Vantage User Guide

The web-based management console for VDURA storage realms

Documentation for PanFS 12.0.0

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

Parts of this product are derived from Berkeley Database software, licensed from Sleepycat Software.

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

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

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

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

© UNIX System Laboratories, Inc.

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

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

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. 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 12.0.0-2026-11-12-11:46:23

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 Dialog 24
6.1.4 Add Storage Dialog 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 29
6.3.1 Volumes List Page 30
6.3.2 Volume Detail Page 31
6.3.3 Delete Volume Confirmation 34
6.3.4 Create Volume Page 35

6.4 Creating and Managing Volumes 38
6.4.1 Create a Volume 39
6.4.2 Delete a Volume 39

7. Realm Nodes 40
7.1 Director Nodes Table 40
7.2 Storage Nodes Table 41
7.3 Node Detail Page 43

8. Hardware 47
8.1 Hardware List Page 47
8.2 Chassis Detail Page 48
8.3 Parent Field on Realm Node Detail 52

9. Tenants 54
9.1 Tenants List Page 54
9.2 Tenant Detail Page 55
9.2.1 Delete Tenant Confirmation 61
9.2.2 Finishing Tenant Setup 62
9.2.3 Networking Tab 64
9.2.4 Create Network Configuration Page 67

9.3 Create Tenant Page 70
9.4 Acting as a Tenant (Context Switching) 71
9.4.1 Give a Tenant Its Own Isolated Storage 77

9.5 Creating and Managing Tenants 77
9.5.1 Create a Tenant 78
9.5.2 Assign a StorageSet to a Tenant 78
9.5.3 Modify a Tenant's StorageSet Assignment 78
9.5.4 Modify a Tenant 78
9.5.5 Delete a Tenant 78
9.5.6 Create a Network Configuration for a Tenant 79

10. Identity and Access 80
10.1 Users List Page 80
10.2 User Detail Page 81
10.2.1 Edit Email Dialog 82
10.2.2 Reset Password Dialog 83
10.2.3 Manage Roles Dialog 84
10.2.4 Delete User Confirmation 85
10.3 Create User Page 85
10.4 Creating and Managing Users 86
10.4.1 Create a User 86
10.4.2 Change a User's Email 87
10.4.3 Reset a User's Password 87
10.4.4 Change a User's Roles 87
10.4.5 Delete a User 87
10.5 Service Accounts List Page 88
10.6 Service Account Detail Page 89
10.6.1 API Tokens Tab 90
10.6.2 Issue Token Dialog 91
10.6.3 Token Issued Dialog 92
10.6.4 Manage Roles Dialog 93
10.6.5 Delete Service Account Confirmation 94
10.7 Create Service Account Page 94
10.8 Creating and Managing Service Accounts 95
10.8.1 Create a Service Account 95
10.8.2 Issue an API Token 95
10.8.3 Revoke an API Token 96
10.8.4 Change a Service Account's Roles 96
10.8.5 Delete a Service Account 96
11. Account, Preferences, and Help 97
11.1 Account Menu and My Account 97
11.2 Theme 100
11.3 Get Help 100
11.4 About 102
11.5 PanActive Manager 103
12. S3 104
12.1 Concepts 104

12.2 Logging In with panctl 105
12.3 Deployment Options 106
12.3.1 Dedicated Tenant (Recommended) 106
12.3.2 Platform Tenant 106
12.3.3 Choices Within a Tenant 106

12.4 Prerequisites 107
12.4.1 Tenant and Storage 107
12.4.2 Tenant Subnet 107
12.4.3 Attaching Client Hosts to the Tenant Subnet 108
12.4.4 Volume 108

12.5 Enabling S3 109
12.5.1 Realm Switch 109
12.5.2 Tenant Configuration 109
12.5.3 Settings Reference 110

12.6 Gateway Instances 111
12.6.1 Choose Nodes 111
12.6.2 Create and Start 112
12.6.3 Endpoints 113
12.6.4 States and Health 113
12.6.5 Stop, Start, Delete 114

12.7 S3 Users 114
12.8 Buckets 115
12.9 TLS Certificates 116
12.9.1 Tenant-Provided Certificate (Recommended) 116
12.9.2 Realm-Generated Certificate (Default) 119

12.10 Named Access (DNS) 120
12.10.1 Named Access with Your Own DNS (External) 120

12.11 Client Access 121
12.11.1 AWS CLI 122
12.11.2 Python (boto3) 123
12.11.3 rclone 123
12.11.4 Connecting to a Realm-Generated Certificate 124

12.12 Supported S3 Features 125
12.12.1 Bucket Policies 126
12.12.2 Bucket ACLs 126

12.12.3 Public Access 127
12.12.4 Browser Uploads (POST Object) 128
12.12.5 Object Lock 129
12.12.6 Not Implemented 131
12.12.7 Server-Side Encryption 131

12.13 S3 and File Access to the Same Data 132
12.14 Monitoring and Logs 132
12.15 Day-to-Day Operations 133
12.16 Removing S3 from a Tenant 134
12.17 S3 on the Platform Tenant 134
12.18 Troubleshooting 135

13. Multi-Tenant Networking (VLAN and VXLAN) 139
13.1 Summary 139
13.2 Network Models 140
13.2.1 VXLAN Overlay and Platform Underlay 140
13.2.2 VLAN-Based Tenants 141
13.2.3 VLAN-Based Tenants with a VXLAN Underlay 141
13.2.4 Topology Matrix 143

13.3 How PanFS Uses VLAN and VXLAN 144
13.4 Configuring Segments in Vantage 145
13.5 Configuring Network Segments with panctl and PanAPI 146
13.6 Changing the Realm Segment Type 148
13.7 Physical Network Requirements 149
13.8 MTU for Tenant Networks 150
13.9 Switch Configuration for Tenant Networks 150
13.9.1 Realm with VXLAN Tenant Networks (the Default) 151
13.9.2 Adding Peers with the Attachments Endpoint 151
13.9.3 Realm with VLAN (802.1Q) Tenant Networks 152

13.10 Setting Up a DirectFlow Client on a Tenant Network 152
13.10.1 Values a DirectFlow Client Needs 152
13.10.2 DirectFlow Client on VXLAN Tenant Networks 153
13.10.3 DirectFlow Client on VLAN Tenant Networks 154
13.10.4 Mounting PanFS over the Tenant Network 154
13.10.5 Restricting the DirectFlow Client to the Tenant Network 155
13.10.6 Checking a DirectFlow Client 155

13.11 RDMA and Tenant Isolation 155

13.12 Switch Checklist 156
13.13 Background: Tagged VLANs and VXLAN 157

14. Troubleshooting 158

1. Introduction

This chapter covers Important Note Regarding VDURA System GUI Options, What Vantage Is, How Vantage Relates to panctl and PanAPI, and Audience and Scope of This Guide.

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, with 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.
- The **Audit Log** panels in Vantage (see 5.3, 9.2, and 10.2) show the same records as `panctl audit list`.
- The **Events** panels in Vantage 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 `panctl` capability, and `panctl` 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 by reviewing health, managing StorageSets and volumes, and managing tenants and users. It walks through every page, panel, button, and dialog in the UI, explaining what each screen is telling you and what each control does.

Four topics that apply to the whole system are covered only in the *VDURA Data Platform User Guide* and are not repeated here. They are the chapters "Security Enhanced Linux in VDURA Systems", "End-to-End Encryption", "Kubernetes CSI Driver for PanFS", and "VDURA System Error Messages".

2. Getting Started

This chapter covers Accessing the UI, Signing In, and Layout Overview.

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.

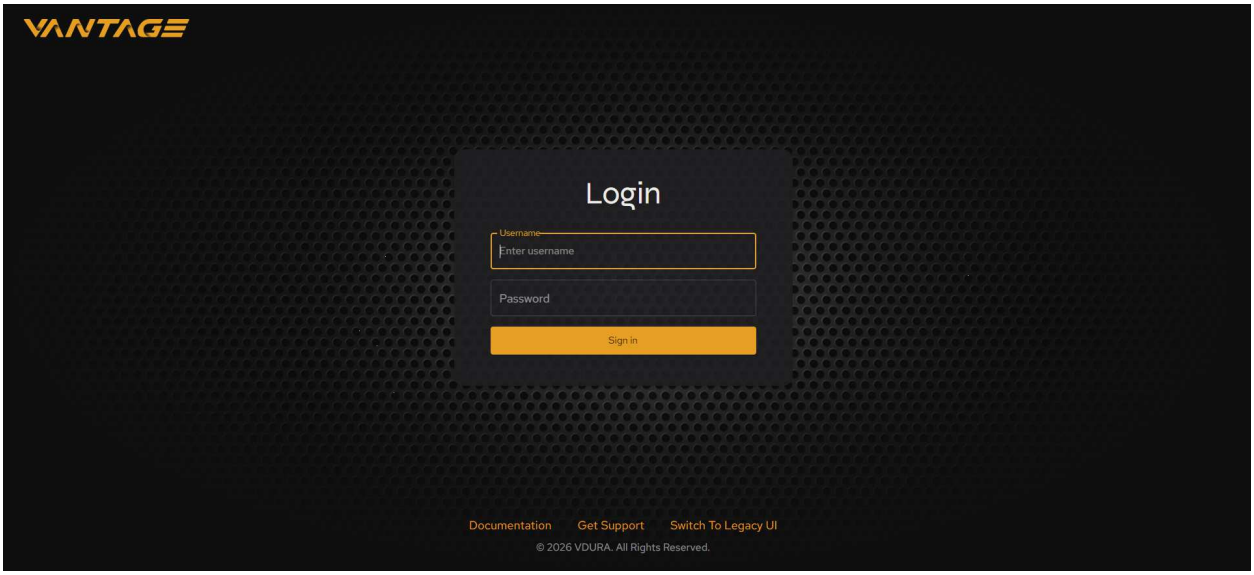


Figure 2.1. The Login page: Vantage wordmark, centered Username/Password form, Sign in button, and footer links.

Vantage uses the same credentials as `panctl` 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 (for example `realm1.example.com`).
- User drop-down menu: Displays the username currently logged in (e.g. "AD" for the built-in `admin` user). Click the drop-down 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.
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 (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).

Icon	Function
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.

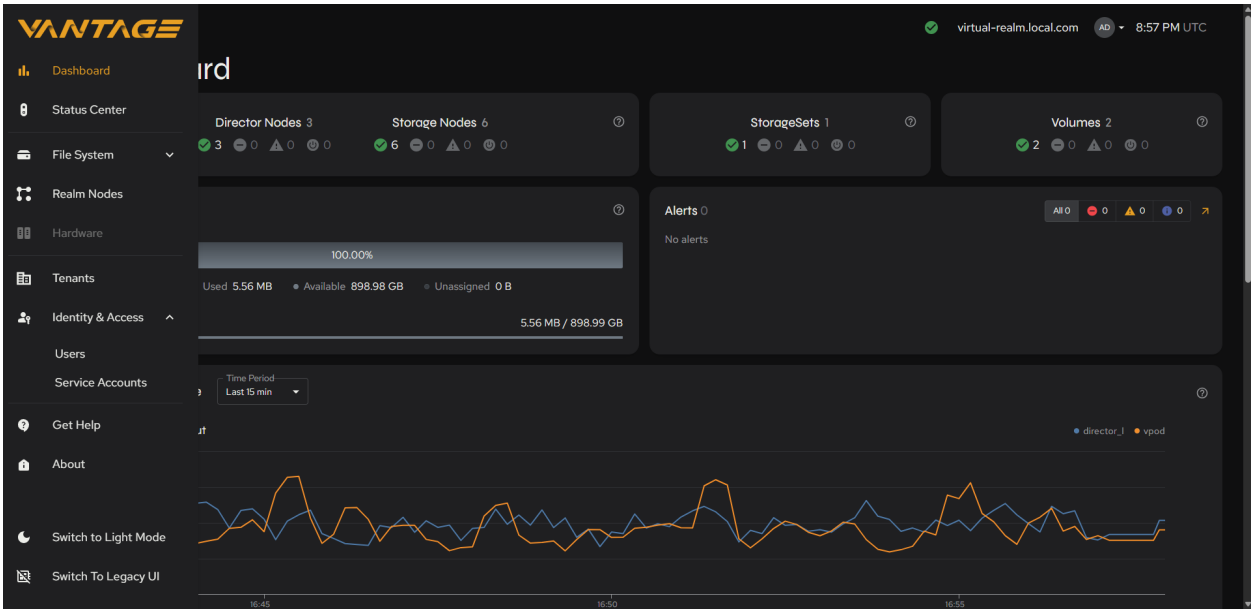


Figure 2.2. Full left sidebar expanded with text labels, including the greyed-out Hardware item and the open Identity & Access menu.

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 and includes a **Rows per page** drop-down menu, 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 (they open on click, not on hover).

3.4 Confirm Before Destroy

Every destructive action in Vantage (deleting a StorageSet, a volume, a user, or a service account) requires a confirmation dialog 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 dialog's own confirm button executes the action immediately, so don't treat the dialog 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.

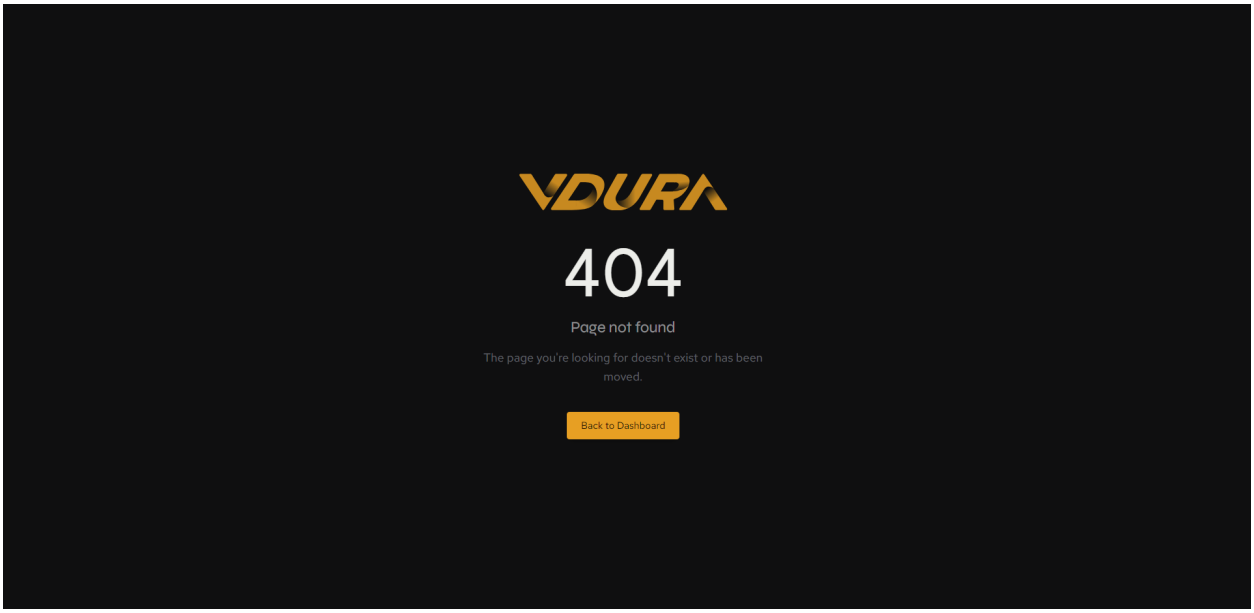


Figure 3.1. The Page not found screen.

4. Dashboard

The Dashboard is the Vantage 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.

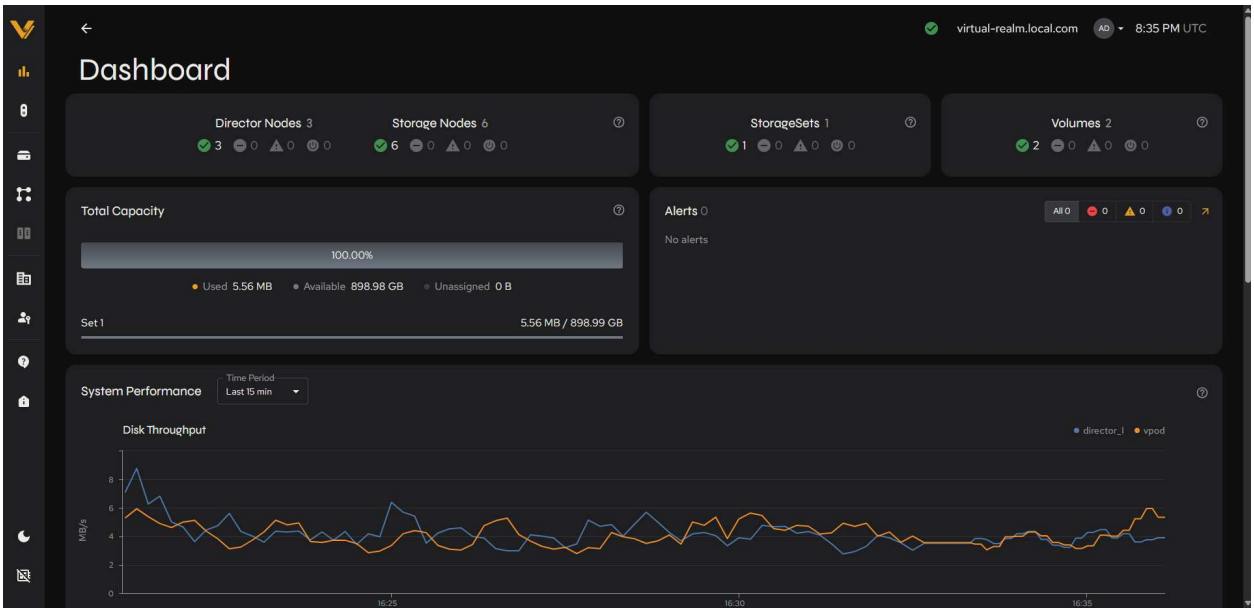


Figure 4.1. The whole Dashboard page, with all four card rows visible.

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.

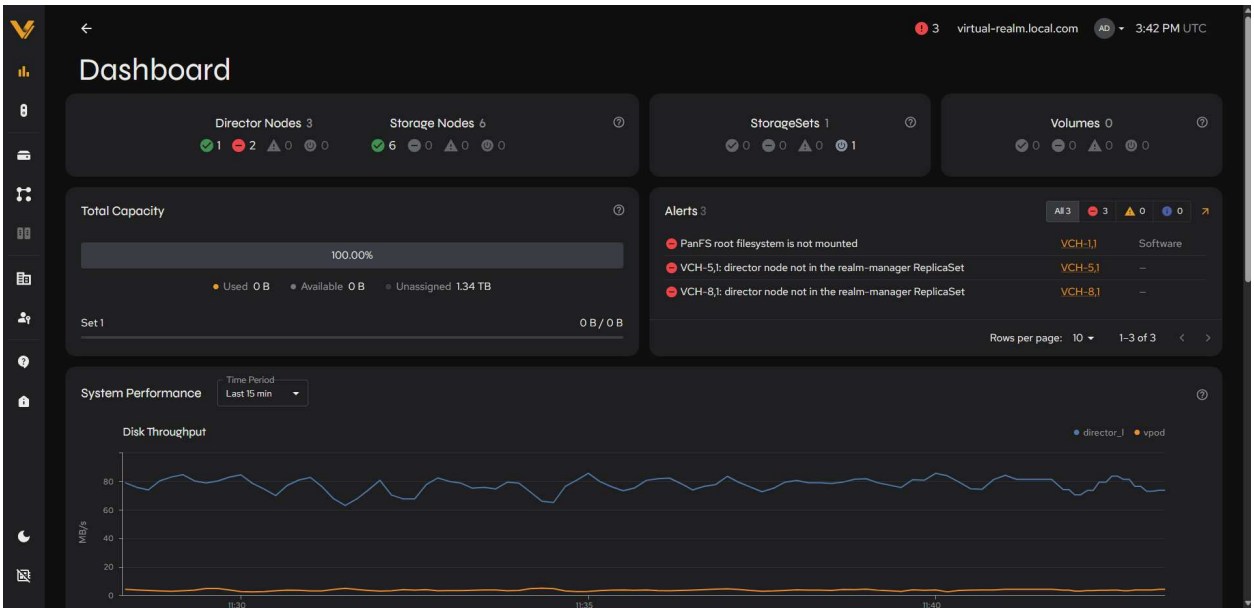


Figure 4.2. Dashboard on a freshly initialized realm: StorageSets pending, Volumes 0, capacity fully Unassigned.

The top bar shows a red alert badge with the alert count, and the Alerts panel lists each active alert. 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** drop-down menu (`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.” 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.

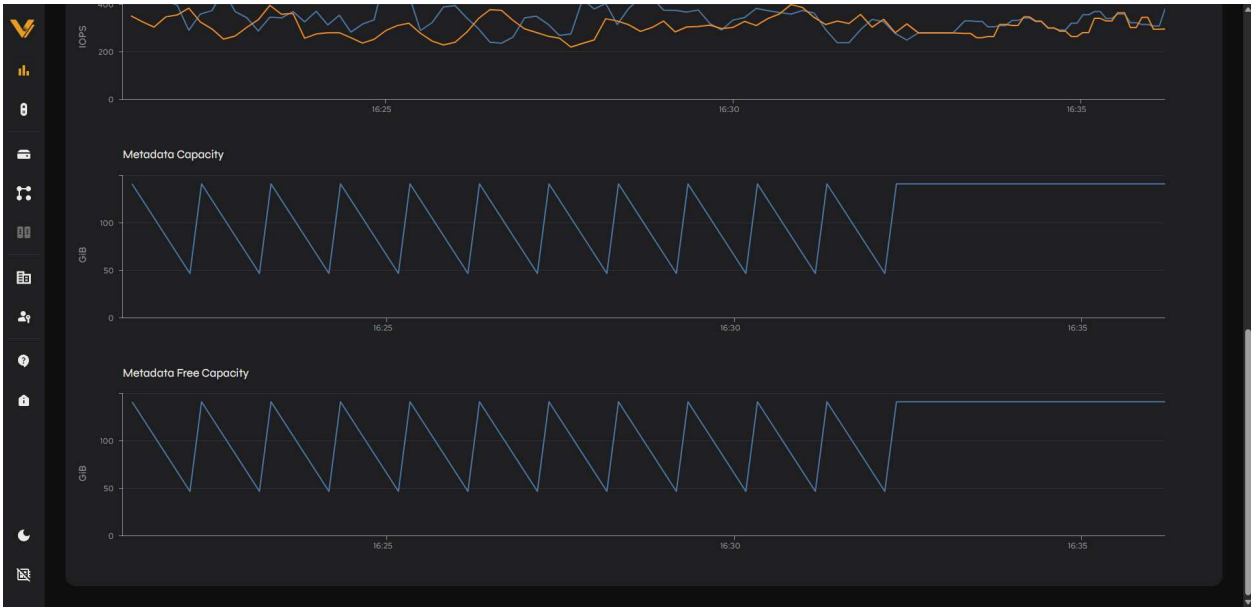


Figure 4.5. The Metadata Capacity and Metadata Free Capacity charts at the bottom of the System Performance card.

> **Note:** A chart that has no data for the selected period shows “No data to display”. On a newly built realm, a chart’s line starts at its first recorded sample, so the earlier part of the time axis is blank.

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**.

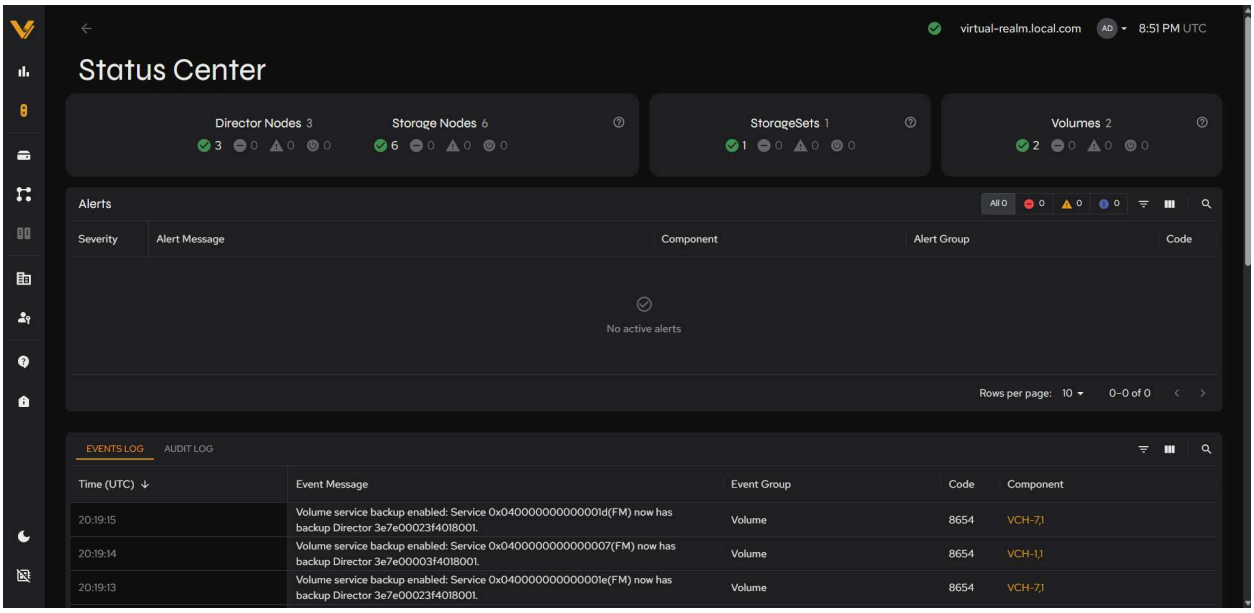


Figure 5.1. Full Status Center page: summary cards + Alerts panel + Events Log/Audit Log panel.

5.1 Alerts Panel

The following table describes each column.

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**.

The toolbar at the top right of the panel has 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). This panel lists the realm's events with 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.

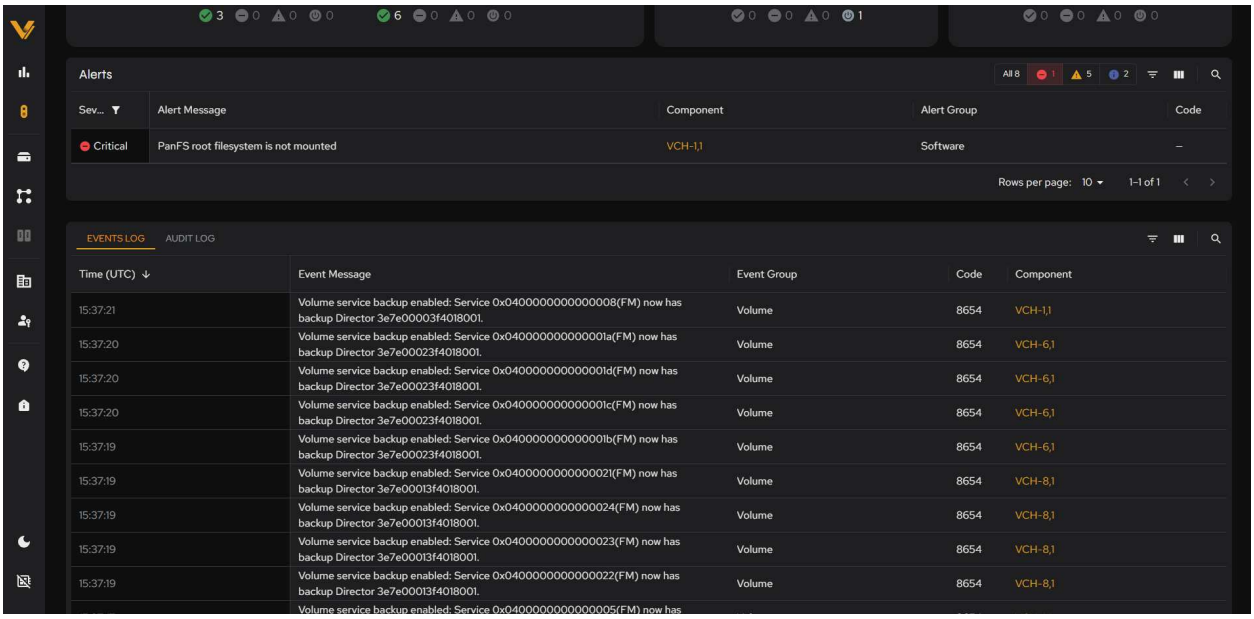


Figure 5.2. Alerts panel filtered to a Critical alert, with the Events Log tab beneath it populated with Volume-component rows.

> **Note:** When the realm has recorded no events, this panel shows an inbox icon and "No events recorded", with no rows. The Audit Log tabs and panels (5.3, and on the Tenant and User detail pages, 9.2 and 10.2) always have rows because viewing a page is itself an audited action.

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.

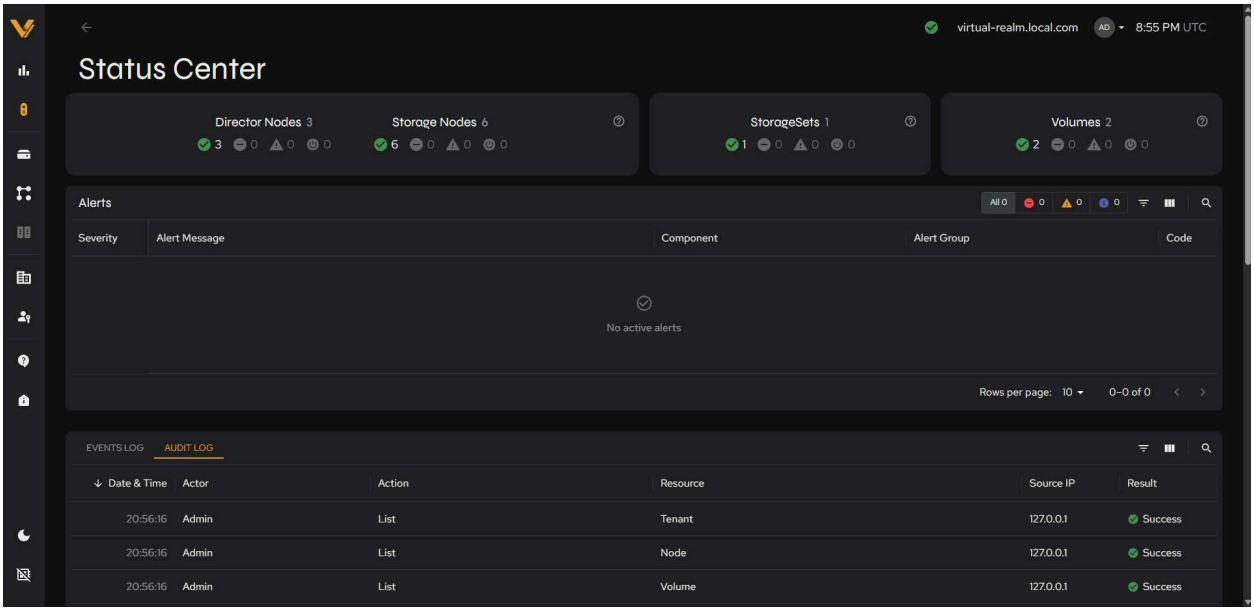


Figure 5.3. The Audit Log tab of the Status Center bottom panel.

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

This section covers StorageSets List Page, StorageSet Detail Page, Edit Spare VPODs Dialog, Add Storage Dialog, Delete StorageSet Confirmation, and Create StorageSet Page.

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.

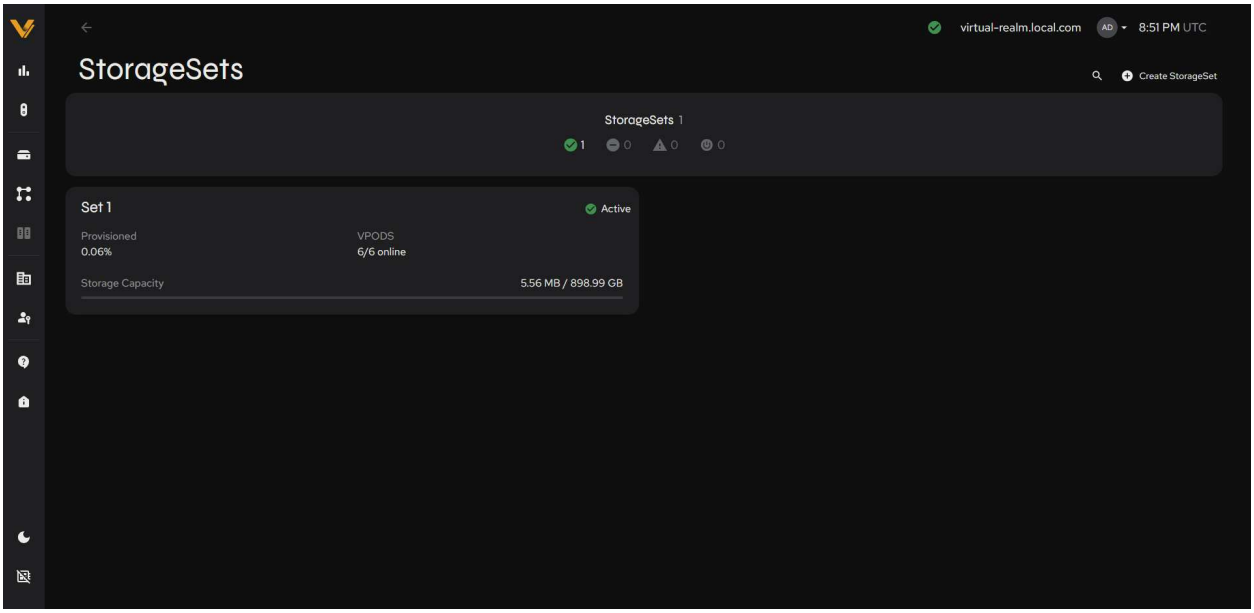


Figure 6.1. StorageSets list page with a single active StorageSet card.

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)

Card element	Description
Provisioned	Percentage of capacity provisioned
VPODS	`x/y online`
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.

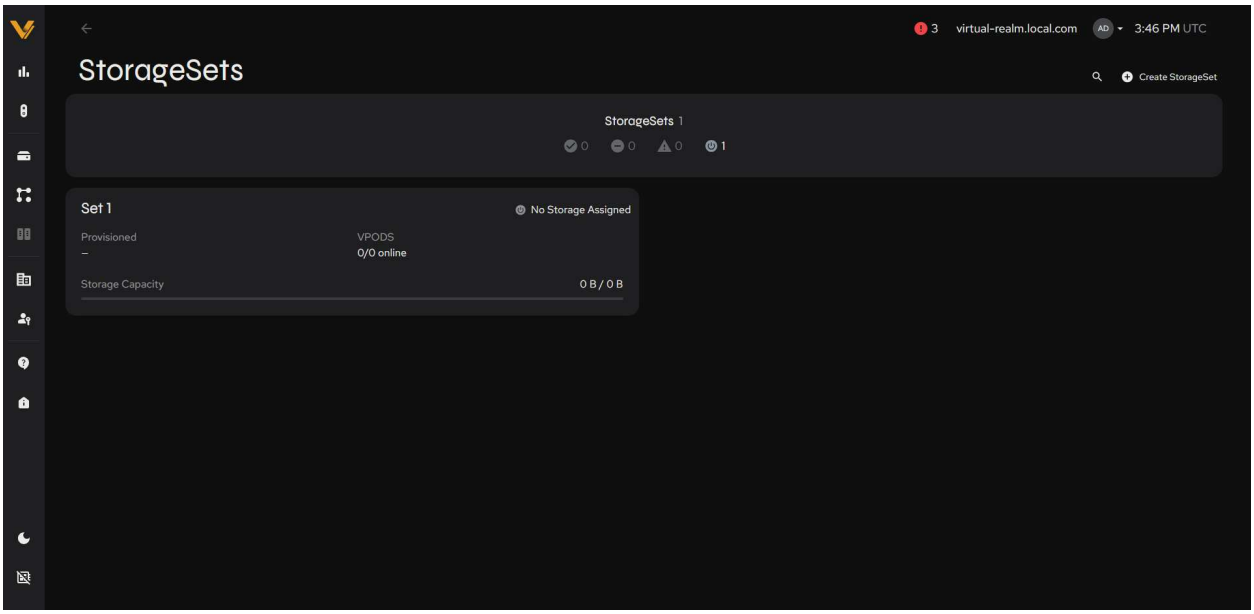


Figure 6.2. StorageSets list on a freshly initialized realm: a single empty shell, no active StorageSets yet.

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.

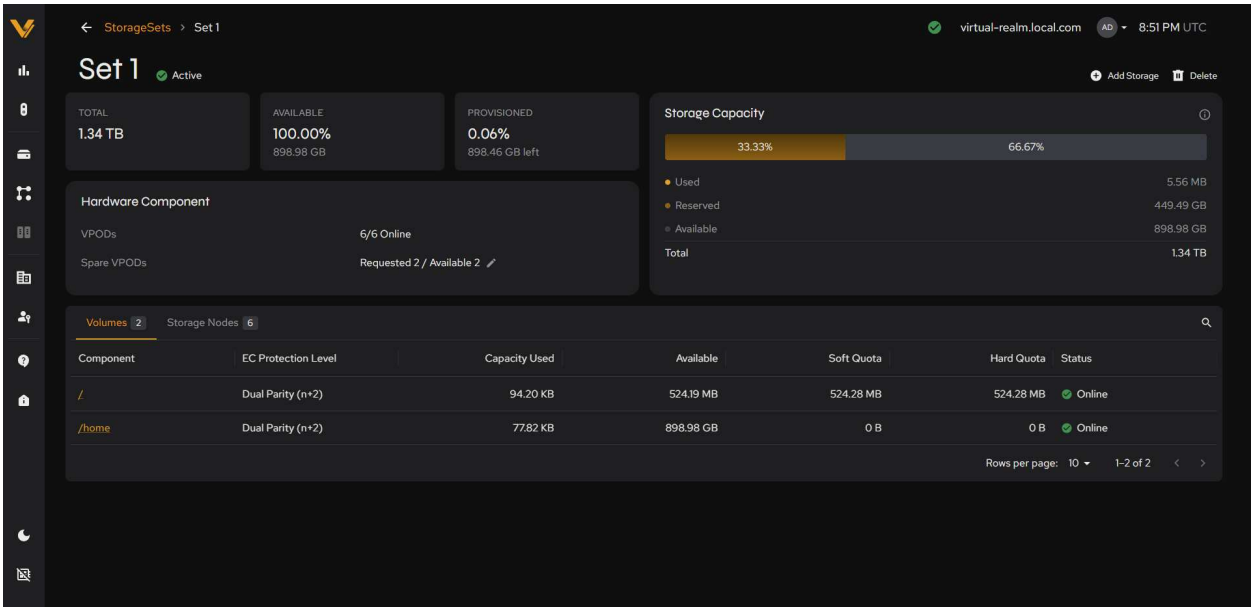


Figure 6.3. Full detail page for the active StorageSet, Volumes tab selected.

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)

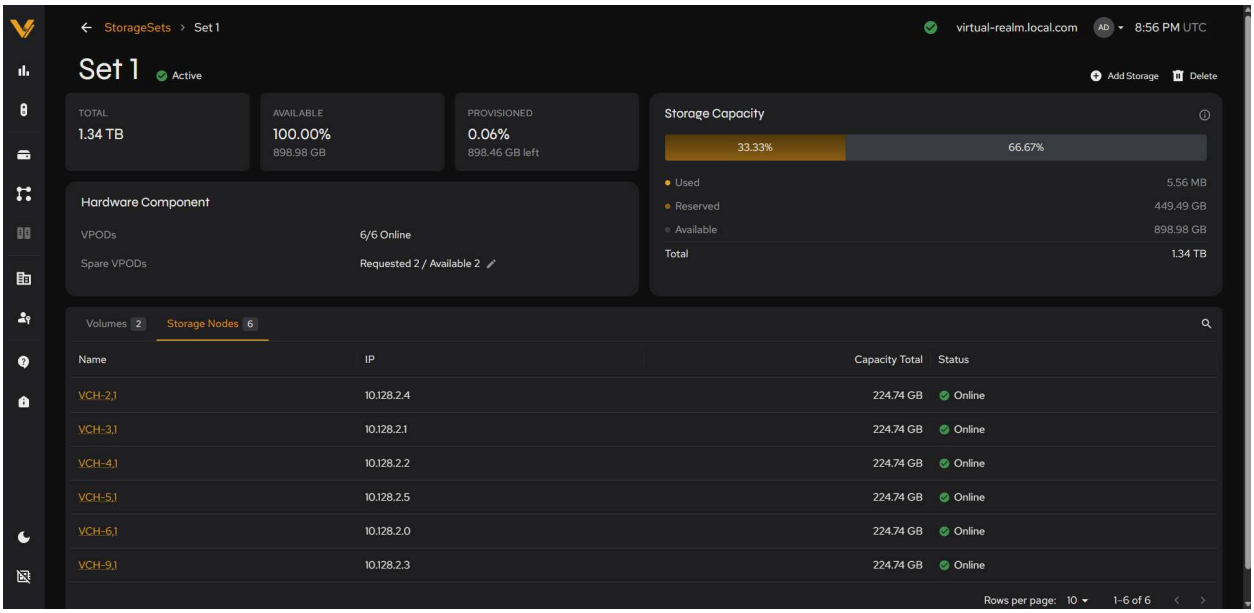


Figure 6.4. Same page with the Storage Nodes tab selected.

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.

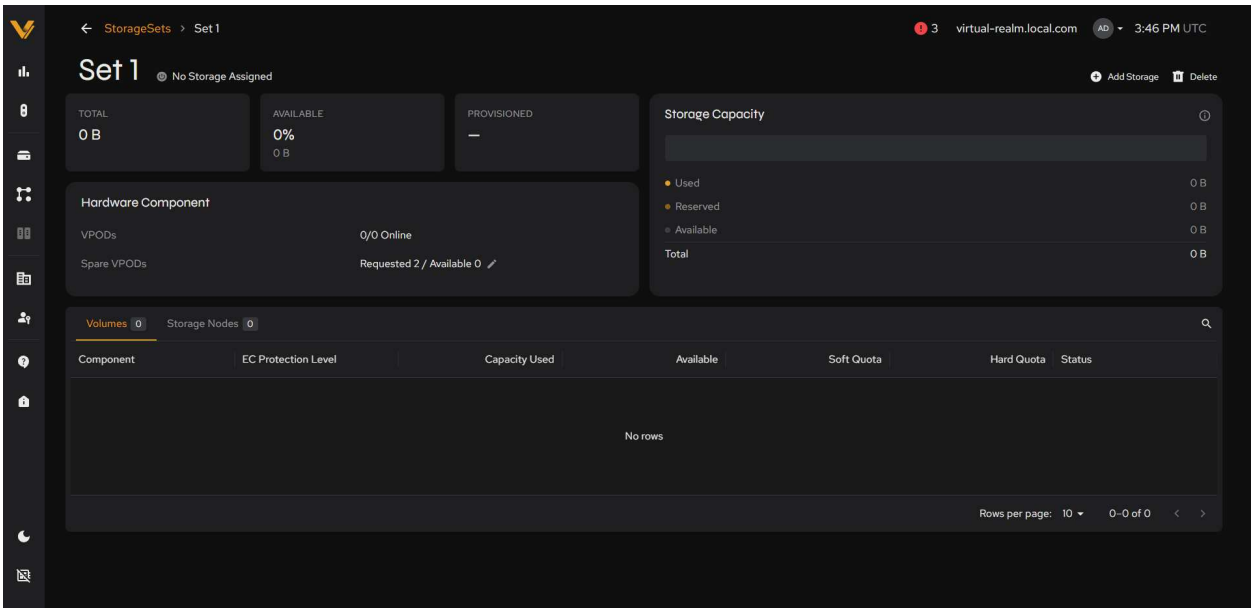


Figure 6.5. Detail page for the empty StorageSet (‘No Storage Assigned’, 0/0 VPODs).

6.1.3 Edit Spare VPODs Dialog

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

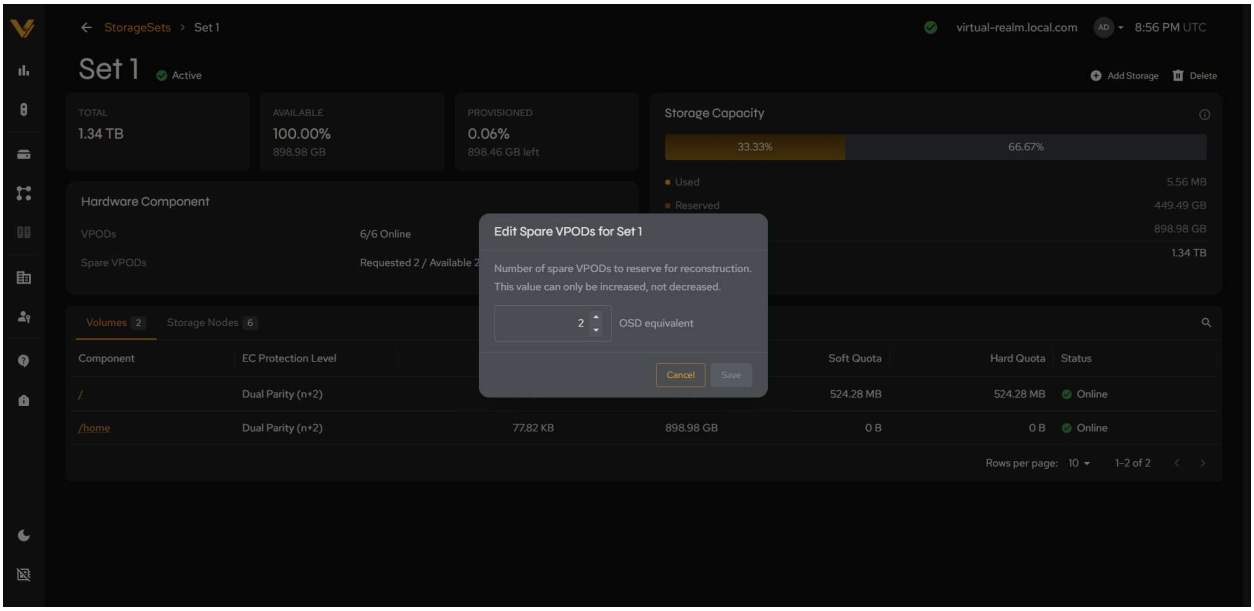


Figure 6.6. The Edit Spare VPODs for Set N dialog open.

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 Dialog

This dialog adds storage to a StorageSet.

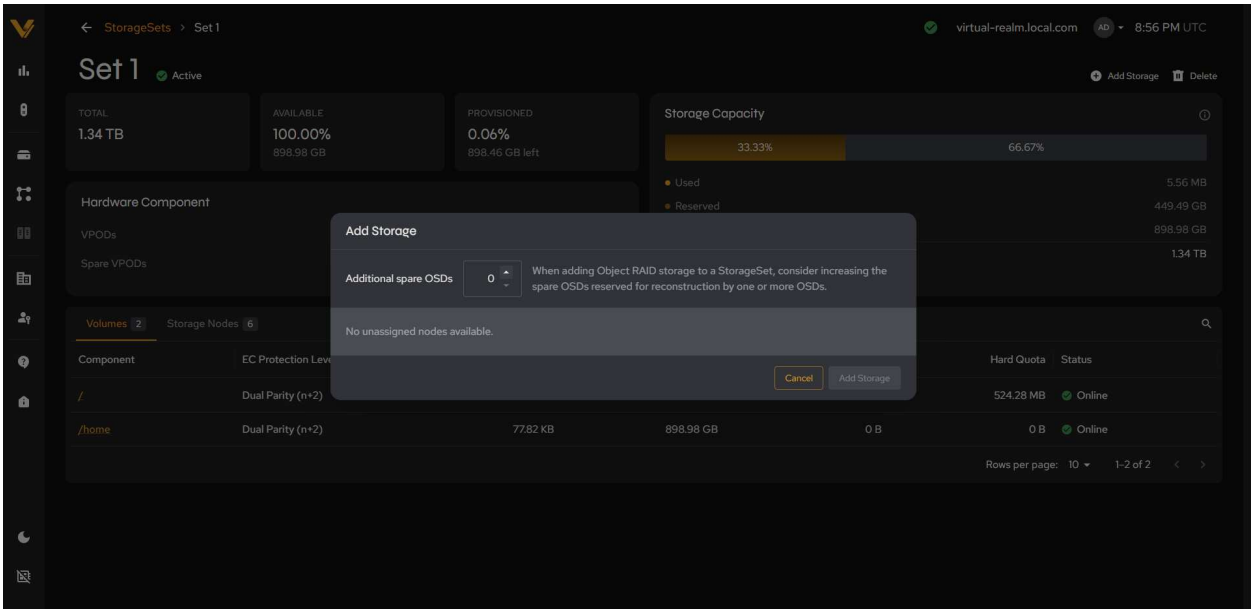


Figure 6.7. The Add Storage dialog showing that there are no unassigned nodes available.

This dialog has an **Additional spare OSDs** field at the top (described below) and, under it, a **Component** list that shows 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.

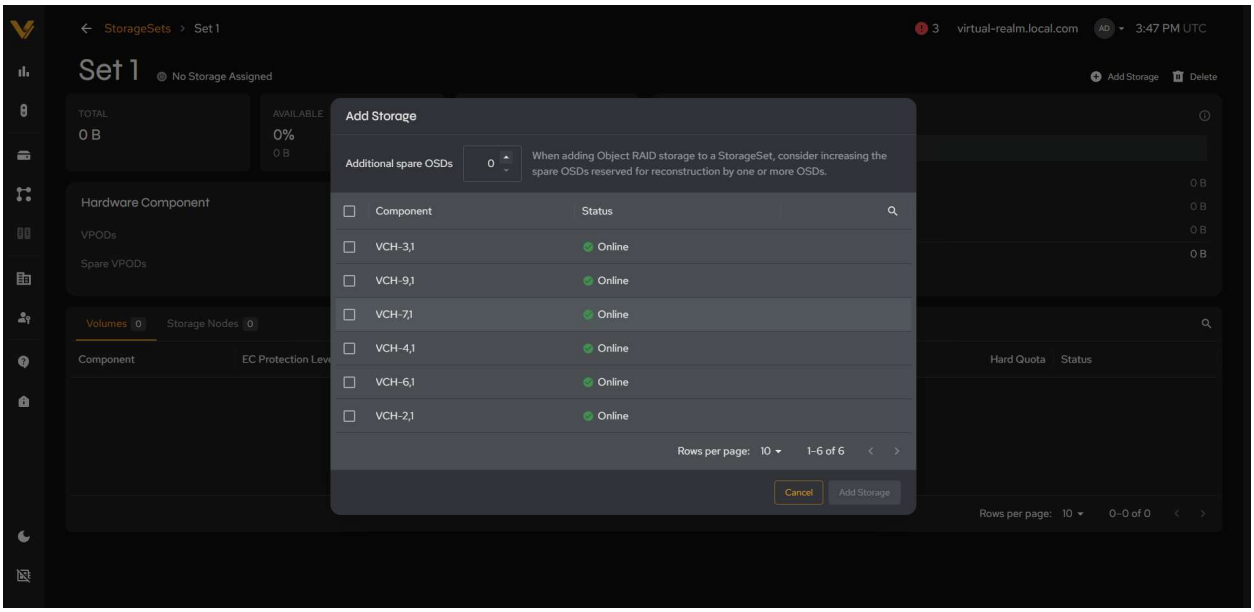


Figure 6.8. The Add Storage dialog's node picker populated with unassigned Storage Nodes.

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, such as 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.

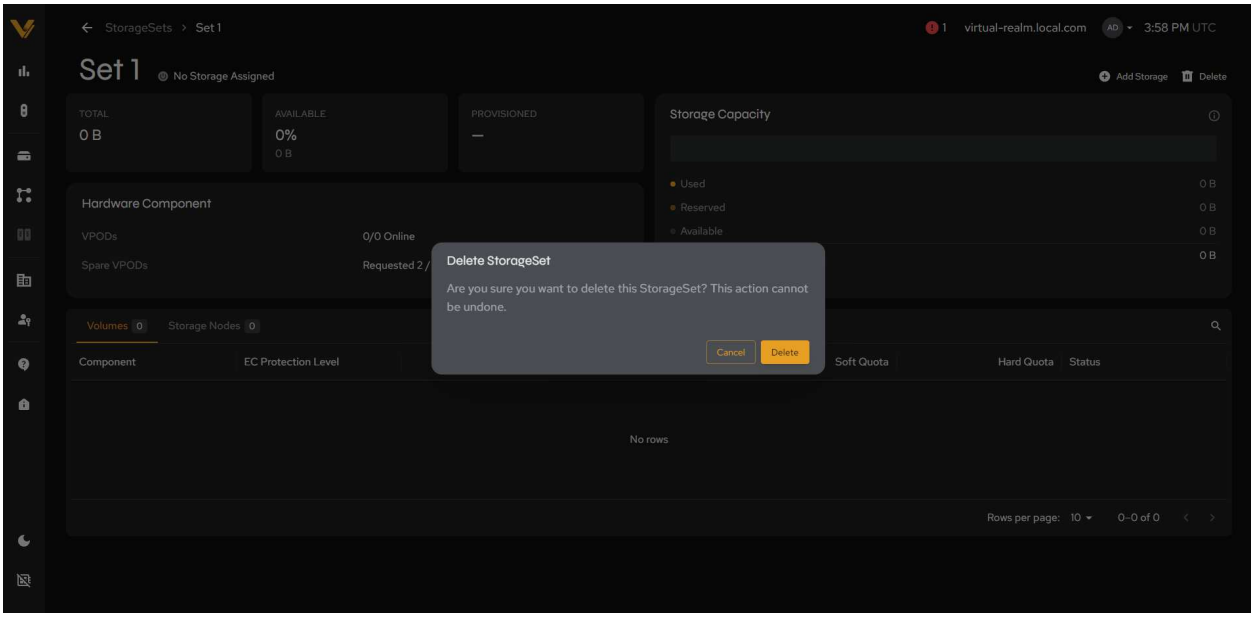


Figure 6.9. The delete confirmation dialog for a StorageSet.

6.1.6 Create StorageSet Page

Use this page to create a new StorageSet.

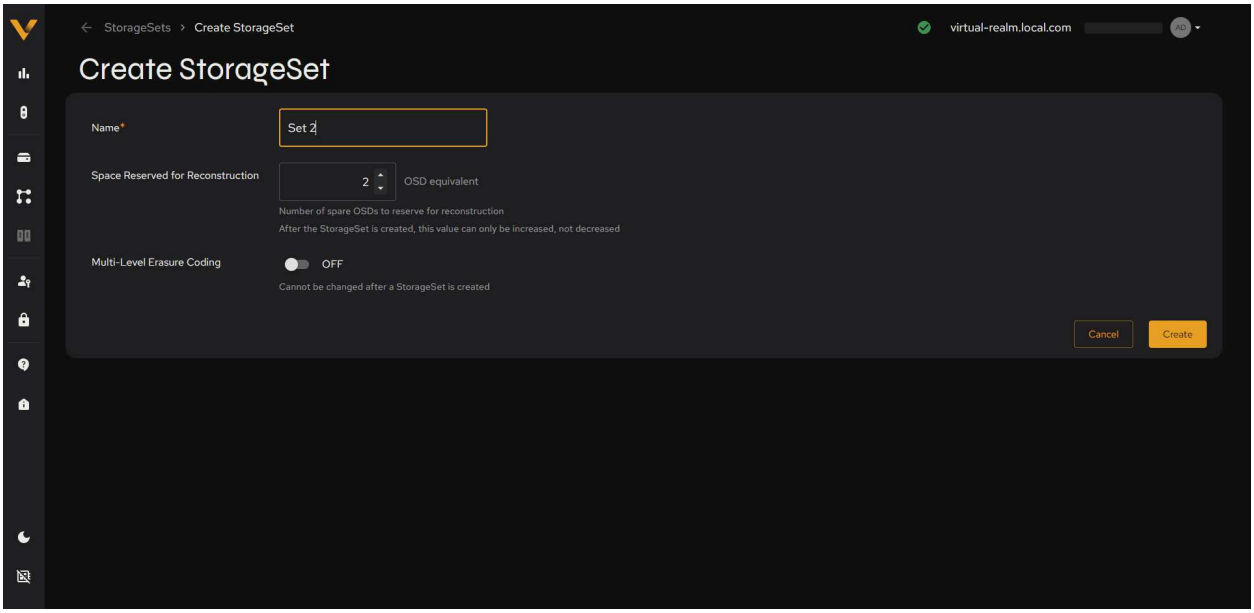


Figure 6.10. The full Create StorageSet form.

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

This section covers Volumes List Page, Volume Detail Page, Delete Volume Confirmation, and Create Volume Page.

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).

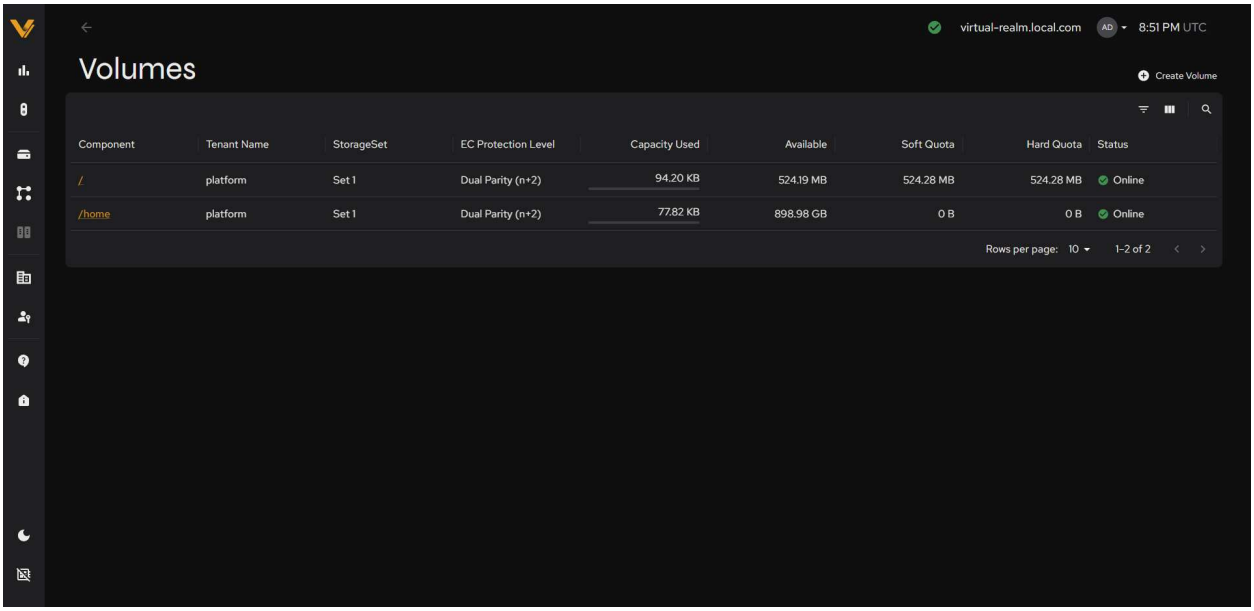


Figure 6.11. The Volumes table with at least two rows.

Column	Description
Component	volume path (e.g. `/`, `/home`); 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.
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.

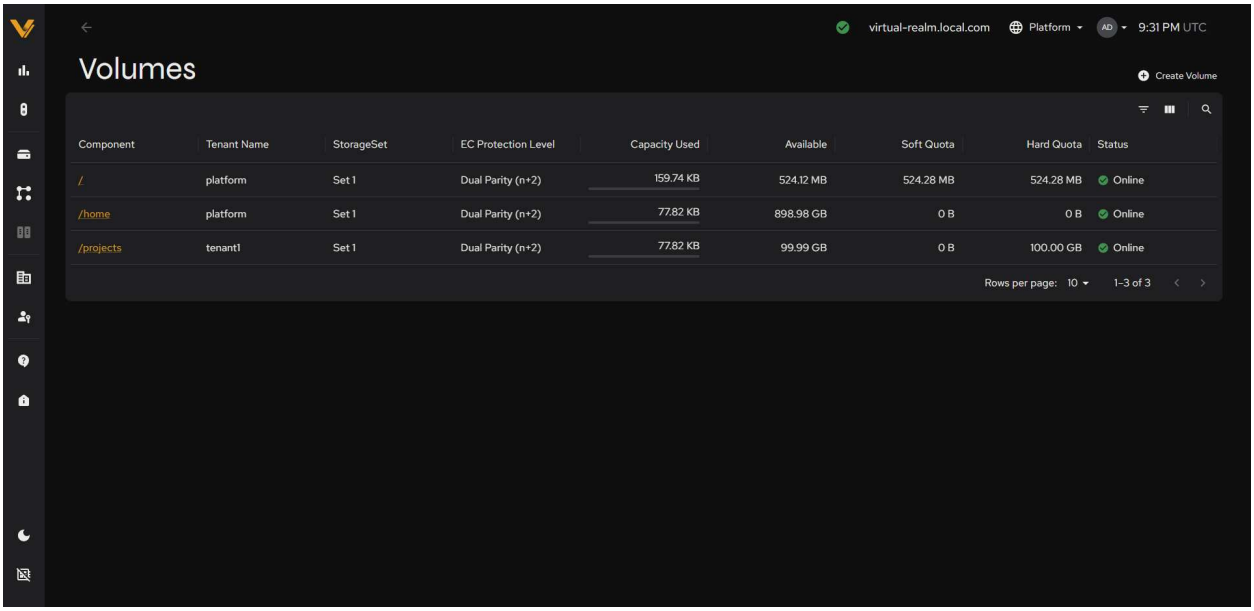


Figure 6.12. The realm-wide Volumes list showing volumes from more than one tenant side by side.

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.

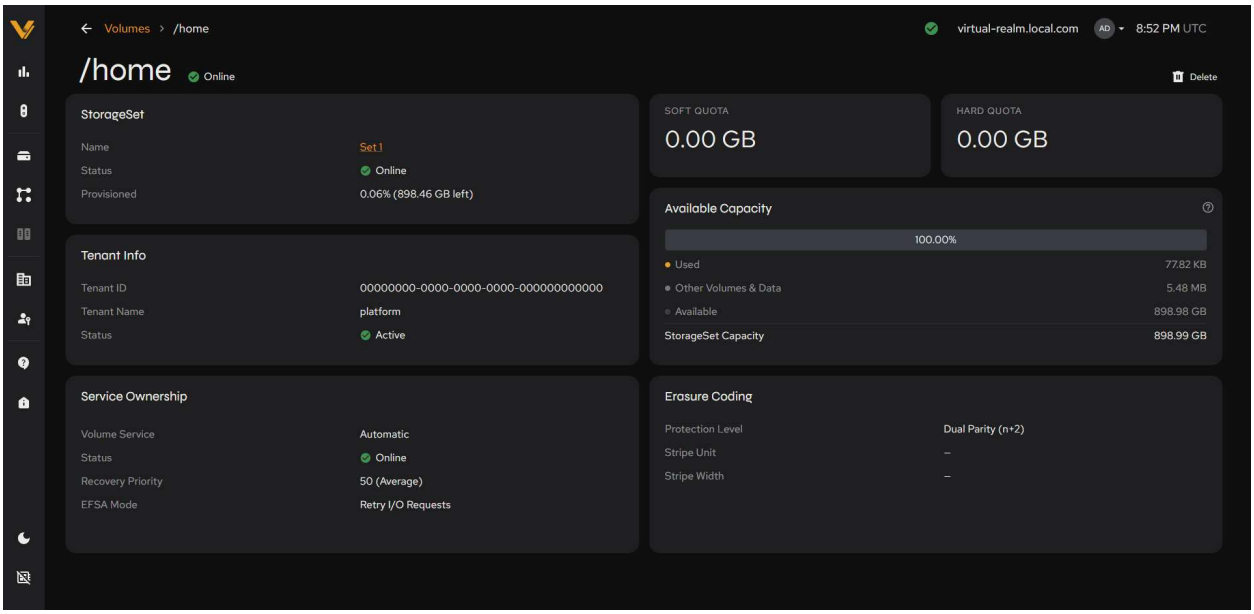


Figure 6.13. Full detail page for one volume.

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")
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")

Card	Field	Meaning
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)”).
Service Ownership	EFSA Mode	Extended File System Availability Mode: the path pursued when an object has exceeded its fault-tolerance level. `retry` blocks and retries I/O requests until the needed storage devices are back online; `file-unavailable` returns I/O errors instead. The system default is `retry` (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.

Card	Field	Meaning
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 dialog confirms deletion of a volume.

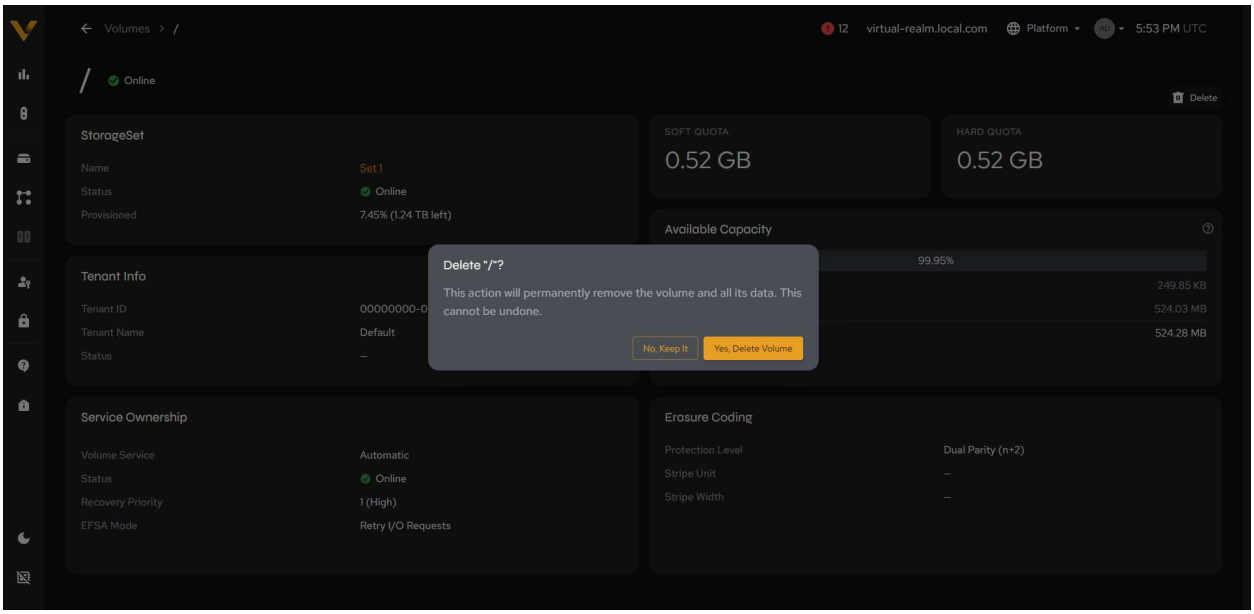


Figure 6.14. The delete confirmation dialog.

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.

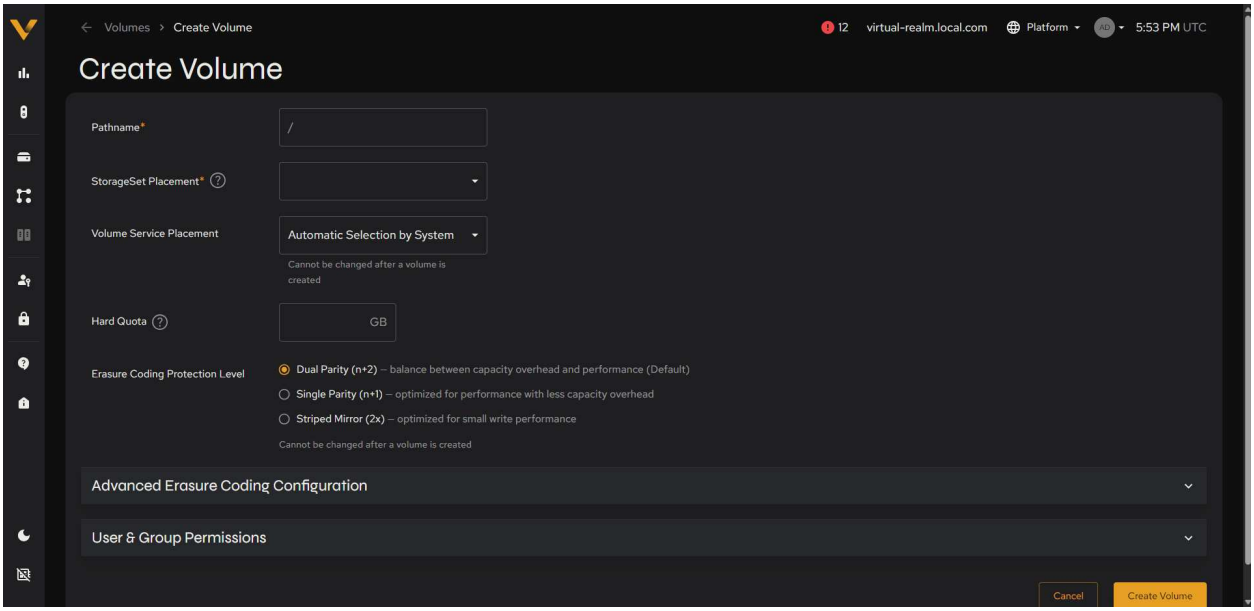


Figure 6.15. Top of the Create Volume form: Pathname through Erasure Coding Protection Level radios.

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 drop-down menu. 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.

Field	Notes
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. Drop-down menu, 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 *(Default)*
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 drop-down menu (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 drop-down menu (default)
User	`root`	`all`
Group	`wheel`	`read-execute`
Others	*(none, implicit)*	`read-execute`

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

Cancel/Create Volume.

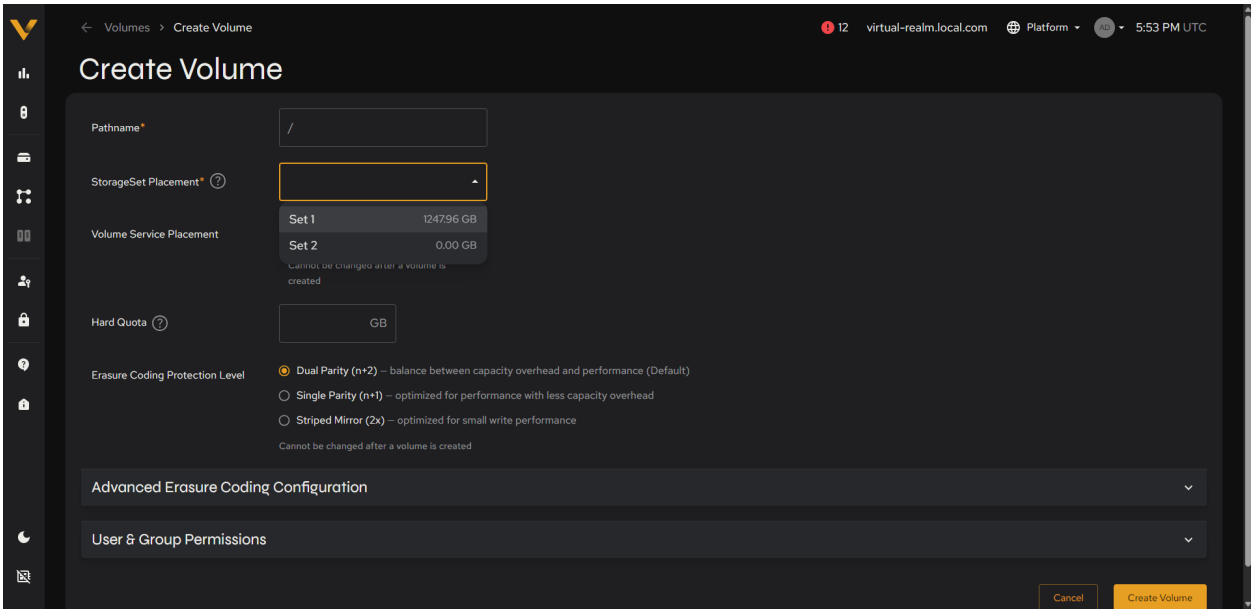


Figure 6.16. The StorageSet Placement drop-down menu open, showing each StorageSet with its available capacity.

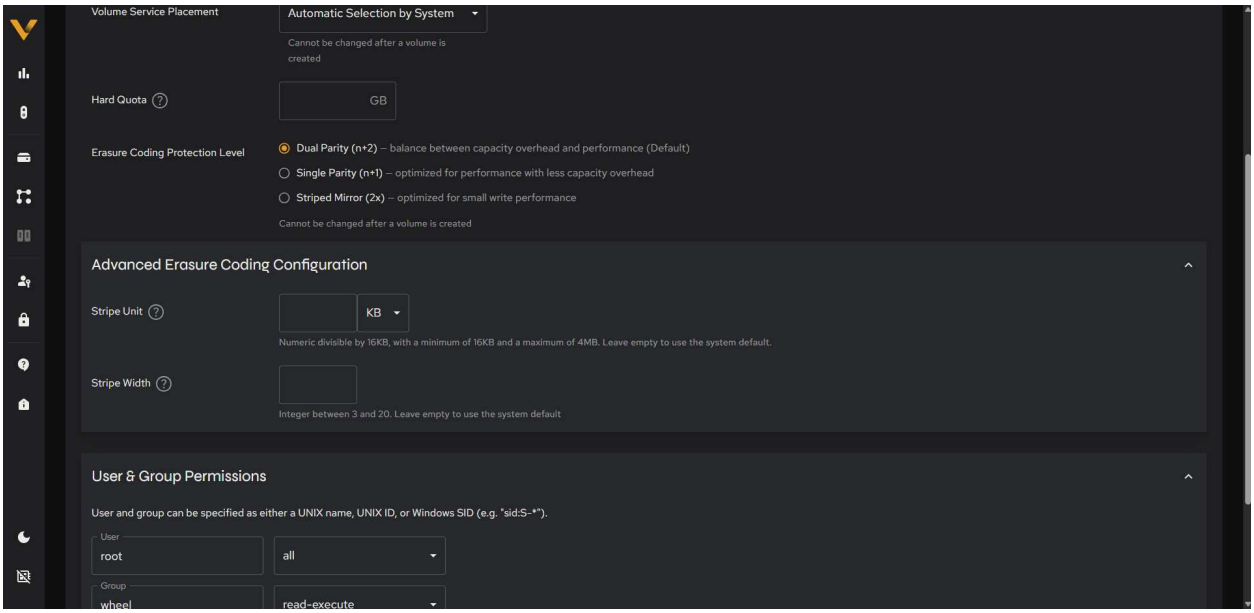


Figure 6.17. Advanced Erasure Coding Configuration section expanded (Stripe Unit/Stripe Width).

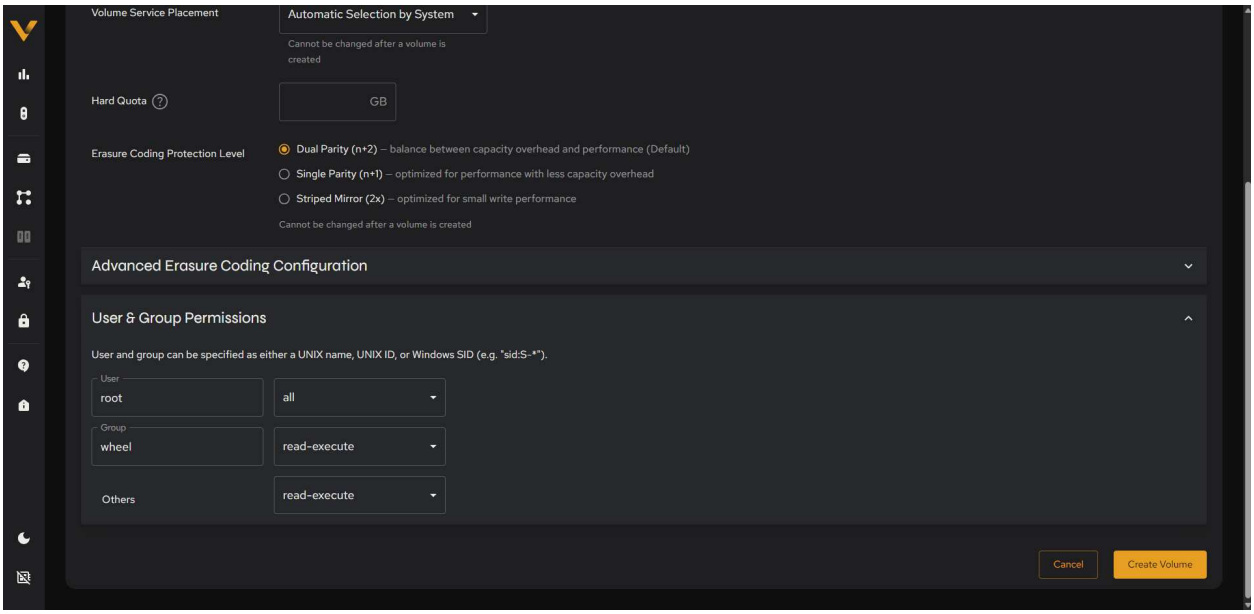


Figure 6.18. User & Group Permissions section expanded, with a permission drop-down menu open.

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. Vantage does not 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.

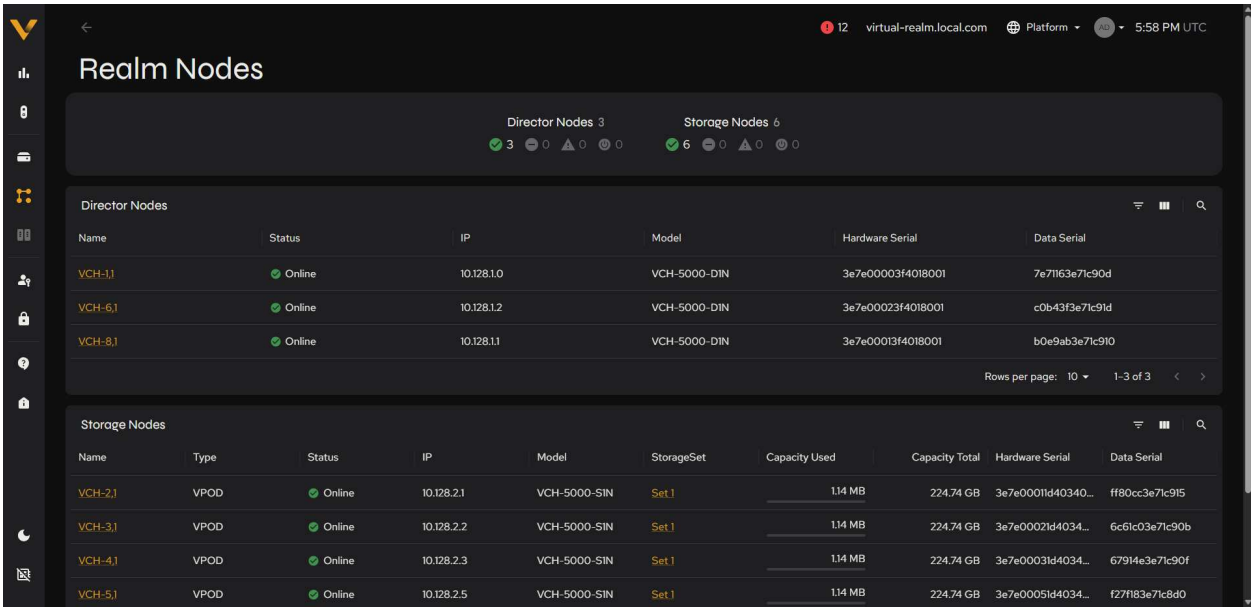


Figure 7.1. Realm Nodes page displaying both the Director Nodes and Storage Nodes tables.

7.1 Director Nodes Table

The following table describes each column.

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 `dataserial` identifier used by `panctl nodes`

7.2 Storage Nodes Table

The following table describes each column.

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 `dataserial` identifier used by `panctl nodes`

Until storage is assigned to a StorageSet, every Storage Node shows `–` in the StorageSet column.

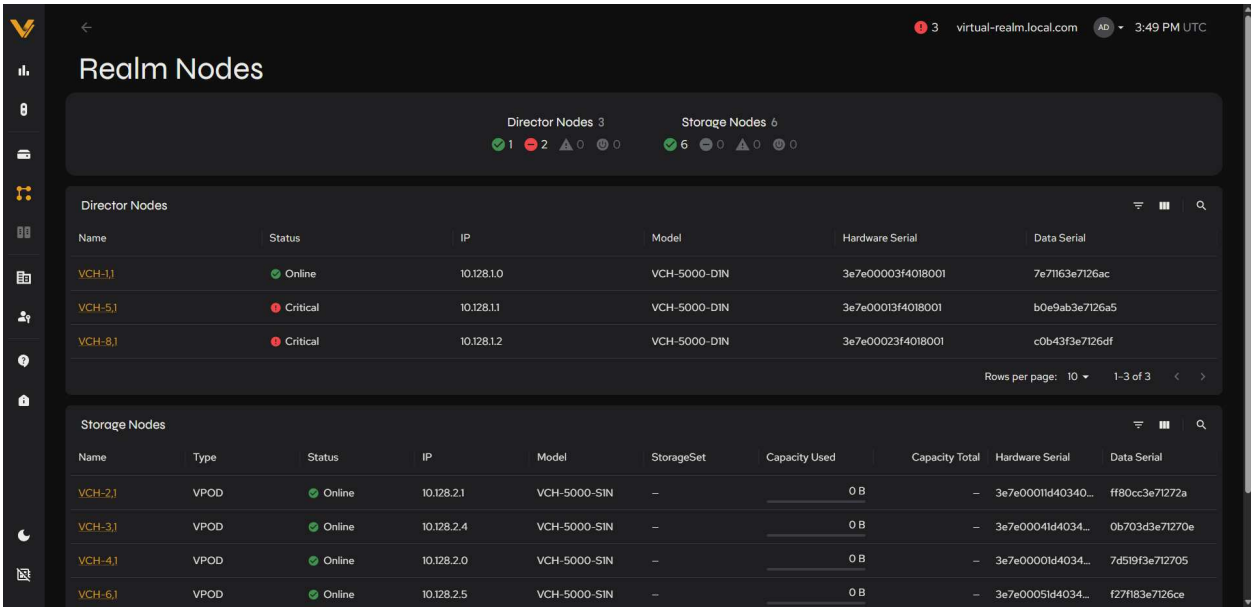


Figure 7.2. Realm Nodes' Storage Nodes table with every StorageSet column showing a dash.

A Director or Storage Node with a problem shows a `Critical` or warning status (see 3.2).

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 VDP nodes.

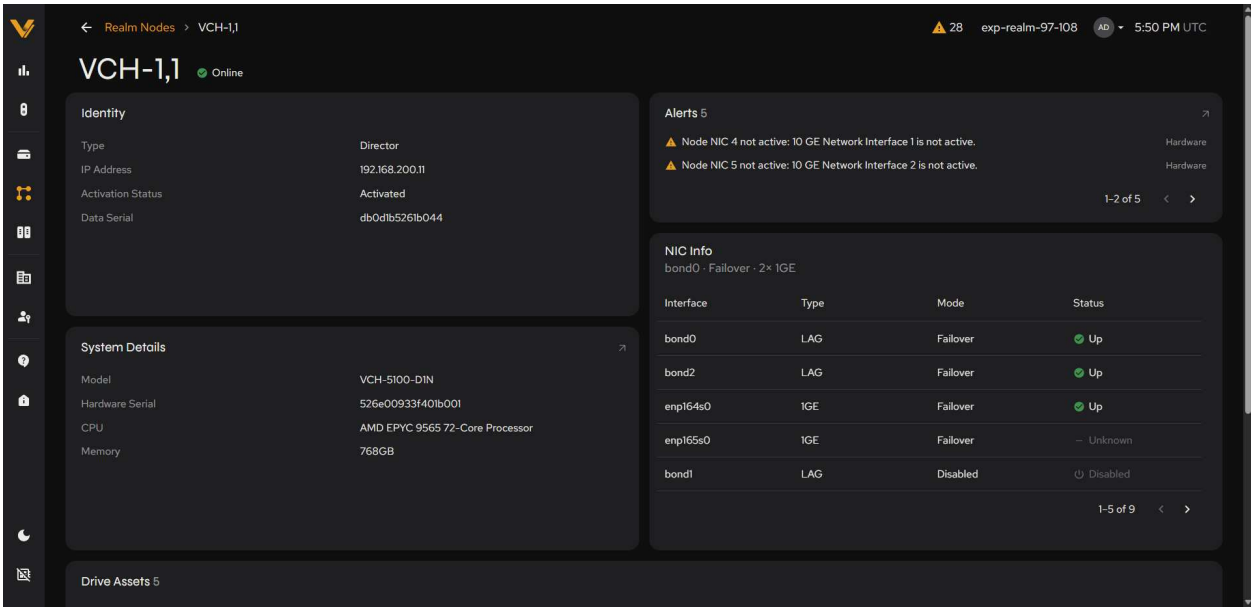


Figure 7.3. Detail page for a Director Node (Identity/System Details/Alerts/NIC Info, with Drive Assets below).

Director Node:

Card	Fields
Identity	Type (`Director`), 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. An arrow in the card's corner (**Open hardware reference**) opens that node's chassis detail page (8.2).
Alerts *(right column)*	Expand icon jumps to Status Center

Card	Fields
NIC Info *(right column)*	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. Status shows **Up**, **Unknown** or **Disabled**. On physical hardware, the bonds in use and their active member interfaces show **Up**, an interface with no link shows **Unknown**, and unused interfaces show **Disabled**.
Drive Assets *(below the cards)*	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 (`VPOD`), 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 *(right column)*	Expand icon jumps to Status Center
Drives Storage *(right column)*	Capacity bar with Used/Available legend
Drive Assets *(below the cards)*	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.

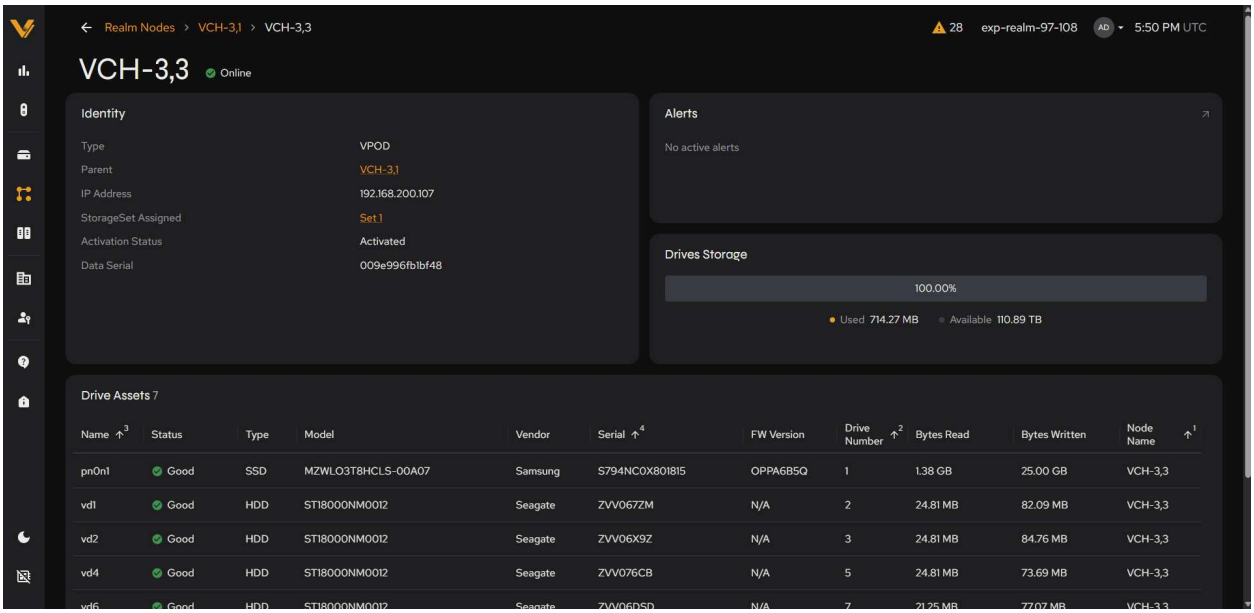


Figure 7.4. Detail page for a Storage Node/VPOD (Identity/Alerts/Drives Storage/Drive Assets).

> **Note:** A VPOD's Identity card shows an additional **Parent** field, and the breadcrumb grows a level to match. See 8.3.

VDP Node:

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 *(right column)*	Expand icon jumps to Status Center
VPODs Storage *(right column)*	Capacity bar with Used/Available legend for the chassis's VPODs
NIC Info *(right column)*	Summary line plus a table: Interface, Type, Mode, Status

Card	Fields
VPODs *(below the cards)*	A table of the chassis's VPOD nodes with the columns Name, Status, Used Capacity, and Total Capacity
Drive Assets *(below the cards)*	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

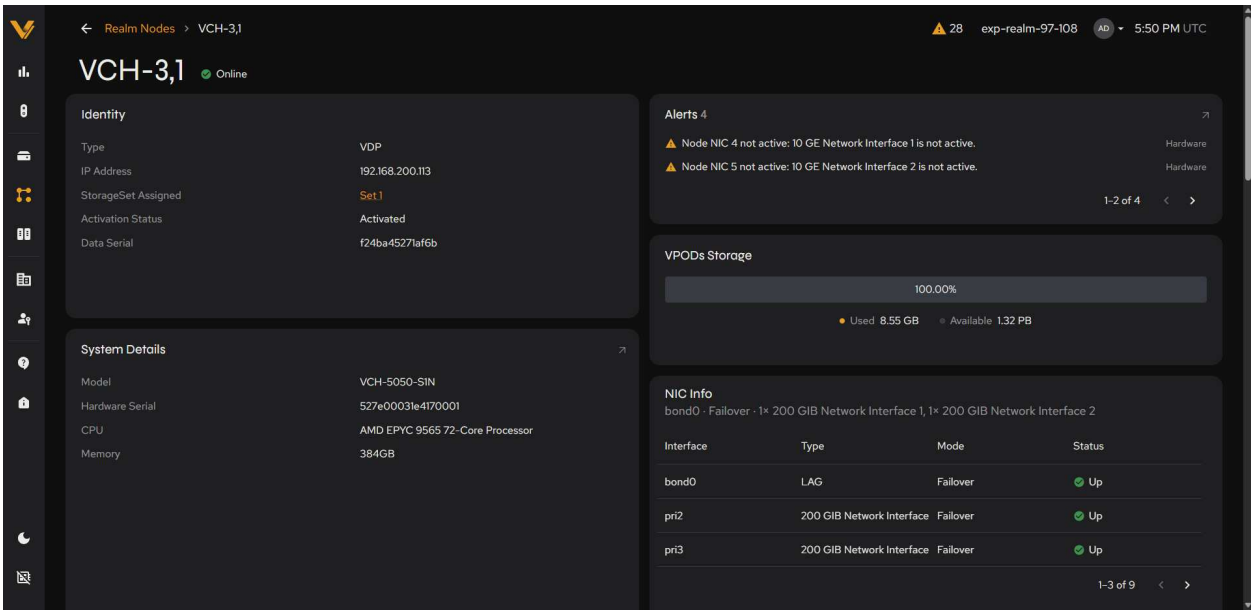


Figure 7.5. Detail page for a VDP Node (Identity/System Details/Alerts/VPODs Storage/NIC Info, with the VPODs and Drive Assets tables below).

8. Hardware

The **Hardware** sidebar entry opens a list of the realm's chassis.

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. `VCH-5050-S1N` (Storage) or `VCH-5100-D1N` (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

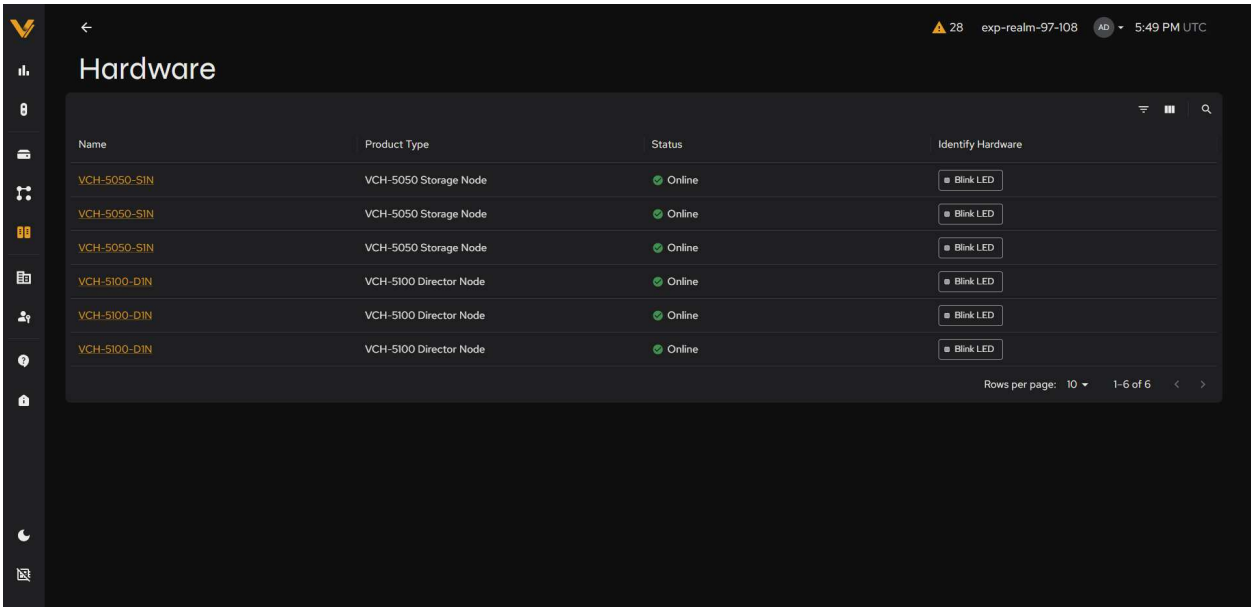


Figure 8.1. Hardware list showing Storage and Director chassis rows.

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.

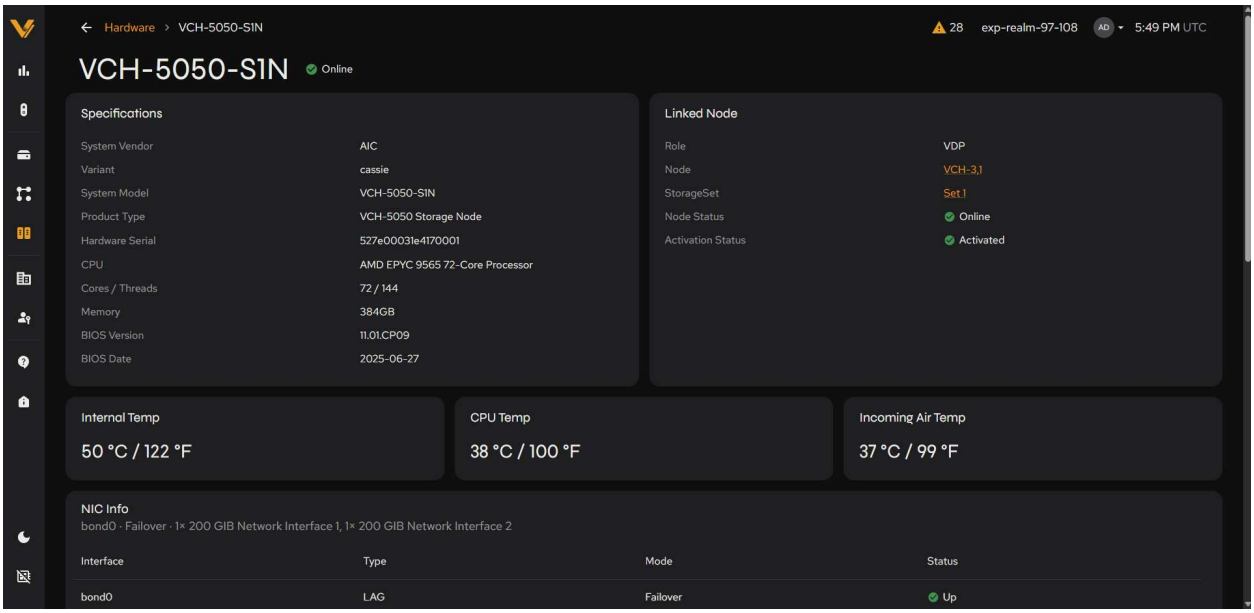


Figure 8.2. Storage chassis detail page (VCH-5050-S1N), top of the page: Specifications, Linked Node, the three temperature cards, and the start of the NIC Info table.

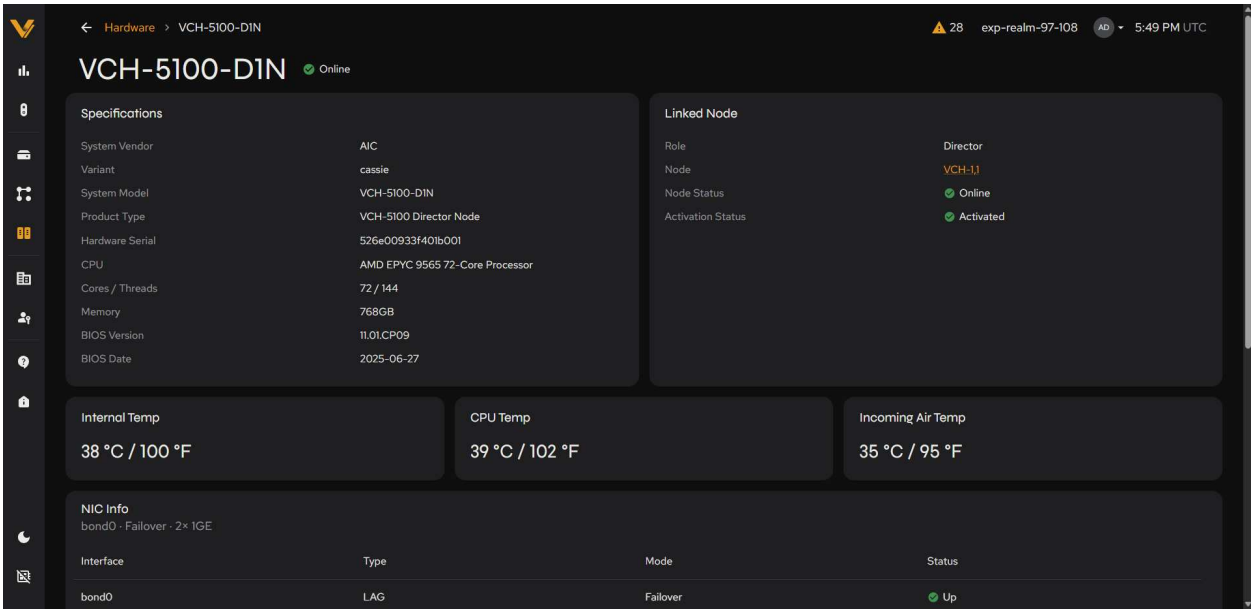


Figure 8.3. Director chassis detail page (VCH-5100-D1N), top of the page. A Director's Linked Node card has no StorageSet row.

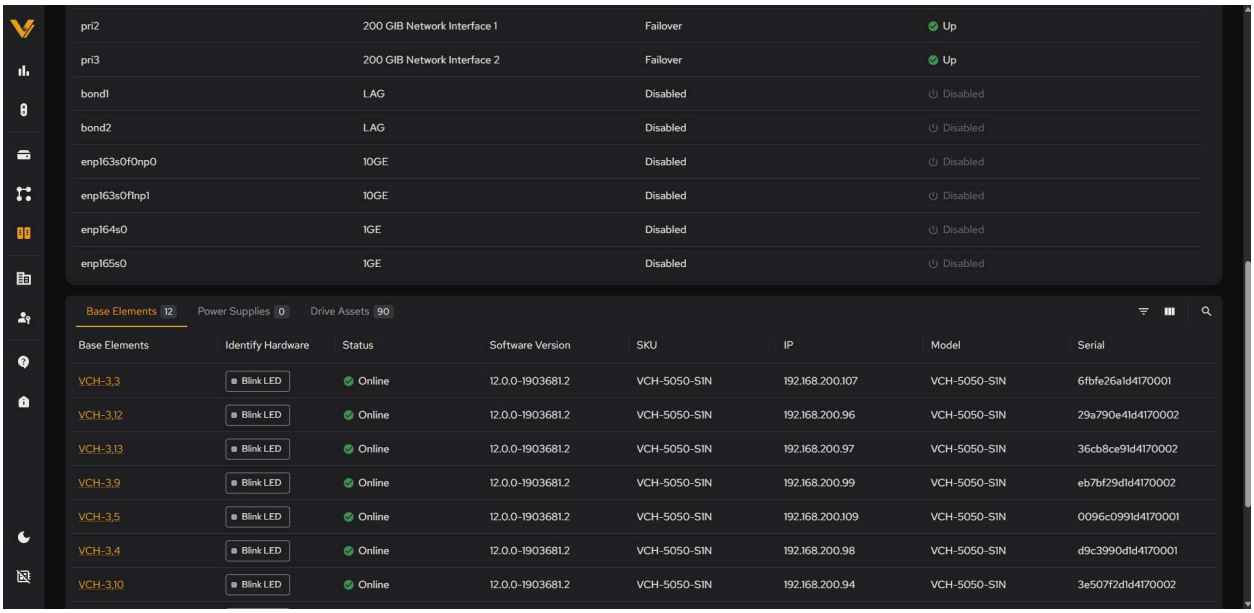


Figure 8.4. Storage chassis detail page scrolled down to the tabs, with the Base Elements tab open below the NIC Info table.

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.”

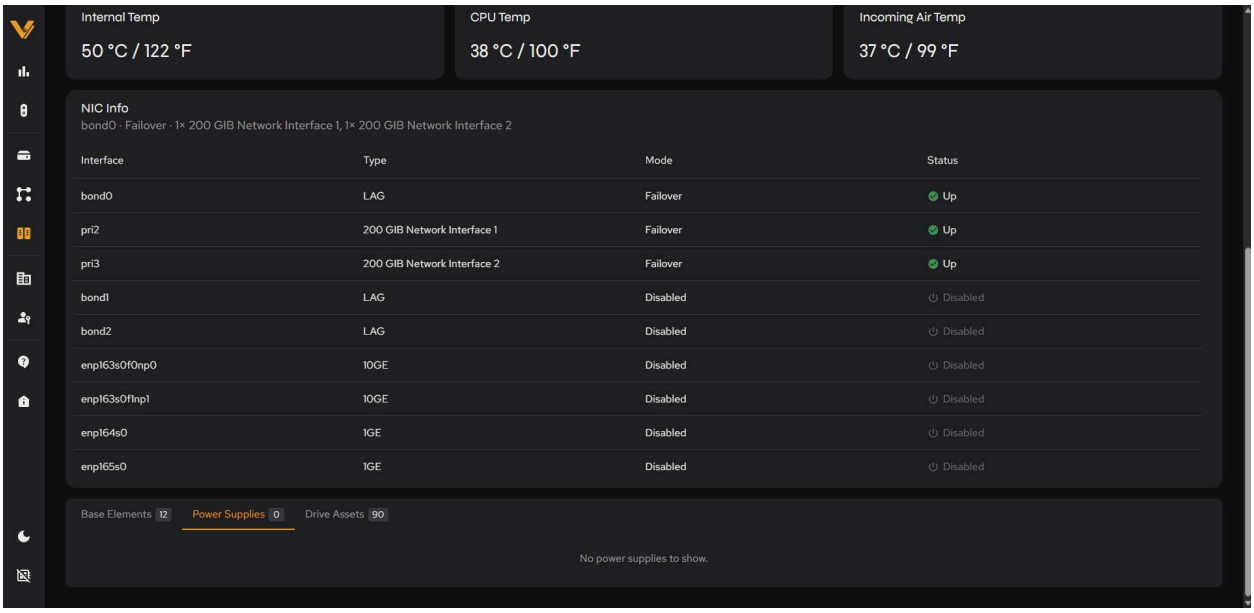


Figure 8.5. Power Supplies tab in its empty state, below the NIC Info table.

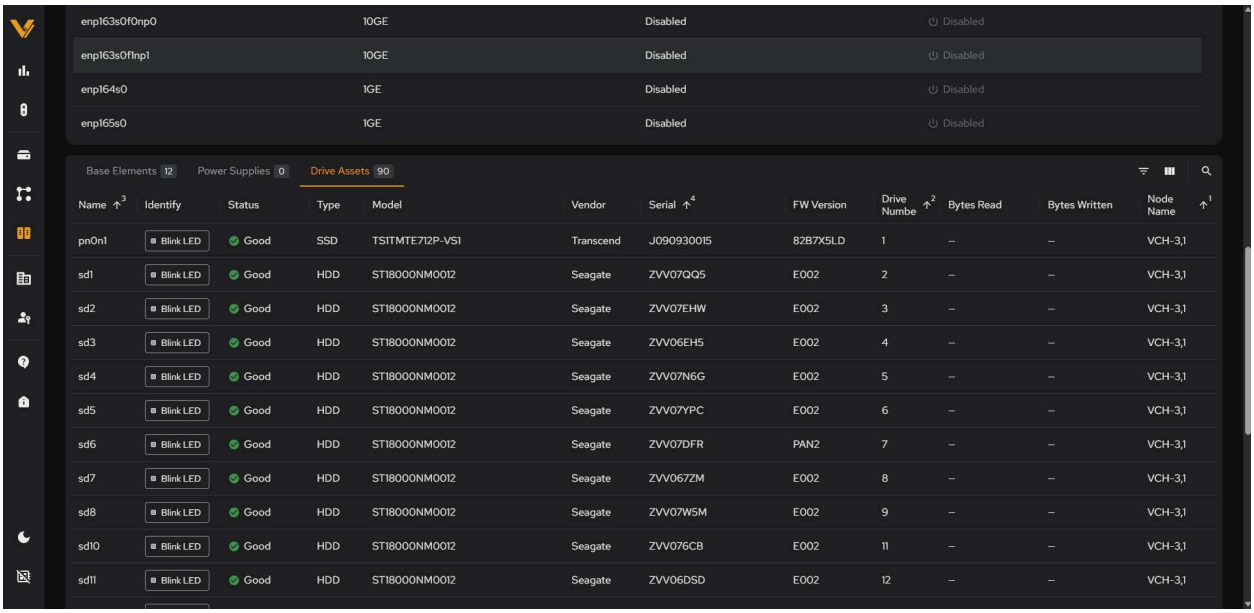


Figure 8.6. Drive Assets tab of a Storage chassis, showing the table's columns and sort-order markers.

8.3 Parent Field on Realm Node Detail

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.

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).

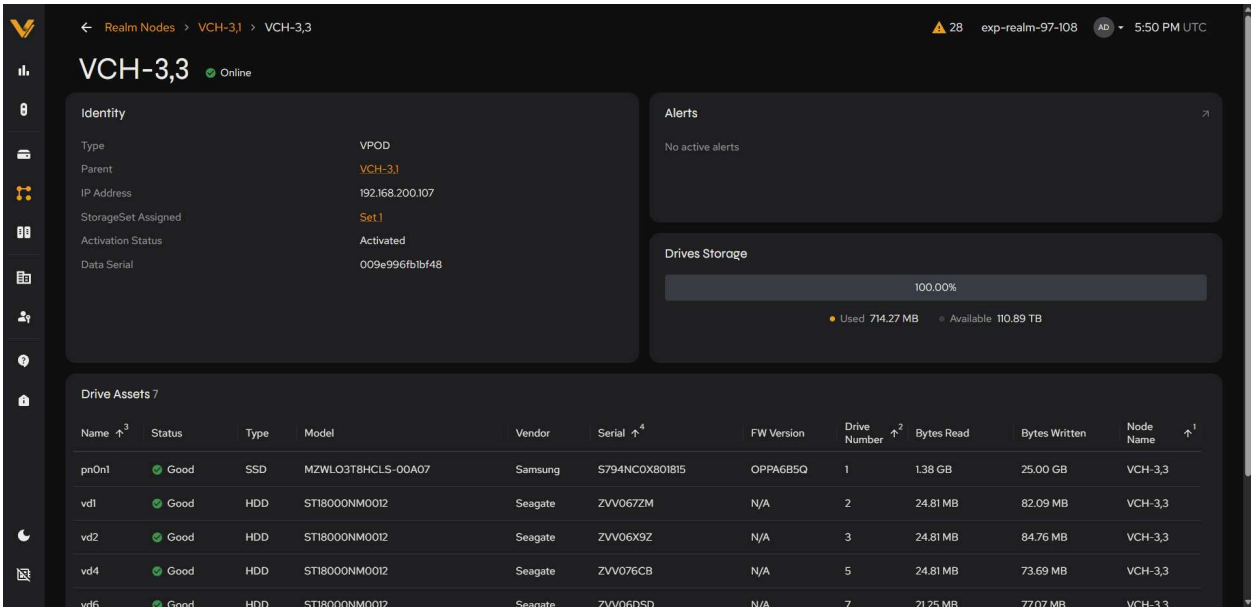


Figure 8.7. Realm Node detail page showing the Parent field and chassis-level breadcrumb.

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.

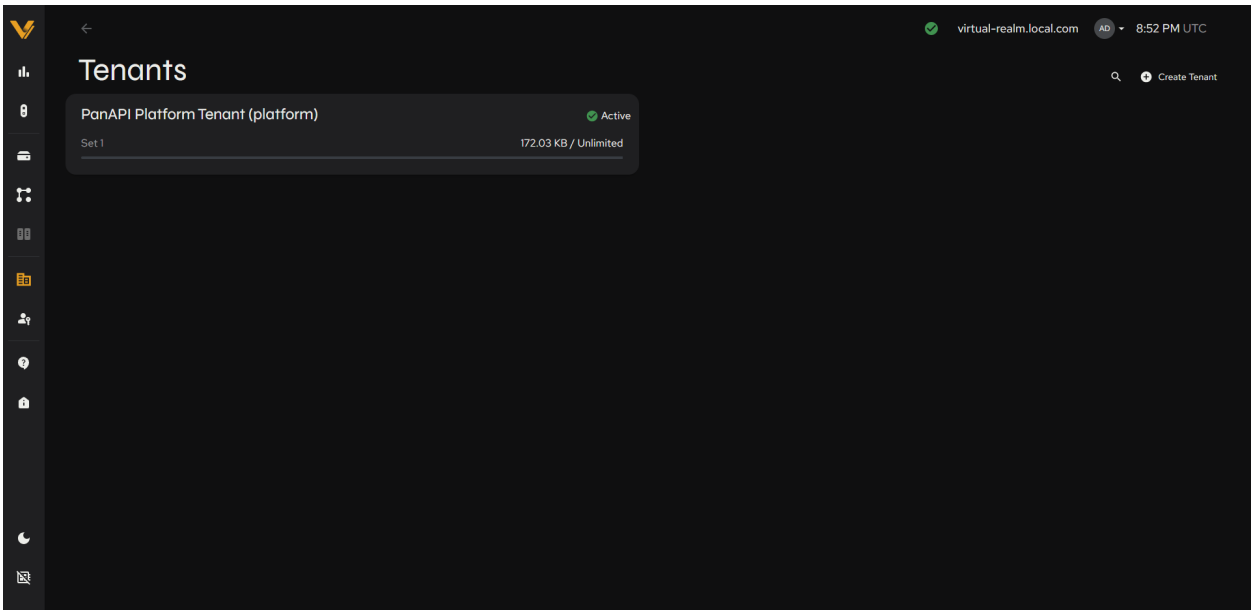


Figure 9.1. The Tenants list with at least one tenant card.

Each card shows:

Card element	Description
Name	e.g. “PanAPI Platform Tenant (platform)”
Status badge	e.g. `Active` or `Setup Incomplete` 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.

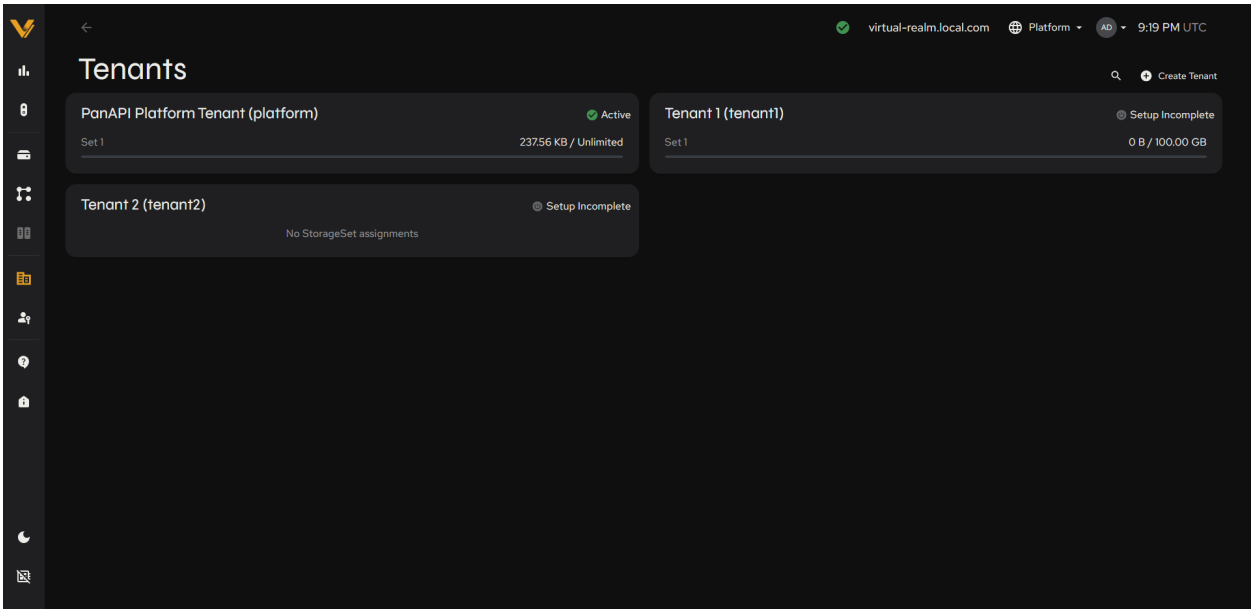


Figure 9.2. The Tenants list with the platform tenant and two newly created tenants.

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).

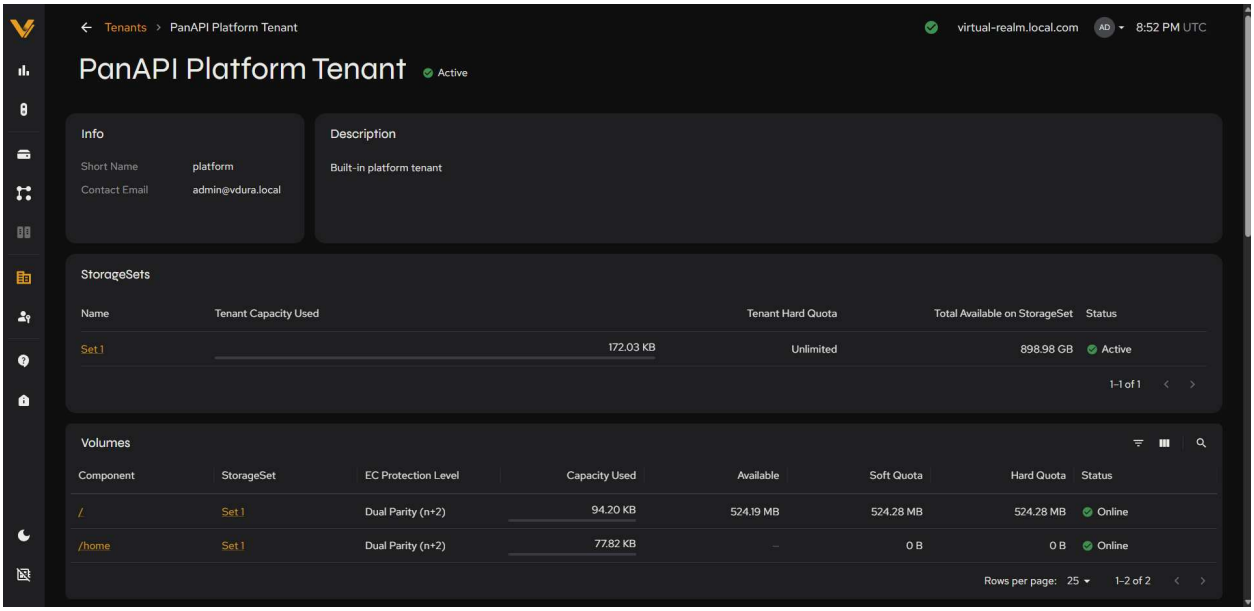


Figure 9.3. Info/Description/StorageSets/Volumes cards.

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.

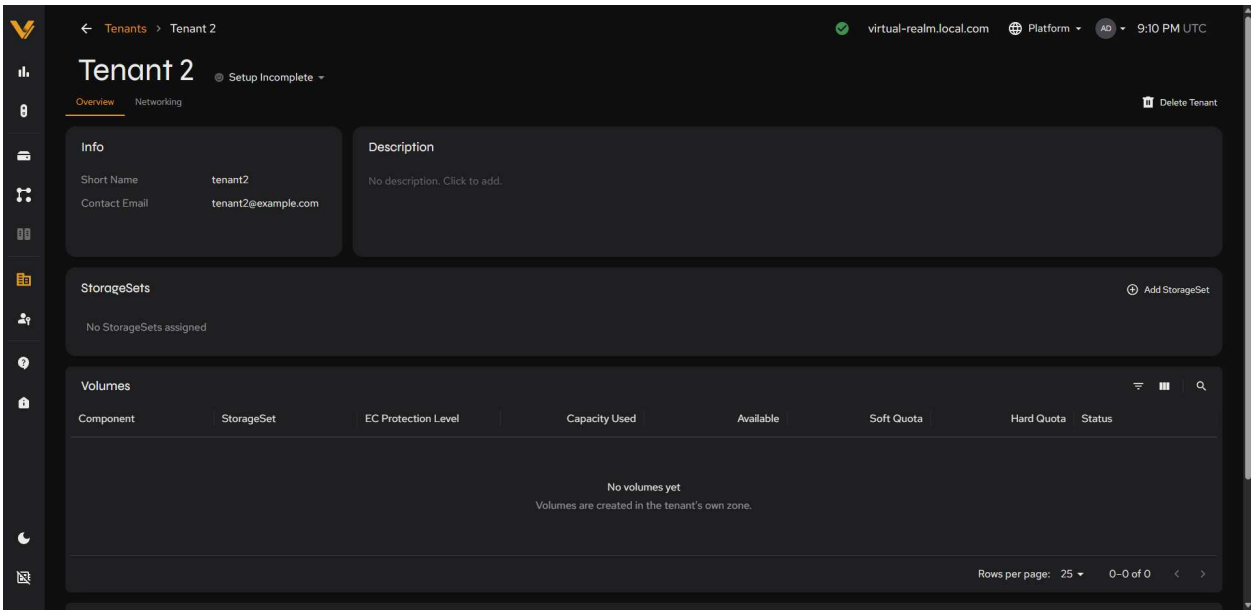


Figure 9.4. A brand-new tenant's detail page: 'No StorageSets assigned' + the + Add StorageSet button.

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

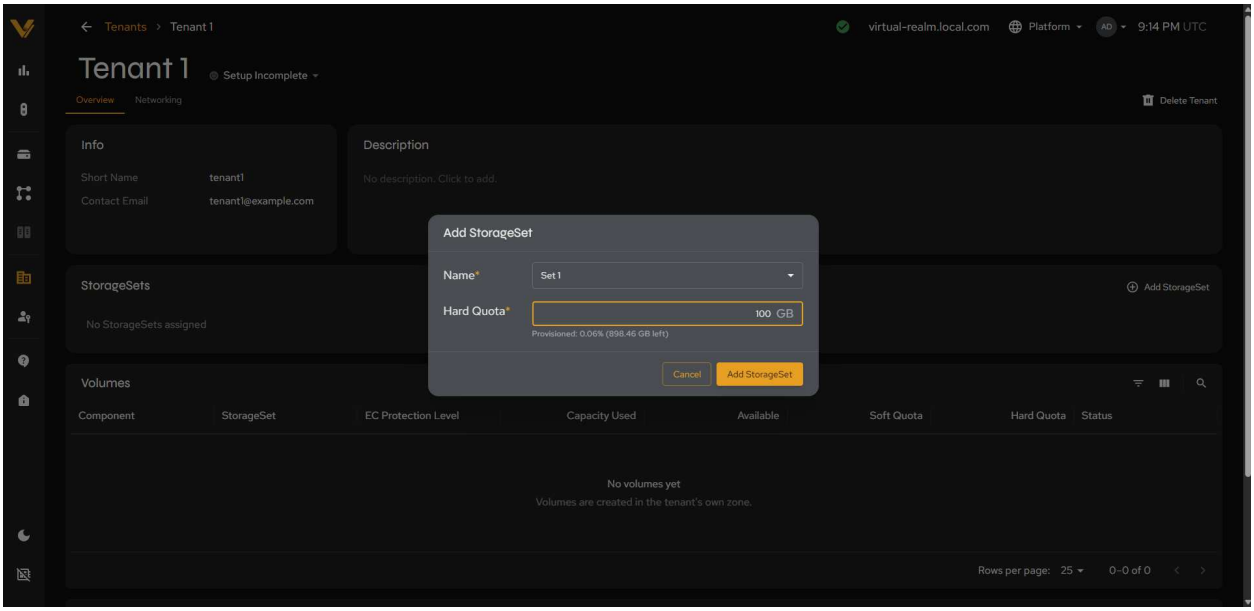


Figure 9.5. The 'Add StorageSet' dialog: Name drop-down menu (with capacities) + required Hard Quota.

Field	Notes
Name*	Drop-down menu 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.

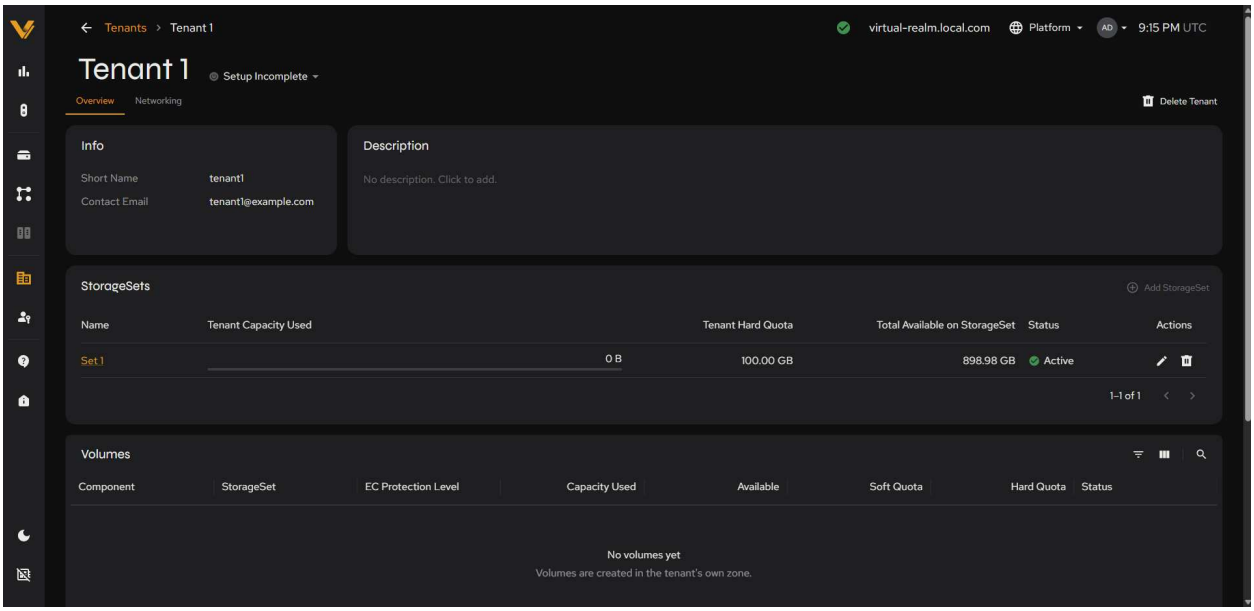


Figure 9.6. The tenant's StorageSets table after assignment: quota shown, plus edit/delete Actions icons.

Volumes table:

Column	Description
Component	Volume path (e.g. `/`, `/home`); 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.

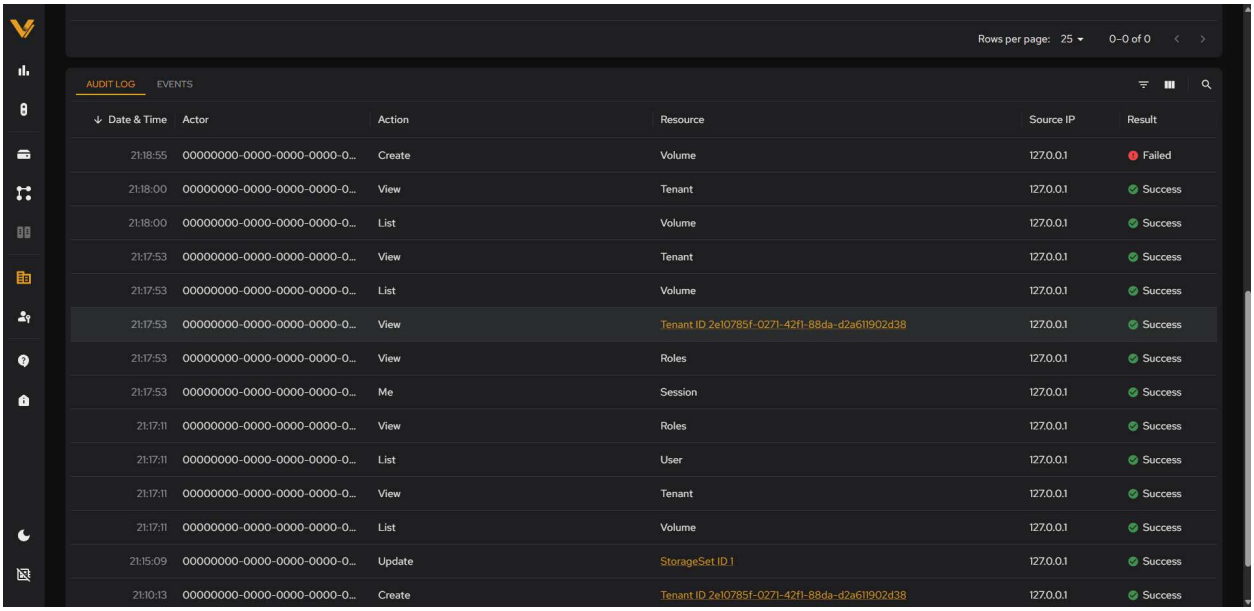


Figure 9.7. The bottom Audit Log tab, populated with rows.

9.2.1 Delete Tenant Confirmation

This dialog 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.

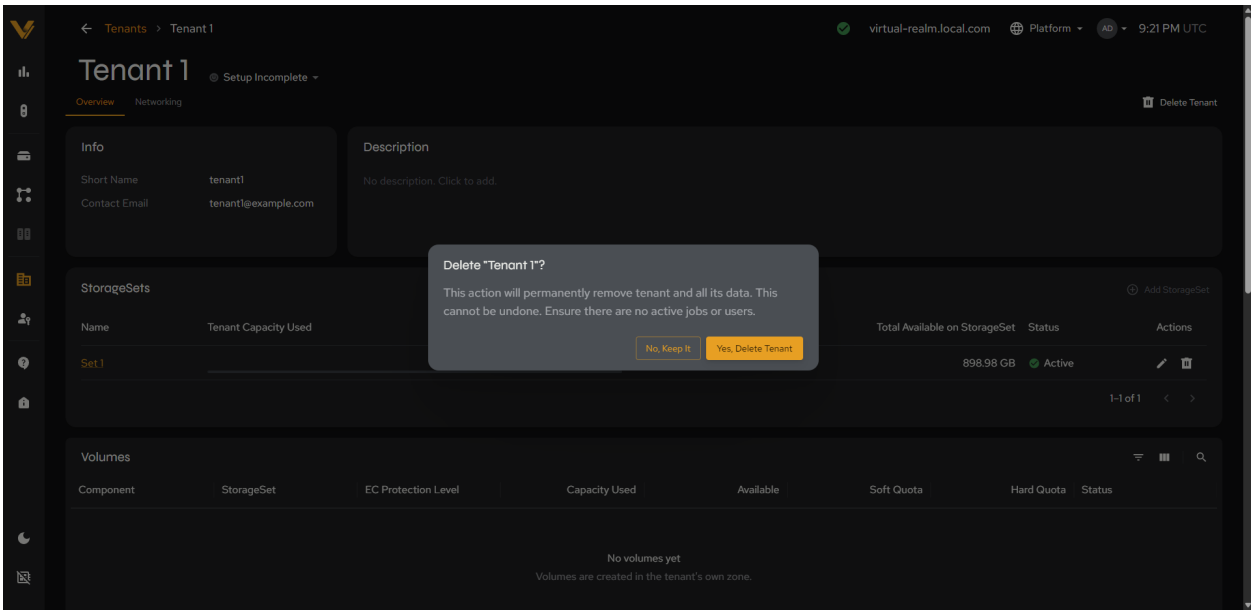


Figure 9.8. The delete confirmation dialog for a tenant.

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.

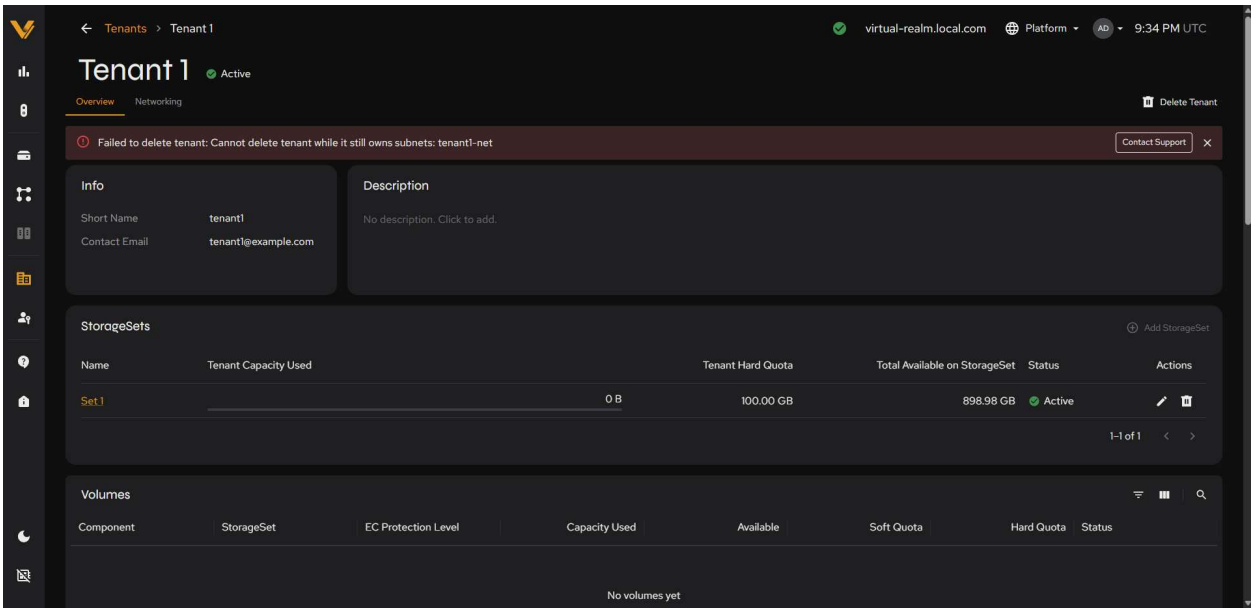


Figure 9.9. The error shown when deleting a tenant that still owns a network configuration.

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).

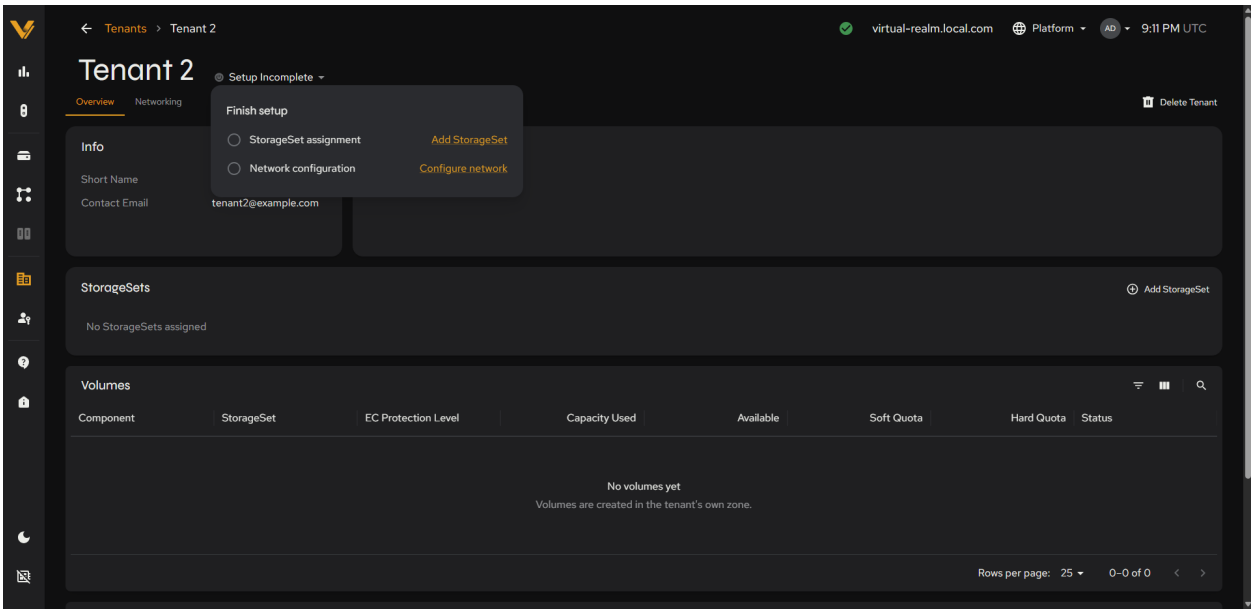


Figure 9.10. The Finish setup checklist opened from the Setup Incomplete chip.

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**.

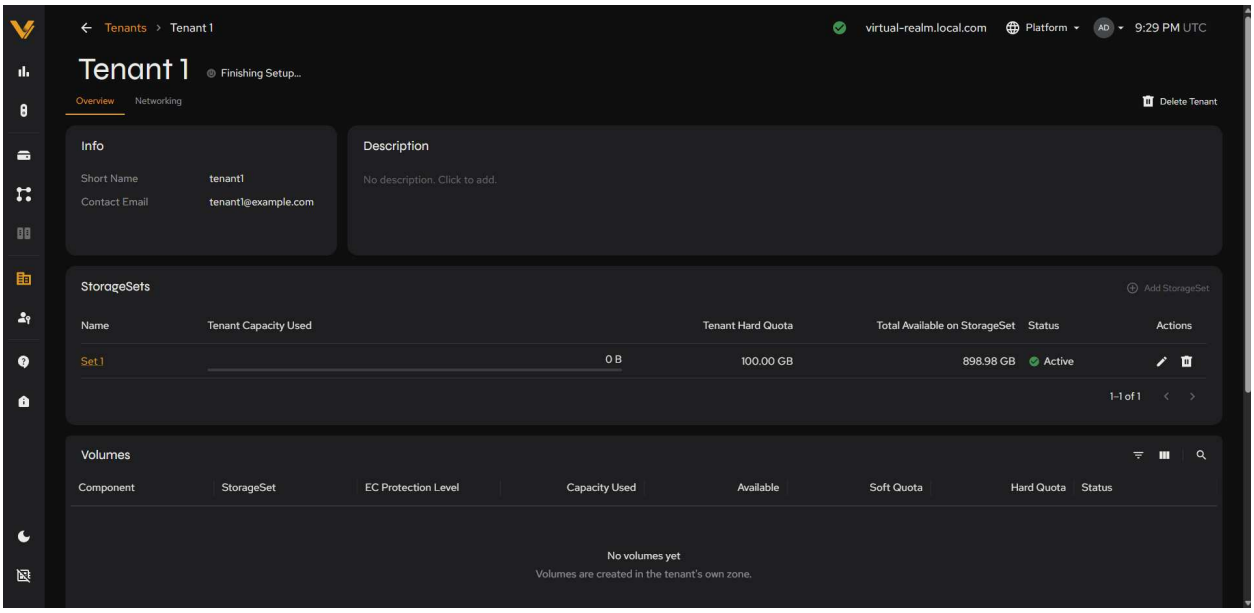


Figure 9.11. A tenant in the Finishing Setup… state.

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.

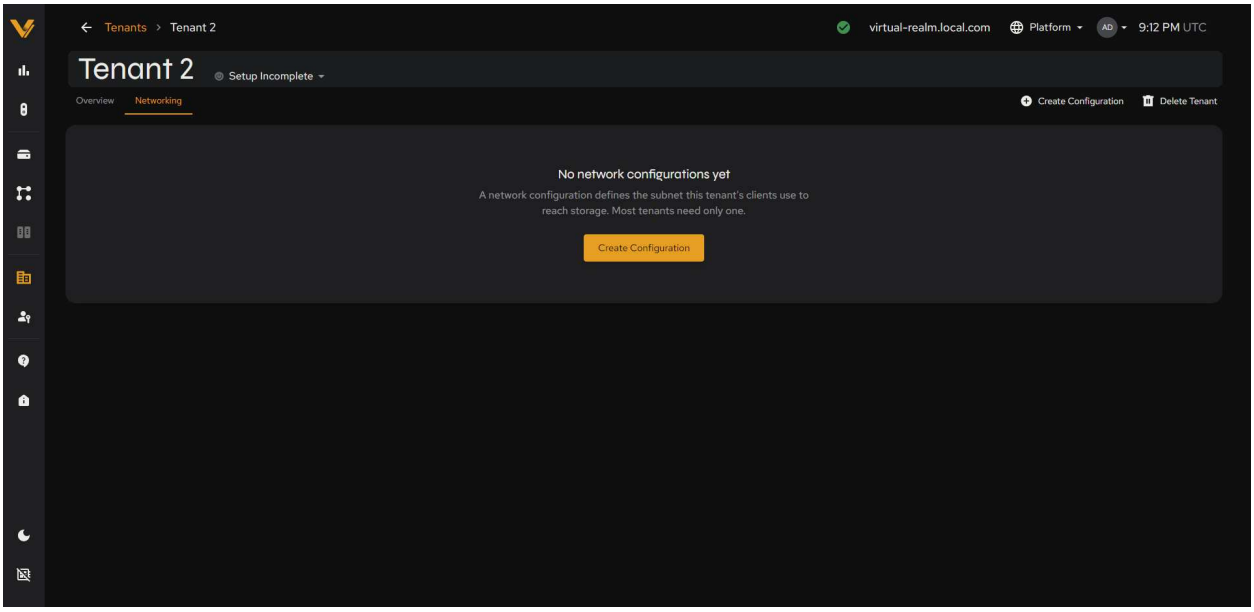


Figure 9.12. The Networking tab of a new tenant, with no network configurations.

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

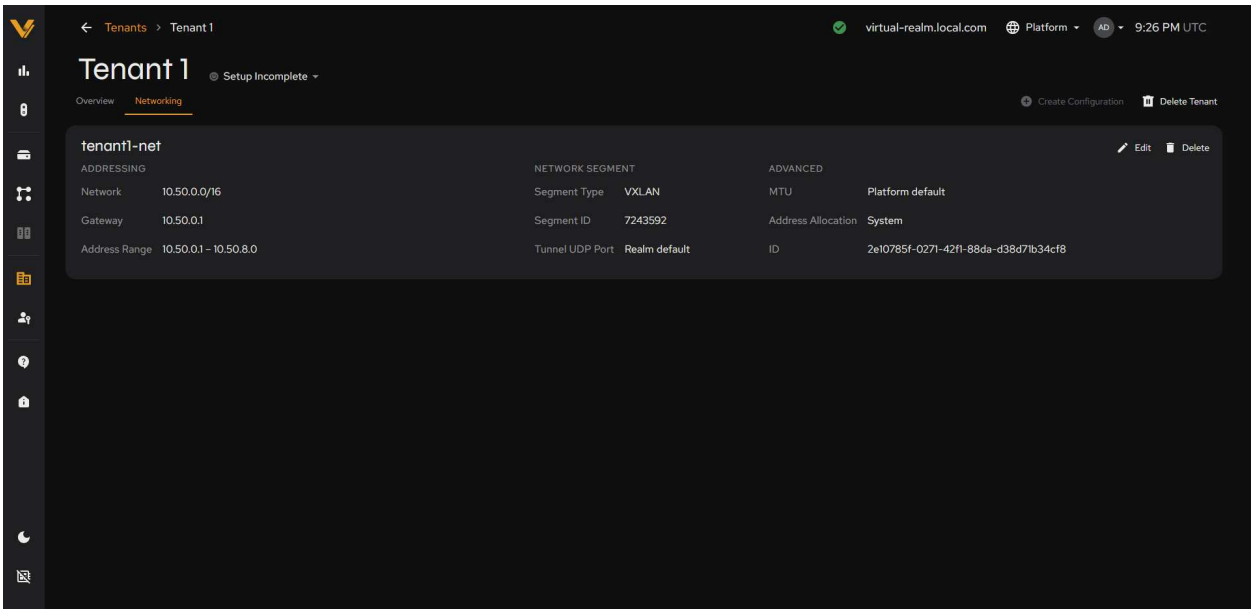


Figure 9.13. The Networking tab showing a network configuration.

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**).

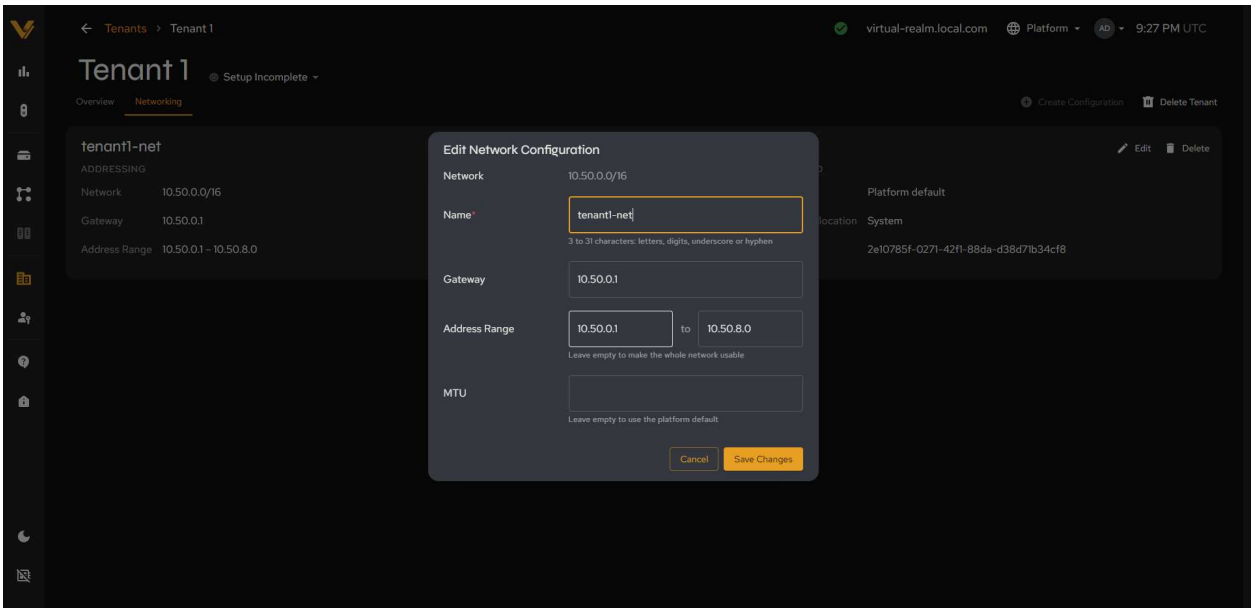


Figure 9.14. The Edit Network Configuration dialog.

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

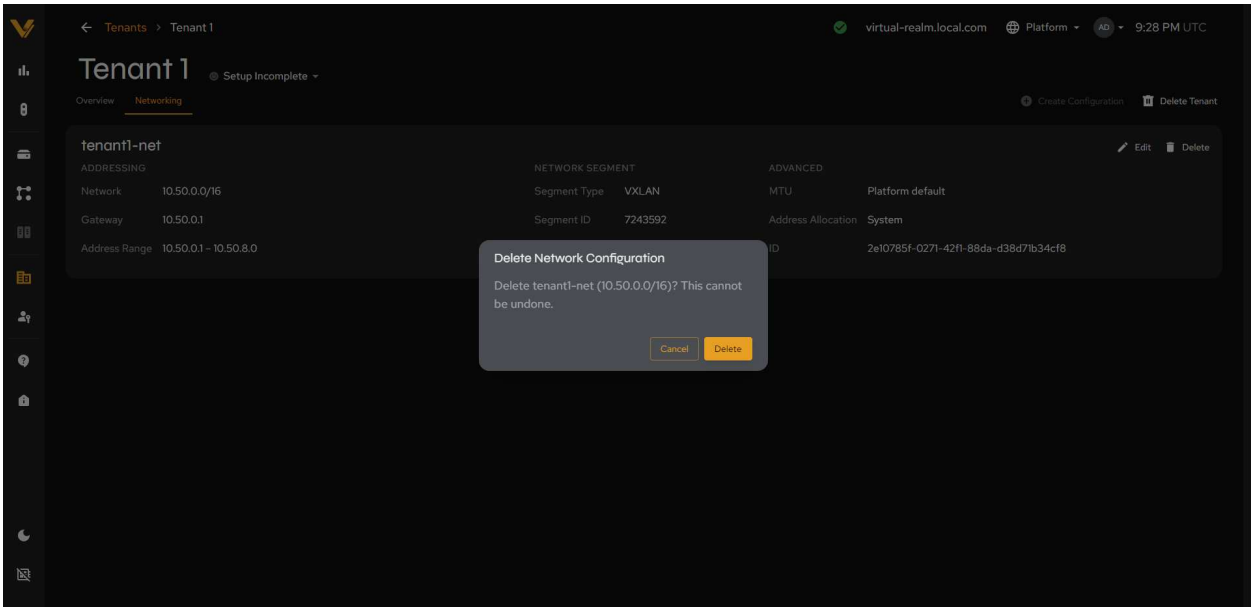


Figure 9.15. The Delete Network Configuration confirmation.

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.”

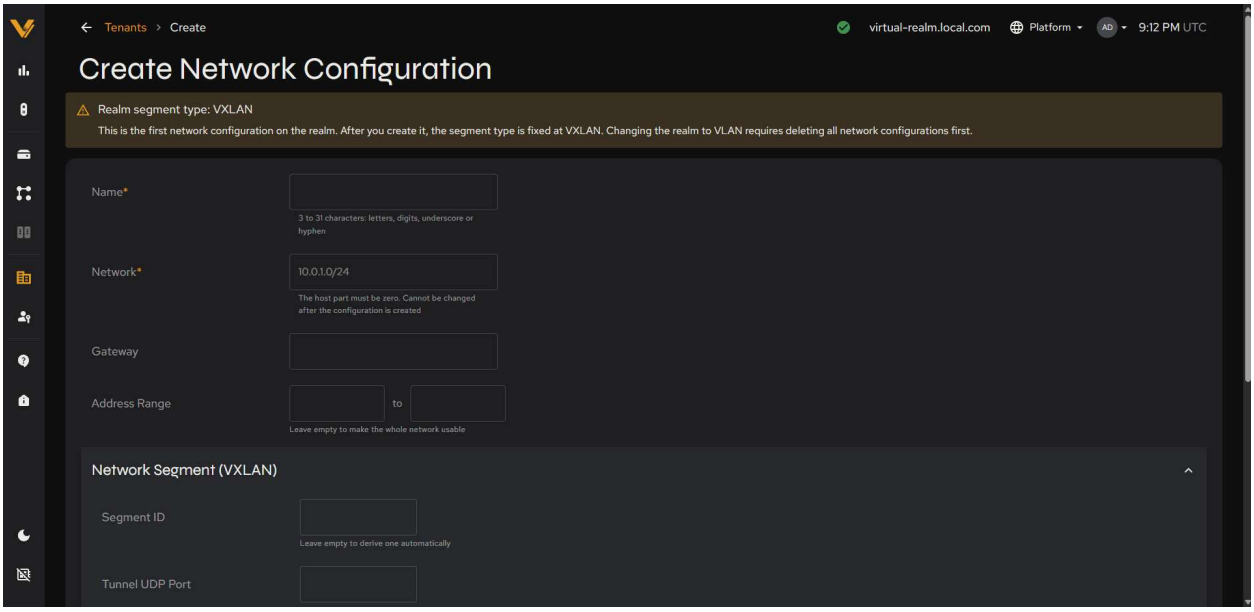


Figure 9.16. The top of the Create Network Configuration form, with the realm segment type banner.

Field	Notes
Name*	Text, 3 to 31 characters, using 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**.

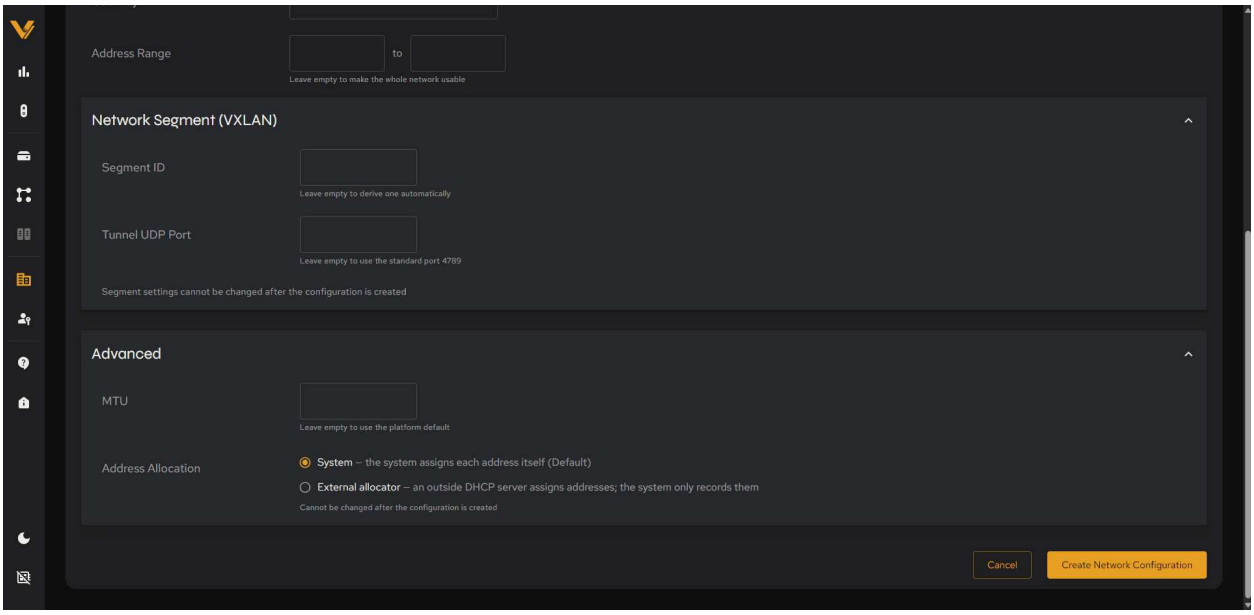


Figure 9.17. The Network Segment and Advanced sections of the Create Network Configuration form.

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.

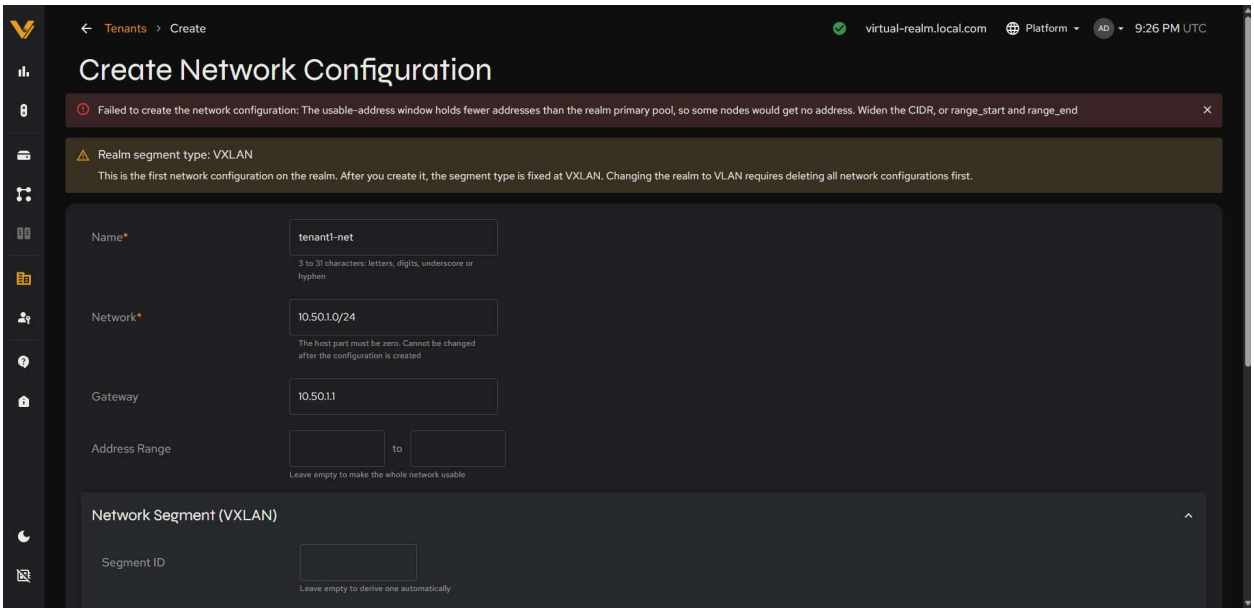


Figure 9.18. The error shown when the network is too small for the realm.

9.3 Create Tenant Page

This page is the form for creating a new tenant.

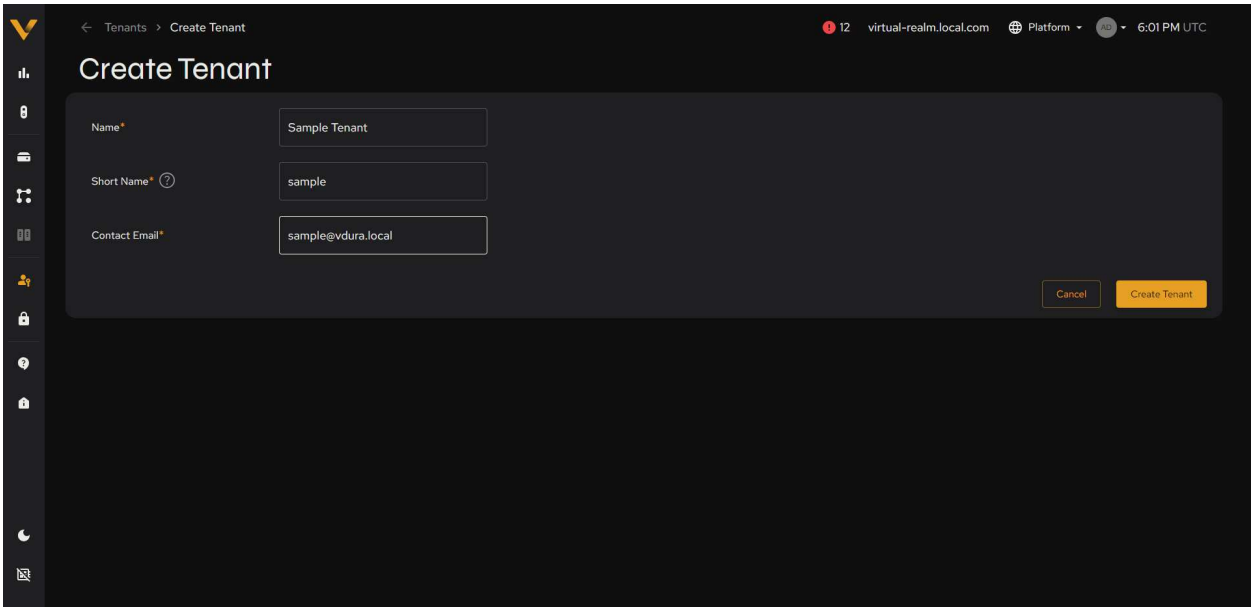


Figure 9.19. The Create Tenant form.

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.

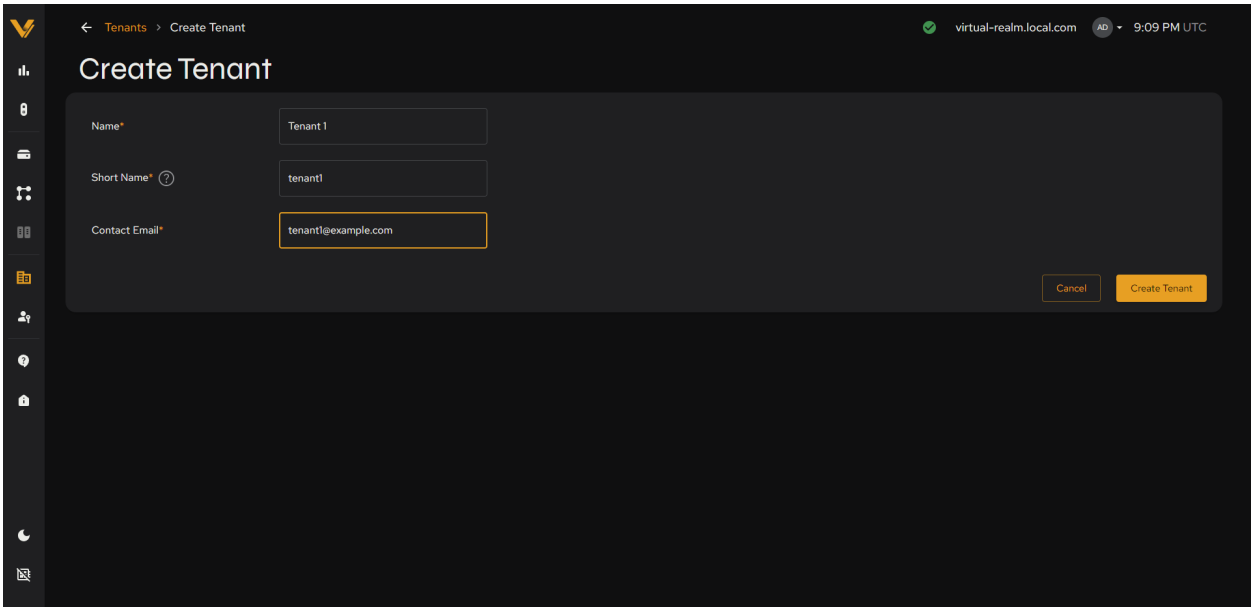


Figure 9.20. The Create Tenant form filled in before submitting.

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** ▾ drop-down menu appears in the top bar, between the status icon and the user drop-down 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 drop-down menu 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 drop-down menu 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, including 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.

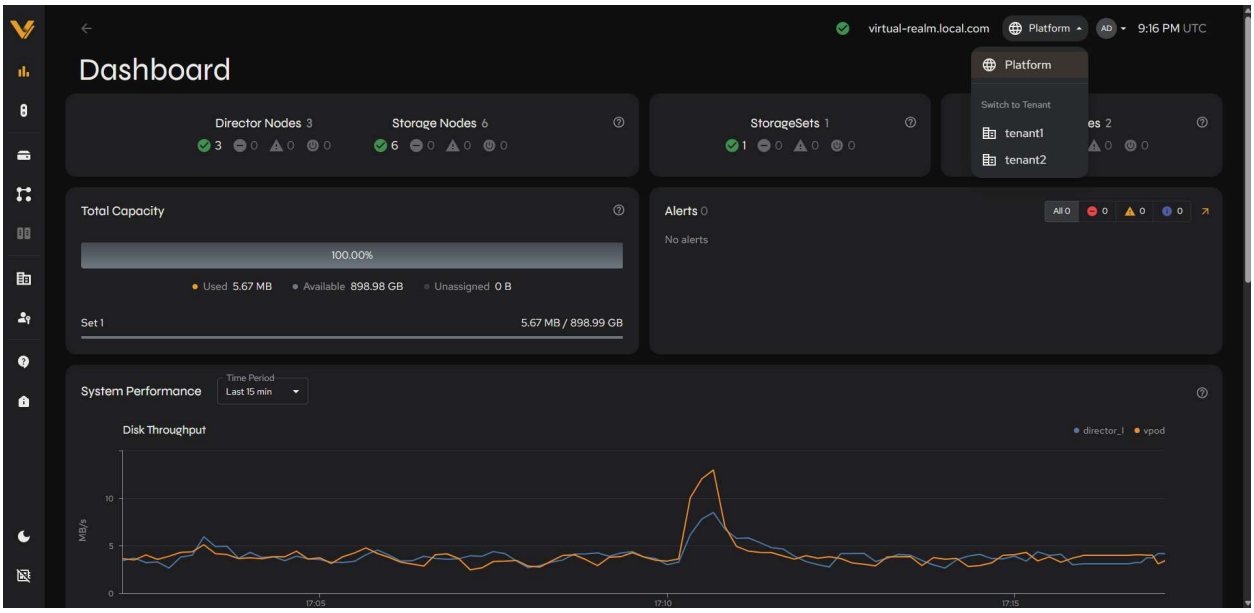


Figure 9.21. The Platform ▾ drop-down menu open, showing Platform plus a Switch to Tenant list.

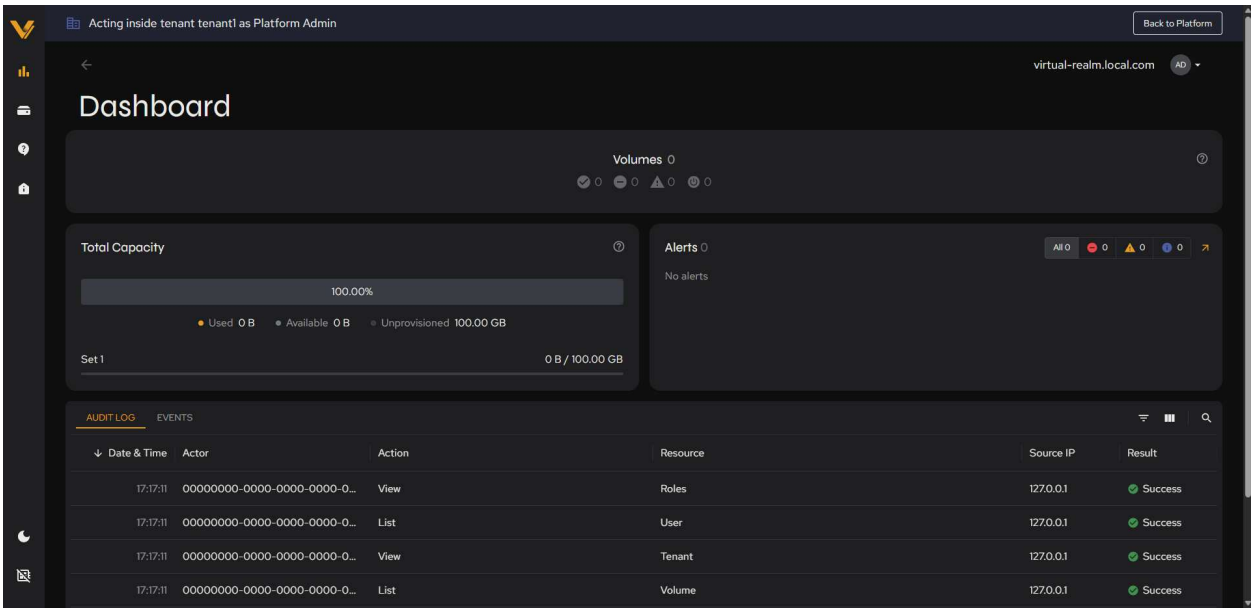


Figure 9.22. The Acting inside tenant `<name>` as Platform Admin banner and reduced sidebar.

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.

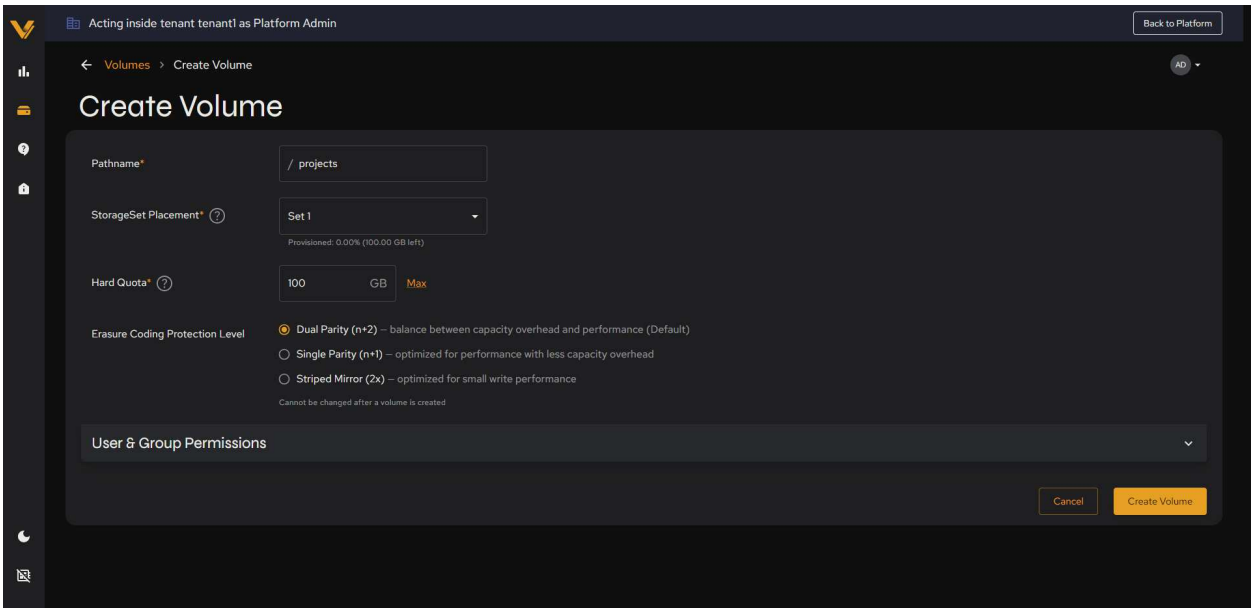


Figure 9.23. Create Volume while acting inside a tenant: required Hard Quota with the Max link, StorageSet labeled with its quota.

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.

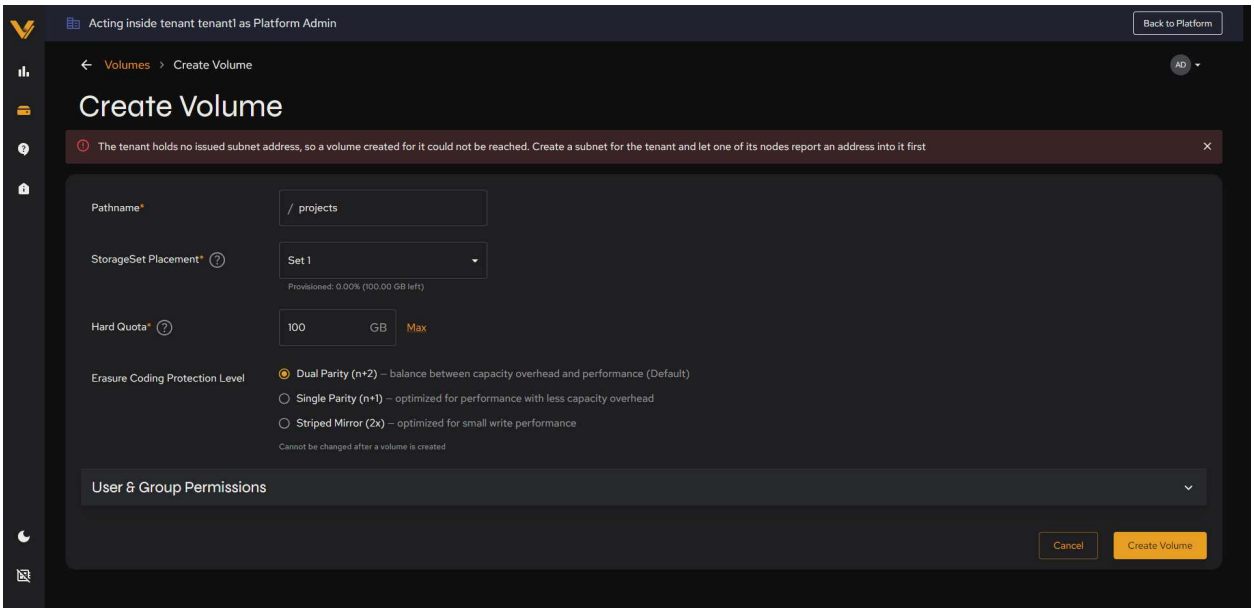


Figure 9.24. The Create Volume error for a tenant with no network configuration.

Volume detail pages also depend on the context. Inside a tenant, a volume's page shows a simpler layout with 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).

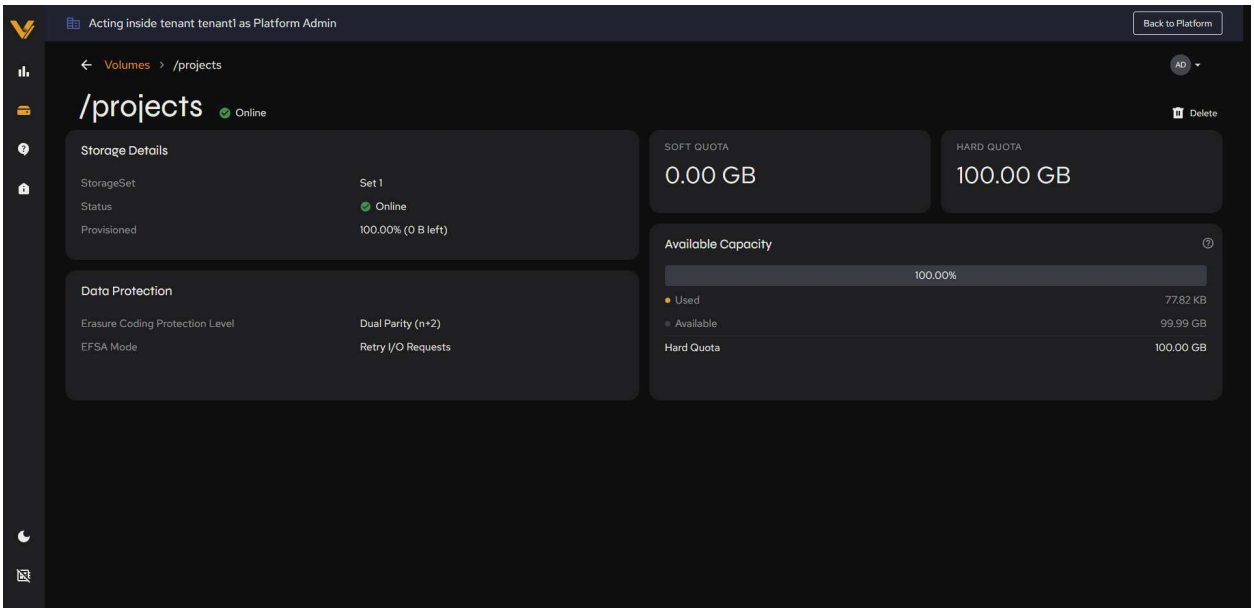


Figure 9.25. A volume's detail page while acting inside its tenant.

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.

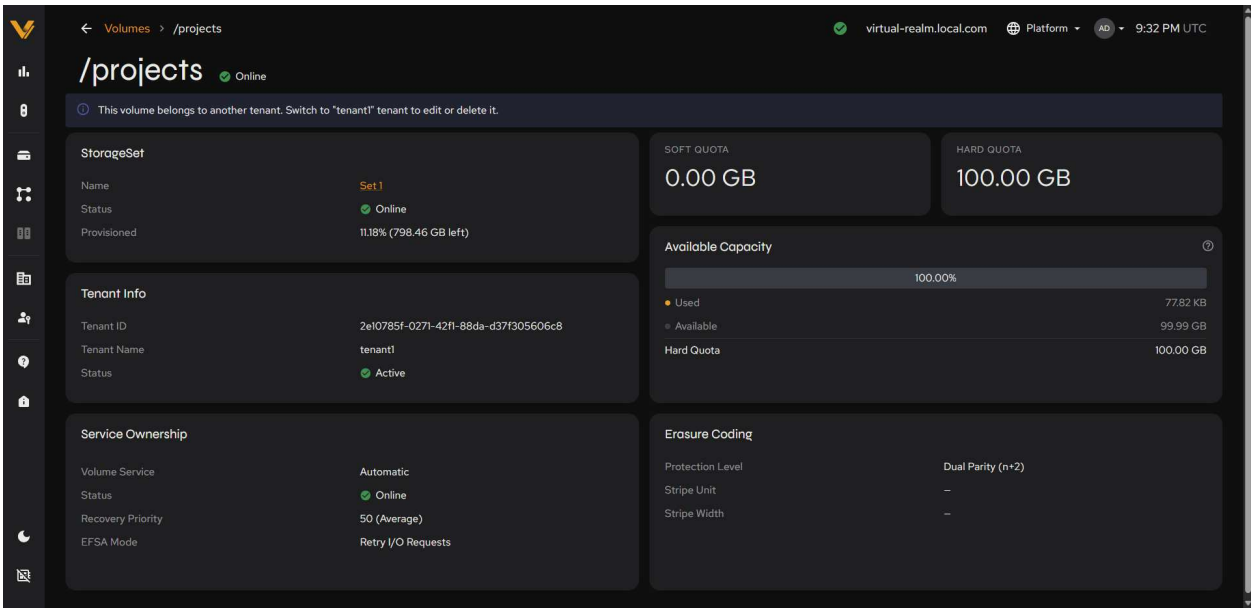


Figure 9.26. A tenant's volume opened from the Platform context, with the belongs-to-another-tenant banner.

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** drop-down menu, 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, `panctl`, 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. Vantage manages users only in the Platform area. Tenant-level users are not fully supported in this release; the API and `panctl` accept them, but Vantage does not manage them, and a tenant has no sign-in of its own here. The 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.

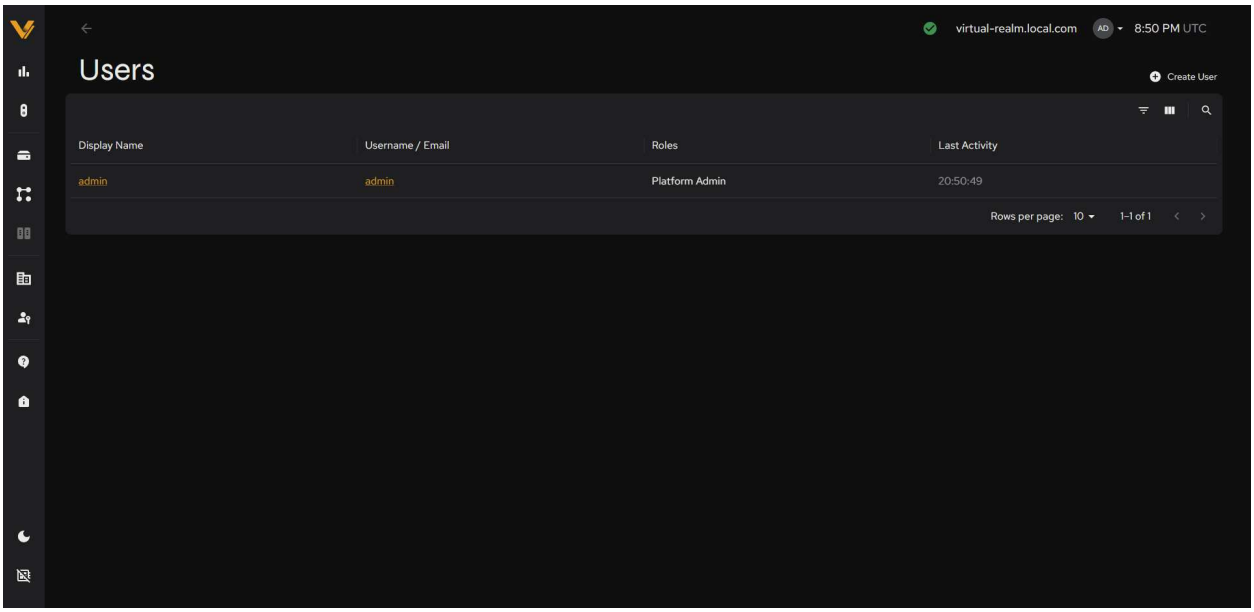


Figure 10.1. The Users table.

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.

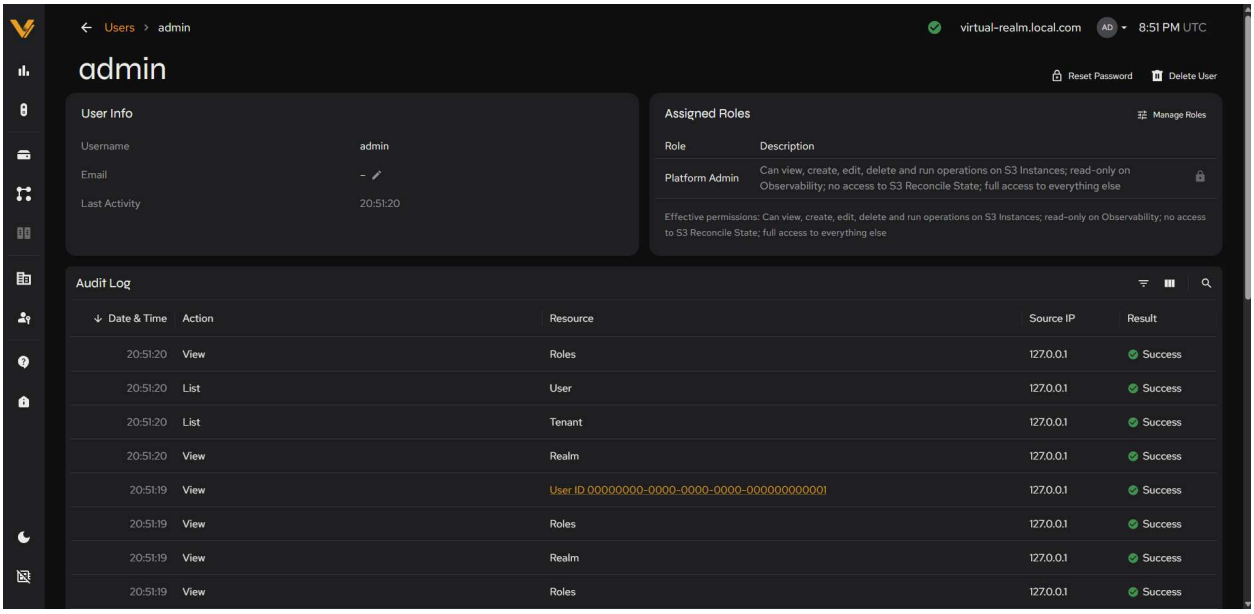


Figure 10.2. Full user detail page (User Info/Assigned Roles/Audit Log).

Card	Fields
User Info	Username, Email (with a pencil icon that opens the Edit Email dialog), 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 Dialog

This dialog 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.

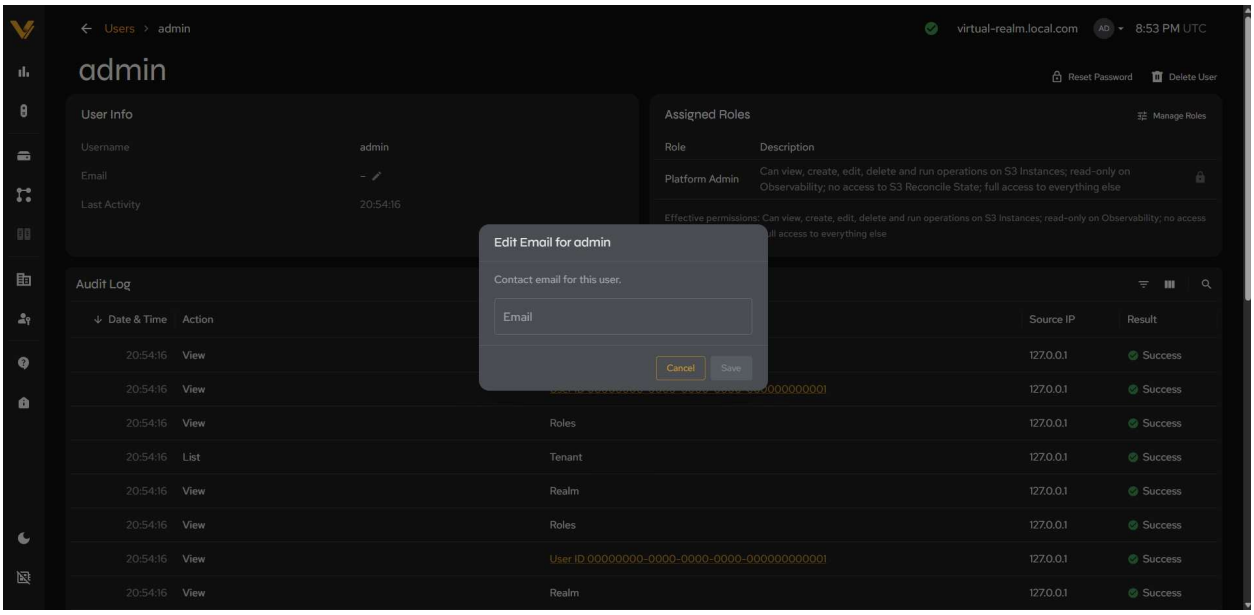


Figure 10.3. The Edit Email dialog.

10.2.2 Reset Password Dialog

This dialog resets a user's password.

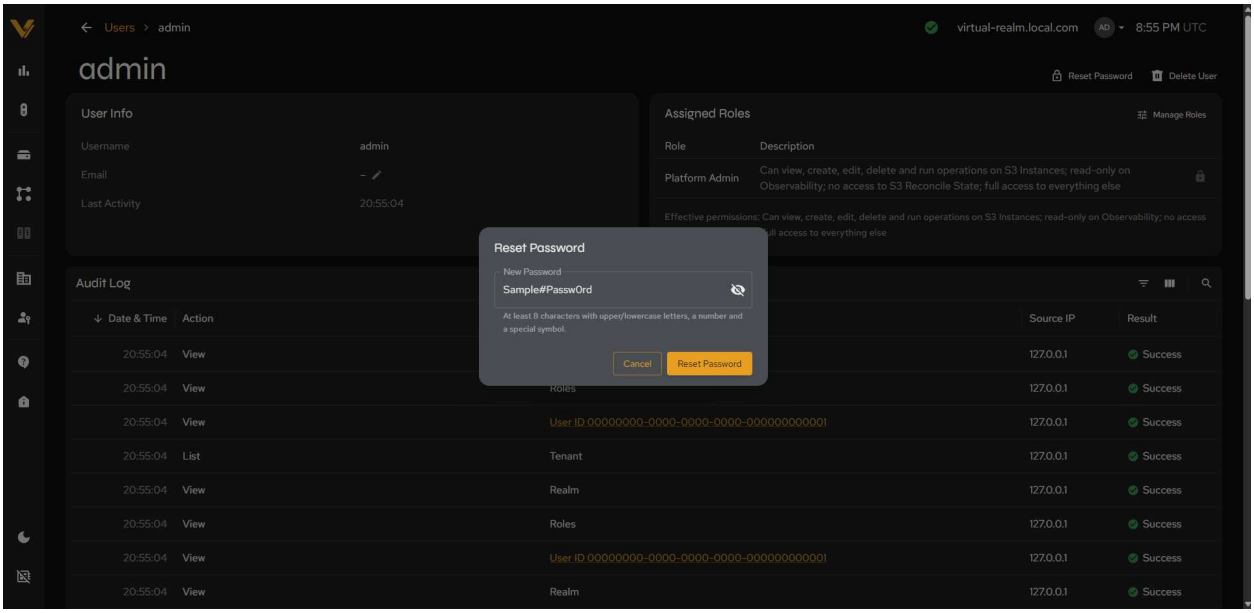


Figure 10.4. The Reset Password dialog, with the password field's show/hide toggled on.

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 Dialog

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

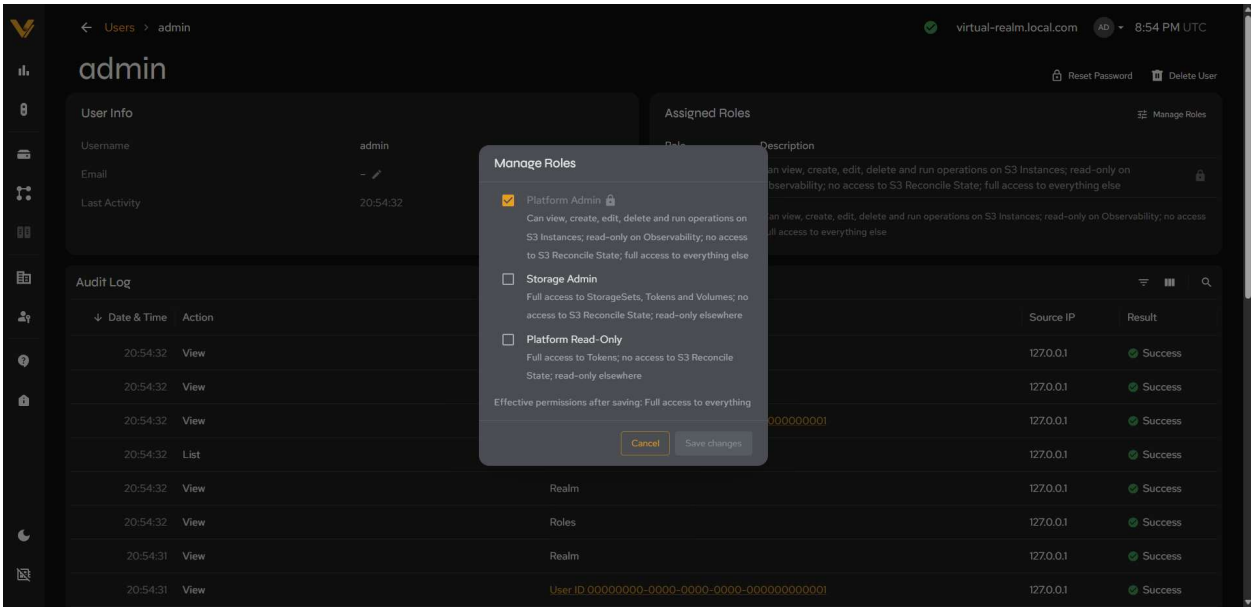


Figure 10.5. The Manage Roles dialog for a user.

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 12); 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. Vantage manages users only in the Platform area; tenant-level users are not fully supported in this release.

10.2.4 Delete User Confirmation

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

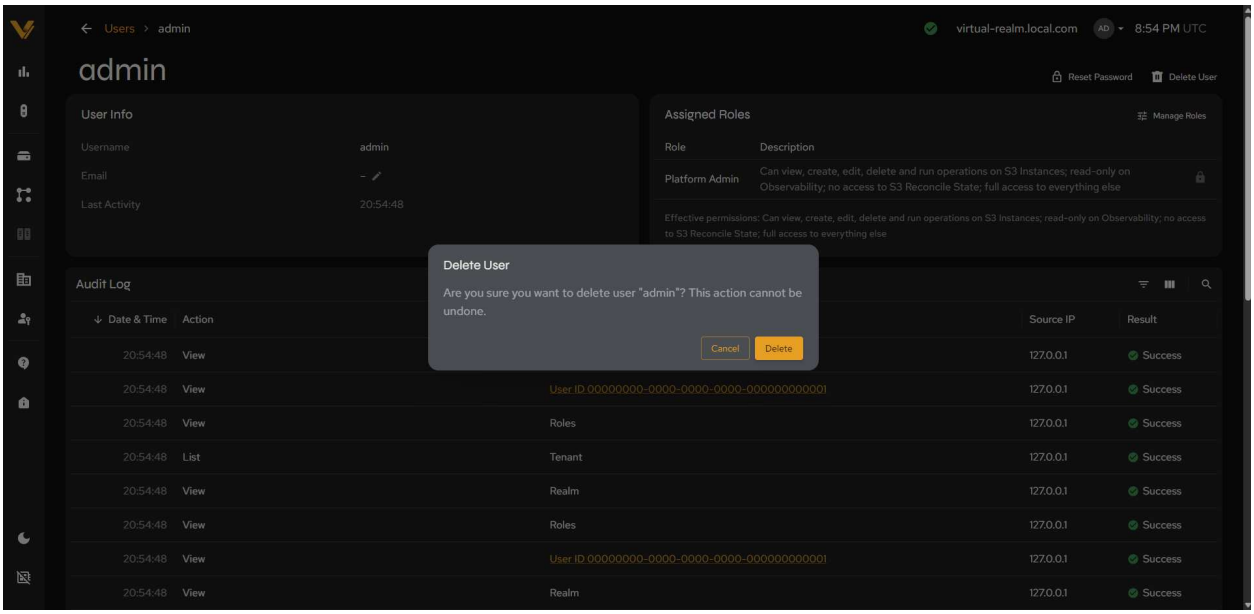


Figure 10.6. The delete confirmation dialog.

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.

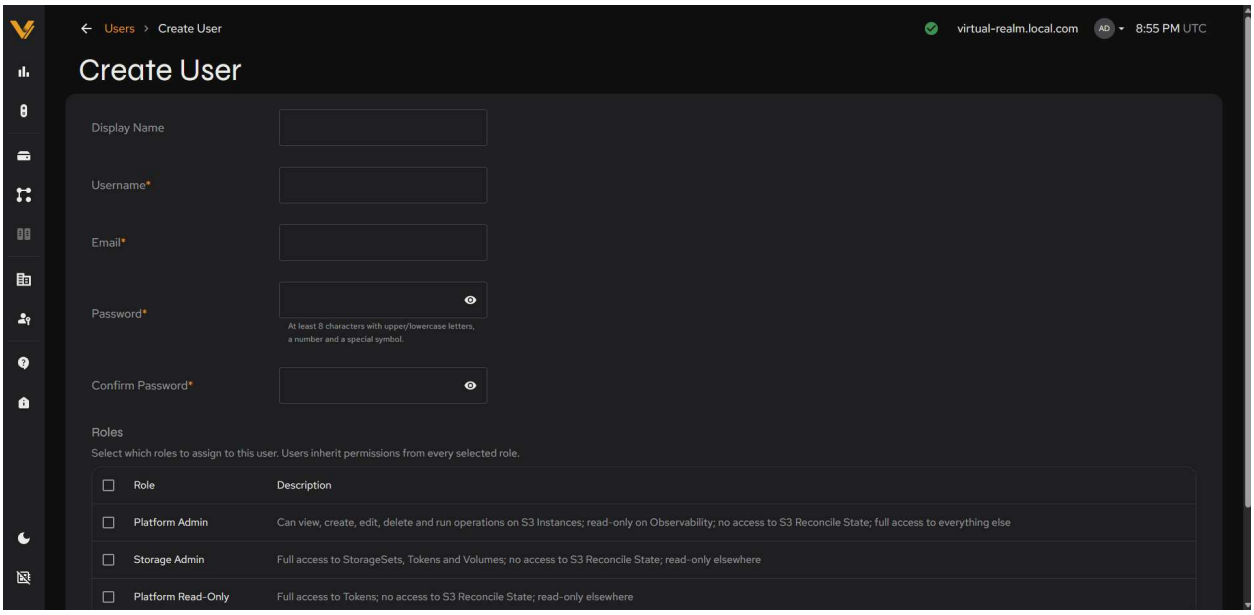


Figure 10.7. The full Create User form, including the Roles checkbox table.

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.

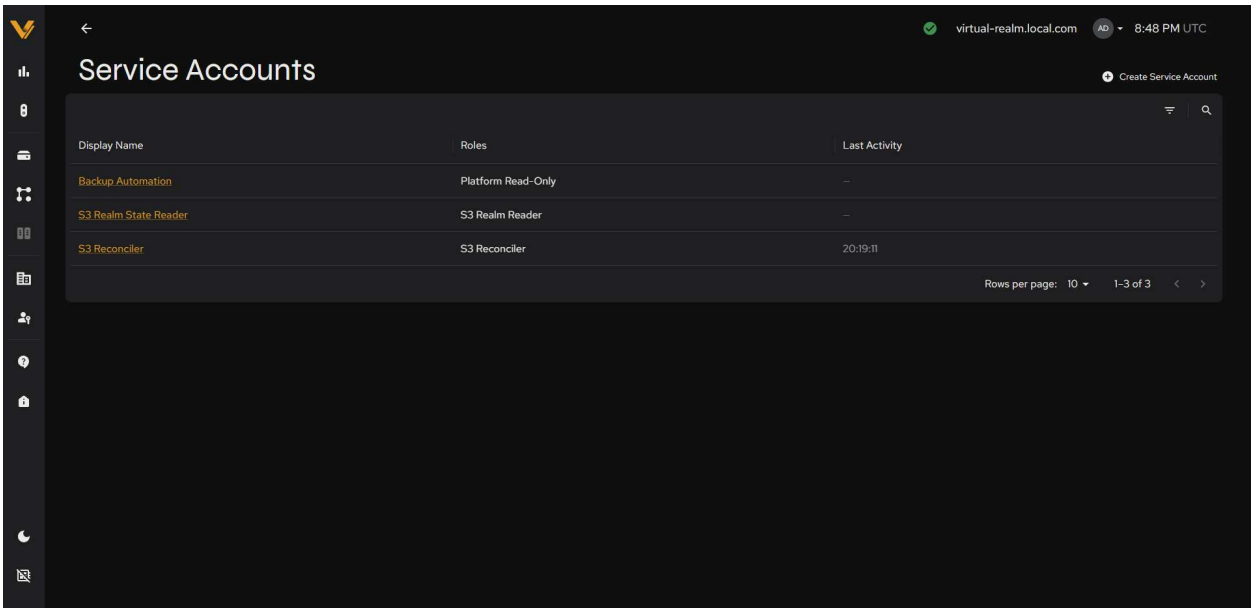


Figure 10.8. The Service Accounts table.

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 12.

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**.

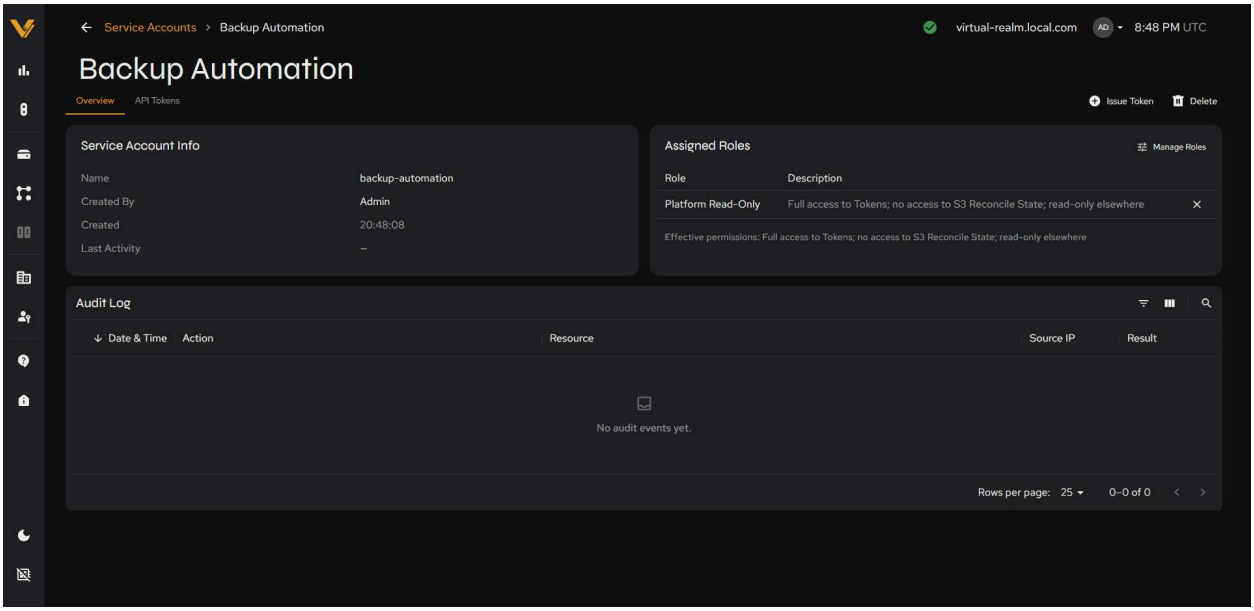


Figure 10.9. Service account detail page, Overview tab (Service Account Info/Assigned Roles/Audit Log).

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.

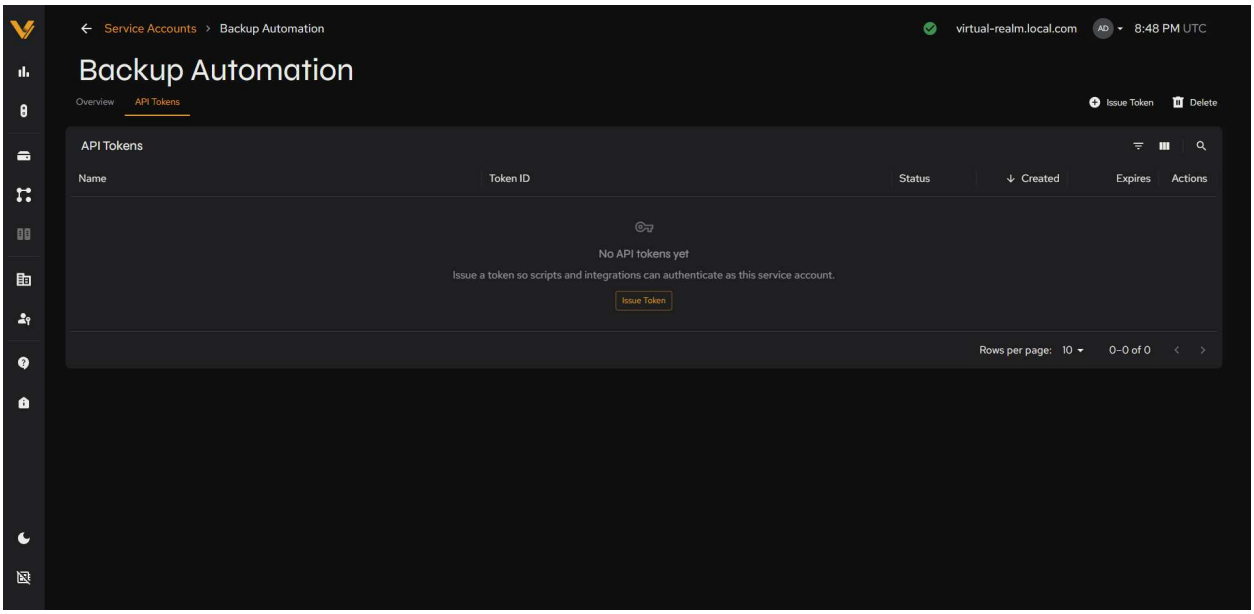


Figure 10.10. The API Tokens tab with no tokens issued.

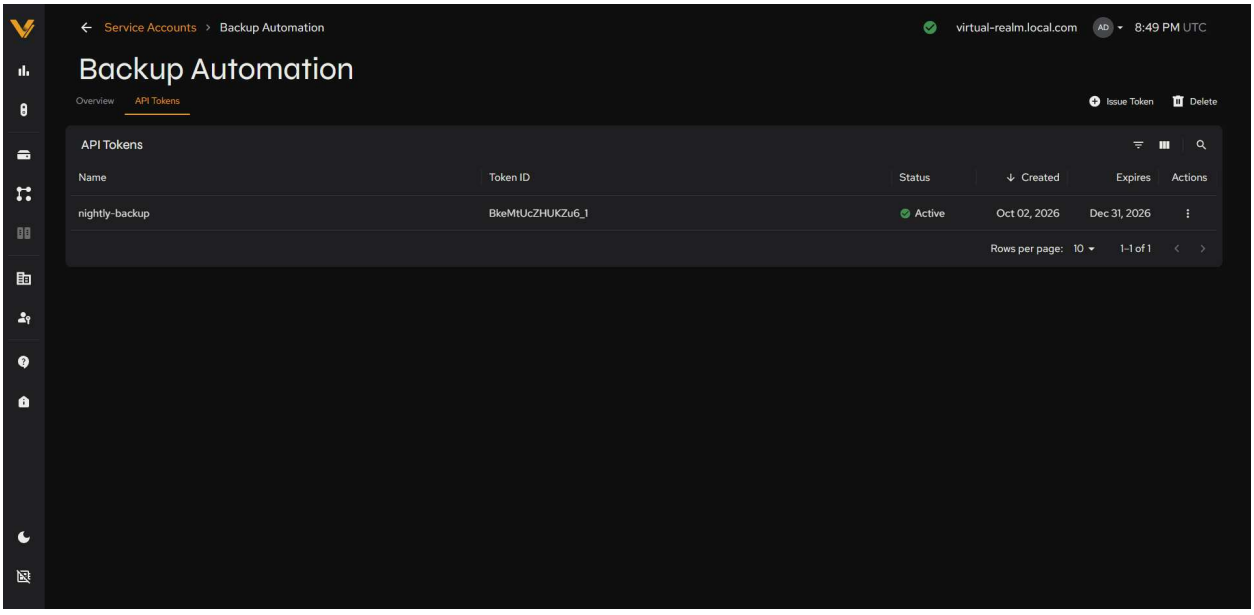


Figure 10.11. The API Tokens tab after a token has been issued.

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 Dialog

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

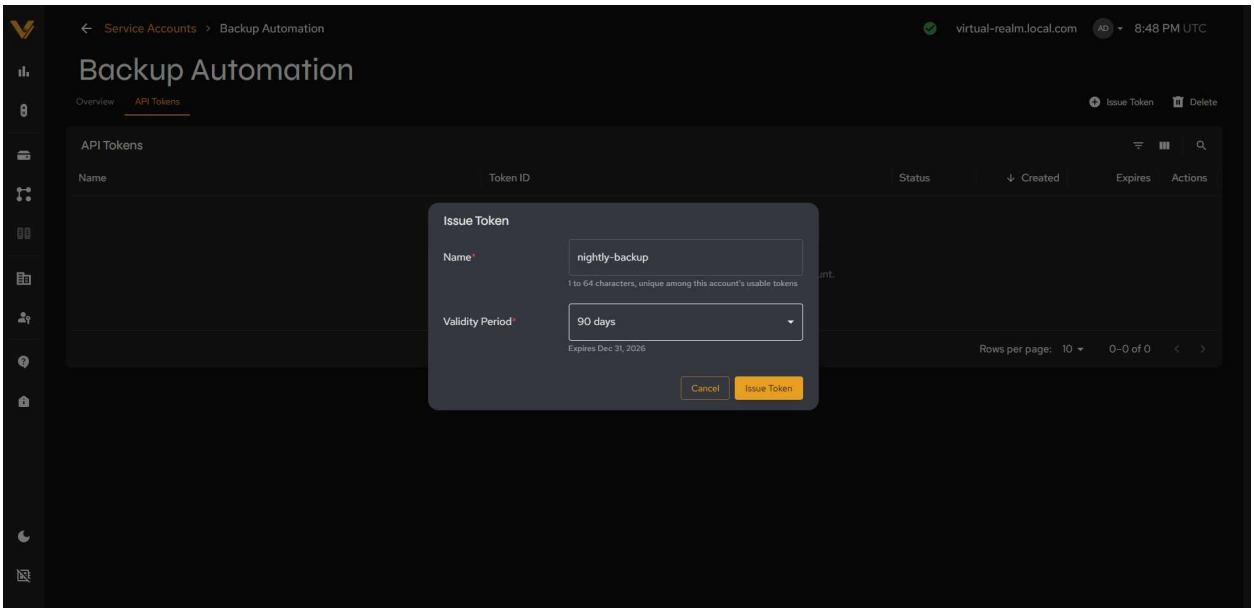


Figure 10.12. The Issue Token dialog.

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

Field	Notes
Validity Period*	Drop-down menu 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 Dialog

After you click **Issue Token**, a “Token Issued” dialog 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 dialog.

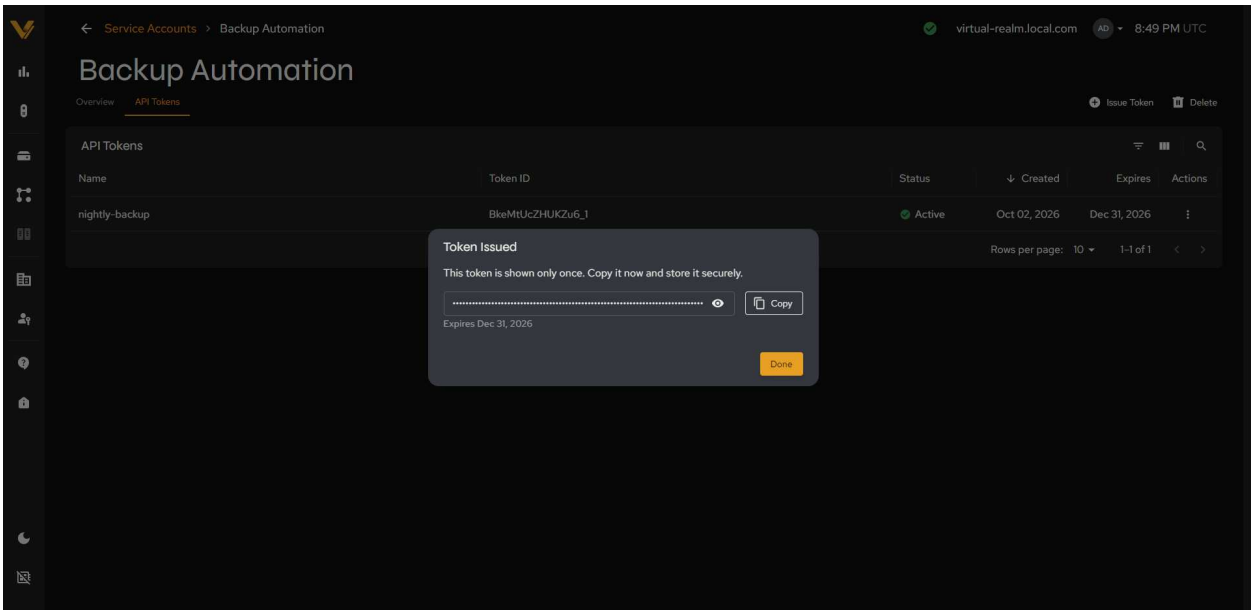


Figure 10.13. The Token Issued dialog.

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

10.6.4 Manage Roles Dialog

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

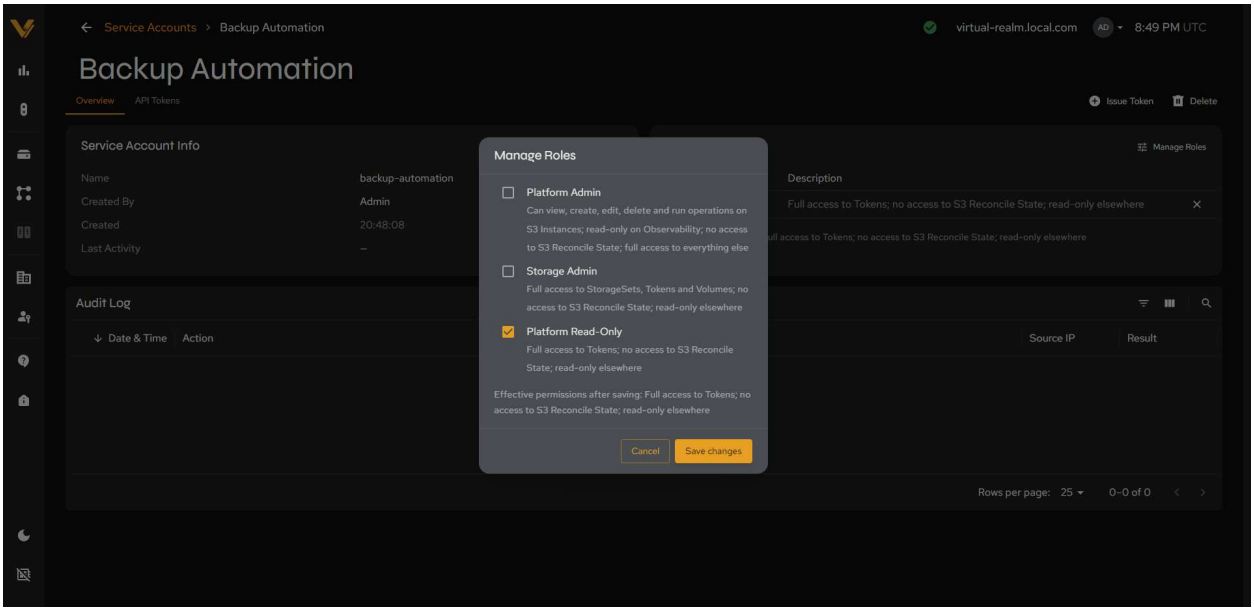


Figure 10.14. The Manage Roles dialog for a service account.

10.6.5 Delete Service Account Confirmation

This dialog 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.”

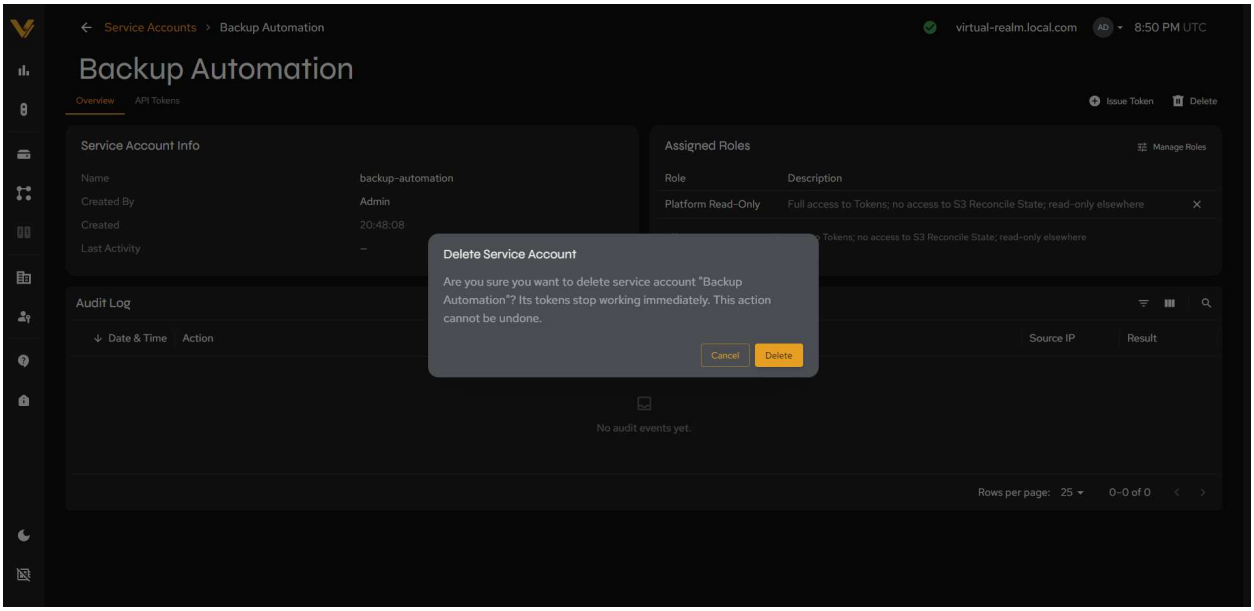


Figure 10.15. The delete confirmation dialog for a service account.

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**.

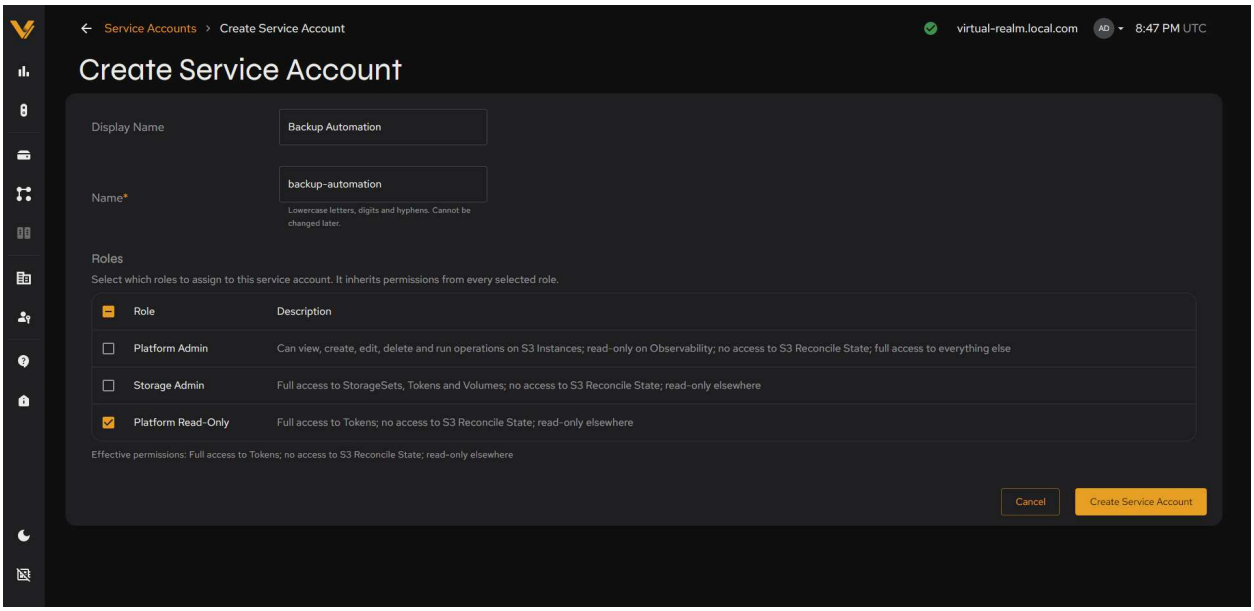


Figure 10.16. The Create Service Account form.

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 dialog, 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

This chapter covers Account Menu and My Account, Theme, Get Help, About, and PanActive Manager.

11.1 Account Menu and My Account

Clicking the user drop-down menu in the top bar opens a small drop-down menu.

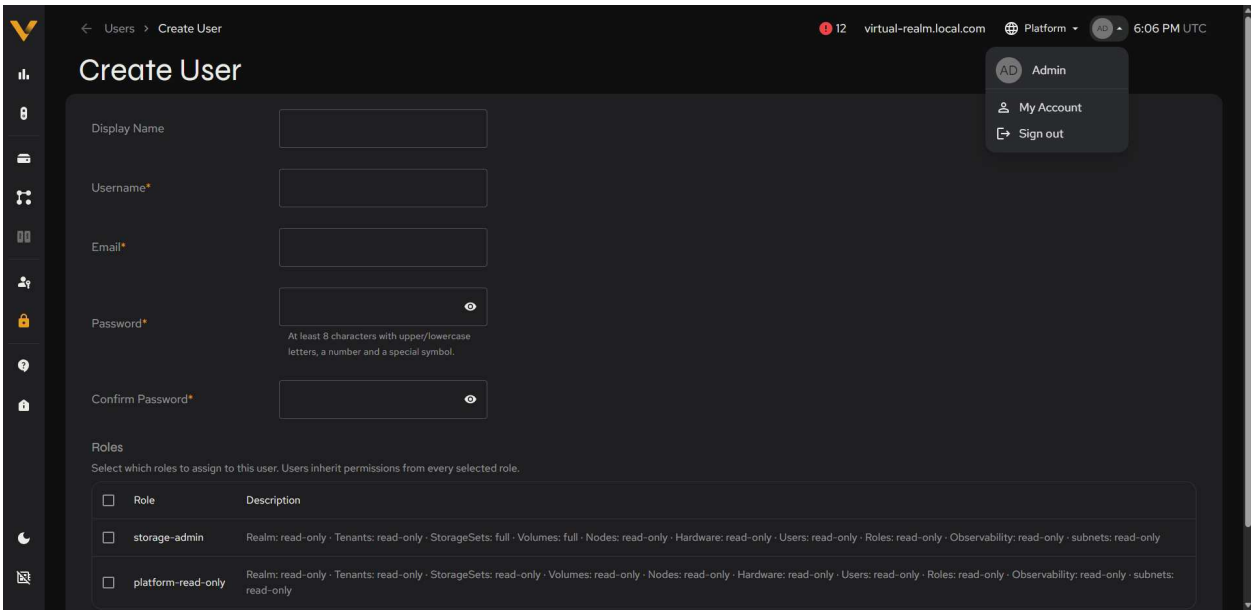


Figure 11.1. The account drop-down menu open (Admin/My Account/Sign out).

The drop-down menu 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.

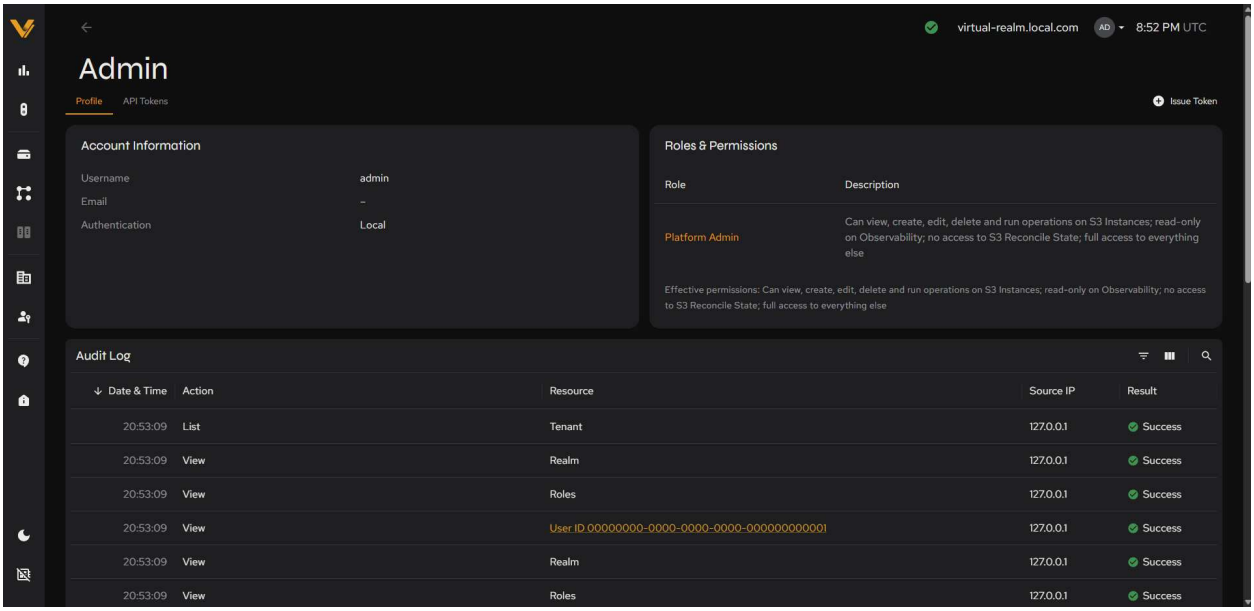


Figure 11.2. The full My Account page.

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** dialog, 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."

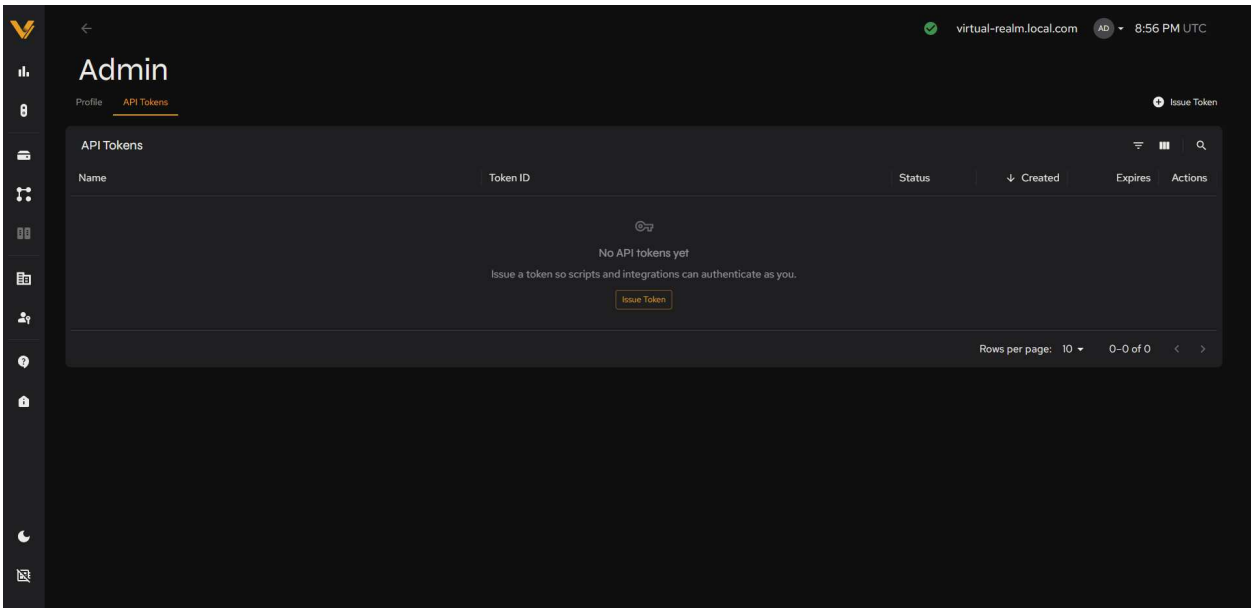


Figure 11.3. The My Account API Tokens tab.

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.

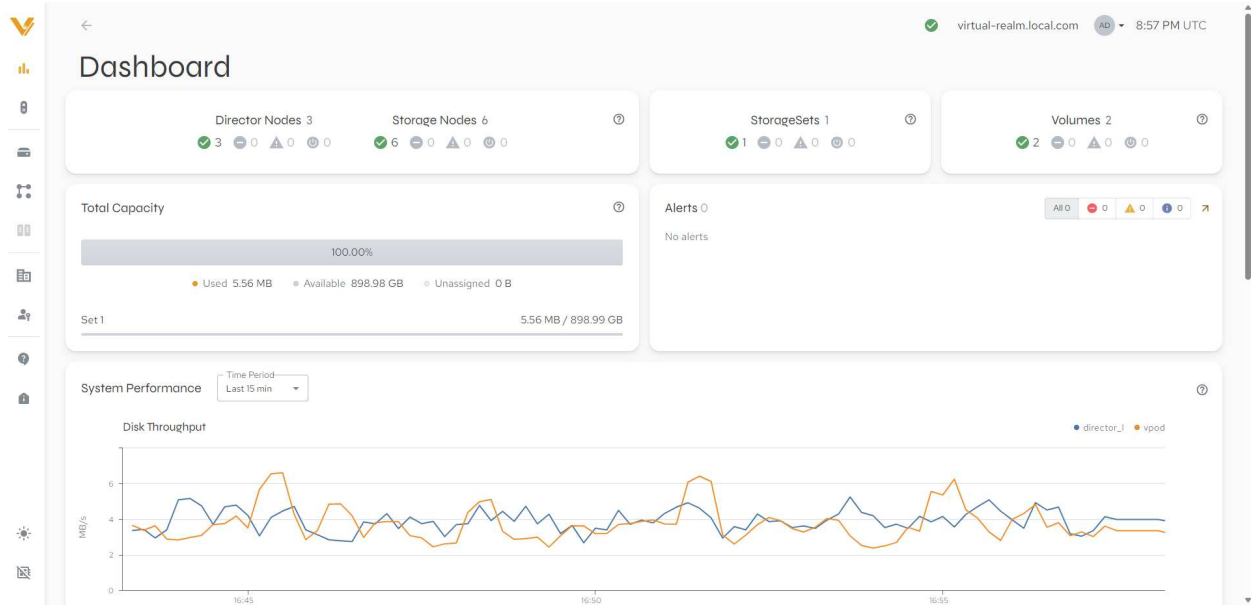


Figure 11.4. The Dashboard (or any page) rendered in light mode, for comparison.

11.3 Get Help

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

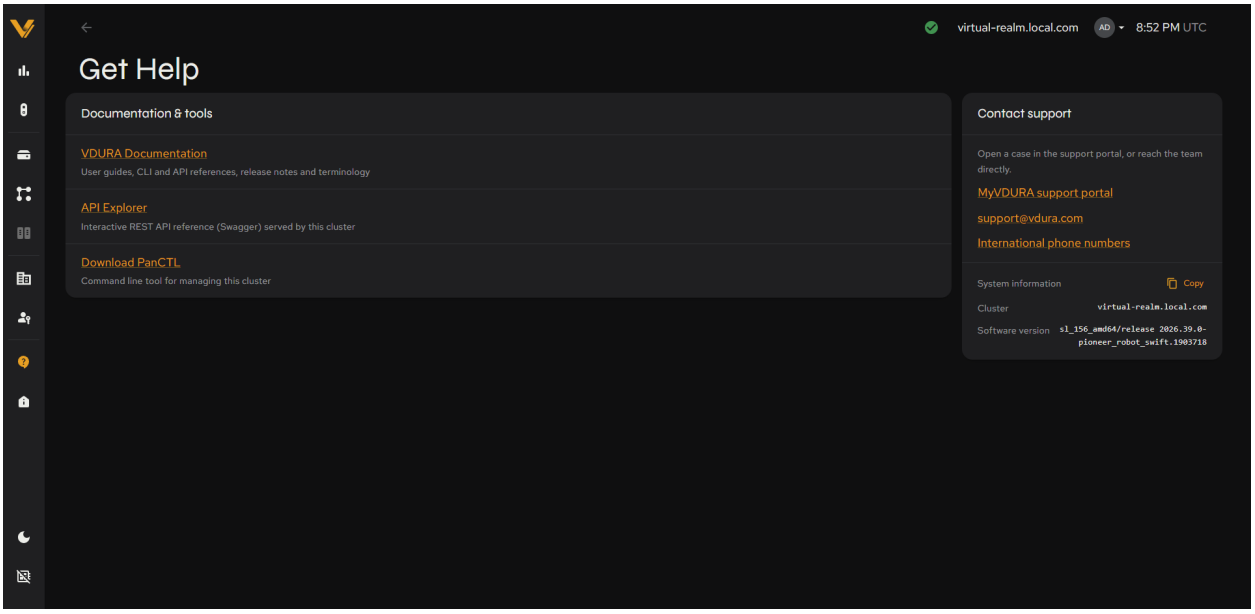


Figure 11.5. The Get Help page.

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 Vantage, titled “About Vantage”.

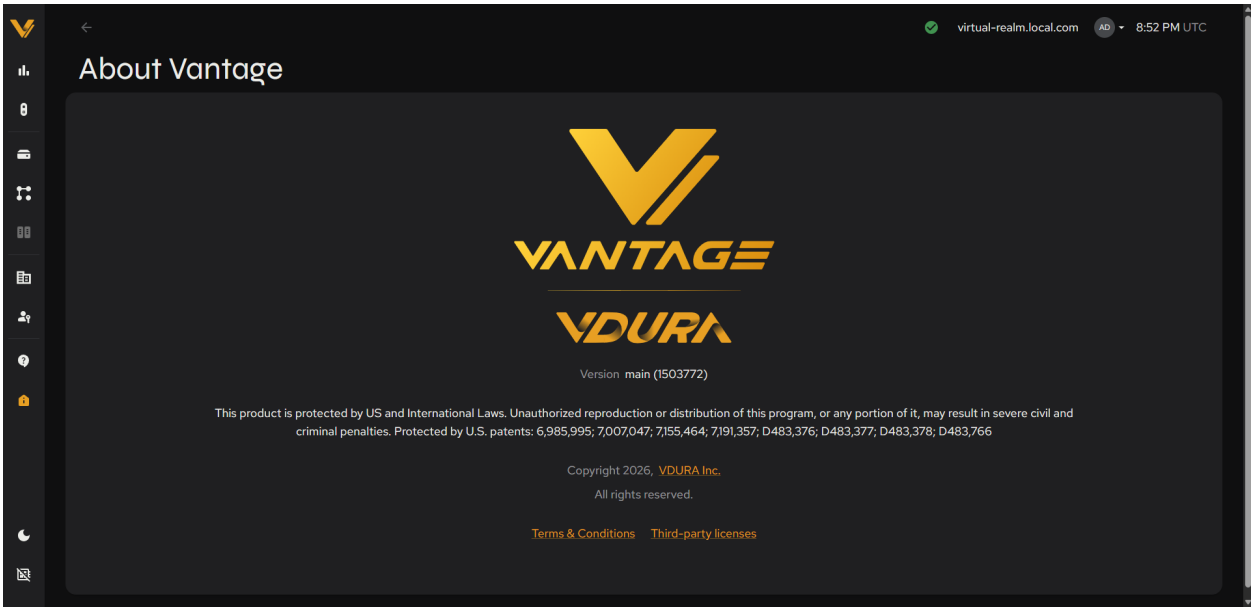


Figure 11.6. The About page.

Item	Description
Logos	The Vantage and VDURA logos
Version	The Vantage UI build version, e.g. `main (1503772)`
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.

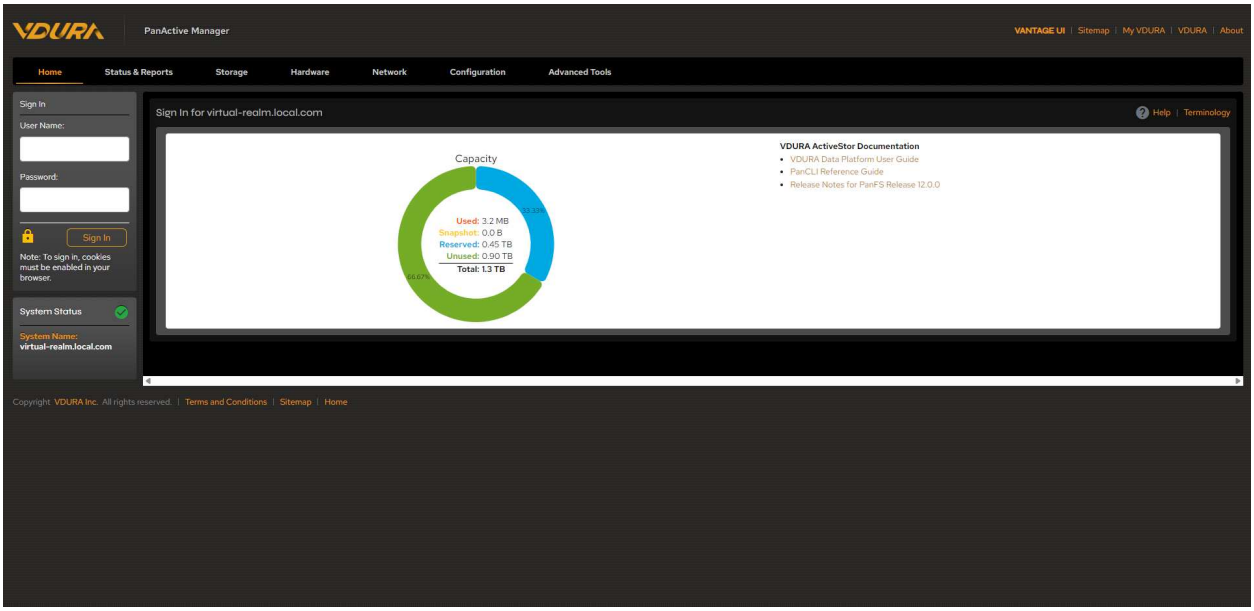


Figure 11.7. The landing and sign-in page of PanActive Manager, showing its own top nav and Sign In form.

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.

> For full coverage of PanActive Manager, see the VDURA Data Platform User Guide.

12. S3

> **Note:** S3 is configured with `panctl` 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 "Gateway Protocols" in "End-to-End Encryption" in the *VDURA Data Platform User Guide*).

This chapter shows a **platform administrator** how to set up, run, and use the PanFS S3 service, including enabling S3, gateway instances, S3 users, buckets, TLS, named access, monitoring, client tools, and day-to-day operations. It describes PanFS 12.0.0.

12.1 Concepts

The following table defines each term.

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, 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.

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

12.2 Logging In with panctl

`panctl` is installed on every Director at `/pan/bin/panctl`. You do not need to run it as root. If `/pan/bin` is not on your `PATH`, type the full path.

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.

To find the Director that hosts the management API, run the following command on a Director. If a line with `LISTEN` appears, that Director hosts the API. If nothing appears, run the command on another Director.

```
ss -ltn | grep ':8000'
```

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
```

12.3 Deployment Options

This section covers Dedicated Tenant (Recommended), Platform Tenant, and Choices Within a Tenant.

12.3.1 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 12.3.2) does not use a tenant subnet.

12.3.2 Platform Tenant

The built-in `platform` tenant can also serve S3 (see 12.17). Its gateways listen on the nodes' own realm addresses, so any host that can reach the realm network can reach them.

12.3.3 Choices Within a Tenant

The following table lists each choice, its options, and the recommendation.

Choice	Options	Recommendation
Instance nodes	Directors	Several Directors, for availability (see 12.6)

Choice	Options	Recommendation
Named access	`off` (default) or `external`	`external` : You publish the names in your own DNS (see 12.10)
TLS certificate	`realm-generated` (default) or `tenant-provided`	`tenant-provided` , so S3 clients can verify the endpoint (see 12.9)

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

12.4 Prerequisites

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

12.4.1 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 --storageset-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`).

12.4.2 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.

12.4.3 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": 8950,
  "segment_type": "vxlan",
  "segment_id": 10100,
  "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. Create the tenant interface on the host with the values from that response, as described in 13.10.2 for a VXLAN tenant network or 13.10.3 for a VLAN tenant network. The segment, MTU, and peer list in the response are the values that section uses.

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

12.4.4 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 --storageset-id 1 --hard-quota
        100GiB
panctl volumes list --tenant acme
```

12.5 Enabling S3

This section covers Realm Switch, Tenant Configuration, and Settings Reference.

12.5.1 Realm Switch

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

12.5.2 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
  MODE:            realm-generated
  CA CERTIFICATE:  -
  CERTIFICATE:     -----BEGIN CERTIFICATE-----\nMIID...\n-----END CERTIFICATE-----\n
  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`.

12.5.3 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
`--enabled`	`false`	Tenant S3 on/off	n/a
`--prefix`	empty	Volume path (or a directory inside a volume) served by S3, e.g. `/s3data`	n/a
`--owner`	`uid:65534`	Default file owner, used for S3 users that have no uid of their own	after restart
`--group`	`gid:65534`	Default file group, used for S3 users that have no gid of their own	after restart
`--dir-mode`	`755`	Mode of directories created through S3 (3–4 octal digits)	after restart
`--obj-mode`	`644`	Mode of files created through S3 (3–4 octal digits)	after restart
`--access-log`	`true`	Per-request access log (see 12.14)	after restart
`--debug-log`	`false`	Accepted, but has no visible effect on gateway logging in this release	after restart

Flag	Default	Meaning	Running instances pick it up
`--metrics-enabled`	`true`	Publishes S3 metrics (see 12.14)	after restart

The TLS and DNS settings are covered in sections 12.9 and 12.10. 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 12.15). 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.

12.6 Gateway Instances

This section covers Choose Nodes, Create and Start, Endpoints, States and Health, and Stop, Start, Delete.

12.6.1 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.

12.6.2 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
```

12.6.3 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`.

12.6.4 States and Health

The following table lists each field and its possible values.

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 12.10.1), so clients are not tied to a single node.

12.6.5 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}'
```

12.7 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 `user` can, and also create its own buckets over S3.
`disabled`	Nothing: The credential is kept, but the gateways reject it (`InvalidAccessKeyId`).

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 12.16).

12.8 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 12.18).

12.9 TLS Certificates

All S3 endpoints use HTTPS. Each tenant has one certificate, used by all of its instances. Certificate changes take effect within seconds, without restarting instances.

12.9.1 Tenant-Provided Certificate (Recommended)

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

Requirements

The following table lists each item and its requirement.

Item	Requirement
Format	PEM. `certificate` holds the server certificate **followed by any intermediate CA certificates**. `private_key` 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.
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 12.6.3). 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 12.9.2). Include them in a tenant-provided certificate if clients use them.

Realm DNS zone	Realm names (tenant `acme`)
not set	`*.s3-acme`
`example.com`	`*.s3-acme.example.com`

Check the zone with `pancli -c "dns zone"` on a Director, and set it with `pancli -c "dns zone 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 12.6.3).
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. Issue it from a certificate authority that your clients trust. The certificate carries the realm's names for zone `example.com`, which section 12.10.1 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.example.com"
cat > s3-acme.ext <<'EOF'
subjectAltName=DNS:s3-acme.example.com,DNS:*.s3-acme.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.example.com:3000 --ca-bundle ca.crt s3 ls
```

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

12.9.2 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 12.11.4). 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
```

12.10 Named Access (DNS)

Named access is off by default, and clients use the IP endpoints. `panctl s3 dns off` turns it off again.

12.10.1 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 `example.com`)	Records	Port
`<iid>.s3-acme.example.com`	one A record per Director hosting that instance	3 0 0 0

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

Check from a client:

```
dig +short gw001.s3-acme.example.com      # 10.240.0.1 10.240.0.2
aws --endpoint-url https://gw001.s3-acme.example.com:3000 --ca-bundle ca.crt s3 ls
aws --endpoint-url https://gw004.s3-acme.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.example.com:3000/<bucket>/<key>`). Bucket names in the host name (virtual-hosted style) are not supported.

12.11 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 12.12.4) with `InvalidArgument` ;
  - presigned URLs with `InvalidRequest` .
- **Addressing:** Path-style.

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

12.11.1 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.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>`.

12.11.2 Python (boto3)

```
import boto3
from botocore.config import Config

s3 = boto3.client(
    "s3",
    endpoint_url="https://gw001.s3-acme.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`).

12.11.3 rclone

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

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

12.11.4 Connecting to a Realm-Generated Certificate

A realm-generated certificate cannot be verified (see 12.9.2), so turn verification off in the client.

> **Caution:** Turning verification off removes the client's protection against a connection to something other than the gateway. Use it only until you install a certificate from a certificate authority that your clients trust (see 12.9).

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

```
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.

12.12 Supported S3 Features

The following table lists each area and what is and is not supported.

Area	Supported	Not supported
Buckets	create, delete (empty), list, head, location (`us-east-1`), tagging, CORS configuration	versioning (`VersioningNotConfigured`), lifecycle, encryption, replication, logging, notifications, public access block (`NotImplemented`)
Objects	put (single PUT up to 5 GiB), get (incl. byte ranges), head, delete, multi-object delete, copy, user metadata (`x-amz-meta-*`), tagging, conditional put (`If-None-Match: *`), checksums (CRC32, SHA-256)	single PUT over 5 GiB (`EntityTooLarge` ; use multipart), S3 Select (`NotImplemented`), 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
Browser uploads (POST)	HTML form POST with a Signature V4 policy (below)	Signature V2 forms (`InvalidArgument`)
Access control	bucket owner, bucket policies, bucket ACLs once the owner turns them on (below)	object ACLs (`NotImplemented`), AWS grantee forms in grants (`InternalError`), 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 (`VersioningNotConfigured`)

12.12.1 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.

12.12.2 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 (`--acl`)	`private` , `public-read` , `public-read-write`	every other canned ACL, `authenticated-read` included (`InvalidArgument`)
Grants (`--grant-read` , `--grant-write` , `--grant-full-control` , …)	a bare access key, e.g. `--grant-read bob`	the AWS forms `id=...` , `uri=...` and `emailAddress=...` (`InternalError`)
Full policy (`--access-control-policy`)	`CanonicalUser` grantees; the `ID` is the S3 user's access key	`Group` grantees (`MalformedACLError`); 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
`public-read`	list the bucket, **read every object**
`public-read-write`	list the bucket, read every object, **upload and delete objects**

12.12.3 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.

12.12.4 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.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.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'`.

12.12.5 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: `AccessDenied` ("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	`--object-lock-mode GOVERNANCE` and `--object-lock-retain-until-date` on upload or `put-object-retention`	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	`put-object-legal-hold ... Status=ON`	`Status=OFF`; cannot be bypassed
Bucket default retention	`put-object-lock-configuration` 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 12.13), locked files can be changed and deleted with ordinary file permissions.

12.12.6 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	`PutBucketLifecycleConfiguration`, `GetBucketLifecycleConfiguration`, `DeleteBucketLifecycle`	Expire objects from a client job. Clean up unfinished multipart uploads with `ListMultipartUploads` and `AbortMultipartUpload`.
Bucket default encryption	`PutBucketEncryption`, `GetBucketEncryption`, `DeleteBucketEncryption`	None. See Server-Side Encryption.
Public access block	`PutPublicAccessBlock`, `GetPublicAccessBlock`, `DeletePublicAccessBlock`	None. See Public Access.
Event notifications	`PutBucketNotificationConfiguration`, `GetBucketNotificationConfiguration`	None.
Server access logging	`PutBucketLogging`, `GetBucketLogging`	Each gateway writes an access log on its node (see 12.14).
Replication	`PutBucketReplication`, `GetBucketReplication`, `DeleteBucketReplication`	None.

12.12.7 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.

12.13 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 12.7);

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 12.4.3) 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.

12.14 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 (its panels and queries are in section 4.6 of the

VDURA PanFS Metrics Reference Guide). To get the dashboard file, request it from your VDURA representative or contact VDURA Technical Support.

- 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.

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

12.15 Day-to-Day Operations

The following table lists each task and the commands that perform it.

Task	Commands
Stop all S3 in the realm	`panctl s3 enabled set false`: Every gateway stops; tenant settings and instance states are kept
Restart all S3 in the realm	`panctl s3 enabled set true`: Instances that were started come back
Apply a setting that needs a restart	`panctl s3 instances stop-all --tenant acme`, then `start-all`
Add capacity or availability	create the same or a new iid on another Director, then start it; update DNS records
Change the prefix	`instances delete-all --yes`, `set-config --enabled false`, `set-config --prefix /new`, `set-config --enabled true`, create and start the instances again
Rotate a user's secret	`panctl fs-users rotate-s3-secret <id> --tenant acme`
Block a user	`panctl fs-users update <id> --tenant acme --s3-profile disabled`
Replace the certificate	upload a new one (see 12.9.1) or `panctl s3 tls renew` for a realm-generated one
Check instance state	`panctl s3 instances list --tenant acme` (`OBSERVED`)

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
```

12.16 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`.

12.17 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 StorageSet 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 enabled set true
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
```

12.18 Troubleshooting

The following table lists each message with its cause and fix.

Message	Cause and fix
`S3 is disabled realm-wide; run `panctl s3 enabled set true` before enabling tenant-level S3.`	Turn on the realm switch first (see 12.5.1).
`name: must match ^[^\s/?#%]+$` or `uid ... less than or equal to 4294967294` (fs-users create)	S3 usernames cannot contain whitespace, `/`, `?`, `#` or `%`; uid and gid go up to 4294967294.
`tenant "default" not found: no tenant has that short name`	The platform tenant's short name is `platform`.
`Storageset assignment 1 not found` (volume create)	Assign the StorageSet to the tenant first (see 12.4.1).
`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 12.4.2).
`S3 configuration not found for this tenant`	Run `panctl s3 create-config` (see 12.5.2).
`S3 configuration already exists for this tenant`	Use `panctl s3 set-config` to change it.

Message	Cause and fix
`Default prefix must be configured before enabling S3`	Pass `--prefix` with or before `--enabled true`.
`S3 prefix '...' does not reference an existing volume of this tenant`	Create the volume first (see 12.4.4) or use a directory inside one.
Buckets of the old prefix still served after a prefix change	Delete and recreate the instances (see 12.5.3).
`Tenant has no subnet ...` (instance create)	Create the tenant subnet first (see 12.4.2).
`Cannot change the S3 prefix while S3 is enabled for this tenant; disable it first`	Stop instances, `--enabled false`, change the prefix, enable again.
`Cannot disable S3 for this tenant while instances are desired=started; stop them first`	`panctl s3 instances stop-all --tenant <t>`.
`Cannot delete S3 config for this tenant while any instances still exist`	`panctl s3 instances delete-all --tenant <t> --yes`, 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 ^[A-Za-z0-9]{5}$`	Use exactly 5 letters or digits.
`confirm=true required` (instance delete)	The iid is on several nodes: Add `--confirm-fanout` or `--dataserial` for one node.
`delete-all did not complete within 60s; ... Stuck instances: ... desired=deleting`	An instance on a Storage Node cannot be removed (see 12.6.1). The other instances are deleted. `panctl` gives up first (`context deadline exceeded`) unless you pass `--request-timeout 120s`.
`No PanFS user in this tenant holds access_key '...'. Create one first.`	`--owner` must be an S3 user's access key (`panctl fs-users list`).
`Tenant has no S3 instances. Bucket create requires at least one running instance`	Create and start an instance (see 12.6), then retry.

Message	Cause and fix
`Tenant has no running S3 instance (1 configured, none started)`	Start an instance and wait for `running`, then retry.
`The S3 gateway refused to create bucket "..." (InvalidBucketName)`	Use a valid S3 bucket name (see 12.8).
`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 (`panctl s3 buckets list`).
Client: `SSL validation failed ... certificate verify failed`	Realm-generated certificate: Skip verification (see 12.11.4). Tenant-provided: Give the client the CA (`--ca-bundle`), include intermediates in the certificate, and add IP SANs when connecting by IP.
Client: `InvalidAccessKeyId`	Wrong access key or the user is `disabled`.
Client: `SignatureDoesNotMatch`	Wrong or rotated secret key.
Client: `InvalidLocationConstraint` on bucket create	Set the client region to `us-east-1`.
Client: `AccessDenied` on a bucket with a policy	The policy does not grant the caller, possibly the owner (see 12.12.1).
Client: `AccessControlListNotSupported` (`PutBucketAcl`)	ACLs are off on the bucket: Set its object ownership to `BucketOwnerPreferred` first (see 12.12.2).
Client: `InternalError` on `PutBucketAcl`	Give grantees as bare access keys of existing S3 users, not `id=` / `uri=` / `emailAddress=` (see 12.12.2).
Client: `AccessDenied ... object protected by object lock`	The object has retention, a legal hold, or the bucket has a default retention (see 12.12.5).
Client: `InvalidRequest ... authorization mechanism ... not supported` on a presigned URL	Create the URL with Signature V4 (boto3: `signature_version="s3v4"`).

Message	Cause and fix
Client: `Connect timeout on endpoint URL`	The client host is not attached to the tenant subnet (see 12.4.3).
Instance stays `OBSERVED unknown` after start	Check it was placed on a Director, not a Storage Node (see 12.6.1).

13. Multi-Tenant Networking (VLAN and VXLAN)

This chapter explains how PanFS keeps each tenant's network traffic separate with VLAN or VXLAN, how to configure a tenant network, and what your switches and DirectFlow Clients need.

> **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.

13.1 Summary

PanFS uses VLAN or VXLAN to isolate traffic between tenants. Each tenant network configuration is placed on a numbered VLAN or VXLAN segment, and a realm runs only one segment type, VLAN or VXLAN, at a time.

- **Where it applies:** Tenant networks only. The platform data network does not support tagged VLANs or VXLAN.
- **Default:** Every realm defaults to the VXLAN segment type. No tenants exist by default, and access to storage follows the normal DirectFlow mount rules.
- **When multi-tenancy is enabled or disabled:** Global mounts are not supported while multi-tenancy is enabled. Unmount any platform tenant mounts before you create the first tenant. Delete client mounts before you delete tenant volumes and remove the tenants.
- **Who chooses:** An administrator. The segment type defaults to VXLAN and can be changed to VLAN through PanAPI before any tenant network configuration exists (see 13.6). Vantage shows the segment type and does not change it.
- **What you set:** A segment ID, plus a tunnel UDP port for VXLAN. Both are fixed for a tenant once its network configuration exists. To change them, delete the tenant's volumes and network configuration and create the configuration again. For VLAN, specify only the VLAN ID. As with VXLAN, the VLAN ID cannot be changed without recreating the configuration.
- **Switches:** On a realm with VXLAN tenant networks, node ports stay untagged access ports, and the switches carry only standard untagged Ethernet frames. The underlay network between clients and storage may be routed; that infrastructure is outside the control of VDURA. On a realm with VLAN tenant networks, the switch ports connected to VDURA Data Platform nodes are 802.1Q trunks. Clients may be on the same subnet or on a different subnet behind a standard layer 3 gateway.

- **RDMA:** RoCEv2 works with both types of tenant network. For VLAN tenant networks, all traffic uses the VLAN tenant network. With VXLAN isolation, the DirectFlow Client establishes TCP connections over the VXLAN interface, and an exchanged key then activates secured RDMA transfers over the platform network.
- **Not available:** InfiniBand is not supported for multi-tenancy or tenant networks.

13.2 Network Models

This section covers VXLAN Overlay and Platform Underlay, VLAN-Based Tenants, VLAN-Based Tenants with a VXLAN Underlay, and Topology Matrix.

13.2.1 VXLAN Overlay and Platform Underlay

Figure 13.1 shows a simple generic pairing of an overlay network and an underlay network. The overlay network is the tenant network, and the underlay network is only the physical transport, which is the platform network. There are multiple overlay networks, one per tenant. The diagram shows two endpoints (VTEPs). In an actual deployment there is an endpoint for every VDURA Data Platform device, and an endpoint for each client when the VTEP is terminated within the Linux client, or on a network switch that hosts the endpoint for a VLAN with several clients that are not VXLAN aware.

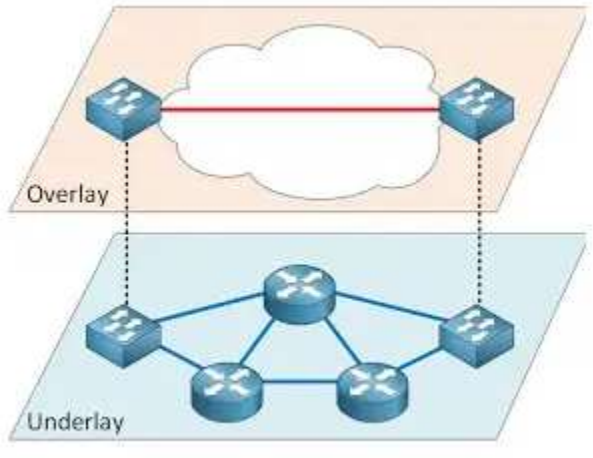


Figure 13.1. A generic overlay network above an underlay network.

The underlay network is invisible to the overlay network, and each overlay network is invisible to the other overlay networks. If the Linux DirectFlow Client is itself the VTEP, the client needs an interface on the platform network, or a routable path to the platform network that uses an interface not associated with the VNI, so that RDMA can work.

13.2.2 VLAN-Based Tenants

Figure 13.2 shows a generic physical network in gray, which is the platform network, with several independent VLANs (10, 20, and 30) that map to tenant VLANs. The clients are not on the platform network.

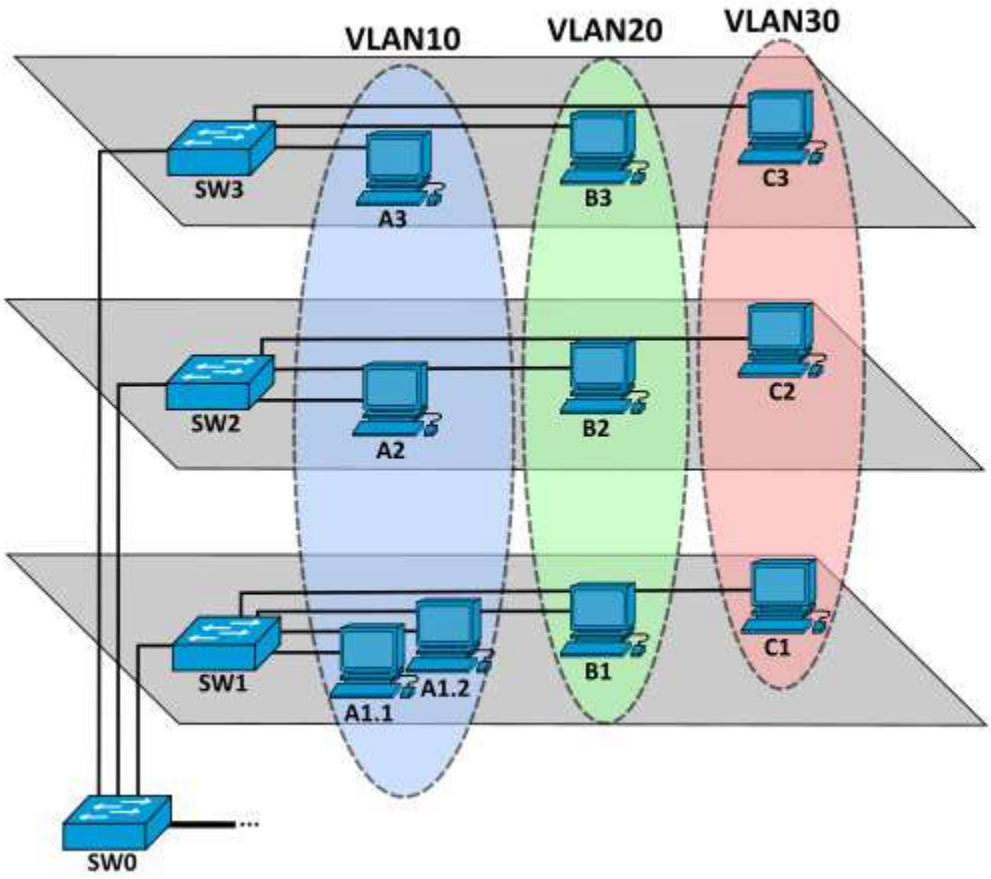


Figure 13.2. A physical network carrying three independent tenant VLANs.

13.2.3 VLAN-Based Tenants with a VXLAN Underlay

A VXLAN underlay network can be invisible to the storage. In that case PanFS implements a standard VLAN-based tenant, and your own infrastructure is configured to tunnel that VLAN through its underlay network (Figure 13.3). The disadvantage of this option is how broadcast and unknown unicast traffic is replicated in the platform network. For sophisticated network environments, this topology lets your network team implement more advanced VTEPs on their own switches, acting as endpoints on both the

VDURA Data Platform side and the client side. The platform network used by the VDURA Data Platform nodes can stay untagged and be mapped to a VNI that is invisible to the storage system. The usual rule still applies; all VDURA Data Platform nodes must be in the same broadcast domain.

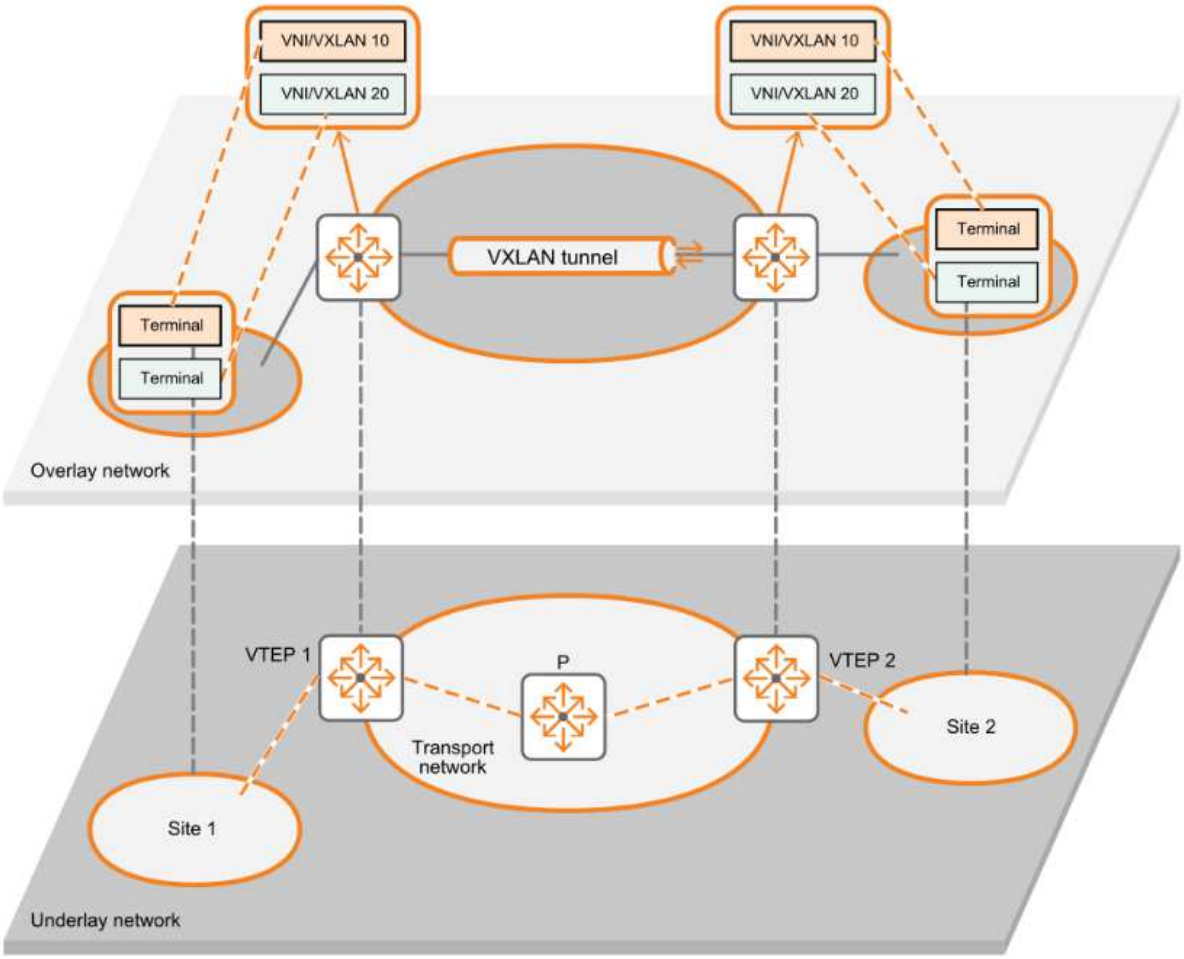


Figure 13.3. A VXLAN tunnel between two sites over an underlay network.

13.2.4 Topology Matrix

The following table summarizes the topologies by realm interfaces, realm switches, client switches, and client interfaces.

Realm Interfaces	Realm Switches	Client Switches	Client Interfaces	Notes
VXLAN VTEP Tenant Untagged Platform	Untagged	Untagged	VXLAN VTEP Tenant Platform Routable	End-to-end platform network connectivity is required for RDMA, in addition to the tenant VNI network.
VXLAN VTEP Tenant Untagged Platform	Untagged	VXLAN VTEP Tenant	Untagged or VLAN Platform Routable	Realm nodes and client switches act as VTEPs. Clients still need a routable path to the platform network for RDMA, which can use VLAN tags on the same physical NIC.
VLAN Tenant Untagged Platform	Untagged or Access Platform Tagged Tenant	Tagged and/or Access	Untagged or Tagged	The VDURA Data Platform nodes must have untagged frame support in the platform network, which may be an access VLAN. Clients need access to the tenant network only.

Realm Interfaces	Realm Switches	Client Switches	Client Interfaces	Notes
VLAN Tenant Untagged Platform	Untagged or Access Platform VXLAN VTEP	VXLAN VTEP Tenant	Untagged or Tagged	The VDURA Data Platform nodes must have untagged frame support in the platform network, which may be an access VLAN. VLAN-to-VNI mapping is used in the realm switches and the client switches. Clients need access to the tenant network only. Because the VTEPs are implemented only in the switches, more sophisticated VXLAN features, such as L2 EVPN with BGP, can be implemented on your side.

13.3 How PanFS Uses VLAN and VXLAN

A tenant reaches storage through one network configuration (called a subnet in PanAPI). A network configuration is a private IP subnet, the address range the tenant uses, placed on a VLAN or VXLAN segment. Each tenant's traffic stays separate from other tenants' traffic.

Setting	What it means
Segment type	`vlan` or `vxlan`. A realm-wide setting. The default is VXLAN. To change it, see 13.6.
Segment ID	A VLAN ID when the realm runs 802.1Q (1 to 4094), or a VNI when it runs VXLAN (1 to 16777215). Leave it empty and the realm assigns an unused one.
Tunnel UDP port	Used only on a realm with VXLAN tenant networks. Leave it empty and the realm applies its default, the standard port 4789.
Underlay and overlay	A host attached to a VXLAN subnet has an underlay address (its address on the platform network, which is also its tunnel endpoint) and an overlay address (its address on the tenant network).

The rules that follow from this:

- **One segment type per realm.** A tenant's network configuration is limited to the segment type the realm already holds. A subnet on a different segment type is refused.
- **Segment IDs are unique across the whole realm, not per tenant.** The identifier space belongs to the physical interface that every tenant on a node shares.
- **Fixed at create.** The segment ID, tunnel UDP port, and address-allocation model cannot be changed afterward. In Vantage, only the name, gateway, address range, and MTU can be edited.
- **A realm with no segment type.** A realm whose segment type setting is missing or unrecognized refuses to create a network configuration.
- **Switching types** requires removing all tenants, volumes, and network configurations, then changing the realm setting (see 13.6).

13.4 Configuring Segments in Vantage

You configure a VLAN or VXLAN segment when you create a tenant's network configuration, on the tenant's **Networking** tab (see 9.2.3 and 9.2.4). The form's banner shows the realm segment type.

1. Open the tenant and select the **Networking** tab. A tenant with no network configuration shows "No network configurations yet".
2. Click **Create Configuration**. A tenant can have only one network configuration, so the button is greyed out afterward.
3. On a realm with no network configuration yet, a banner says the segment type will be fixed at VXLAN once you create the first one. To use VLAN instead, change the realm segment type before you create any configuration (see 13.6).
4. Enter a **Name** (3 to 31 characters, using letters, digits, underscore, or hyphen) and the **Network** in CIDR notation, for example `10.0.1.0/24`. The host part must be zero.
5. Optionally set a **Gateway** and an **Address Range**. Leave the range empty to make the whole network usable.
6. In the **Network Segment** section, optionally set a **Segment ID**, the VXLAN VNI or VLAN ID of the network the DirectFlow Clients connect from. Leave it empty and the realm assigns an unused one. On a realm with VXLAN tenant networks, you can also set a **Tunnel UDP Port**, or leave it empty to use the standard port 4789.
7. In the **Advanced** section, set an **MTU** (empty uses the platform default; see 13.8) and choose **Address Allocation**: **System** (the default) or **External allocator**.
8. Click **Create Network Configuration**.

Two things to expect:

- **The network must be large enough** for all the realm nodes. Otherwise Vantage shows an error that the usable-address window holds fewer addresses than the realm primary pool. Widen the CIDR (a /16 instead of a /24, for example) or the address range, and create the configuration again.
- **The tenant is not ready right away.** It moves from Setup Incomplete to Finishing Setup, which takes about a minute, and then to Active. Volumes can be created only once the tenant is Active.

After creation, the card shows the segment type, segment ID, and tunnel port as read-only fields. **Edit** changes only the name, gateway, address range, and MTU. **Delete** asks you to confirm and cannot be undone. A tenant that still owns a network configuration cannot be deleted. Unmount clients before you remove the configuration.

13.5 Configuring Network Segments with panctl and PanAPI

With panctl, `panctl subnets create` takes `--segment-id` and `--tunnel-udp-port`; both can be set only at creation time (see the VDURA PanCTL User Guide). Through PanAPI, the segment ID and tunnel port are set on the subnet, and the realm segment type is reported by the segment type endpoints (see chapter 12 of the VDURA PanAPI Reference Guide).

Endpoint	Use
`POST /api/v1/tenants/{tenant_id}/subnets`	Create a subnet. Accepts `segment_id` and `tunnel_udp_port`.
`GET /api/v1/tenants/{tenant_id}/subnets/{subnet_id}`	Read a subnet, including `segment_type`.
`PATCH /api/v1/tenants/{tenant_id}/subnets/{subnet_id}`	Edit a subnet. The segment values are refused here because they are fixed at create.
`POST /api/v1/tenants/{tenant_id}/subnets/{subnet_id}/attachments`	Attach a host and add it as a VXLAN peer. The response carries a `tunnel` recipe only when the segment is tunnelled.
`GET /api/v1/subnets/segment-type`, `PUT /api/v1/subnets/segment-type`	Read or set the realm segment type (see 13.6).

A create request:

```
{
  "name": "storage-net",
  "ip_range_base": "192.168.50.0",
  "ip_range_prefix": 24,
  "gateway": "192.168.50.1",
  "segment_id": 10,
  "tunnel_udp_port": 4789
}
```

The fields that matter here:

- `segment_id` : Integer, up to 16777215. Leave it out and the realm assigns an unused identifier; 0 leaves it unassigned.
- `tunnel_udp_port` : Integer, 1 to 65535; 0 uses the realm default. Meaningful only when the realm runs VXLAN.
- `segment_type` : `none` , `vlan` or `vxlan` . Returned, never accepted.
- `mtu` : 0 for the platform default, otherwise 1 to 9216.

Errors you may meet when creating a subnet:

Code	Cause
400	A segment ID was offered while the realm names no segment type, or it falls outside the range that type defines.
409	The name is taken, the segment ID is held by another subnet, the realm already carries subnets on a different segment type, or the realm serves DirectFlow Clients over InfiniBand.
422	The request body fails validation, including segment ID and port bounds.
503	The realm cannot say which segment type applies. Only an operator can fix it.

13.6 Changing the Realm Segment Type

The realm segment type is set through PanAPI. Vantage and panctl show the type of each network configuration and do not change the realm setting. The change is accepted only on a realm with no tenant network configurations, so remove all tenants, with their volumes and network configurations, first.

1. Log in as the Platform Admin. `<realm-address>` is the address of the realm president.

```
API=https://<realm-address>
read -rs -p "Password for admin: " PANPW; echo
TOKEN=$(curl -s -X POST "$API/api/v1/auth/login" \
  -H "Content-Type: application/x-www-form-urlencoded" \
  --data-urlencode "username=admin" \
  --data-urlencode "password=$PANPW" | jq -r .access_token)
unset PANPW
```

2. Show the current segment type.

```
curl -s "$API/api/v1/subnets/segment-type" \
  -H "Authorization: Bearer $TOKEN"
```

Example output: `{"segment_type":"vxlan"}`

3. Confirm that no subnets remain.

```
curl -s "$API/api/v1/subnets" \
  -H "Authorization: Bearer $TOKEN"
```

4. Change the segment type (`vlan` or `vxlan`).

```
curl -s -w "\nHTTP %{http_code}\n" -X PUT "$API/api/v1/subnets/segment-type" \
  -H "Authorization: Bearer $TOKEN" \
  -H "Content-Type: application/json" \
  -d '{"segment_type": "vlan"}'
```

5. Verify the change by repeating step 2.

Response codes for step 4:

Code	Meaning
204	Segment type changed. No body is returned.
409	Refused because a tenant still has a subnet.
422	Invalid value. Only `vlan` and `vxlan` are accepted.
401, 403	Not logged in, or the user lacks the required role.

13.7 Physical Network Requirements

The platform network (also called the realm network) and the underlay network (when VXLAN is in use) do not support tagged VLANs. Storage Nodes, Directors, and DirectFlow Clients send and receive platform traffic untagged.

- **One adapter, several interfaces.** Tenant networks do not layer on top of the platform network; they share the physical network adapter with it. Linux builds one interface on the adapter for the platform network and an additional interface for each tenant network, placed in that tenant's network namespace. A tenant VLAN interface (802.1Q) tags its frames and sends them out of the shared adapter. A tenant VXLAN interface encapsulates its frames and sends them as UDP/IP, using the node's platform network address as its tunnel endpoint.
- **Untagged ports.** In single link, failover, and active/active modes, the port or link aggregation group (LAG) carries the platform network untagged (see section 6.3 of the VDURA V5000 Installation and Service Guide).
- **Port-based VLAN.** The realm can run on a port-based VLAN configured on the switch, which the nodes never see. All Storage Nodes, Directors, and DirectFlow Clients must then be on a network that can reach the tenant network over an IP network. VXLAN tunnel traffic can travel on that VLAN as ordinary IP traffic.
- **Routed fabric for VXLAN tenant networks.** For a realm with VXLAN tenant networks, the expected design is an L2 or L3 fabric with routed IP addressing across the whole platform network. VXLAN tenant traffic then rides over that fabric.
- **Same network.** Ports in a LAG attach to the same VLAN or physical network and sit in the same L2 broadcast domain. The primary network is on the same physical network or VLAN as the storage enclosures.
- **MTU with VLAN tags.** The system supports 1500-byte and 9000-byte IP MTUs on Ethernet. For VLAN interfaces, the node sets the MTU to the underlying interface MTU minus 4 bytes for the VLAN tag. For example, if the underlying interface has an MTU of 9000, the VLAN interface has an MTU of 8996.

- **InfiniBand.** InfiniBand is not supported for multi-tenancy or tenant networks. A node has Ethernet or InfiniBand HPC cards, not both, and a realm that serves DirectFlow Clients over InfiniBand cannot create tenants.

13.8 MTU for Tenant Networks

VXLAN encapsulation adds 50 bytes to every tenant packet, so the tenant MTU must be at least 50 bytes smaller than the underlay MTU. Recommended on-premises values:

- **Switches:** 9216 on all underlay ports, which covers 9000-byte IP packets plus headers and tags with room to spare.
- **Storage Nodes, Directors, and DirectFlow Clients** (the underlay interface and tunnel endpoint): 9000.
- **Tenant VXLAN interfaces and subnet records:** 8950.

On a realm with VLAN tenant networks the tag adds only 4 bytes on the wire, so tenant interfaces can stay at 9000 with the switches at 9216. Things to check:

- **Every endpoint in a tenant subnet must agree:** Directors, OSD nodes, gateways, and DirectFlow Clients. DirectFlow Client address configuration is your responsibility, so set the tenant MTU on each DirectFlow Client's VXLAN interface.
- **Do not rely on path MTU discovery across VXLAN.** ICMP "fragmentation needed" messages are often dropped at the tunnel endpoint. The symptom is mounts that work while large reads and writes hang. A wrong MTU is a silent black hole, not an error.
- **Validate end to end.** `ping -M do -s 8922 <peer tenant IP>` should succeed and `-s 8923` should fail (8950 minus 28 bytes for the IP and ICMP headers).
- **RDMA.** The RoCE MTU is fixed at 1024, 2048 or 4096, and RDMA bulk data does not travel inside the VXLAN tunnel, so the tenant MTU matters for control traffic and for I/O smaller than 64K.

13.9 Switch Configuration for Tenant Networks

What you configure on your switches depends on the realm's segment type. In both cases the platform network stays untagged on the nodes, and PanFS does the tagging or encapsulation itself, using a network namespace and a virtual interface per tenant subnet.

13.9.1 Realm with VXLAN Tenant Networks (the Default)

Every V5000 node is its own VXLAN tunnel endpoint (VTEP). Tenant traffic leaves the node already encapsulated in UDP/IP on the platform network, so when DirectFlow Clients terminate VXLAN themselves the switches see only ordinary IP traffic and need no VXLAN configuration. The switch side needs:

- **Untagged access ports.** Node-facing ports or LAGs stay untagged access ports on the platform network. No trunking is needed.
- **Underlay reachability.** IP connectivity between every node's underlay address and every remote VTEP (DirectFlow Clients, gateways or switch VTEPs), over a flat L2 network or a routed L3 fabric.
- **MTU headroom.** At least 50 bytes more than the tenant MTU on every hop (see 13.8).
- **UDP port 4789 open** (or the subnet's tunnel UDP port) through any access control list (ACL) or firewall in the path.

13.9.2 Adding Peers with the Attachments Endpoint

Why peers are added explicitly. VXLAN has no peer discovery of its own. A tenant subnet still needs broadcast, unknown unicast and multicast (BUM) delivery, for ARP among other things, and the VXLAN standard (RFC 7348) provides it by mapping each VNI to an IP multicast group in the underlay. That mode needs multicast routing (Protocol Independent Multicast, PIM) or Internet Group Management Protocol (IGMP) snooping designed and enabled across the fabric for every VNI. PanFS tenant networks do not use it. The underlay is an ordinary L2 or L3 IP fabric with its usual multicast, but the tenant VXLAN overlay is not mapped to any multicast group, so it carries no multicast and does not depend on your multicast design. It also works on networks that carry no multicast at all.

Without a multicast group, BUM traffic is delivered by head-end (ingress) replication, in which the sending VTEP makes one unicast copy for each known remote VTEP. The attachments endpoint supplies that peer list. Each remote VTEP is added with `POST /api/v1/tenants/{tenant_id}/subnets/{subnet_id}/attachments` (or `panctl subnets attachments create`), and the nodes replicate to the list. PanFS therefore needs no control plane such as BGP Ethernet VPN (EVPN), the other common way to learn the list. How the far end is set up depends on where the tenant's VXLAN terminates:

- **DirectFlow Clients terminate VXLAN themselves** (a Linux VXLAN interface, a data processing unit (DPU) or a hypervisor virtual switch). Add each DirectFlow Client's underlay address as an attachment, and use the `tunnel` recipe in the response to configure the DirectFlow Client with the PanFS node VTEPs as its peers. The switches see underlay traffic only.
- **Your leaf switches are the VTEPs** (bare-metal DirectFlow Clients on ordinary VLANs, for example). Add the leaf VTEP addresses as attachments, map the VNI to the DirectFlow Client VLAN on the leaf, and configure the PanFS node VTEPs as static ingress-replication peers for that VNI on the switch.

- **Fabrics that run EVPN** for the same VNI cannot learn the PanFS nodes through EVPN. Give each PanFS tenant network its own VNI with statically configured remote VTEPs.

13.9.3 Realm with VLAN (802.1Q) Tenant Networks

The nodes tag tenant traffic themselves, so the switch ports must carry the tags:

- **Trunk ports.** Configure each node-facing port or LAG as an 802.1Q trunk, with the platform network as the untagged (native) VLAN and each tenant's segment ID allowed as a tagged VLAN.
- **Extend the tenant VLANs** to the ports facing DirectFlow Clients, as access or trunk ports, so the DirectFlow Clients reach the same broadcast domain.
- **MTU.** Allow for the 4-byte tag; a switch MTU of at least 9022 for 9000-byte packets (see section 9.5 of the VDURA V5000 Installation and Service Guide).
- **Routing.** If tenant traffic must leave the subnet, provide the L3 gateway or switched virtual interface (SVI) for each tenant VLAN at the address set as the subnet's gateway.

13.10 Setting Up a DirectFlow Client on a Tenant Network

PanFS does not configure DirectFlow Client networking. The tenant interface, its address and its MTU are your responsibility. The examples below show the standard Linux commands a DirectFlow Client needs. Names and addresses are examples.

13.10.1 Values a DirectFlow Client Needs

The following table lists each value, where it comes from, and an example.

Value	Where it comes from	Example
Platform (underlay) address	The DirectFlow Client's own address on the platform network	10.128.9.3 on eth0
Segment ID	The tenant's network configuration: VNI or VLAN ID	10100
Overlay address and prefix	Assigned to this DirectFlow Client when it is attached to the tenant network	10.0.1.50/24
Interface MAC (VXLAN)	The DirectFlow Client's attachment record	02:5a:0e:00:01:32
Peer VTEPs (VXLAN)	The realm node platform addresses in the attachment record	10.128.1.2, 10.128.2.5, …

Value	Where it comes from	Example
MTU	The subnet MTU, or the platform MTU minus 50 for VXLAN	8950
Realm manager address	The first Director address on the tenant network	10.0.1.10

13.10.2 DirectFlow Client on VXLAN Tenant Networks

Create a VXLAN interface bound to the platform interface, give it the tenant address, and add one flood entry per realm node so broadcast traffic reaches every peer (head-end replication):

```
ip link add acme type vxlan id 10100 local 10.128.9.3 dev eth0 dstport 4789
ip link set acme address 02:5a:0e:00:01:32 mtu 8950
ip addr add 10.0.1.50/24 dev acme
ip link set acme up
# one entry per realm node VTEP from the attachment record
bridge fdb append 00:00:00:00:00:00 dev acme dst 10.128.1.2
bridge fdb append 00:00:00:00:00:00 dev acme dst 10.128.2.5
sysctl -w net.ipv4.conf.acme.rp_filter=2
```

- **Interface name.** Any Linux interface name up to 15 characters; naming it after the tenant makes it easy to recognize.
- **MAC and MTU.** Use the MAC from the attachment record. If the subnet sets no MTU, use the platform interface MTU minus 50.
- **Reverse-path filter.** Loose mode (`rp_filter=2`) on the tenant interface keeps the kernel from dropping tenant replies.
- **Surviving a reboot.** These commands do not persist. Recreate the interface at boot with the distribution's network configuration (systemd-networkd, NetworkManager or netplan) or a systemd unit that runs them after `network-online.target` .

13.10.3 DirectFlow Client on VLAN Tenant Networks

Create a standard 802.1Q sub-interface on the platform interface:

```
ip link add link eth0 name acme type vlan id 100
ip addr add 10.0.1.50/24 dev acme
ip link set acme mtu 9000 up
```

- The switch port facing the DirectFlow Client must be an 802.1Q trunk, with the platform network untagged and the tenant VLAN tagged.
- No peer list or flood entries are needed; the switch delivers broadcast traffic within the VLAN.

13.10.4 Mounting PanFS over the Tenant Network

Mount with the standard Linux `mount` command. With `-t panfs` it calls the `mount.panfs` helper installed with the DirectFlow Client:

```
mkdir -p /panfs
mount -t panfs -o src-ip=10.0.1.50,rdma-mode=auto panfs://10.0.1.10:global /panfs
```

- **`panfs://<realm manager>:global`:** The realm manager's address on the tenant network, which is the first Director address there. Use one address because the realm tells a mounted DirectFlow Client about realm manager changes.
- **`src-ip=<overlay address>`:** Required. It pins the DirectFlow Client's iSCSI and realm manager sessions to the tenant interface.
- **`rdma-mode=auto`:** RDMA over the platform network when the DirectFlow Client has an RDMA device and the Storage Nodes offer it, TCP otherwise. Use `rdma-mode=off` for a TCP-only mount.
- **A tenant with no volume cannot be mounted yet.** `mount.panfs` reports that it cannot find the object indicated by path "/".

To mount at boot, add an `/etc/fstab` line. Mark it `noauto` and mount it after the tenant interface is up, for example from a systemd unit ordered after the unit that creates the interface. At early boot the Director may not answer yet, so retry:

```
panfs://10.0.1.10:global  /panfs  panfs  noauto,src-ip=10.0.1.50,rdma-mode=auto  0 0
```

13.10.5 Restricting the DirectFlow Client to the Tenant Network

For tighter isolation, a host firewall on the DirectFlow Client can limit traffic to and from the realm node addresses to the paths tenant traffic needs:

- **UDP 4789:** The VXLAN tunnel.
- **UDP 4791:** RoCEv2, the RDMA data path on the platform network. If it is blocked on an RDMA-enabled realm, the DirectFlow Client does not fall back to TCP; it keeps retrying the RDMA connect and I/O hangs.

Reject or drop everything else to and from the realm nodes, and reapply the rules at boot.

13.10.6 Checking a DirectFlow Client

```
ip -d link show acme                                    # vxlan id 10100 ... dstport 4789
bridge fdb show dev acme | grep -c 00:00:00:00:00:00    # one entry per realm node
ping -c 1 -I acme 10.0.1.10                             # realm manager over the tenant network
ping -M do -s 8922 -I acme 10.0.1.10                    # full-size packet, no fragmentation
mount -t panfs                                          # panfs://10.0.1.10:global on /panfs
grep 'panfs://' /etc/fstab                              # carries src-ip=10.0.1.50
```

13.11 RDMA and Tenant Isolation

RoCEv2 works with tenant networks. PanFS splits each DirectFlow Client connection into two paths:

- **Control path over the tenant network.** All control traffic, and every I/O smaller than 64K, travels over the tenant's VXLAN or VLAN segment.
- **Bulk data over kernel RDMA on the platform network.** I/O of 64K and larger moves between Storage Nodes and DirectFlow Clients over RoCEv2, which is not encapsulated in VXLAN, so the VXLAN overhead does not apply on this path.

Tenant separation on the RDMA path comes from the control path. When a DirectFlow Client logs in to a Storage Node over the tenant network, it registers a random per-session secret. Its RDMA connect request must then carry an encrypted proof of that secret for the session it names. The Storage Node recomputes the proof, compares it in constant time, and rejects any request that does not match. A host that cannot reach a tenant's control path cannot open RDMA connections to that tenant's sessions, so the switches need no per-tenant RDMA configuration.

Because RDMA traffic carries no per-packet tenant identity, the network carrying the RDMA path needs:

- **Lossless RoCEv2.** Priority Flow Control (PFC) on the RoCE priority, Explicit Congestion Notification (ECN) marking for Data Center Quantized Congestion Notification (DCQCN), and consistent Differentiated Services Code Point (DSCP) trust end to end.
- **DirectFlow Client-to-client isolation,** with private VLAN or port isolation, so DirectFlow Clients reach Storage Node ports and not each other.
- **ACLs** that allow RoCEv2 (UDP 4791) and RDMA connection management traffic only between DirectFlow Client and Storage Node addresses.
- **Anti-spoofing** on ports facing DirectFlow Clients (dynamic ARP inspection, IP source guard or the equivalent), so one host cannot take over another's flows.

13.12 Switch Checklist

The following table compares the requirements for VXLAN and VLAN tenant networks.

Requirement	VXLAN tenant networks	VLAN tenant networks
Node-facing ports	Untagged access port on the platform network	802.1Q trunk: platform network native, tenant VLANs tagged
MTU	9216 switch, 9000 node, 8950 tenant	At least 9022 switch for 9000-byte packets (9216 recommended)
Tenant segment on the switch	None, unless the switch is the DirectFlow Client VTEP: then VNI-to-VLAN mapping and static ingress-replication peers	Tenant VLANs allowed on node and DirectFlow Client ports
Peers	Each remote VTEP added with the attachments endpoint	Not applicable
Firewall and ACL	UDP 4789 (or the subnet's tunnel port) allowed	Not applicable
Gateway	You provide it, if tenant traffic leaves the subnet	You provide an SVI or router per tenant VLAN
RDMA (if used)	Lossless RoCEv2, DirectFlow Client-to-client isolation, DirectFlow Client-to-storage ACLs, anti-spoofing	Same

13.13 Background: Tagged VLANs and VXLAN

When a network adapter adds a VLAN tag or a VXLAN Network Identifier (VNI) to packets, it acts as a tagging or encapsulation point. That differs fundamentally from a port-based VLAN configured on an Ethernet switch.

- **Port-based (untagged) VLAN.** The switch does all the work. A switch port is assigned to a single VLAN (for example VLAN 10), the adapter sends ordinary untagged frames, and the switch places them in VLAN 10 on arrival. This is how the platform network can sit on a VLAN without the nodes knowing.
- **Adapter-tagged (trunked) VLAN.** The adapter inserts a 4-byte 802.1Q tag into each frame before sending it. The switch port must be a trunk port, or a tagged member of that VLAN, to accept it. One physical adapter can then carry several VLANs at once. This is how VLAN tenant networks are carried.
- **VXLAN.** The adapter acts as a VXLAN tunnel endpoint (VTEP). It wraps the original Ethernet frame in a VXLAN header carrying the 24-bit VNI, then in UDP, IP and a new outer Ethernet header. The physical switches see only ordinary IP traffic and are unaware of the tenant network inside. This is how VXLAN tenant networks are carried.

Feature	Adapter-tagged VLAN	VXLAN (VNI)
Where tagging happens	At the network adapter	At the network adapter (VTEP)
Packet change	Adds a 4-byte 802.1Q tag	Encapsulates the frame in a UDP/IP tunnel (50 bytes)
Switch requirement	Trunk or tagged port	Standard Layer 2 or Layer 3 IP forwarding
Scale limit	4,094 VLANs	16 million+ VNIs

14. Troubleshooting

The following table lists each symptom and its likely cause and fix.

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) shows “No data to display”, or the Events Log shows “No events recorded”.	The realm has not recorded any performance data or events for that period. The chart or the Events Log fills in as the realm records them.
A tenant shows `Setup Incomplete`, 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.
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.