'Leave empty to use the platform default'.
- Address Allocation:** Two radio button options:
 - System** (selected): 'the system assigns each address itself (Default)'
 - External allocator**: 'an outside DHCP server assigns addresses; the system only records them'
 Below these options is the text: 'Cannot be changed after the configuration is created'.

At the bottom right of the form, there are two buttons: 'Cancel' and 'Create Network Configuration'.

Figure 9.17. The Network Segment and Advanced sections of the Create Network Configuration form.

The network must be large enough for the realm's nodes. If it is not, Vantage shows an error banner above the form, for example "Failed to create the network configuration: The usable-address window holds fewer addresses than the realm primary pool, so some nodes would get no address. Widen the CIDR, or range_start and range_end". Widen the **Network** (for example, use a /16 instead of a /24) or the **Address Range** and click **Create Network Configuration** again.

The screenshot shows the 'Create Network Configuration' page in the Vantage interface. At the top, there is a navigation bar with a back arrow, 'Tenants > Create', a status indicator, 'virtual-realm.local.com', a 'Platform' dropdown, and a user profile 'AD' with the time '9:26 PM UTC'. The main heading is 'Create Network Configuration'. Below the heading, a red error banner reads: 'Failed to create the network configuration: The usable-address window holds fewer addresses than the realm primary pool, so some nodes would get no address. Widen the CIDR, or range_start and range_end'. Below the error banner, a yellow warning banner states: 'Realm segment type: VXLAN. This is the first network configuration on the realm. After you create it, the segment type is fixed at VXLAN. Changing the realm to VLAN requires deleting all network configurations first.' The form fields are: 'Name*' with value 'tenant1-net' (3 to 31 characters: letters, digits, underscore or hyphen); 'Network*' with value '10.50.1.0/24' (The host part must be zero. Cannot be changed after the configuration is created); 'Gateway' with value '10.50.11'; 'Address Range' with two empty input boxes separated by 'to' (Leave empty to make the whole network usable); and 'Network Segment (VXLAN)' section with 'Segment ID' (Leave empty to derive one automatically).

Figure 9.18. The error shown when the network is too small for the realm.

9.3 Create Tenant Page

This page is the form for creating a new tenant.

The screenshot shows a web application interface for creating a new tenant. The page title is "Create Tenant". The form contains three required fields: "Name*" (filled with "Sample Tenant"), "Short Name*" (filled with "sample" and a tooltip icon), and "Contact Email*" (filled with "sample@vdera.local"). At the bottom right of the form are "Cancel" and "Create Tenant" buttons. The page has a dark theme and a sidebar on the left.

Figure 9.19. The Create Tenant form.

Field	Notes
Name*	Text
Short Name*	Text, with a (?) tooltip explaining the naming rule: 3 to 10 lowercase alphanumeric characters. Used for volume mount paths. Cannot be changed after creation.
Contact Email*	Contact Email address of tenant.

Cancel/Create Tenant.




Figure 9.20. The Create Tenant form filled in before submitting.

9.4 Acting as a Tenant (Context Switching)

This control only appears once the realm has **more than one tenant**. With only the built-in platform tenant present, there is nothing to switch to and the top bar looks exactly as described in 2.3. As soon as a second tenant exists, a new **Platform ▾** dropdown appears in the top bar, between the status icon and the user dropdown menu.

Clicking it opens a menu with the current context (“Platform”) at the top and a “Switch to Tenant” list underneath, one entry per tenant (by short name, e.g. “tenant1”). Selecting a tenant re-scopes the entire session to that tenant, with several visible effects:

What changes	Platform (admin) context	Inside a tenant
Top bar	Normal	A banner reading “Acting inside tenant <name> as Platform Admin” appears above the normal top bar, with a Back to Platform button. The Platform dropdown and the status icon are no longer shown.

What changes	Platform (admin) context	Inside a tenant
Sidebar	Full: Dashboard, Status Center, File System, Realm Nodes, Hardware, Tenants, Identity & Access	Reduced to only what a tenant would see: Dashboard, Volumes, Get Help, About
Volumes list columns	Includes Tenant Name	Tenant Name column dropped (scope is implicit)
Create Volume's Hard Quota	Optional ("leave empty for unlimited") if the tenant's own StorageSet quota is Unlimited	Required , with a Max link that fills in the tenant's remaining quota on that StorageSet
StorageSet Placement dropdown label	Shows the StorageSet's plain name (e.g. "Set 1")	Shows the StorageSet name labeled with the tenant's own quota (e.g. "Set 1 100.00 GB") rather than the StorageSet's full capacity

Click **Back to Platform** (in the banner) to return to the full administrative view; your platform-admin session is unaffected by having visited a tenant's context.

Inside a tenant, the Dashboard shows only what the tenant owns: a **Volumes** status tile, a **Total Capacity** card (Used, Available, and Unprovisioned capacity, with a usage bar against the tenant's quota for each StorageSet), the **Alerts** panel, and a bottom panel with **Audit Log** and **Events** tabs.




Figure 9.21. The Platform ▾ dropdown open, showing Platform plus a Switch to Tenant list.




Figure 9.22. The Acting inside tenant <name> as Platform Admin banner and reduced sidebar.

Inside a tenant's context, Create Volume shows its tenant-scoped variant: **StorageSet Placement** lists only the StorageSet(s) available to this tenant, labeled with the tenant's own quota rather than the StorageSet's full capacity, and **Hard Quota** is required with a **Max** link that fills in the tenant's full remaining allowance in one click.

The screenshot shows the 'Create Volume' interface within a tenant's context. The breadcrumb navigation at the top indicates 'Volumes > Create Volume'. The form includes the following fields and options:

- Pathname***: A text input field containing '/ projects'.
- StorageSet Placement***: A dropdown menu showing 'Set 1'. Below it, a status bar indicates 'Provisioned: 0.00% (100.00 GB left)'.
- Hard Quota***: A text input field containing '100' followed by a 'GB' unit selector and a 'Max' link.
- Erasure Coding Protection Level**: A section with three radio button options:
 - ☒ **Dual Parity (n+2)** – balance between capacity overhead and performance (Default)
 - ☐ **Single Parity (n+1)** – optimized for performance with less capacity overhead
 - ☐ **Striped Mirror (2x)** – optimized for small write performanceA note below states: 'Cannot be changed after a volume is created'.
- User & Group Permissions**: A dropdown menu at the bottom of the form.

At the bottom right of the form are two buttons: 'Cancel' and 'Create Volume'.

Figure 9.23. Create Volume while acting inside a tenant: required Hard Quota with the Max link, StorageSet labeled with its quota.

A tenant must also be **Active**, with a network configuration (see 9.2.4), before its volumes can be reached. If you click **Create Volume** for a tenant that does not, Vantage shows an error banner reading “The tenant holds no issued subnet address, so a volume created for it could not be reached. Create a subnet for the tenant and let one of its nodes report an address into it first”, and the volume is not created.

The screenshot shows the Vantage web interface for creating a volume. At the top, the breadcrumb navigation is 'Volumes > Create Volume'. A dark red error banner at the top of the form area contains the message: 'The tenant holds no issued subnet address, so a volume created for it could not be reached. Create a subnet for the tenant and let one of its nodes report an address into it first'. Below the banner, the form fields are: 'Pathname*' with a text input containing '/ projects'; 'StorageSet Placement*' with a dropdown menu showing 'Set 1' and a note 'Provisioned: 0.00% (100.00 GB left)'; 'Hard Quota*' with a text input '100' and a unit selector 'GB' with a 'Max' link; and 'Erasure Coding Protection Level' with three radio button options: 'Dual Parity (n+2) - balance between capacity overhead and performance (Default)', 'Single Parity (n+1) - optimized for performance with less capacity overhead', and 'Striped Mirror (2x) - optimized for small write performance'. A note below these options states 'Cannot be changed after a volume is created'. At the bottom of the form is a 'User & Group Permissions' dropdown menu. The bottom right of the interface has 'Cancel' and 'Create Volume' buttons.

Figure 9.24. The Create Volume error for a tenant with no network configuration.

Volume detail pages also depend on the context. Inside a tenant, a volume's page shows a simpler layout: a **Storage Details** card (StorageSet, Status, Provisioned), a **Data Protection** card (Erasure Coding Protection Level, EFSA Mode), the Soft Quota and Hard Quota tiles, and an **Available Capacity** card (Used, Available, and the Hard Quota).




Figure 9.25. A volume's detail page while acting inside its tenant.

From the Platform context, you can open any tenant's volume to look at it (from the realm-wide Volumes list, see 6.3.1), but not change it. The page shows a banner reading "This volume belongs to another tenant. Switch to " <name> " tenant to edit or delete it." and the **Delete** action is not offered. Switch into the owning tenant to edit or delete the volume.




Figure 9.26. A tenant's volume opened from the Platform context, with the belongs-to-another-tenant banner.

9.4.1 Give a Tenant Its Own Isolated Storage

1. **Create a StorageSet** (see 6.1.6) if one doesn't already exist as an empty shell; decide spare-VPOD reservation and Multi-Level Erasure Coding up front, since neither can be loosened or changed the same way later.
2. **Add Storage to it** (see 6.1.4): requires unassigned realm nodes. Expect transient alert noise while VPODs activate; it clears on its own within a minute or two.
3. For the built-in platform tenant, stop here: `/` and `/home` volumes appear automatically. For any other tenant, **create the tenant** (see 9.3), **assign it a StorageSet with a Hard Quota** (see 9.5.2), and **create a network configuration for it** (see 9.2.4) before it can hold any volumes.
4. Wait until the tenant's status is **Active** (see 9.2.2); a tenant cannot hold volumes before then. Then **switch into that tenant's context** (see 9.4) and **create a volume** (see 6.3.4); Hard Quota is required here, and the **Max** link is the fastest way to use the tenant's full allowance.

9.5 Creating and Managing Tenants

Use Vantage to create a new tenant when you need to isolate storage and volumes for a project, team, or customer.

9.5.1 Create a Tenant

1. Click **+ Create Tenant** on the Tenants list page (see [9.1](#)).
2. Enter a **Name** and **Short Name** for the tenant; the Short Name is immutable once the tenant is created.
3. Enter a **Contact Email**.
4. Click **Create Tenant** (see [9.3](#)). The new tenant starts with no StorageSets assigned and no network configuration, shows `Setup Incomplete`, and must be given both before it can hold any volumes (see [9.2.2](#)). If the realm serves its clients over InfiniBand, the realm refuses to create the tenant. See the Note at the start of this chapter.

9.5.2 Assign a StorageSet to a Tenant

1. Open the tenant's detail page (see [9.2](#)).
2. Click **+ Add StorageSet**.
3. Select a StorageSet from the **Name** dropdown, which lists every StorageSet in the realm along with its total capacity.
4. Enter a **Hard Quota**. A helper line shows how much of the StorageSet's capacity is still unprovisioned.
5. Click **Add StorageSet**.

9.5.3 Modify a Tenant's StorageSet Assignment

- To change the quota: click the pencil icon on the StorageSet's row in the StorageSets table.
- To remove the assignment: click the trash icon on the StorageSet's row in the StorageSets table.

9.5.4 Modify a Tenant

- To rename it: click the tenant's name, enter the new name, and click **Save**.
- To add or change its description: click the Description card's placeholder text, enter a description, and click **Save**.

To act on behalf of a tenant, see [9.4](#).

9.5.5 Delete a Tenant

1. If the tenant has a network configuration, open its **Networking** tab and delete the configuration first (see [9.2.3](#)); a tenant that still owns one cannot be deleted. Delete the tenant's volumes as well: switch into the tenant's context (see [9.4](#)), delete each volume (see [6.4.2](#)), then switch back to the Platform context.
2. Open the tenant's detail page (see [9.2](#)).
3. Click  **Delete Tenant** (see [9.2.1](#)).
4. Click **Yes, Delete Tenant** in the confirmation prompt.

9.5.6 Create a Network Configuration for a Tenant

1. Open the tenant's detail page (see [9.2](#)) and select the **Networking** tab (see [9.2.3](#)).
2. Click **Create Configuration**.
3. Enter a **Name** and the **Network** in CIDR notation. Optionally enter a **Gateway** and an **Address Range**.
4. Optionally expand **Network Segment** and **Advanced** to set the Segment ID, Tunnel UDP Port, MTU, and Address Allocation. Leave them empty to accept the defaults.
5. Click **Create Network Configuration** (see [9.2.4](#)). The network, the segment settings, and the address allocation mode cannot be changed afterward. If Vantage reports that the usable-address window holds fewer addresses than the realm primary pool, widen the Network or the Address Range and try again.
6. Wait for the tenant's status to change from **Setup Incomplete** to **Finishing Setup...** to **Active** (see [9.2.2](#)). Once it is **Active**, you can create volumes for it.


10. Identity and Access

Identity and Access groups the two kinds of accounts that can authenticate to the realm: **Users**, who sign in to Vantage with a username and password, and **Service Accounts**, which let scripts and integrations authenticate with API tokens. Both are reached from the **Identity & Access** item in the sidebar (see 2.3), and both are assigned the same roles.

A user is an account for signing in to Vantage, `panct1`, and PanAPI. It does not give access to PanFS data. Its main use is to replace the default admin account: create a dedicated user and give it full or limited permissions. Users exist only in the Platform area. Tenant-level users are not supported in this release, because a tenant has no sign-in of its own; the only way into a tenant is the tenant switcher (see 9.4).

10.1 Users List Page

This page lists every user in the realm. The top-right contains a **+ Create User** button; the table toolbar has filter/column/search icons.



Display Name	Username / Email	Roles	Last Activity
admin	admin	Platform Admin	20:50:49

Figure 10.1. The Users table.

Column	Description
Display Name	Link to user detail
Username/Email	Link to user detail

Column	Description
Roles	The role names assigned to the user, e.g. “Platform Admin”
Last Activity	Timestamp of the user’s most recent activity

10.2 User Detail Page

This page shows full details for a single user. The header shows the username, and the top-right contains the **Reset Password** and **Delete User** actions.




Figure 10.2. Full user detail page (User Info/Assigned Roles/Audit Log).

Card	Fields
User Info	Username, Email (with a pencil icon that opens the Edit Email modal), Last Activity
Assigned Roles	Table of Role + Description, plus a Manage Roles action. An “Effective permissions” line beneath the table summarizes what the user’s roles allow together. A role shown with a lock icon is required and cannot be removed (for example, Platform Admin on the built-in admin account).

The **Audit Log** panel has the same toolbar pattern as Status Center scoped to this user’s own activity.

Column	Description
Date & Time (UTC)	Sortable
Action	e.g. “View”, “List”
Resource	e.g. “User” or a linked ID
Source IP	The IP address the action originated from
Result	e.g. “Success”

10.2.1 Edit Email Modal

This modal changes a user’s contact email. It opens when you click the pencil icon beside the Email field in the User Info card, and is titled “Edit Email for <username>”. It contains a single **Email** field (described as the contact email for this user), with **Cancel** and **Save** buttons.




Figure 10.3. The Edit Email modal.

10.2.2 Reset Password Modal

This modal resets a user’s password.




Figure 10.4. The Reset Password modal, with the password field’s show/hide toggled on.

Field	Notes
New Password*	Required, hidden with a show/hide toggle. A hint beneath the field reads “At least 8 characters with upper/lowercase letters, a number and a special symbol.”

Cancel/Reset Password.

10.2.3 Manage Roles Modal

This modal changes the roles assigned to a user. It opens when you click **Manage Roles** in the Assigned Roles card.




Figure 10.5. The Manage Roles modal for a user.

It lists every assignable role as a checkbox, each with a plain-language description of what the role allows. Roles the user already holds are checked. A line at the bottom, “Effective permissions after saving”, updates as you check and uncheck roles. Click **Save changes** to apply or **Cancel** to discard.

Role	Description
Platform Admin	Can view, create, edit, delete and run operations on S3 Instances (see 14); read-only on Observability and Observability Realm; no access to S3 Reconcile State; full access to everything else. The checkbox is locked (checked and unchangeable) on the built-in admin account.
Storage Admin	Full access to StorageSets, Tokens and Volumes; no access to S3 Reconcile State; read-only elsewhere.
Platform Read-Only	Full access to Tokens; no access to S3 Reconcile State; read-only elsewhere.

The roles listed here apply across the platform; there are no tenant-scoped roles. Users exist only in the Platform area; tenant-level users are not supported in this release.

10.2.4 Delete User Confirmation

This modal confirms deletion of a user. It reads, for example, “Are you sure you want to delete user”admin”? This action cannot be undone.”




Figure 10.6. The delete confirmation modal.

10.3 Create User Page

This page is the form for creating a new user:

Field	Notes
Display Name	Optional text
Username*	Required text
Email*	Required text
Password*	Required, hidden with a show/hide toggle. A hint beneath the field reads “At least 8 characters with upper/lowercase letters, a number and a special symbol.”

Field	Notes
Confirm Password*	Required, hidden with a show/hide toggle
Roles	A checkbox table (Role, Description) with a select-all checkbox in the header. The offered roles are Platform Admin, Storage Admin, and Platform Read-Only (see 10.2.3). Users inherit permissions from every selected role.

The screenshot shows the 'Create User' form in the Vantage interface. The form is dark-themed and includes a sidebar on the left with navigation icons. The top navigation bar shows the user is logged in as 'virtual-realm.local.com' at 8:55 PM UTC. The form fields are: Display Name (optional), Username* (required), Email* (required), Password* (required, with a toggle to show/hide), and Confirm Password* (required, with a toggle to show/hide). Below the password fields is a 'Roles' section with a table of roles. The table has two columns: 'Role' and 'Description'. The roles listed are Platform Admin, Storage Admin, and Platform Read-Only. Each role has a checkbox in the 'Role' column and a description in the 'Description' column. The description for Platform Admin is 'Can view, create, edit, delete and run operations on S3 Instances; read-only on Observability; no access to S3 Reconcile State; full access to everything else'. The description for Storage Admin is 'Full access to StorageSets, Tokens and Volumes; no access to S3 Reconcile State; read-only elsewhere'. The description for Platform Read-Only is 'Full access to Tokens; no access to S3 Reconcile State; read-only elsewhere'.

Figure 10.7. The full Create User form, including the Roles checkbox table.

10.4 Creating and Managing Users

Use Vantage to create new realm admin and observer users and manage their roles.

10.4.1 Create a User

1. Click **+ Create User** on the Users list page (see [10.1](#)).
2. Enter a **Username** and **Email**; a **Display Name** is optional.
3. Enter a **Password** and confirm it. Passwords require at least 8 characters with upper/lowercase letters, a number, and a special symbol.
4. Select one or more **Roles** to assign to the user. Users inherit permissions from every selected role.
5. Click **Create User** (see [10.3](#)).

10.4.2 Change a User's Email

1. Open the user's detail page (see [10.2](#)).
2. Click the pencil icon beside **Email** in the User Info card (see [10.2.1](#)).
3. Enter the new **Email** and click **Save**.

10.4.3 Reset a User's Password

1. Open the user's detail page (see [10.2](#)).
2. Click **Reset Password** (see [10.2.2](#)).
3. Enter a **New Password** meeting the same complexity requirement, then click **Reset Password**.

10.4.4 Change a User's Roles

1. Open the user's detail page (see [10.2](#)).
2. Click **Manage Roles** in the Assigned Roles card (see [10.2.3](#)).
3. Check the roles the user should hold and uncheck any they should not. The "Effective permissions after saving" line shows the combined result.
4. Click **Save changes**.

10.4.5 Delete a User

1. Open the user's detail page (see [10.2](#)).
2. Click **Delete User** (see [10.2.4](#)).
3. Click **Delete** in the confirmation prompt.

10.5 Service Accounts List Page

A service account is a non-human identity for scripts and integrations. It has no password; it authenticates by presenting an API token (see 10.6.1), and it can do only what its assigned roles allow. This page lists every service account in the realm. The top-right contains a **+ Create Service Account** button; the table toolbar has filter and search icons.

Display Name	Roles	Last Activity
Backup Automation	Platform Read-Only	—
S3 Realm State Reader	S3 Realm Reader	—
S3 Reconciler	S3 Reconciler	20:19:11

Figure 10.8. The Service Accounts table.

Column	Description
Display Name	Link to service account detail
Roles	The role names assigned to the service account
Last Activity	Timestamp of the service account’s most recent activity or “—” if it has never been used

The realm includes built-in service accounts used by VDURA services (for example, “S3 Reconciler” and “S3 Realm State Reader”). These carry system roles that are not offered for any account you create. For S3, see 14.

10.6 Service Account Detail Page

This page shows full details for a single service account. The header shows the display name, and the top-right contains the **Issue Token** and **Delete** actions. The page has two tabs, **Overview** and **API Tokens**.




Figure 10.9. Service account detail page, Overview tab (Service Account Info/Assigned Roles/Audit Log).

Card	Fields
Service Account Info	Name (the immutable identifier), Created By, Created, Last Activity
Assigned Roles	Table of Role + Description, plus a Manage Roles action. An “Effective permissions” line beneath the table summarizes what the account’s roles allow together.

The **Audit Log** panel has the same toolbar pattern as Status Center scoped to this service account’s own activity.

Column	Description
Date & Time (UTC)	Sortable
Action	e.g. “View”, “List”
Resource	The resource acted on or a linked ID

Column	Description
Source IP	The IP address the action originated from
Result	e.g. "Success"

10.6.1 API Tokens Tab

This tab lists the API tokens issued to the service account. Until a token is issued it shows an empty state.




Figure 10.10. The API Tokens tab with no tokens issued.




Figure 10.11. The API Tokens tab after a token has been issued.

Column	Description
Name	The name given to the token when it was issued
Token ID	The token’s identifier (not the secret token value)
Status	“Active”, or “Revoked” once the token has been revoked
Created	When the token was issued
Expires	When the token stops working
Actions	A kebab menu whose only action is Revoke

10.6.2 Issue Token Modal

This modal issues a new API token. It opens when you click **Issue Token**.




Figure 10.12. The Issue Token modal.

Field	Notes
Name*	Required text, 1 to 64 characters, unique among the account’s usable tokens

Field	Notes
Validity Period*	Dropdown of 30, 90, 180, or 365 days; the default is 90. A helper line shows the resulting expiration date, e.g. "Expires Dec 31, 2026".

Cancel/Issue Token.

10.6.3 Token Issued Modal

After you click **Issue Token**, a "Token Issued" modal shows the new token once. It reads "This token is shown only once. Copy it now and store it securely." The token is masked by default; click the eye icon to reveal it, and click **Copy** to copy it to the clipboard. The expiration date appears beneath the token. Click **Done** to close the modal.




Figure 10.13. The Token Issued modal.

Important: Vantage cannot display the token again after you close this modal. If it is lost, revoke the token and issue a new one.

10.6.4 Manage Roles Modal

This modal changes the roles assigned to a service account. It works the same way as the user version described in 10.2.3: a checkbox per assignable role (Platform Admin, Storage Admin, Platform Read-Only), a live “Effective permissions after saving” line, and **Cancel/Save changes** buttons.




Figure 10.14. The Manage Roles modal for a service account.

10.6.5 Delete Service Account Confirmation

This modal confirms deletion of a service account. It reads, for example, “Are you sure you want to delete service account”Backup Automation”? Its tokens stop working immediately. This action cannot be undone.”




Figure 10.15. The delete confirmation modal for a service account.

10.7 Create Service Account Page

This page is the form for creating a new service account:

Field	Notes
Display Name	Optional text
Name*	Required. Lowercase letters, digits, and hyphens only. It cannot be changed later.
Roles	A checkbox table (Role, Description) with a select-all checkbox in the header. The offered roles are Platform Admin, Storage Admin, and Platform Read-Only. An “Effective permissions” line below the table updates as you select roles and reads “No access” when none are selected.

Cancel/Create Service Account.

← Service Accounts > Create Service Account

virtual-realm.local.com AD 8:47 PM UTC

Create Service Account

Display Name: Backup Automation

Name*: backup-automation
Lowercase letters, digits and hyphens. Cannot be changed later.

Roles
Select which roles to assign to this service account. It inherits permissions from every selected role.

Role	Description
☐ Platform Admin	Can view, create, edit, delete and run operations on S3 Instances; read-only on Observability; no access to S3 Reconcile State; full access to everything else
☐ Storage Admin	Full access to StorageSets, Tokens and Volumes; no access to S3 Reconcile State; read-only elsewhere
☒ Platform Read-Only	Full access to Tokens; no access to S3 Reconcile State; read-only elsewhere

Effective permissions: Full access to Tokens; no access to S3 Reconcile State; read-only elsewhere

Cancel Create Service Account

Figure 10.16. The Create Service Account form.

10.8 Creating and Managing Service Accounts

Use Vantage to create service accounts for automation, and to issue and revoke the API tokens those scripts use.

10.8.1 Create a Service Account

1. Click **+ Create Service Account** on the Service Accounts list page (see [10.5](#)).
2. Enter a **Name** using lowercase letters, digits, and hyphens; it cannot be changed later. A **Display Name** is optional.
3. Select the **Roles** to assign. Grant only what the automation needs.
4. Click **Create Service Account** (see [10.7](#)).

10.8.2 Issue an API Token

1. Open the service account's detail page (see [10.6](#)).
2. Click **Issue Token** (see [10.6.2](#)).
3. Enter a **Name** for the token and choose a **Validity Period**.
4. Click **Issue Token**.
5. In the Token Issued modal, click **Copy** and store the token securely (see [10.6.3](#)), then click **Done**.

10.8.3 Revoke an API Token

1. Open the service account's detail page and select the **API Tokens** tab (see [10.6.1](#)).
2. Click the kebab menu in the token's **Actions** column.
3. Click **Revoke**.
4. Click **Yes, Revoke Token** in the confirmation prompt. The prompt reads "Revoke <token name> ? This action will permanently revoke the token. Any client still using it will start failing immediately. This cannot be undone." The token stays in the list with the status **Revoked**.

10.8.4 Change a Service Account's Roles

1. Open the service account's detail page (see [10.6](#)).
2. Click **Manage Roles** in the Assigned Roles card (see [10.6.4](#)).
3. Check and uncheck roles as needed, then click **Save changes**.

10.8.5 Delete a Service Account

1. Open the service account's detail page (see [10.6](#)).
2. Click **Delete** (see [10.6.5](#)).
3. Click **Delete** in the confirmation prompt. Its tokens stop working immediately.

11. Account, Preferences, and Help

11.1 Account Menu and My Account

Clicking the user dropdown menu in the top bar opens a small dropdown.




Figure 11.1. The account dropdown open (Admin/My Account/Sign out).

The dropdown shows your display name (e.g. "Admin"), **My Account**, and **Sign out**.

My Account shows your own account details. The page has two tabs, **Profile** and **API Tokens**, and the top-right contains an **Issue Token** action.




Figure 11.2. The full My Account page.

Card	Fields
Account Information	Username, Email, Authentication (e.g. “Local”)
Roles & Permissions	Role + Description (e.g. Role “Platform Admin”, with a plain-language description of what the role allows), followed by an “Effective permissions” line summarizing what your roles allow together. The role name links to its detail.

The **Audit Log** panel has the same toolbar pattern as Status Center scoped to your own activity.

Column	Description
Date & Time (UTC)	Sortable
Action	e.g. “View”, “List”
Resource	e.g. “User” or a linked ID
Source IP	The IP address the action originated from
Result	e.g. “Success”

The **API Tokens** tab lists the personal API tokens you have issued, with the same columns, **Issue Token** modal, and token-shown-once behavior as a service account's tokens (see 10.6.1 through 10.6.3). Until you issue one, it shows a "No API tokens yet" empty state: "Issue a token so scripts and integrations can authenticate as you."




Figure 11.3. The My Account API Tokens tab.

11.2 Theme

The bottom of the sidebar has a **Switch to Light Mode/Switch to Dark Mode** toggle. It re-themes the entire UI instantly.




Figure 11.4. The Dashboard (or any page) rendered in light mode, for comparison.

11.3 Get Help

This page collects documentation links and ways to contact VDURA support.




Figure 11.5. The Get Help page.

The **Documentation & tools** card contains:

Link	Description
VDURA Documentation	Opens the VDURA documentation page on vdura.com: user guides, CLI and API references, release notes, and terminology.
API Explorer	Opens the interactive REST API reference (Swagger) served by this cluster.
Download PanCTL	Opens the page for downloading PanCTL, the command line tool for managing this cluster (see the PanCTL guide).

The **Contact support** card contains:

Item	Description
MyVDURA support portal	Opens the support portal, where you can open a case.
technical-support@vdura.com	Opens an email to the support team.
International phone numbers	Opens the list of VDURA support phone numbers by region.
System information	The cluster name and the software version of the realm. Click Copy to copy both to the clipboard so you can paste them into a support case.

11.4 About

This page shows information about this Vantage build, titled “About Vantage”.




Figure 11.6. The About page.

Item	Description
Logos	The Vantage and VDURA logos
Version	The Vantage UI build version, e.g. <code>main (1503772)</code>
Legal notice	The product protection notice and the list of U.S. patents
Copyright	"Copyright 2026, VDURA Inc."; the company name links to the public VDURA website
Terms & Conditions	A link to the software terms and conditions
Third-party licenses	A link to the third-party license information

11.5 PanActive Manager

Switch To Legacy UI, at the bottom of the sidebar, leaves Vantage entirely and loads the older classic interface at the same realm address.




Figure 11.7. PanActive Manager’s landing/sign-in page, showing its own top nav and Sign In form.

It has its own top navigation (**Home, Status & Reports, Storage, Hardware, Network, Configuration, Advanced Tools**) and its own **separate sign-in form**; it does *not* share the new UI’s session, so you must log in again there with your credentials. A **VANTAGE UI** link in its top bar returns you to the new interface, and your original Vantage session remains intact when you do.

Because it requires separate authentication, this guide does not document PanActive Manager’s internal pages beyond its landing/sign-in screen. For full coverage of that interface, see the VDURA Data Platform User Guide.

12. Security Enhanced Linux in VDURA Systems

This chapter explains the VDURA implementation of Security Enhanced Linux (SELinux).

12.1 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)?”.

12.2 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).

12.3 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.

12.4 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, Active Directory, 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.

12.5 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.

13. 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.

13.1 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.

13.1.1 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.

13.1.2 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.

13.1.3 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.

13.2 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.

13.2.1 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.

13.2.2 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.

13.2.3 KMS Integration

VDURA's encryption implementation implements a KMIP protocol (version 1.4 or better).

13.3 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.

13.4 Gateway Protocols

While VDURA offers the DirectFlow Client for a high-performance, parallel data path, other options exist to access data from other protocols, such as NFS, SMB, and S3. To enable these protocols, Directors act as gateways hosting these services. Each protocol accesses data through a PanFS Client that runs on the Director, called the Gateway Client.

The Gateway Client currently does not support encryption. Therefore, no gateway protocols are supported. Current behavior for non-DirectFlow protocols is as follows:

- **NFS:** If encryption is enabled, NFS support is disabled. NFS access to encrypted volumes can likely be done by re-exporting NFS over a DirectFlow Client that can decrypt data. User experience with NFS behavior will largely depend on the Linux distribution and kernel version.
- **SMB:** If encryption is enabled, SMB support is disabled.
- **S3:** If encryption is enabled, S3 support is disabled.

13.5 Configuration Considerations

Several configuration options must be set to allow a user to manage and control a system with end-to-end encryption enabled.

13.5.1 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 [13.9.1 KMIP Server Configuration](#) below for instructions on using PanCLI commands to configure the KMIP server.

13.5.2 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.

13.5.3 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.

13.5.4 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 [13.9.4 Client Interface Configuration](#) below for mount options to explicitly indicate that an encrypted volume is being mounted.

13.5.5 Secure Status

By querying the current state of the VDURA system via the `syscontrol encryption` and `volume list showcryption` PanCLI commands, the administrator can see what parts of the system are encrypted.

13.5.6 Mounting PanFS

When the new encryption mount option is specified, 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 new encrypted mount option will also result in a failure.

13.6 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.

13.6.1 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.

13.6.2 Lost or Destroyed Keys

If keys are lost or destroyed, then data in an encrypted volume will be permanently inaccessible.

13.6.3 Compromised Keys

No provision is made to deal with compromised keys.

13.6.4 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.

13.6.5 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.

13.7 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.

13.8 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.

13.9 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.

13.9.1 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
```

13.9.2 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 either NFS, CIFS, or S3 are 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
```

13.9.3 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
```

Volume	StorageSet	RAID	E2EE State	Space Used	Soft Quota	Hard 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.

13.9.4 Client Interface Configuration

When mounting the DirectFlow Client, credentials must be set to allow KMS access. 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.

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_III.pem /tmp
cp <path-to-certs>/dfc_III.key /tmp
```

2. Create a KMS config file. 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 = 10.70.25.173
port = 5696
version = 1.4

connect-timeout = 300
reconnect-retries = 5

ca-cert-path = /tmp/RootCertificate.pem
client-cert-path = /tmp/dfc_III.pem
client-key-path = /tmp/dfc_III.key
```

3. Create a directory to mount under. You cannot use `/panfs` because global mounts are not allowed:

```
mkdir /mnt
```

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
```

14. S3

Note: S3 is configured with `panct1` or the API. Vantage does not manage S3 in this release, so this chapter has no screens.

If end-to-end encryption is enabled, S3 support is disabled (see [13.4](#)).

This chapter shows a **platform administrator** how to set up, run and use the PanFS S3 service: enabling S3, gateway instances, S3 users, buckets, TLS, named access, monitoring, client tools and day-to-day operations. It describes PanFS 12.0.0. The commands and example output were run on a live realm.

14.1 Concepts

Term	Meaning
Realm S3 switch	Realm-wide on/off switch. While it is off, no S3 gateway runs anywhere in the realm.
Tenant S3 config	Per-tenant settings include on/off, the prefix (the filesystem path S3 serves), default POSIX ownership and modes, logging, TLS, aggregator, DNS.
Prefix	A tenant volume (or a directory inside one). Each bucket is a directory directly under the prefix, and each object is a file inside it.
Gateway instance	An S3 endpoint running on one node, identified by a 5-character instance ID (iid) . One iid can run on several nodes.
S3 user (fs-user)	A PanFS user that holds an S3 access key and secret key, plus the POSIX uid/gid that S3 requests run as. It never signs in to the management API.
Bucket owner	The S3 user that a bucket belongs to and that has full control of it.
Aggregator	Experimental: An optional extra endpoint on every director, in front of the tenant's running instances. Off by default.

S3 runs **per tenant**. Each tenant has its own prefix, users, buckets, instances, certificate and network endpoints.

14.2 Logging In with panctl

`panctl` is installed on every director at `/pan/bin/panctl`. Run it as root.

The management API listens on port 8000 of **one** director, the one hosting the realm's management services. On that director, `panctl` connects to `http://localhost:8000` by default. On any other director, give that director's address.

The password is the realm `admin` password.

```
# On the director that hosts the management API
printf '%s' "$ADMIN_PASSWORD" | panctl auth login --username admin --password-stdin

# From another director
printf '%s' "$ADMIN_PASSWORD" | panctl auth login http://<management-director-ip>:8000 --
    username admin --password-stdin

panctl auth whoami          # ROLES: platform-admin, IS PLATFORM ADMIN: true
```

The session is saved as a profile in `~/.panctl/profiles.toml`. Use `--profile <name>` to keep several sessions.

Every S3 command in this chapter takes `--tenant <short name or UUID>`. The examples use a tenant whose short name is `acme`.

Confirmation. Commands that delete something ask for confirmation in a terminal. In a script or whenever there is no terminal, they refuse with `this operation requires confirmation; pass --yes to proceed`. The examples pass `--yes`.

API tokens for scripts. A token lets a script run `panctl` without a saved login:

```
TOKEN=$(panctl auth tokens create --name s3-automation --expires-in-days 30 --format 'jsonpath=
    {.token}') # shown once
echo "$TOKEN" | PANCTL_API_SERVER=http://<management-director-ip>:8000 panctl --token-stdin s3
    get-config --tenant acme

panctl auth tokens list
panctl auth tokens revoke <token id> --yes
```

14.3 Deployment Options

14.3.1 Dedicated Tenant (Recommended)

A tenant created for S3 users (for example `acme`) has its own private network, the **tenant subnet**: – The gateways listen on the nodes' addresses in that subnet. – Only hosts attached to the subnet can reach them.

This chapter uses this option throughout. A dedicated tenant needs an Ethernet network. Multi-tenancy is not supported on InfiniBand, so tenants cannot be created on a realm whose clients connect over InfiniBand (see the Note at the start of 9). The platform tenant (see 14.3.2) does not use a tenant subnet.

14.3.2 Platform Tenant

The built-in `platform` tenant can also serve S3 (see 14.17). Its gateways listen on the nodes' own realm addresses, so any host that can reach the realm network can reach them.

14.3.3 Choices Within a Tenant

Choice	Options	Recommendation
Instance nodes	Directors	Several directors, for availability (see 14.6)
Aggregator	<code>off</code> (default) or <code>on</code> (experimental)	<code>off</code> (see 14.10.1)
Named accesses	<code>off</code> (default), <code>external</code> , or <code>panfs</code> (experimental, platform tenant only)	<code>external</code> : You publish the names in your own DNS (see 14.10)
TLS certificate	<code>realm-generated</code> (default) or <code>tenant-provided</code>	<code>tenant-provided</code> , so S3 clients can verify the endpoint (see 14.9)

S3 is always served over HTTPS. Every instance listens on port **3000**.

14.4 Prerequisites

This section gives the minimum non-S3 setup needed for S3. It does not cover these objects' other options.

14.4.1 Tenant and Storage

```
panctl storagesets list                                # note a storage set ID, e.g. 1
panctl tenants create --name acme --display-name "ACME Corporation" --email s3-
admin@acme.example
panctl tenants assign-storagesets --tenant acme --storage-set-id 1 --quota 200GiB
```

The tenant short name (`--name`) is fixed once created. It also appears in S3 host names and certificate names (`s3-acme`).

14.4.2 Tenant Subnet

Instances in a dedicated tenant need a subnet. The realm gives every node an address in it.

```
panctl subnets create --tenant acme --name acme-net --cidr 10.240.0.0/20
panctl subnets list --tenant acme          # ADDR COUNT = number of realm nodes that have an
address
```

Use a range that does not overlap your other networks. It needs enough addresses for every realm node plus every client host.

14.4.3 Attaching Client Hosts to the Tenant Subnet

Every host that will use S3 for this tenant must be attached to the subnet.

1. Register the host. Pass its address on the realm network:

```
panctl subnets attachments create acme-net --tenant acme --underlay-ip 10.128.9.0 --hostname
client-0 --format json
```

The response tells the host how to join:

```
{
  "overlay_ip": "10.240.8.1",
  "overlay_prefix": 20,
  "mtu": 1450,
  "segment_type": "vxlan",
  "segment_id": 12010160,
  "tunnel": {
    "mac": "9a:89:c2:d5:41:81",
    "udp_port": 4789,
    "vteps": ["10.128.1.1", "10.128.1.2", "10.128.2.2", "...one address per realm node..."]
  }
}
```

2. Configure the interface. On a Linux client, use the values from that response:

```
modprobe vxlan
ip link add acme0 type vxlan id 12010160 local 10.128.9.0 dev eth0 dstport 4789
ip link set acme0 address 9a:89:c2:d5:41:81 mtu 1450
ip addr add 10.240.8.1/20 dev acme0
ip link set acme0 up
for v in <every address in tunnel.vteps>; do bridge fdb append 00:00:00:00:00:00 dev acme0 dst $v; done
```

- `dev eth0` is the interface that holds the `--underlay-ip` address.
- These commands do not survive a reboot. Make them persistent with your distribution's network configuration.

3. Check it. `ping 10.240.0.1` should get replies from a director.

14.4.4 Volume

Create the volume that S3 will serve. A volume is created inside the tenant you act in. To act in `acme`, add the global `--as-tenant-admin acme` option: It runs that single command inside `acme` on behalf of the platform administrator.

```
panctl volumes create --as-tenant-admin acme --name /s3data --storage-set-id 1 --hard-quota 100GiB
panctl volumes list --tenant acme
```

14.5 Enabling S3

14.5.1 Realm Switch

```
panctl s3 enabled set true
panctl s3 enabled get
```

14.5.2 Tenant Configuration

Create the tenant's S3 configuration with its prefix and turn it on in one step:

```
panctl s3 create-config --tenant acme --prefix /s3data --enabled true
panctl s3 get-config --tenant acme
```

```
ENABLED:          true
ACCESS LOG:       true
DEBUG LOG:        false
METRICS ENABLED:  true
DEFAULT PREFIX:   /s3data
DEFAULT OWNER:    uid:65534
DEFAULT GROUP:    gid:65534
DEFAULT DIR MODE: 755
DEFAULT OBJ MODE: 644
TLS:              {"ca_certificate":null,"certificate":"-----BEGIN CERTIFICATE-----
...","mode":"realm-generated","private_key":"***"}
AGGREGATOR:       {"enabled":false}
```

Always pass `--prefix` together with `--enabled true`. `create-config --enabled true` without a prefix is accepted, but it leaves S3 enabled with no prefix, so instances cannot serve anything. Recover with `set-config --enabled false`, then `--prefix`, then `--enabled true`.

14.5.3 Settings Reference

Change settings with `panctl s3 set-config --tenant acme <flags>` . Only the flags you pass are changed.

Flag	Default	Meaning	Running instances pick it up
<code>--enabled</code>	false	Tenant S3 on/off	n/a
<code>--prefix</code>	empty	Volume path (or a directory inside a volume) served by S3, e.g. <code>/s3data</code>	n/a
<code>--owner</code>	uid:65534	Default file owner, used for S3 users that have no uid of their own	after restart
<code>--group</code>	gid:65534	Default file group, used for S3 users that have no gid of their own	after restart
<code>--dir-mode</code>	755	Mode of directories created through S3 (3–4 octal digits)	after restart
<code>--obj-mode</code>	644	Mode of files created through S3 (3–4 octal digits)	after restart
<code>--access-log</code>	true	Per-request access log (see 14.14)	after restart
<code>--debug-log</code>	false	Accepted, but has no visible effect on gateway logging in this release	n/a
<code>--metrics-enabled</code>	true	Publishes S3 metrics (see 14.14)	after restart

The TLS, aggregator and DNS settings are covered in sections [14.9](#) and [14.10](#). They take effect without a restart.

When a change needs a restart, `set-config` prints a note:

```
note: the tenant's running S3 instances pick up access_log only after a restart ('panctl s3
instances stop-all' then 'start-all' or stop/start a single instance by iid)
```

Rules:

- **The prefix** must be an existing volume of the tenant or an existing directory inside one.
- **To change the prefix**, delete the tenant's instances, set `--enabled false`, change `--prefix`, set `--enabled true`, then create and start the instances again (see [14.15](#)). The prefix cannot be changed while tenant S3 is enabled. Stopping and starting the instances is not enough: They keep serving the old prefix until they are recreated.
- **To disable tenant S3**, stop all of its instances first.

14.6 Gateway Instances

14.6.1 Choose Nodes

```
panctl nodes list
```

NAME	DATASERIAL	SHELF ID	IP ADDRESS	STATUS	ROLE	...
VCH-1,1	b022963e7165c9	3e7e0000404001001	10.128.1.0	online	qnode	
VCH-5,1	32cdb93e7165fa	3e7e0001404001001	10.128.1.1	online	qnode	
VCH-9,1	409d673e7165cc	3e7e0002404001001	10.128.1.2	online	qnode	
VCH-2,1	94c5393e716630	3e7e00ff404001001	10.128.2.0	online	vosd	
...						

Where instances can run: – Instances run on **directors** (ROLE `qnode`). – **Do not place instances on storage nodes (ROLE `vosd`)**. Such an instance is accepted but never starts. It also cannot be removed, so it blocks [removing S3 from the tenant](#).

Limits: – A tenant can have **at most one instance per node**. – Different tenants can use the same node.

14.6.2 Create and Start

The iid is exactly 5 letters or digits. The same iid can be created on several nodes. Commands that name only the iid then act on every node that holds it.

```
panctl s3 instances create gw001 b022963e7165c9 --tenant acme      # iid, node DATASERIAL
panctl s3 instances create gw001 32cdb93e7165fa --tenant acme      # same iid on a second
director
panctl s3 instances create gw004 409d673e7165cc --tenant acme      # another iid on a third
director
panctl s3 instances start gw001 --tenant acme                      # starts gw001 on both nodes
panctl s3 instances start gw004 --tenant acme
panctl s3 instances list --tenant acme
```

INSTANCE ID	DATA SERIAL	DESIRED	OBSERVED
gw001	32cdb93e7165fa	started	running
gw001	b022963e7165c9	started	running
gw004	409d673e7165cc	started	running

A new instance is `created` and must be started explicitly. It usually reaches `running` within 30 seconds. To wait for it:

```
panctl s3 instances get gw001 --tenant acme --wait-for-observed=running
```

14.6.3 Endpoints

Endpoints appear in the JSON output:

```
panctl s3 instances list --tenant acme --format 'jsonpath={.instances[*].endpoint}'
# https://10.240.0.2:3000 https://10.240.0.1:3000 https://10.240.0.3:3000
```

Each endpoint is the node's address on the tenant subnet, port 3000. `panctl s3 instances get <iid> -tenant acme --format json` shows each placement with its `endpoint`, `observed` state and `health`.

14.6.4 States and Health

Field	Values
DESIRED	created , started , stopped , deleting
OBSERVED	state reported by the node: running , stopped , crashloop , or unknown when the node does not report the instance (observed_reason says why)
health.last_seen_ts	time (epoch seconds) of the node's last report, refreshed about every 30 seconds
health.restart_count_60s	automatic restarts of the gateway in the last 60 seconds
health.last_error	the last error the node reported, for example while the gateway is down; empty once it runs again

Automatic recovery. If a gateway process exits, the node restarts it within seconds, and `restart_count_60s` counts the restart. After several exits in a short time, `OBSERVED` shows `crashloop` and the gateway stays down for about a minute before the node starts it again.

Availability. An instance stays on the node it was created on. For availability, run the tenant's instances on several directors, for example the same iid on two or three directors. Then publish one DNS name per iid with an address record for each director hosting it (see [14.10.2](#)), so clients are not tied to a single node.

14.6.5 Stop, Start, Delete

```
panctl s3 instances stop gw004 --tenant acme
panctl s3 instances start gw004 --tenant acme
panctl s3 instances stop gw001 --dataserial 32cdb93e7165fa --tenant acme # one placement only
panctl s3 instances stop-all --tenant acme
panctl s3 instances start-all --tenant acme
panctl s3 instances delete gw004 --tenant acme --yes
panctl s3 instances delete gw001 --tenant acme --confirm-fanout --yes # the iid is on more
    than one node
panctl s3 instances delete-all --tenant acme --yes
```

Deletion finishes in the background. To wait until an iid is gone:

```
panctl s3 instances list --tenant acme --wait-for 'it == null' \
-o 'jsonpath={.instances[?(@.desired.instance_id=="gw004")].desired.instance_id}'
```

14.7 S3 Users

S3 users hold the S3 credentials. Each user has:

- a **name**, unique in the tenant;
- an **access key**, which defaults to the name;
- a **uid/gid** that files created through S3 are owned by. If omitted, the tenant's `--owner` / `--group` defaults apply at creation time (65534 unless changed);
- an **S3 profile**:

Profile	Can do
user (default)	Work with the buckets it owns. Cannot create buckets; an administrator creates them with this user as owner.
privileged	Everything <code>user</code> can, and also create its own buckets over S3.
disabled	Nothing: The credential is kept, but the gateways reject it (<code>InvalidAccessKeyId</code>).

Create and list:

```
panctl fs-users create --tenant acme --name alice --uid 2001 --gid 2001 --s3-profile privileged
panctl fs-users create --tenant acme --name bob
panctl fs-users list --tenant acme
```

ID	NAME	UID	GID	PROFILE	ACCESS KEY
789b3eb6-3eda-4fe6-a8c3-a72e3997b4dc	alice	2001	2001	privileged	alice
15bc5774-0c37-4f89-87c5-e9fed641189d	bob	65534	65534	user	bob

Read the secret. It is not printed at creation:

```
panctl fs-users s3-secret 789b3eb6-3eda-4fe6-a8c3-a72e3997b4dc --tenant acme # ACCESS KEY
/ SECRET KEY
```

Rotate the secret. The access key stays the same. The new secret is printed once, and the old secret stops working. `panctl` notes that a gateway may keep accepting the old secret for a short time.

```
panctl fs-users rotate-s3-secret <fs_user_id> --tenant acme
```

Disable and re-enable, change identity, delete:

```
panctl fs-users update <fs_user_id> --tenant acme --s3-profile disabled
panctl fs-users update <fs_user_id> --tenant acme --s3-profile user
panctl fs-users update <fs_user_id> --tenant acme --uid 2002 --gid 2002
panctl fs-users delete <fs_user_id> --tenant acme --yes
```

Deleting a user does not delete its buckets: They remain, still listed with that owner. Empty and delete them first (see [14.16](#)).

14.8 Buckets

Create. An administrator creates a bucket for an owner, given by the owner's **access key**. At least one instance of the tenant must be running.

```
panctl s3 buckets create projects --owner alice --tenant acme
panctl s3 buckets list --tenant acme
```

NAME	OWNER
projects	alice

Bucket names follow the S3 rules: 3–63 characters, lowercase letters, digits, `-` and `.`.

Buckets for `user -profile` users. `--owner` accepts a user with any profile. A `user -profile` user cannot create buckets over S3 (`AccessDenied`), so create its buckets this way:

```
panctl s3 buckets create bobs --owner bob --tenant acme
```

As the owner, the user has full control of the bucket over S3: It reads and writes objects, sets the bucket policy and ACLs, and can delete the empty bucket.

Buckets created by S3 users. Users with the `privileged` profile can also create buckets themselves (`aws s3 mb`). These buckets appear in `panctl s3 buckets list` like any other, and the administrator can delete them.

Delete. Any bucket of the tenant can be deleted, whoever created it. The bucket must be empty. Objects are removed over S3.

```
panctl s3 buckets delete projects --tenant acme --yes
```

`panctl s3 buckets` commands are carried out by a running instance of the tenant. Without one they fail (see [14.18](#)).

14.9 TLS Certificates

All S3 endpoints use HTTPS. Each tenant has one certificate, used by all of its instances (and by the aggregator, if enabled). Certificate changes take effect within seconds, without restarting instances.

14.9.1 Tenant-Provided Certificate (Recommended)

You supply a certificate issued by a CA your S3 clients already trust, and clients verify the endpoint normally.

Requirements

Item	Requirement
Format	PEM. <code>certificate</code> holds the server certificate followed by any intermediate CA certificates . <code>private_key</code> holds the key.
Private key	Unencrypted. RSA of 2048 bits or more, ECDSA P-256, P-384 or P-521, or Ed25519. Must match the certificate.

Item	Requirement
Names (SAN)	Every host name and IP address that clients put in their endpoint URLs: Your own DNS names (Using Your Own Host Names), the realm's names (Names), or the endpoint addresses (see 14.6.3), plus the directors' addresses if clients use the aggregator by IP. IP addresses must be listed one by one: a certificate cannot hold an IP wildcard or range.
ca_certificate	Optional. Not needed for clients to verify the endpoint.

The certificate is checked when you submit it. If it does not meet these requirements it is refused with a message, and the current certificate stays in use. For example:

```
tls.certificate carries a RSA-1024 key, which ... Accepted key types: RSA (at least 2048 bits),
ECDSA P-256, P-384 or P-521, or Ed25519.
tls.private_key does not match tls.certificate.
tls.private_key is not a parseable PEM private key: Password was not given but private key is
encrypted
```

If the chain is incomplete: A certificate issued by an intermediate CA is accepted without the intermediate. Clients then fail to verify it, so always append the intermediates.

Names

The realm's names for a tenant come from the tenant short name and the realm DNS zone. A realm-generated certificate always carries them (see [14.9.2](#)). Include them in a tenant-provided certificate if clients use them.

Realm DNS zone	Realm names (tenant acme)
not set	*.s3-acme (plus s3-acme with the aggregator)
lab.example.com	*.s3-acme.lab.example.com (plus s3-acme.lab.example.com with the aggregator)

Check the zone with `pancli -c "dns zone"` on a director, and set it with `pancli -c "dns zone lab.example.com"`.

Without a realm DNS zone, do not rely on the realm names. Clients disagree on a wildcard over a single-label name such as `*.s3-acme`. The AWS CLI v2, boto3 and rclone accept `gw001.s3-acme`, but curl and other tools that use OpenSSL's host name check refuse it with a hostname mismatch. Connect clients by IP address or by your own host names.

Using Your Own Host Names

Clients can use host names of your choice, for example `s3.acme.example.org`, whatever the realm DNS zone:

1. Publish the names in your DNS, each pointing to the endpoint addresses of the instances it should reach (see [14.6.3](#)).
2. Put the names in the certificate's SAN. The certificate needs no other names.
3. Point clients at `https://<your name>:3000`.

Example

1. Issue a certificate. This uses a lab CA; in production, use your organization's CA. The certificate carries the realm's names for zone `lab.example.com`, which section [14.10.2](#) publishes.

```
openssl genpkey -algorithm RSA -pkeyopt rsa_keygen_bits:2048 -out s3-acme.key
openssl req -new -key s3-acme.key -out s3-acme.csr -subj "/CN=s3-acme.lab.example.com"
cat > s3-acme.ext <<'EOF'
subjectAltName=DNS:s3-acme.lab.example.com,DNS:*.s3-acme.lab.example.com
extendedKeyUsage=serverAuth
keyUsage=critical,digitalSignature,keyEncipherment
EOF
openssl x509 -req -in s3-acme.csr -CA ca.crt -CAkey ca.key -CAcreateserial -out s3-acme.crt -
days 397 -extfile s3-acme.ext
```

2. Upload it. If an intermediate CA issued it, use `cat s3-acme.crt intermediate.crt` as the certificate.

```
jq -n --rawfile cert s3-acme.crt --rawfile key s3-acme.key \
  '{tls: {mode: "tenant-provided", certificate: $cert, private_key: $key}}' |
panctl s3 set-config --tenant acme -f -
```

3. Verify from a client that trusts the CA:

```
aws --endpoint-url https://gw001.s3-acme.lab.example.com:3000 --ca-bundle ca.crt s3 ls
```

To replace the certificate, upload a new one the same way.

14.9.2 Realm-Generated Certificate (Default)

If you do not supply a certificate, the realm issues one: - it covers the names above and the tenant's endpoint addresses; - it is valid for one year; - it is signed by the realm's internal root CA.

Clients cannot verify a realm-generated certificate: The realm root CA is not shared outside the realm. Clients must therefore skip certificate verification (see [14.11.4](#)). Use a tenant-provided certificate wherever clients need verified connections.

The certificate takes its names from the realm DNS zone at the time it is issued. After changing the zone, renew it.

Switch back to a realm-generated certificate or renew it before it expires (renewal is not automatic):

```
echo '{"tls":{"mode":"realm-generated"}}' | panctl s3 set-config --tenant acme -f -
panctl s3 tls renew --tenant acme
```

14.10 Aggregator and Named Access (DNS)

14.10.1 Aggregator (Experimental)

The aggregator is **experimental at this stage; keep it off** (the default) for regular use.

```
panctl s3 aggregator on --tenant acme      # experimental
panctl s3 aggregator off --tenant acme     # default
```

With the aggregator on: - **every director** serves the tenant on port **3001** of its tenant-subnet address, for example `https://10.240.0.1:3001`; - requests are forwarded to the tenant's running instances, and stopped instances are skipped; - the per-instance endpoints on port 3000 keep working.

With a tenant-provided certificate, add the names or addresses clients use for the aggregator to the certificate (see [14.9.1](#)).

14.10.2 Named Access with Your Own DNS (external)

To let clients use host names instead of IP addresses, publish the names in your own DNS:

```
panctl s3 dns external --tenant acme
```

PanFS publishes nothing in this mode. You create the records:

Name (tenant acme , zone lab.example.com)	Records	Port
<iid>.s3-acme.lab.example.com	one A record per director hosting that instance	3000
s3-acme.lab.example.com	only with the experimental aggregator: One A record per director	3001

Getting the addresses: - Instance addresses: `panctl s3 instances list --tenant acme --format json ( endpoint field).`

Getting the names right: - The table shows the realm's names (see [Names in 14.9.1](#)). To use them, set the realm DNS zone to the domain you publish them in, and put them in the tenant-provided certificate. - Alternatively, publish host names of your own and put only those in a tenant-provided certificate ([Using Your Own Host Names](#)). The realm DNS zone then does not matter. - A realm-generated certificate covers only the names in the table.

Example with dnsmasq. Records for the realm above, as served by dnsmasq:

```
host-record=gw001.s3-acme.lab.example.com,10.240.0.1
host-record=gw001.s3-acme.lab.example.com,10.240.0.2
host-record=gw004.s3-acme.lab.example.com,10.240.0.3
```

Check from a client:

```
dig +short gw001.s3-acme.lab.example.com      # 10.240.0.1 10.240.0.2
aws --endpoint-url https://gw001.s3-acme.lab.example.com:3000 --ca-bundle ca.crt s3 ls
aws --endpoint-url https://gw004.s3-acme.lab.example.com:3000 --ca-bundle ca.crt s3 ls
```

Keep the records current when you add, move or delete instances.

Clients must use **path-style** addressing (`https://gw001.s3-acme.lab.example.com:3000/<bucket>/<key>`). Bucket names in the host name (virtual-hosted style) are not supported.

14.10.3 Realm-Published Names (panfs, Experimental)

`panctl s3 dns panfs` makes the realm publish the names itself:

- It is **experimental and not recommended for production use**.
- It is available **only for the platform tenant**. Any other tenant is refused with `dns.mode=panfs` is not supported for this tenant ... Use `dns.mode=off` or `dns.mode=external`.
- Clients can use the names only if they resolve through the realm's DNS service.
- A tenant-provided certificate must carry the realm's names (see [Names in 14.9.1](#)), including `s3-<tenant>[.<zone>]` with the aggregator on. Otherwise the change is refused with `tls.certificate` is missing SAN(s) the realm publishes with `dns.mode=panfs`.

`panctl s3 dns off` (the default) turns named access off; clients use the IP endpoints.

The current DNS mode is not shown by `panctl s3 get-config`.

14.11 Client Access

Clients connect to an instance endpoint (port 3000) with an S3 user's access key and secret key.

Client settings: - **Region:** `us-east-1` . Requests signed for another region are accepted, but creating a bucket fails with `InvalidLocationConstraint` when the client asks for another region, as `aws s3 mb` does when the client's region is not `us-east-1` . - **Signature:** Signature Version 4. Signature Version 2 is not supported. Requests signed with it are rejected: - header-signed requests with `AccessDenied` ; - browser upload forms (see [14.12.4](#)) with `InvalidArgument` ; - presigned URLs with `InvalidRequest` . -

Addressing: Path-style.

The examples use the AWS CLI v2 ([installation](#)), boto3 and rclone.

14.11.1 AWS CLI

```
# ~/.aws/credentials

[acme]
aws_access_key_id = alice
aws_secret_access_key = <secret key>

# ~/.aws/config

[profile acme]
region = us-east-1
endpoint_url = https://gw001.s3-acme.lab.example.com:3000
ca_bundle = /etc/pki/acme-ca.crt
```

```
export AWS_PROFILE=acme
aws s3 ls
aws s3 cp report.pdf s3://projects/docs/report.pdf
aws s3 cp big.bin s3://projects/big.bin          # Large files are uploaded in parts
                                                automatically
aws s3 ls s3://projects --recursive --human-readable
aws s3 cp s3://projects/big.bin ./big.bin
aws s3 presign s3://projects/docs/report.pdf --expires-in 3600
aws s3 rm s3://projects/docs/report.pdf
aws s3 mb s3://alice-scratch                    # privileged users only
aws s3 rb s3://alice-scratch                    # bucket must be empty
```

Without a profile, pass the options on each call: `--endpoint-url https://<address>:<port> --ca-bundle <file> .`

14.11.2 Python (boto3)

```
import boto3
from botocore.config import Config

s3 = boto3.client(
    "s3",
    endpoint_url="https://gw001.s3-acme.lab.example.com:3000",
    aws_access_key_id="alice",
    aws_secret_access_key="<secret key>",
    region_name="us-east-1",
    verify="/etc/pki/acme-ca.crt",
    config=Config(signature_version="s3v4", s3={"addressing_style": "path"}),
)

s3.put_object(Bucket="projects", Key="docs/a.txt", Body=b"hello")
url = s3.generate_presigned_url("get_object", Params={"Bucket": "projects", "Key":
    "docs/a.txt"}, ExpiresIn=3600)
```

Set `signature_version="s3v4"` . Without it, boto3 creates presigned URLs that the gateway rejects (`InvalidRequest`).

14.11.3 rclone

```
# ~/.config/rclone/rclone.conf
[acme]
type = s3
provider = Other
access_key_id = alice
secret_access_key = <secret key>
endpoint = https://gw001.s3-acme.lab.example.com:3000
force_path_style = true
region = us-east-1
```

```
rclone --ca-cert /etc/pki/acme-ca.crt ls acme:
rclone --ca-cert /etc/pki/acme-ca.crt copy ./data acme:projects/data
```

14.11.4 Connecting to a Realm-Generated Certificate

A realm-generated certificate cannot be verified (see [14.9.2](#)), so turn verification off in the client:

Client	How
AWS CLI	<code>--no-verify-ssl</code> on every command. A warning (<code>InsecureRequestWarning</code>) is printed for each request.
botocore	<code>verify=False</code> in <code>boto3.client(...)</code> ; to silence the warnings, <code>urllib3.disable_warnings(urllib3.exceptions.InsecureRequestWarning)</code>
rcloone	<code>--no-check-certificate</code>

```
aws --endpoint-url https://10.240.0.1:3000 --no-verify-ssl s3 ls
```

```
import boto3, urllib3
urllib3.disable_warnings(urllib3.exceptions.InsecureRequestWarning)
s3 = boto3.client("s3", endpoint_url="https://10.240.0.1:3000", verify=False, region_name="us-east-1")
```

Use a current boto3 release. With the older boto3/botocore packaged by SUSE Linux Enterprise 15 (boto3 1.26, botocore 1.29), `verify=False` was ignored and connections still failed certificate verification.

14.12 Supported S3 Features

Area	Supported	Not supported
Buckets	create, delete (empty), list, head, location (<code>us-east-1</code>), tagging, CORS configuration	versioning (<code>VersioningNotConfigured</code>), lifecycle, encryption, replication, logging, notifications, public access block (<code>NotImplemented</code>)

Area	Supported	Not supported
Objects	put (single PUT up to 5 GiB), get (incl. byte ranges), head, delete, multi-object delete, copy, user metadata (<code>x-amz-meta-*</code>), tagging, conditional put (<code>If-None-Match: *</code>), checksums (CRC32, SHA-256)	single PUT over 5 GiB (<code>EntityTooLarge</code> ; use multipart), S3 Select (<code>NotImplemented</code>), server-side encryption (headers ignored; below)
Multipart upload	create, upload part (minimum 5 MiB except the last part), complete, abort, list uploads	None
Listing	ListObjects v1 and v2, delimiters and prefixes, ListBuckets (lists the caller's own buckets)	None
Presigned URLs	GET and PUT, Signature V4, up to 7 days	longer than 7 days
Browse uploads (POST)	HTML form POST with a Signature V4 policy (below)	Signature V2 forms (<code>InvalidArgument</code>)
Access control	bucket owner, bucket policies, bucket ACLs once the owner turns them on (below)	object ACLs (<code>NotImplemented</code>), AWS grantee forms in grants (<code>InternalError</code>), Block Public Access (below)
Object lock	GOVERNANCE and COMPLIANCE retention, bucket default retention, legal hold (below)	versioned object lock: A lock bucket is not versioned (<code>VersioningNotConfigured</code>)

14.12.1 Bucket Policies

Once a bucket has a policy, the policy alone decides access, including the owner's. A policy that grants only public read locks the owner out of the bucket. Always include a statement for the owner:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {"Sid": "PublicRead", "Effect": "Allow", "Principal": "*",
     "Action": ["s3:GetObject"], "Resource": ["arn:aws:s3:::pubread/*"]},
    {"Sid": "OwnerAll", "Effect": "Allow", "Principal": {"AWS": ["alice"]},
     "Action": ["s3:*"], "Resource": ["arn:aws:s3:::pubread", "arn:aws:s3:::pubread/*"]}
  ]
}
```

With this policy anyone can read objects anonymously, anonymous listing is refused, and the owner keeps full access. Deleting the policy makes the bucket private again.

14.12.2 Bucket ACLs

ACLs are off on a new bucket: Its object ownership is `BucketOwnerEnforced`, and `PutBucketAcl` is refused with `AccessControlListNotSupported`. The bucket owner turns ACLs on by changing the object ownership to `BucketOwnerPreferred`:

```
aws s3api put-bucket-ownership-controls --bucket shared --ownership-controls 'Rules=
    [{ObjectOwnership=BucketOwnerPreferred}]'
aws s3api put-bucket-acl --bucket shared --acl public-read      # anyone can list the bucket
                        and read its objects
aws s3api put-bucket-acl --bucket shared --grant-read bob      # S3 user bob can list and read
aws s3api put-bucket-acl --bucket shared --acl private         # owner only again
aws s3api get-bucket-acl --bucket shared
```

To set an ACL when creating a bucket, pass the object ownership in the same request. Without it, the request is refused with `InvalidBucketAclWithObjectOwnership`, as on AWS:

```
aws s3api create-bucket --bucket shared2 --acl public-read --object-ownership
    BucketOwnerPreferred
```

What `PutBucketAcl` accepts:

Form	Accepted	Refused
Canned ACL (<code>--acl</code>)	<code>private</code> , <code>public-read</code> , <code>public-read-write</code>	every other canned ACL, <code>authenticated-read</code> included (<code>InvalidArgument</code>)
Grants (<code>--grant-read</code> , <code>--grant-write</code> , <code>--grant-full-control</code> , ...)	a bare access key, e.g. <code>--grant-read bob</code>	the AWS forms <code>id=...</code> , <code>uri=...</code> and <code>emailAddress=...</code> (<code>InternalError</code>)
Full policy (<code>--access-control-policy</code>)	<code>CanonicalUser</code> grantees; the ID is the S3 user's access key	<code>Group</code> grantees (<code>MalformedACLError</code>); use a canned ACL for public access

A grantee that is not an S3 user of the tenant fails with `InternalError` . The AWS CLI retries `InternalError` a few times before it reports it.

A bucket ACL covers the bucket's objects too. Objects have no ACLs of their own: `GetObjectAcl` and `PutObjectAcl` return `NotImplemented` , and an `x-amz-acl` header on an upload is ignored. A grant on the bucket therefore applies to every object in it. On AWS, a bucket ACL controls listing and object ACLs control reading.

Canned ACL on the bucket	Anonymous requests that then succeed
<code>public-read</code>	list the bucket, read every object
<code>public-read-write</code>	list the bucket, read every object, upload and delete objects

14.12.3 Public Access

A bucket owner can make the bucket public with a public ACL ([above](#)) or with a bucket policy for `"Principal": "*" (above)`. Every owner can, including a `user` -profile user that owns a bucket created for it with `panctl s3 buckets create --owner` .

The platform administrator has no setting that prevents this. There is no Block Public Access: `PutPublicAccessBlock` returns `NotImplemented` ([below](#)), and there is no tenant- or realm-wide equivalent. To limit public buckets, limit who owns buckets: Give the `privileged` profile only to users who may create buckets, and choose the `--owner` of buckets you create.

14.12.4 Browser Uploads (POST Object)

A browser can upload straight to a bucket with an HTML form that POSTs to

`https://<endpoint>/<bucket>` . The form must be signed with Signature Version 4. boto3 creates the signed fields:

```
post = s3.generate_presigned_post(
    "projects", "uploads/photo.jpg", ExpiresIn=3600,
    Conditions=[["starts-with", "$key", "uploads/"], ["content-length-range", 1, 10485760]],
)
# post["url"]: https://gw001.s3-acme.lab.example.com:3000/projects
# post["fields"]: key, policy, x-amz-algorithm, x-amz-credential, x-amz-date, x-amz-signature
```

The form sends every field of `post["fields"]` , then the `file` field last. To test it without a browser:

```
curl --cacert /etc/pki/acme-ca.crt \
  -F key=uploads/photo.jpg -F policy=<policy> -F x-amz-algorithm=AWS4-HMAC-SHA256 -F x-amz-credential=<credential> \
  -F x-amz-date=<date> -F x-amz-signature=<signature> -F file=@photo.jpg \
  https://gw001.s3-acme.lab.example.com:3000/projects
```

- A successful upload returns `204` . With the field `success_action_status` set to `201` , it returns `201` and an XML body. With `success_action_redirect` , it redirects (`303`) to that URL, with the bucket, key and ETag added to the query string.
- A file larger than the `content-length-range` maximum is refused with `EntityTooLarge` . A field that breaks another condition, such as `starts-with` , is refused with `AccessDenied` .
- `${filename}` in the key is not replaced with the uploaded file's name: The object is stored under the literal key.
- A form signed with Signature Version 2 (fields `AWSAccessKeyId` and `signature`) is rejected with `InvalidArgument: Bucket POST must contain a field named 'x-amz-algorithm'` .

14.12.5 Object Lock

Object lock works **without versioning**. A bucket created with object lock is not versioned:

`GetBucketVersioning` returns `VersioningNotConfigured`. On AWS, object lock always turns versioning on. Here each key has a single copy, and a lock protects that copy:

Request on a protected object	Result	On AWS (versioned)
PUT to the same key	refused: <code>AccessDenied</code> ("Access Denied because object protected by object lock.")	succeeds and adds a new version
DELETE	refused with the same error	succeeds and adds a delete marker
DELETE with GOVERNANCE bypass	deletes the object and its data	deletes only the named version

Applications that need versioned object lock cannot use these buckets, for example backup software that checks for versioning or writes new versions of locked keys.

```
aws s3api create-bucket --bucket records --object-lock-enabled-for-bucket
aws s3api put-object --bucket records --key r1.txt --body r1.txt \
  --object-lock-mode GOVERNANCE --object-lock-retain-until-date 2026-12-31T00:00:00Z
aws s3api get-object-retention --bucket records --key r1.txt
aws s3api put-object-legal-hold --bucket records --key r2.txt --legal-hold Status=ON
aws s3api put-object-lock-configuration --bucket records \
  --object-lock-configuration 'ObjectLockEnabled=Enabled,Rule={DefaultRetention={Mode=GOVERNANCE,Days=30}}'
aws s3api delete-object --bucket records --key r1.txt --bypass-governance-retention
```

Protections are enforced as soon as they are set:

Protection	Set with	Ends when
GOVERNANCE retention	<code>--object-lock-mode GOVERNANCE</code> and <code>--object-lock-retain-until-date</code> on upload or <code>put-object-retention</code>	the date passes; can be bypassed (below)

Protection	Set with	Ends when
COMPLIANCE retention	the same with COMPLIANCE	the date passes; cannot be bypassed
Legal hold	<code>put-object-legal-hold ... Status=ON</code>	<code>Status=OFF</code> ; cannot be bypassed
Bucket default retention	<code>put-object-lock-configuration</code> with a rule	the period passes, or the rule is removed

- **Bucket default retention** protects every object in the bucket, including objects uploaded before it was set. Its period counts from when the default was set, not from each object's upload. `get-object-retention` does not report it for an object without retention of its own (`NoSuchObjectLockConfiguration`). Remove the rule with `put-object-lock-configuration --object-lock-configuration ObjectLockEnabled=Enabled`.
- **Bypassing GOVERNANCE retention** needs the `--bypass-governance-retention` option (header `x-amz-bypass-governance-retention: true`) **and** a bucket policy that grants the caller `s3:BypassGovernanceRetention`. The bucket owner has no implicit right to bypass: With no such policy, the delete is refused. Once the policy grants it, the caller can also overwrite GOVERNANCE-protected objects without the option. The policy decides all access to the bucket, so include the owner's other permissions in it ([above](#)).

Object lock is not a tamper-proof (WORM) guarantee:

- `put-object-retention` accepts an earlier retain-until date, for COMPLIANCE as well as GOVERNANCE, and needs no bypass permission. Once the new date passes, the object can be deleted. A COMPLIANCE default retention can likewise be removed from the bucket.
- It applies only to S3 requests. Over the filesystem (see [14.13](#)), locked files can be changed and deleted with ordinary file permissions.

14.12.6 Not Implemented

The following bucket features are not available. Their API calls return `501 NotImplemented`, so a client that tries to configure them gets an error and nothing is applied.

Feature	API calls	Instead
Lifecycle rules	<code>PutBucketLifecycleConfiguration</code> , <code>GetBucketLifecycleConfiguration</code> , <code>DeleteBucketLifecycle</code>	Expire objects from a client job. Clean up unfinished multipart uploads with <code>ListMultipartUploads</code> and <code>AbortMultipartUpload</code> .
Bucket default encryption	<code>PutBucketEncryption</code> , <code>GetBucketEncryption</code> , <code>DeleteBucketEncryption</code>	None. See Server-Side Encryption .
Public access block	<code>PutPublicAccessBlock</code> , <code>GetPublicAccessBlock</code> , <code>DeletePublicAccessBlock</code>	None. See Public Access .
Event notifications	<code>PutBucketNotificationConfiguration</code> , <code>GetBucketNotificationConfiguration</code>	None.
Server access logging	<code>PutBucketLogging</code> , <code>GetBucketLogging</code>	Each gateway writes an access log on its node (see 14.14).
Replication	<code>PutBucketReplication</code> , <code>GetBucketReplication</code> , <code>DeleteBucketReplication</code>	None.

14.12.7 Server-Side Encryption

Data is stored unencrypted. The `x-amz-server-side-encryption*` headers on uploads (SSE-S3, SSE-KMS and SSE-C, for example `aws s3 cp --sse AES256`) are **accepted and ignored**: The upload succeeds, the file under the prefix holds the plain data, and an SSE-C object can be read back without its key. Do not rely on them.

14.13 S3 and File Access to the Same Data

S3 data is ordinary PanFS data under the tenant's prefix:

S3	Filesystem (prefix /s3data)
bucket projects	directory /s3data/projects
object projects/docs/a.txt	file /s3data/projects/docs/a.txt
key prefix docs/	directory docs

Ownership. Files and directories created through S3 get the first of these that applies:

1. object metadata `x-amz-meta-uid` , `x-amz-meta-gid` , `x-amz-meta-mode` on the upload (uid/gid 0 is not accepted);
2. the S3 user's uid/gid (set when the user is created; see [14.7](#));

and mode `--obj-mode` (files) and `--dir-mode` (directories). A bucket created with `panctl s3 buckets create` is owned by its owner's uid/gid.

Reading the data as files. A host attached to the tenant subnet (see [14.4.3](#)) can mount the tenant volume. The director list comes from `panctl subnets mount-info acme-net --tenant acme` :

```
mount -t panfs -o "rmlist=(10.240.0.1;10.240.0.2;10.240.0.3)" panfs://10.240.0.1/s3data
/mnt/acme
ls -l /mnt/acme/projects/docs
```

```
-rw-r--r-- 1 2001 2001 11 Oct 12 11:40 small.txt
-rw----- 1 3001 3001 1 Oct 12 11:40 owned.txt      # uploaded with x-amz-meta-uid/gid/mode
```

Files written over the filesystem: - Files and subdirectories written inside an existing bucket appear as objects. Objects that did not come through S3 have no ETag. - A new top-level directory under the prefix does **not** become a bucket.

14.14 Monitoring and Logs

Metrics. With `--metrics-enabled true` (default), the gateways publish request counts, errors and throughput. VDURA does not ship Grafana. To view them, import the **VersityGW S3 Metrics** sample dashboard into a Grafana instance that you run yourself (see the PanFS Metrics Reference Guide). - In

your Grafana, open the dashboard (path `/d/panfs-s3/versitygw-s3-metrics`). - Select the tenant by its **UUID** (`panctl tenants list`), then the instance. - Panels: Request rate (success and failed), error ratio, throughput, objects created and removed, failures by action and HTTP status.

Access log. With `--access-log true` (default), every request is logged on the node that served it, in `/var/log/versitygw/access_<pid>.log` , in S3 server access log format. With the experimental aggregator on, its periodic health checks also appear there, as `403 AccessDenied` lines for `/` .

The access log setting takes effect after the instances are restarted.

14.15 Day-to-Day Operations

Task	Commands
Stop all S3 in the realm	<code>panctl s3 enabled set false</code> : Every gateway stops; tenant settings and instance states are kept
Restart all S3 in the realm	<code>panctl s3 enabled set true</code> : Instances that were started come back
Apply a setting that needs a restart	<code>panctl s3 instances stop-all --tenant acme</code> , then <code>start-all</code>
Add capacity or availability	create the same or a new iid on another director, then start it; update DNS records
Change the prefix	<code>instances delete-all --yes</code> , <code>set-config --enabled false</code> , <code>set-config --prefix /new</code> , <code>set-config --enabled true</code> , create and start the instances again
Rotate a user's secret	<code>panctl fs-users rotate-s3-secret <id> --tenant acme</code>
Block a user	<code>panctl fs-users update <id> --tenant acme --s3-profile disabled</code>
Replace the certificate	upload a new one (see 14.9.1) or <code>panctl s3 tls renew</code> for a realm-generated one
Check instance state	<code>panctl s3 instances list --tenant acme (OBSERVED)</code>

When realm S3 is turned off while tenants still have S3 enabled, `panctl` warns:

```
WARNING:                1 tenant(s) still have S3 enabled at the tenant level; realm-wide
disable does not turn them off.
TENANTS WITH S3 ENABLED:  acme
```

14.16 Removing S3 from a Tenant

1. Remove all objects and buckets.

- Users empty their buckets over S3. They can also delete their own buckets over S3.
- The administrator deletes the remaining empty buckets with `panctl s3 buckets delete <name> -tenant acme --yes`.

2. Delete the instances, then the configuration:

```
panctl s3 instances delete-all --tenant acme --yes
panctl s3 instances list --tenant acme --wait-for 'it == null' -o 'jsonpath=
    {instances[*].desired.instance_id}'
panctl s3 set-config --tenant acme --enabled false
panctl s3 delete-config --tenant acme --yes
```

3. Delete the S3 users with `panctl fs-users delete <fs_user_id> --tenant acme --yes`.

14.17 S3 on the Platform Tenant

The built-in platform tenant (short name `platform`) can serve S3 with the same commands, using `--tenant platform`.

Differences from a dedicated tenant:

- **No subnet or client attachment.** Endpoints are the nodes' own realm addresses (e.g. `https://10.128.1.0:3000`), reachable from any host on the realm network.

- **Volume.** The volume is created without `--as-tenant-admin`, at the realm root. The platform tenant needs a storage set assignment first:

```
panctl tenants assign-storagesets --tenant platform --storageset-id 1 --quota 100GiB
panctl volumes create --name /s3platform --storageset-id 1 --hard-quota 50GiB
panctl s3 create-config --tenant platform --prefix /s3platform --enabled true
panctl fs-users create --tenant platform --name ops --s3-profile privileged
panctl s3 instances create pgw01 b022963e7165c9 --tenant platform
panctl s3 instances start pgw01 --tenant platform
panctl s3 buckets create ops-data --owner ops --tenant platform
```

- **DNS.** Named access mode `panfs` (experimental) is available (see [14.10.3](#)).

14.18 Troubleshooting

Message	Cause and fix
S3 is disabled realm-wide; run <code>`panctl s3 enabled set true`</code> before enabling tenant-level S3.	Turn on the realm switch first (see 14.5.1).
name: must match <code>^[^\s/?#%]+\$</code> or uid ... less than or equal to 4294967294 (fs-users create)	S3 user names cannot contain whitespace, <code>/</code> , <code>?</code> , <code>#</code> or <code>%</code> ; uid and gid go up to 4294967294.
tenant "default" not found: no tenant has that short name	The platform tenant's short name is <code>platform</code> .
Storageset assignment 1 not found (volume create)	Assign the storage set to the tenant first (see 14.4.1).
The tenant holds no issued subnet address, so a volume created for it could not be reached (volume create)	Create the tenant subnet first (see 14.4.2).
S3 configuration not found for this tenant	Run <code>panctl s3 create-config</code> (see 14.5.2).
S3 configuration already exists for this tenant	Use <code>panctl s3 set-config</code> to change it.

Message	Cause and fix
Default prefix must be configured before enabling S3	Pass <code>--prefix</code> with or before <code>--enabled true</code> .
S3 prefix <code>'...'</code> does not reference an existing volume of this tenant	Create the volume first (see 14.4.4) or use a directory inside one.
Buckets of the old prefix still served after a prefix change	Delete and recreate the instances (see 14.5.3).
Tenant has no subnet ... (instance create)	Create the tenant subnet first (see 14.4.2).
Cannot change the S3 prefix while S3 is enabled for this tenant; disable it first	Stop instances, <code>--enabled false</code> , change the prefix, enable again.
Cannot disable S3 for this tenant while instances are desired=started; stop them first	<code>panctl s3 instances stop-all --tenant <t></code> .
Cannot delete S3 config for this tenant while any instances still exist	<code>panctl s3 instances delete-all --tenant <t> --yes</code> , wait for the list to empty.
Tenant already has instance gw001 on node ... At most one instance per tenant per node	Use another node.
instance_id: must match <code>^[A-Za-z0-9]{5}\$</code>	Use exactly 5 letters or digits.
confirm=true required (instance delete)	The iid is on several nodes: Add <code>--confirm-fanout</code> or <code>--dataserial</code> for one node.
delete-all did not complete within 60s; ... Stuck instances: ... desired=deleting	An instance on a storage node cannot be removed (see 14.6.1). The other instances are deleted. <code>panctl</code> gives up first (context deadline exceeded) unless you pass <code>--request-timeout 120s</code> .
No PanFS user in this tenant holds access_key <code>'...'</code> . Create one first.	<code>--owner</code> must be an S3 user's access key (<code>panctl fs-users list</code>).
Tenant has no S3 instances. Bucket create requires at least one running instance	Create and start an instance (see 14.6), then retry.

Message	Cause and fix
Tenant has no running S3 instance (1 configured, none started)	Start an instance and wait for <code>running</code> , then retry.
The S3 gateway refused to create bucket "..." (InvalidBucketName)	Use a valid S3 bucket name (see 14.8).
Bucket "..." is not empty. Remove its objects over the S3 API before deleting it.	Delete the objects first.
S3 bucket "..." not found. (panctl delete)	The tenant has no bucket with that name (<code>panctl s3 buckets list</code>).
<code>dns.mode=panfs</code> is not supported for this tenant	Use <code>external</code> (see 14.10.2).
<code>tls.certificate</code> is missing SAN(s) the realm publishes with <code>dns.mode=panfs ...</code>	Add the listed names to the certificate or use <code>external</code> (see 14.10.3).
Client: SSL validation failed ... certificate verify failed	Realm-generated certificate: Skip verification (see 14.11.4). Tenant-provided: Give the client the CA (<code>--ca-bundle</code>), include intermediates in the certificate, and add IP SANs when connecting by IP.
Client: InvalidAccessKeyId	Wrong access key or the user is <code>disabled</code> .
Client: SignatureDoesNotMatch	Wrong or rotated secret key.
Client: InvalidLocationConstraint on bucket create	Set the client region to <code>us-east-1</code> .
Client: AccessDenied on a bucket with a policy	The policy does not grant the caller, possibly the owner (see 14.12.1).
Client: AccessControlListNotSupported (<code>PutBucketAcl</code>)	ACLs are off on the bucket: Set its object ownership to <code>BucketOwnerPreferred</code> first (see 14.12.2).
Client: InternalError on <code>PutBucketAcl</code>	Give grantees as bare access keys of existing S3 users, not <code>id=</code> / <code>uri=</code> / <code>emailAddress=</code> (see 14.12.2).

Message	Cause and fix
Client: AccessDenied ... object protected by object lock	The object has retention, a legal hold, or the bucket has a default retention (see 14.12.5).
Client: InvalidRequest ... authorization mechanism ... not supported on a presigned URL	Create the URL with Signature V4 (boto3: <code>signature_version="s3v4"</code>).
Client: Connect timeout on endpoint URL	The client host is not attached to the tenant subnet (see 14.4.3).
Instance stays OBSERVED unknown after start	Check it was placed on a director, not a storage node (see 14.6.1).

15. Kubernetes CSI Driver for PanFS

This chapter explains how to configure and use the Container Storage Interface (CSI) with PanFS.

15.1 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.

15.2 Installation, Configuration, and Diagnostics

The most recent files and documentation can be found at the links below.

Important: A GitHub Personal Access Token (PAT) is required for Kubernetes to authenticate with the GitHub container registry. Follow the instructions described in [GitHub's documentation](#) to create a PAT with at least read permissions.

Basic information to get started: [panfs-container-storage-interface-oss](#)

Resource	Link
PanFS CSI Driver Architecture Overview	Overview.md
CSI Driver Installation - Quick Start	README quick-start section
PanFS CSI Driver & StorageClass Deployment Guide	deploy/k8s
PanFS CSI Driver Diagnostic Commands Reference	Diagnostic.md
PanFS CSI Driver Upgrade Guide	Upgrade.md
Kernel Module Management (KMM) Guide	KMM.md
PanFS CSI Driver Volume Encryption Guide	Encryption.md
SELinux Configuration Reference for Kubernetes CSI Drivers	SELinux.md
PanFS CSI Driver Usage Guide	Usage-Guide.md

16. Troubleshooting

Symptom	Likely Cause/Fix
Browser shows a certificate warning on first visit.	Expected: the realm uses a self-signed certificate by default. Click through it once per browser; it will not reappear on that browser.
System Performance chart (Dashboard) or an Events table spins forever and never loads rows.	No telemetry/events have been generated yet in that realm; not a UI bug. The Audit Log tab (a different feed) on the same page loads correctly, confirming the page itself is working.
A tenant shows <code>Setup Incomplete</code> , or Create Volume fails with "The tenant holds no issued subnet address..."	The tenant is missing a setup step. Click the status chip to see the Finish setup checklist, then assign a StorageSet and create a network configuration for the tenant (see 9.2.2 and 9.2.4).
"Failed to delete tenant: Cannot delete tenant while it still owns subnets".	Delete the tenant's network configuration on its Networking tab (and any volumes), then delete the tenant (see 9.5.5).
"Failed to create the network configuration: The usable-address window holds fewer addresses than the realm primary pool..."	The network is too small for the realm's nodes. Widen the Network CIDR or the Address Range (see 9.2.4).
A page shows "404 Page not found" ("The page you're looking for doesn't exist or has been moved.").	The address does not match a Vantage page, for example a mistyped URL or an old bookmark. Click Back to Dashboard , then navigate with the sidebar.
"Add Storage" on a StorageSet says "No unassigned nodes available."	Every node in the realm is already claimed by a StorageSet. Free up nodes (remove them from their current StorageSet) or add new hardware/VPODs to the realm before retrying.

Symptom	Likely Cause/Fix
Signed in to Vantage but PanActive Manager still asks you to sign in.	Expected: the two UIs keep separate sessions. Sign in to PanActive Manager with the same credentials; it will not disturb your Vantage session.
Hardware sidebar item does nothing when clicked.	Expected in a virtual realm: it's disabled by design ("Hardware is unavailable for virtual realms"). It should be active on physical VDURA hardware.
A burst of critical/warning alerts (e.g. Director-Node software processes reported down) appears in the top bar right after adding storage to a StorageSet, then clears on its own within a minute or two.	Normal, transient churn while VPODs activate and Director-Node services restabilize; not necessarily a real problem. Confirm it clears rather than assuming it's benign every time.
"Add StorageSet"/Create Volume won't let a tenant use a StorageSet.	The built-in platform tenant gets implicit, unlimited access to every StorageSet; any other tenant needs an explicit + Add StorageSet action with a Hard Quota on its own detail page first (see 9.2) before it can place volumes there.

17. VDURA System Error Messages

This chapter provides a numerical list of VDURA system error messages.

Error messages are categorized as follows:

Category	Description
Info	Error recovery messages or configuration changes are typically logged to the Events page (see 5.2 Events Panel).
Warning	Developing problems that might impair system availability or data integrity are logged to the Events page and emailed to system administrators.
Critical	Problems that are an imminent risk to system availability or data integrity are logged to the Events page and emailed to system administrators.

The following table lists error messages and PanActive Manager error handling. %%SERIAL%% , %%SHELFID%% , and similar tokens are placeholders filled in with the actual serial number, shelf ID, or other identifying value at runtime.

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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.	

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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 lag 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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.	
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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.	
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.

Error Code	Category	Error Description	Recommended Action
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.
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.

Error Code	Category	Error Description	Recommended Action
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 %%CLOCKSKREW%% seconds exceeds %%MAXCLOCKSKREW%% 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 %%NUMMSTHRESH OLD%% 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 %%EMAILADDRESS%%S%% 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_NOD E_DRIVES_MISMATC H	Place drives into proper location.
9903	Warning	PAN_MGMT_EV_NOD E_DRIVES_UNEXPEC TED	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.

This guide documents Vantage as of version 12.0.0. Screen labels, tooltips, and available options may change slightly in later releases.